



HAL
open science

State and parameter learning with PARIS particle Gibbs

Gabriel Victorino Cardoso, Yazid Janati, Sylvain Le Corff, Éric Moulines,
Jimmy Olsson

► **To cite this version:**

Gabriel Victorino Cardoso, Yazid Janati, Sylvain Le Corff, Éric Moulines, Jimmy Olsson. State and parameter learning with PARIS particle Gibbs. 2023. hal-03919565

HAL Id: hal-03919565

<https://hal.science/hal-03919565>

Preprint submitted on 3 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

State and parameter learning with PARIS particle Gibbs

Gabriel Cardoso[†], Yazid Janati El Idrissi[‡], Sylvain Le Corff^{*}, Éric Moulines[†], and Jimmy Olsson[†]

[†]CMAP, École Polytechnique, Institut Polytechnique de Paris, Palaiseau.

[‡]Samovar, Télécom SudParis, département CITI, TIPIC, Institut Polytechnique de Paris, Palaiseau.

^{*}LPSM, Sorbonne Université, UMR CNRS 8001, 4 Place Jussieu, 75005 Paris.

[†]Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden.

Abstract

Non-linear state-space models, also known as general hidden Markov models, are ubiquitous in statistical machine learning, being the most classical generative models for serial data and sequences in general. The particle-based, rapid incremental smoother (PARIS) is a sequential Monte Carlo (SMC) technique allowing for efficient online approximation of expectations of additive functionals under the smoothing distribution in these models. Such expectations appear naturally in several learning contexts, such as likelihood estimation (MLE) and Markov score climbing (MSC). PARIS has linear computational complexity, limited memory requirements and comes with non-asymptotic bounds, convergence results and stability guarantees. Still, being based on self-normalised importance sampling, the PARIS estimator is biased. Our first contribution is to design a novel additive smoothing algorithm, the Parisian particle Gibbs (PPG) sampler, which can be viewed as a PARIS algorithm driven by conditional SMC moves, resulting in bias-reduced estimates of the targeted quantities. We substantiate the PPG algorithm with theoretical results, including new bounds on bias and variance as well as deviation inequalities. Our second contribution is to apply PPG in a learning framework, covering MLE and MSC as special examples. In this context, we establish, under standard assumptions, non-asymptotic bounds highlighting the value of bias reduction and the implicit Rao–Blackwellization of PPG. These are the first non-asymptotic results of this kind in this setting. We illustrate our theoretical results with numerical experiments supporting our claims.

1. Introduction

Sequential Monte Carlo (SMC) *methods*, or *particle filters*, are simulation-based approaches used for the online approximation of posterior distributions in the context of Bayesian inference in state space models. In nonlinear *hidden Markov models* (HMM), they have been successfully applied for approximating online the typically intractable posterior distributions of sequences of unobserved states $(X_{s_1}, \dots, X_{s_2})$ given observations $(Y_{t_1}, \dots, Y_{t_2})$ for $0 \leq s_1 \leq s_2$ and $0 \leq t_1 \leq t_2$. Standard SMC methods use Monte Carlo samples generated recursively by means of sequential importance sampling and resampling steps. A particle filter approximates the flow of marginal posteriors by a sequence of occupation measures associated with a sequence $\{\xi_t^i\}_{i=1}^N$, $t \in \mathbb{N}$, of Monte Carlo samples, each *particle* ξ_t^i being a random draw in the state space of the hidden process. Particle filters revolve around two operations: a *selection step* duplicating/discarding particles with large/small importance weights, respectively, and a *mutation step* evolving randomly the selected particles in the state space. Applying alternately and iteratively selection and mutation results in swarms of particles being both temporally and spatially dependent. The joint state posteriors of an HMM can also be interpreted as laws associated with a certain kind of Markovian backward dynamics; this interpretation is useful,

for instance, when designing backward-sampling-based particle algorithms for nonlinear smoothing [Douc et al., 2011, Del Moral et al., 2010].

Throughout the years, several convergence results as the number N of particles tends to infinity have been established; see, *e.g.*, [Del Moral, 2004, Douc and Moulines, 2008, Cappé et al., 2005] and the references therein. In addition, a number of non-asymptotic results have been established, including time-uniform bounds on the SMC L_p error and bias as well as bounds describing the propagation of chaos among the particles. Extensions to the backward-sampling-based particle algorithms can also be found for instance in [Douc et al., 2011, Del Moral et al., 2010, Dubarry and Le Corff, 2013].

In this paper, we focus on the problem of recursively computing smoothed expectations $\eta_{0:t}h_t = \mathbb{E}[h_t(X_{0:t}) \mid Y_{0:t}]$ for additive functionals h_t in the form

$$h_t(x_{0:t}) := \sum_{s=0}^{t-1} \tilde{h}_s(x_{s:s+1}), \quad (1.1)$$

where $X_{0:n}$ and $Y_{0:n}$ denote vectors of states and observations (see below for precise definitions). Such expectations appear frequently in the context of maximum-likelihood parameter estimation in nonlinear HMMs, for instance, when computing the score function (the gradient of the log-likelihood function) or the Expectation Maximization intermediate quantity; see [Cappé, 2001, Andrieu and Doucet, 2003, Poyiadjis et al., 2005, Cappé, 2011, Poyiadjis et al., 2011]. The *particle-based, rapid incremental smoother* (PARIS) proposed in [Olsson and Westerborn, 2017] is tailored for solving online this additive smoothing problem. When the transition density of the latent states is lower and upper bounded, this algorithm can be shown to have a linear computational complexity in the number N of particles and limited memory requirements. An interesting feature of the PARIS, which samples on-the-fly from the backward dynamics induced by the particle filter, is that it requires two or more backward draws per particle to cope with the degeneracy of the sampled trajectories and remain numerically stable in the long run, with an asymptotic variance that grows only linearly with time.

In this paper, we introduce a method to reduce the bias of the PARIS estimator of $\eta_{0:t}h_t$. The idea is to mix—by introducing a *conditional PARIS algorithm*—the PARIS algorithm with a backward-sampling-based version of the *particle Gibbs sampler* [Andrieu et al., 2010b, Lindsten et al., 2014a, Chopin and Singh, 2015a, Del Moral et al., 2016, Del Moral and Jasra, 2018]. This leads to a batch mode *PARIS particle Gibbs (PPG) sampler*, which we furnish with an upper bound of the bias that decreases inversely proportionally to the number N of particles and exponentially fast with the particle Gibbs iteration index (under the assumption that the particle Gibbs sampler is uniformly ergodic).

As an application we consider the problem of likelihood maximization with stochastic gradient. In this specific context, where the smoothing estimator is employed repeatedly to produce mean-field estimates, controlling the bias becomes critical. Thus, it is natural to aim at minimizing the bias for a fixed computational budget, provided that the variance does not explode. For this reason, bias reduction in stochastic simulation has been the subject of extensive research during the last decades [Jacob et al., 2020, Glynn and Rhee, 2014]. The present paper contributes to this line of research. In particular, we show that stochastic approximation (SA) with PPG achieves a $\mathcal{O}(\log(n)/\sqrt{n})$ rate, where n is the number of SA steps. This improves on a previous result of [?], which establishes the almost sure convergence (to a stationary point of the likelihood) of an SA *Expectation Maximization* (EM) algorithm based on particle Gibbs with *ancestor sampling* (PGAS).

The paper is structured as follows. In Section 2, we recall the hidden Markov model framework, the particle filter and the PARIS algorithm. In Section 3, we lay out the PPG algorithm and present the first central result of this paper, an upper bound on the bias of our estimator as a function of the number of particles and the iteration index of the Gibbs algorithm. In addition, we provide an upper bound on the mean-squared error (MSE). In Section 4, we undertake the learning problem and present the second result of this paper, a $\mathcal{O}(\log(n)/\sqrt{n})$ non-asymptotic bound on the expectation of the squared gradient norm taken at a random index K . In Section 5.1, we illustrate our results through numerical experiments. All the proofs are collected in the supplementary material.

Notation. For a given measurable space (X, \mathcal{X}) , where \mathcal{X} is a countably generated σ -algebra, we denote by $F(\mathcal{X})$ the set of bounded $\mathcal{X}/\mathcal{B}(\mathbb{R})$ -measurable functions on X . For any $h \in F(\mathcal{X})$, we let $\|h\|_\infty := \sup_{x \in X} |h(x)|$ and $\text{osc}(h) := \sup_{(x, x') \in X^2} |h(x) - h(x')|$ denote the supremum and oscillator norms of h , respectively. Let $M(\mathcal{X})$ be the set of σ -finite measures on (X, \mathcal{X}) and $M_1(\mathcal{X}) \subset M(\mathcal{X})$ the probability measures. For any $h \in F(\mathcal{X})$ and $\mu \in M(\mathcal{X})$ we write $\mu(f) = \int h(x)\mu(dx)$. For a Markov kernel K from (X, \mathcal{X}) to another measurable space (Y, \mathcal{Y}) , we define the measurable function $Kh : X \ni x \mapsto \int h(y)K(x, dy)$. The composition μK is a probability measure on (Y, \mathcal{Y}) such that $\mu K : \mathcal{X} \ni A \mapsto \int \mu(dx)K(x, dy)\mathbb{1}_A(y)$. For all sequences $\{a_u\}_{u \in \mathbb{Z}}$ and $\{b^u\}_{u \in \mathbb{Z}}$, and all $s \leq t$ we write $a_{s:t} = \{a_s, \dots, a_t\}$ and $b^{s:t} = \{b^s, \dots, b^t\}$.

2. Background

2.1 Hidden Markov models

Hidden Markov models consist of an unobserved state process $\{X_t\}_{t \in \mathbb{N}}$ and observations $\{Y_t\}_{t \in \mathbb{N}}$, where, at each time $t \in \mathbb{N}$, the unobserved state X_t and the observation Y_t are assumed to take values in some general measurable spaces (X_t, \mathcal{X}_t) and (Y_t, \mathcal{Y}_t) , respectively. It is assumed that $\{X_t\}_{t \in \mathbb{N}}$ is a Markov chain with transition kernels $\{M_{t+1}\}_{t \in \mathbb{N}}$ and initial distribution η_0 . Given the states $\{X_t\}_{t \in \mathbb{N}}$, the observations $\{Y_t\}_{t \in \mathbb{N}}$ are assumed to be independent and such that for all $t \in \mathbb{N}$, the conditional distribution of the observation Y_t depends only on the current state X_t . This distribution is assumed to admit a density $g_t(X_t, \cdot)$ with respect to some reference measure. In the following we assume that we are given a fixed sequence $\{y_t\}_{t \in \mathbb{N}}$ of observations and define, abusing notations, $g_t(\cdot) = g_t(\cdot, y_t)$ for each $t \in \mathbb{N}$. We denote, for $0 \leq s \leq t$, $X_{s:t} := \prod_{u=s}^t X_u$ and $\mathcal{X}_{s:t} := \bigotimes_{u=s}^t \mathcal{X}_u$. Consider the unnormalized transition kernel

$$Q_s : X_s \times \mathcal{X}_{s+1} \ni (x, A) \mapsto g_s(x)M_s(x, A) \quad (2.2)$$

and let

$$\gamma_{0:t} : \mathcal{X}_{0:t} \ni A \mapsto \int \mathbb{1}_A(x_{0:t})\eta_0(dx_0) \prod_{s=0}^{t-1} Q_s(x_s, dx_{s+1}). \quad (2.3)$$

Using these quantities, we may define the *joint-smoothing* and *predictor distributions* at time $t \in \mathbb{N}$ as

$$\eta_{0:t} : \mathcal{X}_{0:t} \ni A \mapsto \frac{\gamma_{0:t}(A)}{\gamma_{0:t}(X_{0:t})}, \quad (2.4)$$

$$\eta_t : \mathcal{X}_t \ni A \mapsto \eta_{0:t}(X_{0:t-1} \times A), \quad (2.5)$$

respectively. It can be shown (see [Cappé et al., 2005, Section 3]) that $\eta_{0:t}$ and η_t are the conditional distributions of $X_{0:t}$ and X_t given $Y_{0:t-1}$ respectively, evaluated at $y_{0:t-1}$. Unfortunately, these distributions, which are vital in Bayesian smoothing and filtering as they enable the estimation of hidden states through the observed data stream, are available in a closed form only in the cases of linear Gaussian models or models with finite state spaces; see [Cappé et al., 2009] for a comprehensive coverage.

2.2 Particle filters

For most models of interest in practice, the joint smoothing and predictor distributions are intractable, and so are also any expectation associated with these distributions. Still, such expectations can typically be efficiently estimated using *particle methods*, which are based on the predictor recursion $\eta_{t+1} = \eta_t Q_t / \eta_t g_t$. At time t , if we assume that we have at hand a consistent particle approximation of η_t , formed by N random draws $\{\xi_i^t\}_{i=1}^N$, so-called *particles*, in X_t and given by $\eta_t^N = N^{-1} \sum_{i=1}^N \delta_{\xi_i^t}$, plugging η_t^N into the recursion tying η_{t+1} and η_t yields the mixture $\eta_t^N Q_t$, from which a sample of N new particles can be drawn in order to construct η_{t+1}^N . To do so, we sample, for all $1 \leq i \leq N$, ancestor

indices $\alpha_t^i \sim \text{Categorical}(\{g_t(\xi_t^\ell)\}_{\ell=1}^N)$ and then propagate $\xi_{t+1}^i \sim M_t(\xi_t^{\alpha_t^i}, \cdot)$. This procedure, which is initialized by sampling the initial particles $\{\xi_0^i\}_{i=1}^N$ independently from η_0 , describes the particle filter with multinomial resampling and produces consistent estimators such that for every $h \in \mathbb{F}(\mathcal{X}_t)$, $\eta_t^N(h)$ converges almost surely to $\eta_t(h)$ as the number N of particles tends to infinity.

This procedure can also be extended to produce particle approximations of the joint-smoothing distributions $\{\eta_{0:t}\}_{t \in \mathbb{N}}$. Note that the successive ancestor selection steps described previously generates an ancestor line for each terminal particle ξ_t^i , which we denote by $\xi_{0:t}^i$. It can then be easily shown that $\eta_{0:t}^N = N^{-1} \sum_{i=1}^N \delta_{\xi_{0:t}^i}$ forms a particle approximation of the joint-smoothing distribution $\eta_{0:t}$. However, it is well known that the same selection operation also depletes the ancestor lines, since, at each step, two different particles are likely to originate from the same parent in the previous generation. Thus, eventually, all the particles end up having a large portion of their initial ancestry in common. This means that in practice, this naive approach, which we refer to as the *poor man's smoother*, suffers generally from high variance when used for estimating joint-smoothing expectations of objective functionals depending on the whole state trajectory.

2.3 Backward smoothing and the PARIS algorithm

We now discuss how to avoid the problem of particle degeneracy relative to the smoothing problem by means of so-called *backward sampling*. While this line of research has broader applicability, we restrict ourselves for the sake of simplicity to the case of *additive state functionals* in the form

$$h_t(x_{0:t}) := \sum_{s=0}^{t-1} \tilde{h}_s(x_{s:s+1}), \quad x_{0:t} \in \mathcal{X}_{0:t}. \quad (2.6)$$

Appealingly, using the poor man's smoother described in the previous section, smoothing of additive functionals can be performed online alongside the particle filter by letting, for each s ,

$$\eta_{0:s}^N h_s := N^{-1} \sum_{i=1}^N \beta_s^i, \quad (2.7)$$

where the statistics $\{\beta_s^i\}_{i=1}^N$ satisfy the recursion

$$\beta_{s+1}^i = \beta_s^{\alpha_s^i} + \tilde{h}_s(\xi_s^{\alpha_s^i}, \xi_{s+1}^i), \quad (2.8)$$

where α_s^i is, as described, the ancestor at time s of particle ξ_{s+1}^i .

As mentioned above, the previous estimator suffers from high variance when s is relatively large with respect to N . However, assume now that the model is *fully dominated* in the sense that each state process kernel M_s has a transition density m_s with respect to some reference measure; then, interestingly, it is easily seen that the conditional probability that $\alpha_s^i = j$ given the offspring ξ_{s+1}^i and the ancestors $\{\xi_s^\ell\}_{\ell=1}^N$ is given by

$$\mathbf{\Lambda}_s(i, j) := \frac{\omega_s^j m_s(\xi_s^j, \xi_{s+1}^i)}{\sum_{\ell=1}^N \omega_s^\ell m_s(\xi_s^\ell, \xi_{s+1}^i)}. \quad (2.9)$$

Here $\mathbf{\Lambda}_s$ forms a backward Markov transition kernel on $\llbracket 1, N \rrbracket \times \llbracket 1, N \rrbracket$. Using this observation, we may avoid completely the particle-path degeneracy of the poor man's smoother by simply replacing the naive update (2.8) by the Rao-Blackwellized counterpart

$$\beta_{s+1}^i = \sum_{j=1}^N \mathbf{\Lambda}_s(i, j) \{\beta_s^j + \tilde{h}_s(\xi_s^j, \xi_{s+1}^i)\}. \quad (2.10)$$

This approach, proposed in [Del Moral et al., 2010], avoids elegantly the path degeneracy as it eliminates the ancestral connection between the particles by means of averaging. Furthermore, it is entirely

online since at step s only the particle populations $\xi_s^{1:N}$ and $\xi_{s+1}^{1:N}$ are needed to perform the update. Still, a significant drawback is the overall $\mathcal{O}(N^2)$ complexity for the computation of $\beta_t^{1:N}$, since the calculation of each β_{s+1}^i in (2.10) involves the computation of N^2 terms, which can be prohibitive when the number N of particles is large. Thus, in [Olsson and Westerborn, 2017], the authors propose to sample $M \ll N$ conditionally independent indices $\{J_s^{i,j}\}_{j=1}^M$ from the distribution $\mathbf{\Lambda}_s(i, \cdot)$ and to update the statistics according to

$$\beta_{s+1}^i = M^{-1} \sum_{j=1}^M \left(\beta_s^{J_s^{i,j}} + \tilde{h}_s(\xi_s^{J_s^{i,j}}, \xi_{s+1}^i) \right). \quad (2.11)$$

If the transition density m_s is uniformly bounded from above and below, an accept-reject approach allows the sampling-based update (2.11) to be performed for $i \in \llbracket 1, N \rrbracket$ at an $\mathcal{O}(N(M+1))$ overall complexity if a pre-initialized multinomial sampler is used. A key aspect of this approach is that the number M of sampled indices at each step can be very small; indeed, for any fixed $M \geq 2$, the algorithm, which is referred to as the **PARIS**, can be shown to be stochastically stable with an $\mathcal{O}(t)$ variance (see [Olsson and Westerborn, 2017, Section 1] for details), and setting M to 2 or 3 yields typically fully satisfying results.

The **PARIS** estimator can be viewed as an alternative to the FFBSm, rather than the FFBSi. Even if the **PARIS** and FFBSi are both randomised versions of the FFBSm estimator, the **PARIS** is of a fundamentally different nature than the FFBSi. The **PARIS** approximates the forward-only FFBSm online in the context of additive functionals by approximating each updating step by additional Monte Carlo sampling. The sample size M is an accuracy parameter that determines the precision of this approximation, and by increasing M the statistical properties of the **PARIS** approaches those of the forward-only FFBSm. On the other hand, as shown in [Douc et al., 2011, Corollary 9], the asymptotic variance of FFBSi is always larger than that of the FFBSm, with a gap given by the variance of the state functional under the joint-smoothing distribution. Thus, we expect, especially in the case of a low signal-to-noise ratio, the **PARIS** to be more accurate than the FFBSi for a given computational budget. Another important reason to focus on the **PARIS** estimator rather than the FFBSi is the appealing online properties of the latter, whose interplay with and relevance to the particle MCMC methodology is to be explored. Our results can be naturally extended to the FFBSi and PGAS but since the **PARIS** has a practical edge, we chose to center our contribution around it although the main idea behind our paper is more general.

3. PARIS particle Gibbs

3.1 Particle Gibbs methods

The *conditional particle filter* (CPF) introduced in [Andrieu et al., 2010a] serves the basis of a particle-based MCMC algorithm targeting the joint-smoothing distribution $\eta_{0:t}$. Let $\ell \in \mathbb{N}^*$ be an iteration index and $\zeta_{0:t}[\ell]$ a conditional path used at iteration ℓ of the CPF to construct a particle approximation of $\eta_{0:t}$ as follows. At step $s \in \llbracket 1, t \rrbracket$ of the CPF, a randomly selected particle, with uniform probability $1/N$, is set to $\zeta_s[\ell]$, whereas the remaining $N-1$ particles are all drawn from the mixture $\eta_{s-1}^N Q_{s-1}$. At the final step, a new particle path $\zeta_{0:t}[\ell+1]$ is drawn either:

- by selecting randomly, again with uniform probability $1/N$, a genealogical trace from the ancestral tree of the particles $\{\xi_s^{1:N}\}_{s=0}^t$ produced by the CPF, as in the vanilla particle Gibbs sampler;
- or by generating the path by means of backward sampling, *i.e.*, by drawing indices $J_{0:t}$ backwards in time according to $J_t \sim \text{Categorical}(\{1/N\}_{i=1}^N)$ and, conditionally to J_{s+1} , $J_s \sim \mathbf{\Lambda}_s(J_{s+1}, \cdot)$, $s \in \llbracket 0, t-1 \rrbracket$, and letting $\zeta_{0:t}[\ell+1] = (\xi_0^{J_0}, \dots, \xi_t^{J_t})$ (where the transition kernels $\{\mathbf{\Lambda}_s\}_{s=0}^t$, defined by (2.9), are induced by the particles produced by the CPF), as proposed in [Whiteley, 2010].

The theoretical properties of the different versions of the particle Gibbs sampler are well studied [Singh et al., 2017, Chopin and Singh, 2015b, Andrieu et al., 2018]. In short, the produced conditional paths $(\zeta_{0:t}[\ell])_{\ell \in \mathbb{N}}$ form a Markov chain whose marginal law converges geometrically fast in total variation to the target distribution $\eta_{0:t}$. As it is the case for smoothing algorithms, the vanilla particle Gibbs sampler suffers from bad mixing due to particle path degeneracy while its backward-sampling counterpart exhibits superior performance as t increases [Lee et al., 2020].

3.2 The PPG algorithm

Remarkably, in order for the standard particle Gibbs samplers to output a single conditional path, a whole particle filter is run and then discarded, resulting in significant waste of computational work. Thus, we now introduce a variant of the PARIS algorithm, coined the PARIS particle Gibbs (PPG), in which the conditional path of particle Gibbs with backward sampling is merged with the intermediate particles, ensuring less computational waste and reduced bias with respect to the vanilla PARIS.

In the following we let $t \in \mathbb{N}$ be a fixed time horizon, and describe in detail how the PPG approximates iteratively $\eta_{0:t}h_t$, where h_t is an additive functional in the form (2.6). Using a given conditional path $\zeta_{0:t}[\ell - 1]$ as input, the ℓ -th iteration of the PPG outputs a many-body system $\mathbf{v}_t[\ell] = ((\xi_{0:t}^1, \beta_t^1), \dots, (\xi_{0:t}^N, \beta_t^N))$ comprising N backward particle paths $\{\xi_{0:t}^i\}_{i=1}^N$ with associated PARIS statistics $\{\beta_t^i\}_{i=1}^N$. This is the so-called *conditional PARIS update* detailed in Algorithm 1. After this, an updated conditional path is selected with probability $1/N$ among the N particle paths $\{\xi_{0:t}^i\}_{i=1}^N$ and used as input in the next conditional PARIS operation. At each iteration, the produced statistics $\{\beta_t^i\}_{i=1}^N$ provide an approximation of $\eta_{0:t}h_t$ according to (2.7). The overall algorithm is summarized in Algorithm 2. The function CPF_s describes one step of the conditional particle filter and is given in the supplementary material. In addition, the PPG algorithm defines a Markov chain with Markov transition kernel denoted by \mathbb{K}_t and detailed in (A.41).

Algorithm 1 One conditional PARIS update (CondPaRIS)

Input: $\{(\xi_{0:s}^i, \beta_s^i)\}_{i=1}^N, \zeta_{s+1}, \tilde{h}_{s-1}$

Result: $\{(\xi_{0:s+1}^i, \beta_{s+1}^i)\}_{i=1}^N$

- 1 draw $\xi_{s+1}^{1:N} \sim \text{CPF}_s(\zeta_{s+1}, \xi_s^{1:N})$
 - for** $i \leftarrow 1$ **to** N **do**
 - 2 draw $\{J_s^{i,\ell}\}_{\ell=1}^M \sim \Lambda(i, \cdot)^{\otimes M}$
 - 3 set $\beta_{s+1}^i \leftarrow M^{-1} \sum_{\ell=1}^M \left(\beta_s^{i, J_s^{i,\ell}} + \tilde{h}_s(\xi_s^{i, J_s^{i,\ell}}, \xi_{s+1}^i) \right)$
 - 4 set $\xi_{0:s+1}^i \leftarrow (\xi_{0:s}^{i, J_s^{i,1}}, \xi_{s+1}^i)$
-

Algorithm 2 One iteration of PPG

Input: Initial path $\zeta_{0:t}, \{\tilde{h}_s\}_{s=0}^{t-1}$

Result: $\{\beta_t^i\}_{i=1}^N, \zeta'_{0:t}$

- 5 draw $\xi_0^{1:N} \sim \text{CPF}_0(\zeta_0)$
 - 6 set $\beta_0^i \leftarrow 0$ for $i \in \llbracket 1, N \rrbracket$
 - 7 **for** $s \leftarrow 0$ **to** $t - 1$ **do**
 - 8 set $\{(\xi_{0:s+1}^i, \beta_{s+1}^i)\}_{i=1}^N \leftarrow \text{CondPaRIS}(\{(\xi_{0:s}^i, \beta_s^i)\}_{i=1}^N, \zeta_{s+1}, \tilde{h}_s)$
 - 9 draw $\zeta'_{0:t} \sim N^{-1} \sum_{i=1}^N \delta_{\xi_{0:t}^i}$
-

As performing k steps of the PPG results in k many-body systems, it is natural to consider the

following *roll-out estimator* which combines the backward statistics from step $k_0 < k$ to k :

$$\Pi_{(k_0, k), N}(h_t) = [N(k - k_0)]^{-1} \sum_{\ell=k_0+1}^k \sum_{j=1}^N \beta_t^j[\ell]. \quad (3.12)$$

The total number of particles used in this estimator is $C = (N - 1)k$ per time step. We denote by $v = (k - k_0)/k$ the ratio of the number of particles used in the estimator to the total number of sampled particles.

We now state the first main results of the present paper, in the form of theoretical bounds on the bias and mean-squared error (MSE) of the roll-out estimator (3.12). These results are obtained under the following *strong mixing* assumptions, which are now standard in the literature (see [Del Moral, 2004, Douc and Moulines, 2008, Del Moral, 2013, Del Moral et al., 2016]). It is crucial for obtaining quantitative bounds for particle smoothing algorithms, see [Olsson and Westerborn, 2017] or [Gloaguen et al., 2022] but also for the coupled conditional backward sampling particle filter [Lee et al., 2020].

A 3.1 (strong mixing). For every $s \in \mathbb{N}$ there exist $\underline{\tau}_s, \bar{\tau}_s, \underline{\sigma}_s$, and $\bar{\sigma}_s$ in \mathbb{R}_+^* such that

- (i) $\underline{\tau}_s \leq g_s(x_s) \leq \bar{\tau}_s$ for every $x_s \in \mathbf{X}_s$,
- (ii) $\underline{\sigma}_s \leq m_s(x_s, x_{s+1}) \leq \bar{\sigma}_s$ for every $(x_s, x_{s+1}) \in \mathbf{X}_{s:s+1}$.

Under **A 3.1**, define, for every $s \in \mathbb{N}$,

$$\rho_s := \max_{m \in \llbracket 0, s \rrbracket} \frac{\bar{\tau}_m \bar{\sigma}_m}{\underline{\tau}_m \underline{\sigma}_m} \quad (3.13)$$

and, for every $N \in \mathbb{N}^*$ and $t \in \mathbb{N}$ such that $N > N_t := (1 + 5\rho_t^2/2) \vee 2t(1 + \rho_t^2)$,

$$\kappa_{N,t} := 1 - \frac{1 - (1 + 5t\rho_t^2/2)/N}{1 + 4t(1 + 2\rho_t^2)/N}. \quad (3.14)$$

Note that $\kappa_{N,t} \in (0, 1)$ for all N and t as above.

Theorem 1. *Assume **A 3.1**. Then for every $t \in \mathbb{N}$, $M \in \mathbb{N}^*$, $\xi \in \mathbf{M}_1(\mathcal{X}_{0:t})$, $k_0 \in \mathbb{N}^*$, $k > k_0$ and $N \in \mathbb{N}^*$ such that $N > N_t$,*

$$\begin{aligned} |\mathbb{E}_\xi[\Pi_{(k_0, k), N}(h_t)] - \eta_{0:t} h_t| &\leq \sigma_{bias} \\ \mathbb{E}_\xi \left[\left(\Pi_{(k_0, k), N}(h_t) - \eta_{0:t} h_t \right)^2 \right] &\leq \sigma_{mse}^2, \end{aligned} \quad (3.15)$$

where

$$\begin{aligned} \sigma_{bias} &:= \frac{\mathbf{c}_t^{bias} \kappa_{t,N}^{k_0} \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty}{(k - k_0)(1 - \kappa_{t,N})N}, \\ \sigma_{mse}^2 &:= \frac{\left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2}{N(k - k_0)} \left(\mathbf{c}_t^{mse} + \frac{2\mathbf{c}_t^{cov}}{N^{1/2}(1 - \kappa_{t,N})} \right) \end{aligned}$$

and \mathbf{c}_t^{bias} , \mathbf{c}_t^{mse} and \mathbf{c}_t^{cov} are constants that do not depend on N and \mathbb{E}_ξ denotes the expectation under the law of the Markov chain formed by the PPG when initialized according to ξ .

The proof is provided in the supplementary material. Importantly, (3.15) provides a bound on the bias of the roll-out estimator that decreases exponentially with the burn-in period k_0 and is inversely proportional to the number N of particles. This means that we can improve the bias of the PARIS estimator with a better allocation of the computational resources.

4. Parameter learning with PPG

We now turn to parameter learning using PPG and gradient-based methods. We set the focus on learning the parameter θ of a function $V(\theta)$ whose gradient is the smoothed expectation of an additive functional $s_{0:t,\theta}$ in the form (2.6). Algorithm 4 defines a stochastic approximation (SA) scheme where the noise forms a parameter dependent Markov chain with associated invariant measure π_θ . We follow the approach of [Karimi et al., 2019] to establish a non-asymptotic bound over the mean field $h(\theta) := \pi_\theta s_{0:t,\theta}$. Such a setting encompasses for instance the following estimation procedures.

- (1) *Score ascent.* In the case of fully dominated HMMs, we are often interested in optimizing the log-likelihood of the observations given by $V(\theta) = \log \int \gamma_{0:t,\theta}(dx_{0:t})$. By applying *Fisher's identity*, we may express its gradient as a smoothed expectation of an additive functional according to

$$\begin{aligned} \nabla_\theta V(\theta) &= \int \nabla_\theta \log \gamma_{0:t}(x_{0:t}) \eta_{0:t,\theta}(dx_{0:t}), \\ &= \int \sum_{\ell=0}^{t-1} s_{\ell,\theta}(x_\ell, x_{\ell+1}) \eta_{0:t,\theta}(dx_{0:t}), \end{aligned}$$

where $s_{\ell,\theta} : \mathcal{X}_{\ell:\ell+1} \ni (x, x') \mapsto \nabla_\theta \log \{g_{\ell,\theta}(x)m_{\ell,\theta}(x, x')\}$ and $s_{0:t,\theta} := \sum_{\ell=0}^{t-1} s_{\ell,\theta}$.

- (2) *Inclusive KL surrogates.* Inspired by [Naesseth et al., 2020], we may consider the problem of learning a surrogate model for $\eta_{0:t,\theta}$ in the form $q_\phi(x_{0:t}) = q_\phi(x_0) \prod_{\ell=0}^{t-1} q_\phi(x_{\ell+1}, x_\ell)$ by minimizing $V(\phi) = \text{KL}(\eta_{0:t,\theta}, q_\phi)$.

Algorithm 3 Gradient estimation with roll-out PPG (GdEst)

Input: $\theta, \zeta_{0:t}[0], \{s_{\ell,\theta}\}_{\ell=0}^{t-1}$, number k of PPG iterations, burn-in k_0 .

Result: $\beta_t^{1:N}[k_0 : k], \zeta_{0:t}[k]$

```

10 for  $\ell \leftarrow 0$  to  $k - 1$  do
11   run  $(\tilde{\beta}_t^{1:N}[\ell + 1], \zeta_{0:t}[\ell + 1]) \leftarrow \text{PPG}(\theta; \zeta_{0:t}[\ell], \{s_{\ell,\theta}\}_{\ell=0}^{t-1})$ 
12   if  $\ell \geq k_0 - 1$  then
13     set  $\beta_t^{1:N}[\ell + 1] = \tilde{\beta}_t^{1:N}[\ell + 1]$ 

```

Algorithm 4 Score ascent with PPG.

Input: $\theta_0, \zeta_{0:t}[0]$, number k of PPG iterations, burn-in k_0 , number of SA iterations n , learning-rate sequence $\{\gamma_\ell\}_{\ell \in \mathbb{N}}$.

Result: θ_n

```

14 for  $i \leftarrow 0$  to  $n - 1$  do
15   run  $(\beta_t^{1:N}[k_0 : k], \zeta_{0:t}[i + 1]) \leftarrow \text{GdEst}(\theta_i, \zeta_{0:t}[i], \{s_{\ell,\theta_i}\}_{\ell=0}^{t-1}, k, k_0)$ 
16   set  $\Pi_{(k_0,k),N}(s_{0:t,\theta_i}) = (N(k - k_0))^{-1} \sum_{\ell=k_0}^{k-1} \sum_{j=1}^N \beta_t^j[\ell]$ 
17   set  $\theta_{i+1} \leftarrow \theta_i + \gamma_{i+1} \Pi_{(k_0,k),N}(s_{0:t,\theta_i})$ 

```

Note that Algorithm 3 defines a (collapsed) Markov kernel $\mathbb{P}_{\theta,t}$ defining for each path $\zeta_{0:t}$ a measure $\mathbb{P}_{\theta,t}(\zeta_{0:t}, d(\tilde{\zeta}_{0:t}, \tilde{\beta}_t^{1:N}[k_0 : k]))$ over the extended space of paths and sufficient statistics. Note that by evaluating the function $b_t^{1:N}[k_0 : k] \mapsto [N(k - k_0)]^{-1} \sum_{\ell=k_0+1}^k \sum_{j=1}^N b_t^j[\ell]$ at a realisation of this kernel gives the roll-out estimator whose properties are analysed in Theorem 1. The Markov kernel $\mathbb{P}_{\theta,t}$ is detailed in (B.72).

The following assumptions, are vital when analysing the convergence of Algorithm 4.

- A 4.1.** (i) The function $\theta \mapsto V(\theta)$ is L^V -smooth.

- (ii) The function $\theta \mapsto \eta_{0:t,\theta}$ is L^η -Lipschitz in total variation distance.
- (iii) For each path $\zeta_{0:t} \in \mathbf{X}_{0:t}$, the function

$$\theta \mapsto K_{\theta,t}(\zeta_{0:t}, d\tilde{\zeta}_{0:t}) \quad (4.16)$$

is L_1^P -Lipschitz in total variation distance, where $K_{\theta,t}$ is path-marginalized Markov transition kernel associated with the PPG algorithm when the model is parameterized by θ , see (A.41).

- (iv) For each path $\zeta_{0:t} \in \mathbf{X}_{0:t}$, the function

$$\theta \mapsto \mathbb{P}_{\theta,t} \Pi_{k_0-1,k,N}(s_{0:t,\theta})(\zeta_{0:t}) \quad (4.17)$$

is L_2^P -Lipschitz in total variation distance.

In the case of score ascent we check, in Appendix B, that these assumptions hold if the strong mixing assumption **A 3.1** is satisfied uniformly in θ , and with additional assumptions on the model. We are now ready to state a bound on the mean field $h(\theta)$ for Algorithm 4.

Theorem 2. *Assume **A 3.1** uniformly in θ and **A 4.1** and suppose that the stepsizes $\{\gamma_{\ell+1}\}_{\ell \in \llbracket 0, n-1 \rrbracket}$ satisfy $\gamma_{\ell+1} \leq \gamma_\ell$, $\gamma_\ell < a\gamma_{\ell+1}$, $\gamma_\ell - \gamma_{\ell+1} < a'\gamma_\ell^2$ and $\gamma_1 \leq 0.5(L^V + C_h)$ for some $a > 0$, $a' > 0$ and all $n \in \mathbb{N}$. Then,*

$$\mathbb{E} [\|h(\theta_\varpi)\|^2] \leq 2 \frac{V_{0,n} + C_{0,n} + C_{0,\gamma} \sum_{k=0}^n \gamma_{k+1}^2}{\sum_{k=0}^n \gamma_{k+1}}, \quad (4.18)$$

where $V_{0,n} = \mathbb{E} [V(\theta) - V(\theta_n)]$ and

$$C_{0,n} := \gamma_1 h(\theta_0) C_0 + \sigma_{bias} (\gamma_1 - \gamma_{n+1} + 1) \delta_{k,N,t}^{-1}, \quad (4.19)$$

$$C_{0,\gamma} := \sigma_{mse}^2 L^V + \sigma_{mse} C_1 + \sigma_{mse} \sigma_{bias} \left(L^V + \frac{C_2}{1 - \kappa_{N,t}} \right) \delta_{k,N,t}^{-1} + \sigma_{bias} L^V \delta_{k,N,t}^{-1}, \quad (4.20)$$

$$C_h := \left(C_1 + \sigma_{bias} \frac{C_2}{(1 - \kappa_{N,t}) \delta_{k,N,t}} \right) [(a+1)/2 + a\sigma_{mse}] + (L^V + a' + 1) \sigma_{bias} \delta_{k,N,t}^{-1}, \quad (4.21)$$

$$C_1 = L_2^P \left[1 + \kappa_{N,t}^k \delta_{k,N,t}^{-1} \right] + L^V \quad (4.22)$$

$$C_2 = L_1^P \delta_{k,N,t}^{-1} + L^\eta \kappa_{N,t}^k. \quad (4.23)$$

where C_0 is independent of σ_{bias} , σ_{mse} , N and where $\delta_{k,N,t} = 1 - \kappa_{N,t}^k$.

Theorem 2 establishes not only the convergence of Algorithm 4, but also illustrates the impact of the bias and the variance of the PPG on the convergence rate.

Remark 1. Under additional assumptions on the model (cf Appendix B), if we consider $\gamma_1 \leq 0.5(L^V + C_h)$, $\gamma_\ell = \gamma_1 \ell^{-1/2}$ for all $\ell \in \llbracket 1, n \rrbracket$, then $\sum_{k=0}^n \gamma_{k+1}^2 / \sum_{k=0}^n \gamma_{k+1} \sim \log n / \sqrt{n}$, showing that $\mathbb{E} [\|h(\theta_\varpi)\|^2]$ is $\mathcal{O}(\log n / \sqrt{n})$, where the leading constant depends on σ_{bias} and σ_{mse} .

Remark 1 establishes the rate of convergence of Algorithm 4. In principle we could try to optimize the parameters k, k_0 and N of the algorithm using these bounds, but one of the main challenges with this approach is the determination of the mixing rate, which is underestimated by $\kappa_{N,t}$. Still, our bound provides interesting information of the role of both bias and MSE.

5. Numerics

In this section, we focus on the numerical analysis of the two main results of the paper, namely the bias and MSE bounds of the roll-out estimator established in Theorem 1 and the efficiency of using PPG for learning in the framework developed in Section 4. For the latter, we will restrict ourselves to the case

of parameter learning via score ascent. In this setting, the competing method that corresponds most closely to the one presented here consists of using, as presented in Algorithm 5, a standard particle Gibbs sampler Π_θ instead of the PPG. One of the most common such samplers is the *particle Gibbs with ancestor sampling* (PGAS) presented in [Lindsten et al., 2014b]. In [Lindholm and Lindsten, 2018], the PGAS is used for parameter learning in HMMs via the Expectation Maximization (EM) algorithm.

Algorithm 5 Score ascent with particle Gibbs kernel.

Data: $\zeta_{0:t}[0]$, θ_0 , number k of paths per trajectory, burn-in k_0 , number n of SA iterations, learning-rate sequence $\{\gamma_\ell\}_{\ell \in \mathbb{N}}$, $\Pi_\theta(\zeta_{0:t}, d\tilde{\zeta}_{0:t})$ a Markov kernel targeting $\eta_{0:t}$.

Result: θ_n

```

18 for  $i \leftarrow 0$  to  $n - 1$  do
19   for  $j \leftarrow 0$  to  $k - 1$  do
20     sample  $\tilde{\zeta}_{0:t}[j+1] \sim \Pi_\theta(\tilde{\zeta}_{0:t}[j], \cdot)$ 
21     set  $\theta_{i+1} \leftarrow \theta_i + \frac{\gamma_{i+1}}{k-k_0} \sum_{\ell=k_0+1}^k s_{0:t, \theta_i}(\tilde{\zeta}_{0:t}[\ell])$ 
22     set  $\zeta_{0:t}[i+1] = \tilde{\zeta}_{0:t}[k]$ 

```

5.1 PPG

Linear Gaussian state-space model (LGSSM). We first consider a linear Gaussian HMM

$$X_{m+1} = AX_m + Q\epsilon_{m+1}, \quad Y_m = BX_m + R\zeta_m, \quad m \in \mathbb{N}, \quad (5.24)$$

where $\{\epsilon_m\}_{m \in \mathbb{N}^*}$ and $\{\zeta_m\}_{m \in \mathbb{N}}$ are sequences of independent standard normally distributed random variables, independent of X_0 . The coefficients A , Q , B , and R are assumed to be known and equal to 0.97, 0.60, 0.54, and 0.33, respectively. Using this parameterisation, we generate, by simulation, a record of $t = 999$ observations.

In this setting, we aim at computing smoothed expectations of the state one-lag covariance $h_t(x_{0:t}) := \sum_{m=0}^{t-1} x_m x_{m+1}$. In the linear Gaussian case, the *disturbance smoother* (see [Cappé et al., 2005, Algorithm 5.2.15]) provides the exact values of the smoothed sufficient statistics, which allows us to study the bias of the estimator for a given computational budget C . Figure 1 displays, for three different total budgets C , the distribution of estimates of $\eta_{0:n} h_n$ using the PARIS as well as three different configurations of the PPG corresponding to $k \in \{2, 4, 10\}$ (and $N = C/k$) with $k_0 = k/2$ and $k_0 = k/4$. The reference value is shown as a red-dashed line and the mean value of each distribution is shown as a black-dashed line. Each boxplot is based on 1000 independent replicates of the corresponding estimator. We observe that in this example, all configurations of the PPG are less biased than the equivalent PARIS estimator. The illustration of the bounds from Theorem 1 is postponed to Appendix D.1.

5.2 Score ascent

LGSSM. We consider the LGSSM with state and observation spaces being \mathbb{R}^5 . We assume that the parameters R and Q are known and consider the inference of $\theta = (A, B)$ on the basis of a simulated sequence of $n = 999$ observations. In this setting, the M-step of the EM algorithm can be solved exactly with the disturbance smoother [Cappé et al., 2005, Chapter 11]. The parameter obtained by this procedure (denoted θ_{mle}) is the reference value for any likelihood maximization algorithm. Table 1 shows the L_2 distance between the singular values of θ_{mle} and those of the parameters obtained by Algorithm 4 and Algorithm 5. The CLT confidence intervals were obtained on the basis of 25 replicates. The configurations respect a given particle budget $kN = C = 1024$. The choice of keeping $k_0 = k/2$ is a heuristic rule to achieve a good bias–variance trade-off, but other combinations of k_0 and k may lead to better performance for different problems. We analyse this for the LGSSM in Appendix D.2. All settings are the same for both algorithms and are described in Appendix D.2. The PPG achieves




Figure 1: PARIS and PPG outputs for the LGSSM for $C = 500$, yellow boxes correspond to PPG outputs produced using $k \in \{50, 20, 10, 5\}$ iterations and $N \in \{C/50, C/20, C/10, C/5\}$ particles. The image on the left corresponds to taking $k_0 = k/2$ and the one on the right to $k_0 = k/4$.

consistently a smaller distance to θ_{mle} . Figure 2 displays, for each estimator and configuration, the evolution of the distance to the MLE estimator as a function of the iteration index.




Figure 2: Distance to the MLE estimator as a function of the iteration step for PGAS and PPG with different parameters while keeping the particle budget fixed for LGSSM for 25 different seeds.

CRNN. We consider now the problem of inference in a non-linear HMM and in particular the chaotic recurrent neural network introduced by [Zhao et al., 2021]. We use the same setting as in the original paper. The state and observation equations are

$$\begin{aligned} X_{m+1} &= X_m + \tau^{-1} \Delta(-X_m + \gamma W \tanh(X_m)) + \epsilon_{m+1}, \\ Y_m &= BX_m + \zeta_m, \quad m \in \mathbb{N}, \end{aligned}$$

where $\{\epsilon_m\}_{m \in \mathbb{N}^*}$ is a sequence of 20-dimensional independent multivariate Gaussian random variables with zero mean and covariance $0.01\mathbf{I}$ and $\{\zeta_m\}_{m \in \mathbb{N}}$ is a sequence of independent random variables

Table 1: Distance to θ_{MLE} for each configuration in the LGSSM case.

Algorithm	N	k_0	k	D_{mle}
PGAS	32	32	64	0.793 ± 0.048
PGAS	64	16	32	0.751 ± 0.052
PGAS	128	8	16	0.633 ± 0.054
PGAS	256	4	8	0.580 ± 0.049
PPG	32	32	64	0.358 ± 0.038
PPG	64	16	32	0.373 ± 0.031
PPG	128	8	16	0.355 ± 0.043
PPG	256	4	8	0.351 ± 0.042

Table 2: Per configuration negative loglikelihood for the CRNN model.

Algorithm	N	k_0	k	NLL
PGAS	32	16	32	31364.932 ± 173.708
PGAS	64	8	16	31083.408 ± 380.527
PGAS	128	4	8	30264.836 ± 265.880
PPG	32	16	32	22291.971 ± 47.683
PPG	64	8	16	22314.537 ± 25.028
PPG	128	4	8	22353.416 ± 39.443

where each component is distributed independently according to a Student’s t-distribution with scale 0.1 and 2 degrees of freedom.

In this case, the natural metric used to evaluate the different estimators is the negative log likelihood (NLL). We use the unbiased estimator of the likelihood given by the mean of the log weights produced by a particle filter [Douc et al., 2014, Section 12.1] using $N = 10^4$ particles. Table 2 shows the results obtained for 25 different replications for several different configurations of PPG and PGAS, while keeping total budget of particles fixed. Further numerical details are given in Appendix D.2. We observe that PPG achieves the a considerably lower NLL than PGAS in all configurations.

6. Conclusion

We have presented a new algorithm, referred to as PPG as well as bounds on its bias and MSE in Theorem 1. We then propose a way of using PPG in a learning framework and derive a non-asymptotic bound over the gradient of the updates when doing score ascent with the PPG with explicit dependence on the bias and MSE of the estimator. We provide numerical simulations to support our claims, and we show that our algorithm outperforms the current competitors in the two different examples analysed.

Contents

1	Introduction	1
2	Background	3
2.1	Hidden Markov models	3
2.2	Particle filters	3
2.3	Backward smoothing and the PARIS algorithm	4
3	PARIS particle Gibbs	5
3.1	Particle Gibbs methods	5
3.2	The PPG algorithm	6
4	Parameter learning with PPG	8
5	Numerics	9
5.1	PPG	10
5.2	Score ascent	10
6	Conclusion	12
A	PPG	14
A.1	Many-body Feynman–Kac models	14
A.2	Backward interpretation of Feynman–Kac path flows	15
A.3	Conditional dual processes and particle Gibbs	16
A.4	The PARIS algorithm	17
A.5	Proof of Theorem 1	21
A.6	Proofs of intermediate results	24
A.6.1	Proof of Proposition 1	24
A.6.2	Proof of Theorem 3	24
A.6.3	Proof of Theorem 5	25
A.6.4	Proof of Proposition 2	29
A.6.5	Proof of Theorem 7	31
A.6.6	Proof of Proposition 4	34
A.6.7	Proof of Proposition 5	35
B	Learning with PPG	36
B.1	Non-asymptotic bound	36
B.2	Application to Theorem 2	38
B.2.1	Verification of the assumptions of Theorem 8	38
B.2.2	Proof of Theorem 2	42
B.3	Conditions on the model to verify A 4.1	42
C	Lipschitz properties	44
C.1	Lipschitz continuity of \mathbb{P}_θ ,	44
C.1.1	$\theta \mapsto \mathbb{C}_{m,\theta}$ is Lipschitz.	46
C.1.2	$\theta \mapsto \mathbb{B}_{t,\theta}(\mathbf{x}_{0:t}, \cdot)$ is Lipschitz	48
C.1.3	$\theta \mapsto \int \mathbb{S}_{t,\theta}(\mathbf{x}_{0:t}, d\mathbf{b}_t)\mu(\mathbf{b}_t)(\text{id})$ is Lipschitz	49
C.2	Lipschitz properties of Markov Kernels	51
D	Additional numerical results	53
D.1	PPG	53
D.2	Learning	53

A. PPG

In this section, we develop the theoretical framework necessary to establish Theorem 1. We recall the notions of *Feynman–Kac models*, *many-body Feynman–Kac models*, *backward interpretations*, and *conditional dual processes*. Our presentation follows closely [Del Moral et al., 2016] but with a different and hopefully more transparent definition of the many-body extensions. We restate (in Theorem 3 below) a duality formula of [Del Moral et al., 2016] relating these concepts. This formula provides a foundation for the *particle Gibbs sampler* described in Algorithm 2.

Notations. Let (Z, \mathcal{Z}) be a measurable space and L another possibly unnormalised transition kernel on $Y \times Z$. Define, with K as above,

$$KL : X \times Z \ni (x, A) \mapsto \int L(y, A) K(x, dy)$$

and

$$K \otimes L : X \times (Y \otimes Z) \ni (x, A) \mapsto \iint \mathbb{1}_A(y, z) K(x, dy) L(y, dz),$$

whenever these are well defined. This also defines the \otimes products of a kernel K on $X \times Y$ and a measure ν on X as well as of a kernel L on $Y \times X$ and a measure μ on Y as the measures

$$\nu \otimes K : X \otimes Y \ni A \mapsto \iint \mathbb{1}_A(x, y) K(x, dy) \nu(dx),$$

$$L \otimes \mu : X \otimes Y \ni A \mapsto \iint \mathbb{1}_A(x, y) L(y, dx) \mu(dy).$$

A.1 Many-body Feynman–Kac models

In the following, we assume that all random variables are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The distribution flow $\{\eta_m\}_{m \in \mathbb{N}}$ defined in eq. (2.4) is intractable in general, but can be approximated by random samples $\boldsymbol{\xi}_m = \{\xi_m^i\}_{i=1}^N$, $m \in \mathbb{N}$, referred to as *particles*, where $N \in \mathbb{N}^*$ is a fixed Monte Carlo sample size and each particle ξ_m^i is an X_m -valued random variable. Such particle approximation is based on the recursion $\eta_{m+1} = \Phi_m(\eta_m)$, $m \in \mathbb{N}$, where Φ_m denotes the mapping

$$\Phi_m : M_1(\mathcal{X}_m) \ni \eta \mapsto \frac{\eta Q_m}{\eta g_m} \tag{A.25}$$

taking on values in $M_1(\mathcal{X}_{m+1})$. In order to describe recursively the evolution of the particle population, let $m \in \mathbb{N}$ and assume that the particles $\boldsymbol{\xi}_m$ form a consistent approximation of η_m in the sense that $\mu(\boldsymbol{\xi}_m)h$, where $\mu(\boldsymbol{\xi}_m) := N^{-1} \sum_{i=1}^N \delta_{\xi_m^i}$, with δ_x denotes the Dirac measure located at x , is the occupation measure formed by $\boldsymbol{\xi}_m$, which serves as a proxy for $\eta_m h$ for all η_m -integrable test functions h . Under general conditions, $\mu(\boldsymbol{\xi}_m)h$ converges in probability to η_m with $N \rightarrow \infty$; see [Del Moral, 2004, Chopin and Papaspiliopoulos, 2020] and references therein. Then, in order to generate an updated particle sample approximating η_{m+1} , new particles $\boldsymbol{\xi}_{m+1} = \{\xi_{m+1}^i\}_{i=1}^N$ are drawn conditionally independently given $\boldsymbol{\xi}_m$ according to

$$\xi_{m+1}^i \sim \Phi_m(\mu(\boldsymbol{\xi}_m)) = \sum_{\ell=1}^N \frac{g_m(\xi_m^\ell)}{\sum_{\ell'=1}^N g_m(\xi_m^{\ell'})} M_m(\xi_m^\ell, \cdot), \quad i \in \llbracket 1, N \rrbracket.$$

Since this process of particle updating involves sampling from the mixture distribution $\Phi_m(\mu(\boldsymbol{\xi}_m))$, it can be naturally decomposed into two substeps: *selection* and *mutation*. The selection step consists of randomly choosing the ℓ -th mixture stratum with probability $g_m(\xi_m^\ell) / \sum_{\ell'=1}^N g_m(\xi_m^{\ell'})$ and the mutation step consists of drawing a new particle ξ_{m+1}^i from the selected stratum $M_m(\xi_m^\ell, \cdot)$. In

[Del Moral et al., 2016], the term *many-body Feynman–Kac models* is related to the law of process $\{\xi_m\}_{m \in \mathbb{N}}$. For all $m \in \mathbb{N}$, let $\mathbf{X}_m := \mathbf{X}_m^N$ and $\mathcal{X}_m := \mathcal{X}_m^{\otimes N}$; then $\{\xi_m\}_{m \in \mathbb{N}}$ is an inhomogeneous Markov chain on $\{\mathbf{X}_m\}_{m \in \mathbb{N}}$ with transition kernels

$$\mathbf{M}_m : \mathbf{X}_m \times \mathcal{X}_{m+1} \ni (\mathbf{x}_m, A) \mapsto \Phi_m(\mu(\mathbf{x}_m))^{\otimes N}(A)$$

and initial distribution $\eta_0 = \eta_0^{\otimes N}$. Now, denote $\mathbf{X}_{0:n} := \prod_{m=0}^n \mathbf{X}_m$ and $\mathcal{X}_{0:n} := \otimes_{m=0}^n \mathcal{X}_m$. In the following, we use a bold symbol to stress that a quantity is related to the many-body process. The *many-body Feynman–Kac path model* refers to the flows $\{\gamma_m\}_{m \in \mathbb{N}}$ and $\{\eta_m\}_{m \in \mathbb{N}}$ of the unnormalised and normalised, respectively, probability distributions on $\{\mathcal{X}_{0:m}\}_{m \in \mathbb{N}}$ generated by (2.4) and (2.3) for the Markov kernels $\{\mathbf{M}_m\}_{m \in \mathbb{N}}$, the initial distribution η_0 , the potential functions

$$\mathbf{g}_m : \mathbf{X}_m \ni \mathbf{x}_m \mapsto \mu(\mathbf{x}_m)g_m = \frac{1}{N} \sum_{i=1}^N g_m(x_m^i), \quad m \in \mathbb{N},$$

and the corresponding unnormalised transition kernels

$$\mathbf{Q}_m : \mathbf{X}_m \times \mathcal{X}_{m+1} \ni (\mathbf{x}_m, A) \mapsto \mathbf{g}_m(\mathbf{x}_m)\mathbf{M}_m(\mathbf{x}_m, A), \quad m \in \mathbb{N}.$$

A.2 Backward interpretation of Feynman–Kac path flows

Suppose that each kernel Q_n , $n \in \mathbb{N}$, defined in (2.2), has a transition density q_n with respect to some dominating measure $\lambda_{n+1} \in \mathcal{M}(\mathcal{X}_{n+1})$. Then for $n \in \mathbb{N}$ and $\eta \in \mathcal{M}_1(\mathcal{X}_n)$ we may define the *backward kernel*

$$\overleftarrow{Q}_{n,\eta} : \mathcal{X}_{n+1} \times \mathcal{X}_n \ni (x_{n+1}, A) \mapsto \frac{\int \mathbb{1}_A(x_n)q_n(x_n, x_{n+1})\eta(dx_n)}{\int q_n(x'_n, x_{n+1})\eta(dx'_n)}. \quad (\text{A.26})$$

Now, denoting, for $n \in \mathbb{N}^*$,

$$B_n : \mathbf{X}_n \times \mathcal{X}_{0:n-1} \ni (x_n, A) \mapsto \int \cdots \int \mathbb{1}_A(x_{0:n-1}) \prod_{m=0}^{n-1} \overleftarrow{Q}_{m,\eta_m}(x_{m+1}, dx_m), \quad (\text{A.27})$$

we may state the following—now classical—*backward decomposition* of the Feynman–Kac path measures, a result that plays a pivotal role in this paper.

Proposition 1. *For every $n \in \mathbb{N}^*$ it holds that $\gamma_{0:n} = \gamma_n \otimes B_n$ and $\eta_{0:n} = \eta_n \otimes B_n$.*

Although the decomposition in Proposition 1 is well known (see, *e.g.*, [Del Moral et al., 2010, Del Moral et al., 2016]), we provide a proof in Appendix A.6.1 for completeness. Using the backward decomposition, a particle approximation of a given Feynman–Kac path measure $\eta_{0:n}$ is obtained by first sampling, in an initial forward pass, particle clouds $\{\xi_m\}_{m=0}^n$ from $\eta_0 \otimes \mathbf{M}_0 \otimes \cdots \otimes \mathbf{M}_{n-1}$ and then sampling, in a subsequent backward pass, for instance N conditionally independent paths $\{\tilde{\xi}_{0:n}^i\}_{i=1}^N$ from $\mathbb{B}_n(\xi_0, \dots, \xi_n, \cdot)$, where

$$\mathbb{B}_n : \mathbf{X}_{0:n} \times \mathcal{X}_{0:n} \ni (\mathbf{x}_{0:n}, A) \mapsto \int \cdots \int \mathbb{1}_A(x_{0:n}) \left(\prod_{m=0}^{n-1} \overleftarrow{Q}_{m,\mu(\mathbf{x}_m)}(x_{m+1}, dx_m) \right) \mu(\mathbf{x}_n)(dx_n) \quad (\text{A.28})$$

is a Markov kernel describing the time-reversed dynamics induced by the particle approximations generated in the forward pass. Here and in the following we use blackboard notation to denote kernels related to many-body path spaces. Finally, $\mu(\{\tilde{\xi}_{0:n}^i\}_{i=1}^N)h$ is returned as an estimator of $\eta_{0:n}h$ for any $\eta_{0:n}$ -integrable test function h . This algorithm is in the literature referred to as the *forward-filtering backward-simulation (FFBSi) algorithm* and was introduced in [Godsill et al., 2004]; see also [Cappé et al., 2007, Douc et al., 2011]. More precisely, given the forward particles $\{\xi_m\}_{m=0}^n$, each path

$\tilde{\xi}_{0:n}^i$ is generated by first drawing $\tilde{\xi}_n^i$ uniformly among the particles ξ_n in the last generation and then drawing, recursively,

$$\tilde{\xi}_m^i \sim \overleftarrow{Q}_{m,\mu(\xi_m)}(\tilde{\xi}_{m+1}^i, \cdot) = \sum_{j=1}^N \frac{q_m(\xi_m^j, \tilde{\xi}_{m+1}^i)}{\sum_{\ell=1}^N q_m(\xi_m^\ell, \tilde{\xi}_{m+1}^i)} \delta_{\xi_m^j}(\cdot), \quad (\text{A.29})$$

i.e., given $\tilde{\xi}_{m+1}^i, \tilde{\xi}_m^i$ is picked at random among the ξ_m according to weights proportional to $\{q_m(\xi_m^j, \tilde{\xi}_{m+1}^i)\}_{j=1}^N$. Note that in this basic formulation of the FFBSi algorithm, each backward-sampling operation (A.29) requires the computation of the normalising constant $\sum_{\ell=1}^N q_m(\xi_m^\ell, \tilde{\xi}_{m+1}^i)$, which implies an overall quadratic complexity of the algorithm. Still, this heavy computational burden can be eased by means of an effective accept-reject technique discussed in Appendix A.4.

A.3 Conditional dual processes and particle Gibbs

The *dual process* associated with a given Feynman–Kac model (2.4–2.3) and a given trajectory $\{z_n\}_{n \in \mathbb{N}}$, where $z_n \in \mathbf{X}_n$ for every $n \in \mathbb{N}$, is defined as the canonical Markov chain with kernels

$$\mathbf{M}_n \langle z_{n+1} \rangle : \mathbf{X}_n \times \mathcal{X}_{n+1} \ni (\mathbf{x}_n, A) \mapsto \frac{1}{N} \sum_{i=0}^{N-1} \left(\Phi_n(\mu(\mathbf{x}_n))^{\otimes i} \otimes \delta_{z_{n+1}} \otimes \Phi_n(\mu(\mathbf{x}_n))^{\otimes (N-i-1)} \right) (A), \quad (\text{A.30})$$

for $n \in \mathbb{N}$, and initial distribution

$$\boldsymbol{\eta}_0 \langle z_0 \rangle := \frac{1}{N} \sum_{i=0}^{N-1} \left(\eta_0^{\otimes i} \otimes \delta_{z_0} \otimes \eta_0^{\otimes (N-i-1)} \right). \quad (\text{A.31})$$

As clear from (A.30–A.31), given $\{z_n\}_{n \in \mathbb{N}}$, a realisation $\{\xi_n\}_{n \in \mathbb{N}}$ of the dual process is generated as follows. At time zero, the process is initialised by inserting z_0 at a randomly selected position in the vector ξ_0 while drawing independently the remaining components from η_0 . Then, given ξ_n at step n , z_{n+1} is inserted at a randomly selected position in ξ_{n+1} while drawing independently the remaining components from $\Phi_n(\mu(\xi_n))$.

In order to describe compactly the law of the conditional dual process, we define the Markov kernel

$$\mathbb{C}_n : \mathbf{X}_{0:n} \times \mathcal{X}_{0:n} \ni (z_{0:n}, A) \mapsto \boldsymbol{\eta}_0 \langle z_0 \rangle \otimes \mathbf{M}_0 \langle z_1 \rangle \otimes \cdots \otimes \mathbf{M}_{n-1} \langle z_n \rangle (A).$$

The following result elegantly combines the underlying model (2.4–2.3), the many-body Feynman–Kac model, the backward decomposition, and the conditional dual process.

Theorem 3 ([Del Moral et al., 2016]). *For all $n \in \mathbb{N}$,*

$$\mathbb{B}_n \otimes \boldsymbol{\gamma}_{0:n} = \boldsymbol{\gamma}_{0:n} \otimes \mathbb{C}_n. \quad (\text{A.32})$$

In [Del Moral et al., 2016], each state ξ_n of the many-body process maps an outcome ω of the sample space Ω into an *unordered set* of N elements in \mathbf{X}_n . However, we have chosen to let each ξ_n take on values in the standard *product space* \mathbf{X}_n^N for two reasons: first, the construction of [Del Moral et al., 2016] requires sophisticated measure-theoretic arguments to endow such unordered sets with suitable σ -fields and appropriate measures; second, we see no need to ignore the index order of the particles as long as the Markovian dynamics (A.30–A.31) of the conditional dual process is symmetrised over the particle cloud. Therefore, in Appendix A.6.2, we include our own proof of duality (A.32) for completeness. Note that the measure (A.32) on $\mathcal{X}_{0:n} \otimes \mathcal{X}_{0:n}$ is unnormalised, but since the kernels \mathbb{B}_n and \mathbb{C}_n are both Markovian, normalising the identity with $\boldsymbol{\gamma}_{0:n}(\mathbf{X}_{0:n}) = \boldsymbol{\gamma}_{0:n}(\mathbf{X}_{0:n})$ yields immediately

$$\mathbb{B}_n \otimes \boldsymbol{\eta}_{0:n} = \boldsymbol{\eta}_{0:n} \otimes \mathbb{C}_n. \quad (\text{A.33})$$

Since the two sides of (A.33) provide the full conditionals, it is natural to choose a data-augmentation approach and sample the target (A.33) using a two-stage deterministic-scan Gibbs sampler [Andrieu et al., 2010b, Chopin and Singh, 2015a]. More specifically, assume that we have generated a state $(\boldsymbol{\xi}_{0:n}[\ell], \zeta_{0:n}[\ell])$ comprising a dual process with associated path on the basis of $\ell \in \mathbb{N}$ iterations of the sampler; then the next state $(\boldsymbol{\xi}_{0:n}[\ell+1], \zeta_{0:n}[\ell+1])$ is generated in a Markovian fashion by sampling first $\boldsymbol{\xi}_{0:n}[\ell+1] \sim \mathbb{C}_n(\zeta_{0:n}[\ell], \cdot)$ and then sampling $\zeta_{0:n}[\ell+1] \sim \mathbb{B}_n(\boldsymbol{\xi}_{0:n}[\ell+1], \cdot)$. After arbitrary initialisation (and the discard of possible burn-in iterations), this procedure produces a Markov trajectory $\{(\boldsymbol{\xi}_{0:n}[\ell], \zeta_{0:n}[\ell])\}_{\ell \in \mathbb{N}}$, and under weak additional technical conditions this Markov chain admits (A.33) as its unique invariant distribution. In such a case, the Markov chain is ergodic [Douc et al., 2018, Chapter 5], and the marginal distribution of the conditioning path $\zeta_{0:n}[\ell]$ converges to the target distribution $\eta_{0:n}$. Therefore, for every $h \in \mathbb{F}(\mathcal{X}_{0:n})$,

$$\lim_{L \rightarrow \infty} \frac{1}{L} \sum_{\ell=1}^L h(\zeta_{0:n}[\ell]) = \eta_{0:n} h, \quad \mathbb{P}\text{-a.s.}$$

A.4 The PARIS algorithm

In the following, we assume that we are given a sequence $\{h_n\}_{n \in \mathbb{N}}$ of *additive state functionals* as in (2.6). This problem is particularly relevant in the context of maximum-likelihood-based parameter estimation in general state-space models, *e.g.*, when computing the *score-function*, *i.e.* the gradient of the log-likelihood function, via the Fisher identity or when computing the intermediate quantity of the *Expectation Maximization* (EM) *algorithm*, in which case $\eta_{0:n}$ and h_n correspond to the joint state posterior and an element of some sufficient statistic, respectively; see [Cappé and Moulines, 2005, Douc et al., 2011, Del Moral et al., 2010, Poyiadjis et al., 2011, Olsson and Westerborn, 2017] and the references therein. Interestingly, as noted in [Cappé, 2011, Del Moral et al., 2010], the backward decomposition allows, when applied to additive state functionals, a forward recursion for the expectations $\{\eta_{0:n} h_n\}_{n \in \mathbb{N}}$. More specifically, using the forward decomposition $h_{n+1}(x_{0:n+1}) = h_n(x_{0:n}) + \tilde{h}_n(x_n, x_{n+1})$ and the backward kernel B_{n+1} defined in (A.27), we may write, for $x_{n+1} \in \mathcal{X}_{n+1}$,

$$\begin{aligned} B_{n+1} h_{n+1}(x_{n+1}) &= \int \overleftarrow{Q}_{n, \eta_n}(x_{n+1}, dx_n) \int \left(h_n(x_{0:n}) + \tilde{h}_n(x_n, x_{n+1}) \right) B_n(x_n, dx_{0:n-1}) \\ &= \overleftarrow{Q}_{n, \eta_n}(B_n h_n + \tilde{h}_n)(x_{n+1}), \end{aligned} \quad (\text{A.34})$$

which by Proposition 1 implies that

$$\eta_{0:n+1} h_{n+1} = \eta_{n+1} \overleftarrow{Q}_{n, \eta_n}(B_n h_n + \tilde{h}_n). \quad (\text{A.35})$$

Since the marginal flow $\{\eta_n\}_{n \in \mathbb{N}}$ can be expressed recursively via the mappings $\{\Phi_n\}_{n \in \mathbb{N}}$, (A.35) provides, in principle, a basis for online computation of $\{\eta_{0:n} h_n\}_{n \in \mathbb{N}}$. To handle the fact that the marginals are generally intractable we may, following [Del Moral et al., 2010], plug particle approximations $\mu(\boldsymbol{\xi}_{n+1})$ and $\overleftarrow{Q}_{n, \mu(\boldsymbol{\xi}_n)}$ (see (A.29)) of η_{n+1} and $\overleftarrow{Q}_{n, \mu(\eta_n)}$, respectively, into the recursion (A.35). More precisely, we proceed recursively and assume that at time n we have at hand a sample $\{(\boldsymbol{\xi}_n^i, \beta_n^i)\}_{i=1}^N$ of particles with associated statistics, where each statistic β_n^i serves as an approximation of $B_n h_n(\boldsymbol{\xi}_n^i)$; then evolving the particle cloud according to $\boldsymbol{\xi}_{n+1} \sim \mathbf{M}_n(\boldsymbol{\xi}_n, \cdot)$ and updating the statistics using (A.34), with $\overleftarrow{Q}_{n, \eta_n}$ replaced by $\overleftarrow{Q}_{n, \mu(\boldsymbol{\xi}_n)}$, yields the particle-wise recursion

$$\beta_{n+1}^i = \sum_{\ell=1}^N \frac{q_n(\boldsymbol{\xi}_n^\ell, \boldsymbol{\xi}_{n+1}^i)}{\sum_{\ell'=1}^N q_n(\boldsymbol{\xi}_n^{\ell'}, \boldsymbol{\xi}_{n+1}^i)} \left(\beta_n^\ell + \tilde{h}_n(\boldsymbol{\xi}_n^\ell, \boldsymbol{\xi}_{n+1}^i) \right), \quad i \in \llbracket 1, N \rrbracket, \quad (\text{A.36})$$

and, finally, the estimator

$$\mu(\boldsymbol{\beta}_n)(\text{id}) = \frac{1}{N} \sum_{i=1}^N \beta_n^i \quad (\text{A.37})$$

of $\eta_{0:n}h_n$, where $\beta_n := (\beta_n^1, \dots, \beta_n^N)$, $i \in \llbracket 1, N \rrbracket$. The procedure is initialised by simply letting $\beta_0^i = 0$ for all $i \in \llbracket 1, N \rrbracket$. Note that (A.37) provides a particle interpretation of the backward decomposition in Proposition 1. This algorithm is a special case of the *forward-filtering backward-smoothing* (FFBSm) *algorithm* (see [Andrieu and Doucet, 2003, Godsill et al., 2004, Douc et al., 2011, Särkkä, 2013]) for additive functionals satisfying (2.6). It allows for online processing of the sequence $\{\eta_{0:n}h_n\}_{n \in \mathbb{N}}$, but has also the appealing property that only the current particles ξ_n and statistics β_n need to be stored. However, since each update (A.36) requires the summation of N terms, the scheme has an overall *quadratic* complexity in the number of particles, leading to a computational bottleneck in applications to complex models that require large particle sample sizes N .

In order to detour the computational burden of this forward-only implementation of FFBSm, the PARIS algorithm [Olsson and Westerborn, 2017] updates the statistics β_n by replacing each sum (A.36) by a Monte Carlo estimate

$$\beta_{n+1}^i = \frac{1}{M} \sum_{j=1}^M \left(\tilde{\beta}_n^{i,j} + \tilde{h}_n(\tilde{\xi}_n^{i,j}, \xi_{n+1}^i) \right), \quad i \in \llbracket 1, N \rrbracket, \quad (\text{A.38})$$

where $\{(\tilde{\xi}_n^{i,j}, \tilde{\beta}_n^{i,j})\}_{j=1}^M$ are drawn randomly among $\{(\xi_n^i, \beta_n^i)\}_{i=1}^N$ with replacement, by assigning $(\tilde{\xi}_n^{i,j}, \tilde{\beta}_n^{i,j})$ the value of $(\xi_n^\ell, \beta_n^\ell)$ with probability $q_n(\xi_n^\ell, \xi_{n+1}^i) / \sum_{\ell'=1}^N q_n(\xi_n^{\ell'}, \xi_{n+1}^i)$, and the Monte Carlo sample size $M \in \mathbb{N}^*$ is supposed to be much smaller than N (say, less than 5). Formally,

$$\{(\tilde{\xi}_n^{i,j}, \tilde{\beta}_n^{i,j})\}_{j=1}^M \sim \left(\sum_{\ell=1}^N \frac{q_n(\xi_n^\ell, \xi_{n+1}^i)}{\sum_{\ell'=1}^N q_n(\xi_n^{\ell'}, \xi_{n+1}^i)} \delta_{(\xi_n^\ell, \beta_n^\ell)} \right)^{\otimes M}, \quad i \in \llbracket 1, N \rrbracket.$$

The resulting procedure, summarised in Algorithm 1, allows for online processing with constant memory requirements, since it only needs to store the current particle cloud and the estimated auxiliary statistics at each iteration. Moreover, in the case where the Markov transition densities of the model can be uniformly bounded, *i.e.* when there exists, for every $n \in \mathbb{N}$, an upper bound $\bar{\sigma}_n > 0$ such that for all $(x_n, x_{n+1}) \in \mathbf{X}_n \times \mathbf{X}_{n+1}$, $m_n(x_n, x_{n+1}) \leq \bar{\sigma}_n$ (a weak assumption satisfied for most models of interest), a sample $(\tilde{\xi}_n^{i,j}, \tilde{\beta}_n^{i,j})$ can be generated by drawing, with replacement and until acceptance, candidates $(\tilde{\xi}_n^{i,*}, \tilde{\beta}_n^{i,*})$ from $\{(\xi_n^i, \beta_n^i)\}_{i=1}^N$ according to the normalised particle weights $\{g_n(\xi_n^\ell) / \sum_{\ell'} g_n(\xi_n^{\ell'})\}_{\ell=1}^N$, obtained as a by-product in the generation of ξ_{n+1} , and accepting the same with probability $m_n(\tilde{\xi}_n^{i,*}, \xi_{n+1}^i) / \bar{\sigma}_n$. As this sampling procedure bypasses completely the calculation of the normalising constant $\sum_{\ell'=1}^N q_n(\xi_n^{\ell'}, \xi_{n+1}^i)$ of the targeted categorical distribution, it yields an overall $\mathcal{O}(MN)$ complexity of the algorithm as a whole; see [Douc et al., 2011] for details.

Increasing M improves the accuracy of the algorithm at the cost of additional computational complexity. As shown in [Olsson and Westerborn, 2017], there is a qualitative difference between the cases $M = 1$ and $M \geq 2$, and it turns out that the latter is required to keep PARIS numerically stable. More precisely, in the latter case, it can be shown that the PARIS estimator $\mu(\beta_n)$ satisfies, as N tends to infinity while M is held fixed, a central limit theorem (CLT) at the rate \sqrt{N} and with an n -normalised asymptotic variance of order $\mathcal{O}(1 - 1/(M - 1))$. As clear from this bound, using a large M only yields a waste of computational work, and setting M to 2 or 3 typically works well in practice.

We now introduce the *Parisian particle Gibbs* (PPG) *algorithm*. For all $t \in \mathbb{N}^*$, let $\mathbf{Y}_t := \mathbf{X}_{0:t} \times \mathbb{R}$ and $\mathcal{Y}_t := \mathcal{X}_{0:t} \otimes \mathcal{B}(\mathbb{R})$. Moreover, let $\mathbf{Y}_0 := \mathbf{X}_0 \times \{0\}$ and $\mathcal{Y}_0 := \mathcal{X}_0 \otimes \{\{0\}, \emptyset\}$. An element of \mathbf{Y}_t will always be denoted by $y_t = (x_{0:t|t}, b_t)$. The Parisian particle Gibbs sampler comprises, as a key ingredient, a *conditional PARIS step*, which updates recursively a set of \mathbf{Y}_t -valued random variables $v_t^i := (\xi_{0:t|t}^i, \beta_t^i)$, $i \in \llbracket 1, N \rrbracket$. Let $(\mathbf{v}_t)_{t \in \mathbb{N}}$ denote the corresponding many-body process, each $\mathbf{v}_t := \{(\xi_{0:t|t}^i, \beta_t^i)\}_{i=1}^N$ taking on values in the space $\mathbf{Y}_t := \mathbf{Y}_t^N$, which we furnish with a σ -field $\mathcal{Y}_t := \mathcal{Y}_t^{\otimes N}$. The space \mathbf{Y}_0 and the corresponding σ -field \mathcal{Y}_0 are defined accordingly. For every $t \in \mathbb{N}$, we write $\xi_{0:t|t}$ for the collection $\{\xi_{0:t|t}^i\}_{i=1}^N$ of paths in \mathbf{v}_t , and $\xi_{t|t}$ for the collection $\{\xi_{t|t}^i\}_{i=1}^N$ of end points of the same.

In the following, we let $t \in \mathbb{N}$ be a fixed time horizon, and describe in detail how the PPG approximates $\eta_{0:t}h_t$ iteratively. In short, at each iteration ℓ , the PPG produces, given an input conditional

path $\zeta_{0:t}[\ell]$, a many-body system $\mathbf{v}_t[\ell + 1]$ by means of a series of conditional PARIS operations; then, an updated path $\zeta_{0:t}[\ell + 1]$, serving as input at the next iteration, is generated by picking one of the paths $\xi_{0:t}[\ell + 1]$ in $\mathbf{v}_t[\ell + 1]$ at random. At each iteration, the produced statistics β_t in \mathbf{v}_t provides an approximation of $\eta_{0:t}h_t$ according to (A.37).

More precisely, given the path $\zeta_{0:t}[\ell]$, the conditional PARIS operations are executed as follows. In the initial step, $\xi_{0|0}[\ell + 1]$ are drawn from $\eta_0\langle\zeta_0[\ell]\rangle$ defined in (A.31) and $v_0^i[\ell + 1] \leftarrow (\xi_{0|0}^i[\ell + 1], 0)$ for all $i \in \llbracket 1, N \rrbracket$; then, recursively for $m \in \llbracket 0, t \rrbracket$, assuming access to $\mathbf{v}_m[\ell + 1]$,

- (1) we generate an updated particle cloud $\xi_{m+1}[\ell + 1] \sim \mathbf{M}_m\langle\zeta_{m+1}[\ell]\rangle(\xi_{m|m}[\ell + 1], \cdot)$,
- (2) we pick at random, for each $i \in \llbracket 1, N \rrbracket$, an ancestor path with associated statistics $(\tilde{\xi}_{0:m}^{i,1}[\ell + 1], \tilde{\beta}_m^{i,1}[\ell + 1])$ among $\mathbf{v}_m[\ell + 1]$ by drawing

$$(\tilde{\xi}_{0:m}^{i,1}[\ell + 1], \tilde{\beta}_m^{i,1}[\ell + 1]) \sim \sum_{s=1}^N \frac{q_m(\xi_{m|m}^s[\ell + 1], \xi_{m+1}^i[\ell + 1])}{\sum_{s'=1}^N q_m(\xi_{m|m}^{s'}[\ell + 1], \xi_{m+1}^i[\ell + 1])} \delta_{v_m^s[\ell + 1]}, \quad i \in \llbracket 1, N \rrbracket,$$

- (3) we draw, with replacement, $M - 1$ ancestor particles and associated statistics $\{(\tilde{\xi}_m^{i,j}[\ell + 1], \tilde{\beta}_m^{i,j}[\ell + 1])\}_{j=2}^M$ at random from $\{(\xi_{m|m}^s[\ell + 1], \beta_m^s[\ell + 1])\}_{s=1}^N$ according to

$$\{(\tilde{\xi}_m^{i,j}[\ell + 1], \tilde{\beta}_m^{i,j}[\ell + 1])\}_{j=2}^M \sim \left(\sum_{s=1}^N \frac{q_m(\xi_{m|m}^s[\ell + 1], \xi_{m+1}^i[\ell + 1])}{\sum_{s'=1}^N q_m(\xi_{m|m}^{s'}[\ell + 1], \xi_{m+1}^i[\ell + 1])} \delta_{(\xi_{m|m}^s[\ell + 1], \beta_m^s[\ell + 1])} \right)^{\otimes(M-1)},$$

- (4) we set, for all $i \in \llbracket 1, N \rrbracket$, $\xi_{0:m+1|m+1}^i[\ell + 1] \leftarrow (\tilde{\xi}_{0:m}^{i,1}[\ell + 1], \xi_{m+1}^i[\ell + 1])$ and $v_{m+1}^i[\ell + 1] \leftarrow (\xi_{0:m+1|m+1}^i[\ell + 1], \beta_{m+1}^i[\ell + 1])$, where

$$\beta_{m+1}^i[\ell + 1] \leftarrow M^{-1} \sum_{j=1}^M \left(\tilde{\beta}_m^{i,j}[\ell + 1] + \tilde{h}_m(\tilde{\xi}_m^{i,j}[\ell + 1], \xi_{m+1}^i[\ell + 1]) \right).$$

This conditional PARIS procedure is summarised in Algorithm 1.

Once the set of trajectories and associated statistics $\mathbf{v}_t[\ell + 1]$ is formed by means of n recursive conditional PARIS updates, an updated path $\zeta_{0:t}[\ell + 1]$ is drawn from $\mu(\xi_{0:t}[\ell + 1])$. A full sweep of the PPG is summarised in Algorithm 2.

The following Markov kernels will play an instrumental role in the following. For a given path $\{z_m\}_{m \in \mathbb{N}}$, the conditional PARIS update in Algorithm 1 defines an inhomogeneous Markov chain on the spaces $\{(\mathbf{Y}_m, \mathbf{Y}_m)\}_{m \in \mathbb{N}}$ with kernels

$$\mathbf{Y}_m \times \mathbf{Y}_{m+1} \ni (\mathbf{y}_m, A) \mapsto \int \mathbf{M}_m\langle z_{m+1} \rangle(\mathbf{x}_{m|m}, d\mathbf{x}_{m+1}) \mathbf{S}_m(\mathbf{y}_m, \mathbf{x}_{m+1}, A), \quad m \in \mathbb{N},$$

where

$$\begin{aligned} \mathbf{S}_m : \mathbf{Y}_m \times \mathbf{X}_{m+1} \times \mathbf{Y}_{m+1} \ni (\mathbf{y}_m, \mathbf{x}_{m+1}, A) & \tag{A.39} \\ \mapsto \int \cdots \int \mathbb{1}_A \left(\left\{ \left((\tilde{x}_{0:m}^{i,1}, x_{m+1}^i), \frac{1}{M} \sum_{j=1}^M (\tilde{\beta}_m^{i,j} + \tilde{h}_m(\tilde{x}_m^{i,j}, x_{m+1}^i)) \right) \right\}_{i=1}^N \right) & \\ \times \prod_{i=1}^N \left(\sum_{\ell=1}^N \frac{q_m(x_{m|m}^\ell, x_{m+1}^i)}{\sum_{\ell'=1}^N q_m(x_{m|m}^{\ell'}, x_{m+1}^i)} \delta_{y_m^\ell} (d(\tilde{x}_{0:m}^{i,1}, \tilde{b}_m^{i,1})) \right) & \\ \times \left(\sum_{\ell=1}^N \frac{q_m(x_{m|m}^\ell, x_{m+1}^i)}{\sum_{\ell'=1}^N q_m(x_{m|m}^{\ell'}, x_{m+1}^i)} \delta_{(x_{m|m}^\ell, b_m^\ell)} \right)^{\otimes(M-1)} & (d(\tilde{x}_m^{i,2}, \tilde{b}_m^{i,2}, \dots, \tilde{x}_m^{i,M}, \tilde{b}_m^{i,M})). \end{aligned}$$

In addition, we introduce the joint law

$$\mathbb{S}_t : \mathbf{X}_{0:t} \times \mathcal{Y}_t \ni (\mathbf{x}_{0:t}, A) \mapsto \int \cdots \int \mathbb{1}_A(\mathbf{y}_t) \mathcal{S}_0(\mathbf{J}\mathbf{x}_0, \mathbf{x}_1, d\mathbf{y}_1) \prod_{m=1}^{t-1} \mathcal{S}_m(\mathbf{y}_m, \mathbf{x}_{m+1}, d\mathbf{y}_{m+1}), \quad (\text{A.40})$$

where we have defined $\mathbf{J} := \text{Id}_N \otimes (0, 1)^\top$.

The kernel \mathbb{S}_t can be viewed as a *superincumbent sampling kernel* describing the distribution of the output \mathbf{v}_t generated by a sequence of PARIS iterates when the many-body process $\{\boldsymbol{\xi}_m\}_{m=0}^t$ associated with the underlying SMC algorithm is given. This allows us to describe alternatively the PPG as follows: given $\zeta_{0:t}[\ell]$, draw $\boldsymbol{\xi}_{0:t}[\ell+1] \sim \mathbb{C}_t(\zeta_{0:t}[\ell], \cdot)$; then, draw $\mathbf{v}_t[\ell+1] \sim \mathbb{S}_t(\boldsymbol{\xi}_{0:t}[\ell+1], \cdot)$ and pick a trajectory $\zeta_{0:t}[\ell+1]$ from $\boldsymbol{\xi}_{0:t}[\ell+1]$ at random. The following proposition, which will be instrumental in the coming developments, establishes that the conditional distribution of $\zeta_{0:t}[\ell+1]$ given $\boldsymbol{\xi}_{0:t}[\ell+1]$ coincides, as expected, with the particle-induced backward dynamics \mathbb{B}_t .

Proposition 2. *For all $t \in \mathbb{N}^*$, $N \in \mathbb{N}^*$, $\mathbf{x}_{0:t} \in \mathbf{X}_{0:t}$, and $h \in \mathcal{F}(\mathcal{X}_{0:t})$,*

$$\int \mathbb{S}_t(\mathbf{x}_{0:t}, d\mathbf{y}_t) \mu(\mathbf{x}_{0:t|t}) h = \mathbb{B}_t h(\mathbf{x}_{0:t}).$$

Finally, we define the Markov kernel induced by the PPG as well as the extended probability distribution targeted by the same. For this purpose, we introduce the extended measurable space $(\mathbf{E}_t, \boldsymbol{\mathcal{E}}_t)$ with

$$\mathbf{E}_t := \mathbf{Y}_t \times \mathbf{X}_{0:t}, \quad \boldsymbol{\mathcal{E}}_t := \mathcal{Y}_t \otimes \mathcal{X}_{0:t}.$$

The PPG described in Algorithm 2 defines a Markov chain on $(\mathbf{E}_t, \boldsymbol{\mathcal{E}}_t)$ with Markov transition kernel

$$\mathbb{K}_t : \mathbf{E}_t \times \boldsymbol{\mathcal{E}}_t \ni (\mathbf{y}_t, z_{0:t}, A) \mapsto \iiint \mathbb{1}_A(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t}) \mathbb{C}_t(z_{0:t}, d\tilde{\mathbf{x}}_{0:t}) \mathbb{S}_t(\tilde{\mathbf{x}}_{0:t}, d\tilde{\mathbf{y}}_t) \mu(\tilde{\mathbf{x}}_{0:t|t})(d\tilde{z}_{0:t}). \quad (\text{A.41})$$

Note that the values of \mathbb{K}_t defined above do not depend on \mathbf{y}_t , but only on $(z_{0:t}, A)$. For any given initial distribution $\xi \in \mathcal{M}_1(\mathcal{X}_{0:t})$, let \mathbb{P}_ξ be the distribution of the canonical Markov chain induced by the kernel \mathbb{K}_t and the initial distribution ξ . In the special case where $\xi = \delta_{z_{0:t}}$ for some given path $z_{0:t} \in \mathbf{X}_{0:t}$, we use the short-hand notation $\mathbb{P}_{\delta_{z_{0:t}}} = \mathbb{P}_{z_{0:t}}$. In addition, denote by

$$K_t : \mathbf{X}_{0:t} \times \mathcal{X}_{0:t} \ni (z_{0:t}, A) \mapsto \iiint \mathbb{1}_A(\tilde{z}_{0:t}) \mathbb{C}_t(z_{0:t}, d\tilde{\mathbf{x}}_{0:t}) \mathbb{S}_t(\tilde{\mathbf{x}}_{0:t}, d\tilde{\mathbf{y}}_t) \mu(\tilde{\mathbf{x}}_{0:t|t})(d\tilde{z}_{0:t}) \quad (\text{A.42})$$

the path-marginalised version of \mathbb{K}_t . By Proposition 2 it holds that $K_t = \mathbb{C}_t \mathbb{B}_t$, which shows that K_t coincides with the Markov transition kernel of the backward-sampling-based particle Gibbs sampler discussed in Appendix A.3. It is also possible to specify the invariant distribution of \mathbb{K}_t .

Proposition 3. *For all $t \in \mathbb{N}^*$, it holds that*

$$\eta_{0:t} \mathbb{C}_t \mathbb{S}_t \mathbb{K}_t = \eta_{0:t} \mathbb{C}_t \mathbb{S}_t. \quad (\text{A.43})$$

Proof. Let $f \in \mathcal{M}(\mathbf{E}_t^{\otimes(k-k_0)})$.

$$\begin{aligned} & \int f(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t}) \eta_{0:t}(dz_{0:t}) \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d(\mathbf{y}_t, z'_{0:t})) \mathbb{K}_t(z'_{0:t}, \mathbf{y}_t, d(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t})) \\ &= \int f(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t}) \eta_{0:t}(dz_{0:t}) \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d(\mathbf{y}_t, z'_{0:t})) \mathbb{C}_t \mathbb{S}_t(z'_{0:t}, d(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t})) \\ &= \int f(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t}) \eta_{0:t}(dz_{0:t}) K_t(z_{0:t}, dz'_{0:t}) \mathbb{C}_t \mathbb{S}_t(z'_{0:t}, d(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t})) \\ &= \int f(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t}) \eta_{0:t}(dz'_{0:t}) \mathbb{C}_t \mathbb{S}_t(z'_{0:t}, d(\tilde{\mathbf{y}}_t, \tilde{z}_{0:t})). \end{aligned}$$

□

Finally, in order to prepare for the statement of our theoretical results on the PPG we need to introduce the following Feynman–Kac path model *with a frozen path*. More precisely, for a given path $z_{0:t} \in \mathcal{X}_{0:t}$, define, for every $m \in \llbracket 0, t-1 \rrbracket$, the unnormalised kernel

$$Q_m \langle z_{m+1} \rangle : \mathcal{X}_m \times \mathcal{X}_{m+1} \ni (x_m, A) \mapsto \left(1 - \frac{1}{N}\right) Q_m(x_m, A) + \frac{1}{N} g_m(x_m) \delta_{z_{m+1}}(A)$$

and the initial distribution $\eta_0 \langle z_0 \rangle : \mathcal{X}_0 \ni A \mapsto (1 - 1/N)\eta_0(A) + \delta_{z_0}(A)/N$. Given these quantities, define, for $m \in \llbracket 0, t \rrbracket$, $\gamma_m \langle z_{0:m} \rangle := \eta_0 \langle z_0 \rangle Q_0 \langle z_1 \rangle \cdots Q_{m-1} \langle z_m \rangle$ along with the normalised counterpart $\eta_m \langle z_{0:m} \rangle := \gamma_m \langle z_{0:m} \rangle / \gamma_m \langle z_{0:m} \rangle \mathbb{1}_{\mathcal{X}_{0:m}}$. Finally, we introduce, for $m \in \llbracket 0, t \rrbracket$, the kernels

$$B_m \langle z_{0:m-1} \rangle : \mathcal{X}_m \times \mathcal{X}_{0:m-1} \ni (x_m, A) \mapsto \int \cdots \int \mathbb{1}_A(x_{0:m-1}) \prod_{m=0}^{t-1} \overleftarrow{Q}_{m, \eta_m \langle z_{0:m} \rangle}(x_{m+1}, dx_m),$$

as well as the path model $\eta_{0:m} \langle z_{0:m} \rangle := B_m \langle z_{0:m-1} \rangle \otimes \eta_m \langle z_{0:m} \rangle$.

A.5 Proof of Theorem 1

We start by establishing bias, MSE and covariance bounds for a fixed iteration of the PPG estimator.

Theorem 4. *Assume A 3.1. Then for every $t \in \mathbb{N}$ there exist c_t^{bias} , c_t^{mse} , and c_t^{cov} in \mathbb{R}_+^* such that for every $M \in \mathbb{N}^*$, $\xi \in \mathbf{M}_1(\mathcal{X}_{0:t})$, $\ell \in \mathbb{N}^*$, $s \in \mathbb{N}^*$, and $N \in \mathbb{N}^*$ such that $N > N_t$,*

$$|\mathbb{E}_\xi [\mu(\beta_t[\ell])(\text{id})] - \eta_{0:t} h_t| \leq c_t^{bias} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right) N^{-1} \kappa_{N,t}^\ell, \quad (\text{A.44})$$

$$\mathbb{E}_\xi \left[(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t)^2 \right] \leq c_t^{mse} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2 N^{-1}, \quad (\text{A.45})$$

$$|\mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t) (\mu(\beta_t[\ell+s])(\text{id}) - \eta_{0:t} h_t)]| \leq c_t^{cov} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2 N^{-3/2} \kappa_{N,t}^s. \quad (\text{A.46})$$

The constants c_t^{bias} , c_t^{mse} , and c_t^{cov} are explicitly given in the proof. Since the focus of this paper is on the dependence on N and the index ℓ , we have made no attempt to optimise the dependence of these constants on t in our proofs; still, we believe that it is possible to prove, under the stated assumptions, that this dependence is linear. The proof of the bound in Theorem 4 is based on four key ingredients. The first is the following unbiasedness property of the PARIS under the many-body Feynman–Kac path model.

Theorem 5. *For every $t \in \mathbb{N}$, $N \in \mathbb{N}^*$, and $\ell \in \mathbb{N}^*$,*

$$\mathbb{E}_{\eta_{0:t}} [\mu(\beta_t[\ell])(\text{id})] = \int \eta_{0:t} \mathbb{C}_t \mathbb{S}_t(d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) = \int \eta_{0:t} \mathbb{S}_t(d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) = \eta_{0:t} h_t.$$

The proof of Theorem 5 is postponed to Appendix A.6.3. The second ingredient of the proof of Theorem 4 is the uniform geometric ergodicity of the particle Gibbs with backward sampling established in [Del Moral and Jasra, 2018].

Theorem 6. *Assume A 3.1. Then, for every $t \in \mathbb{N}$, $(\mu, \nu) \in \mathbf{M}_1(\mathcal{X}_{0:t})^2$, $\ell \in \mathbb{N}^*$, and $N \in \mathbb{N}^*$ such that $N > 1 + 5\rho_t^2/2$, $\|\mu K_t^\ell - \nu K_t^\ell\|_{\text{TV}} \leq \kappa_{N,t}^\ell$, where $\kappa_{N,t}$ is defined in (3.14).*

As a third ingredient, we require the following uniform exponential concentration inequality of the conditional PARIS with respect to the frozen-path Feynman–Kac model defined in the previous section.

Theorem 7. For every $t \in \mathbb{N}$ there exist $c_t > 0$ and $d_t > 0$ such that for every $M \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, $N \in \mathbb{N}^*$, and $\varepsilon > 0$,

$$\int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) \mathbb{1} \{ |\mu(\mathbf{b}_t)(\text{id}) - \eta_{0:t}\langle z_{0:t} \rangle h_t| \geq \varepsilon \} \leq c_t \exp \left(- \frac{d_t N \varepsilon^2}{2(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty)^2} \right).$$

Theorem 7, whose proof is postponed to Appendix A.6.5, implies, in turn, the following conditional variance bound.

Proposition 4. For every $t \in \mathbb{N}$, $M \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, and $N \in \mathbb{N}^*$,

$$\int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) |\mu(\mathbf{b}_t)(\text{id}) - \eta_{0:t}\langle z_{0:t} \rangle h_t|^2 \leq \frac{c_t}{d_t} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2 N^{-1}.$$

Using Proposition 4, we deduce, in turn, the following bias bound, whose proof is postponed to Appendix A.6.7.

Proposition 5. For every $t \in \mathbb{N}$ there exists $\bar{c}_t^{\text{bias}} > 0$ such that for every $M \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, and $N \in \mathbb{N}^*$,

$$\left| \int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) - \eta_{0:t}\langle z_{0:t} \rangle h_t \right| \leq \bar{c}_t^{\text{bias}} N^{-1} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right).$$

A fourth and last ingredient in the proof of Theorem 4 is the following bound on the discrepancy between additive expectations under the original and frozen-path Feynman–Kac models. This bound is established using novel results in [Gloaguen et al., 2022]. More precisely, since for every $m \in \mathbb{N}$, $(x, z) \in \mathbf{X}_m^2$, $N \in \mathbb{N}^*$, and $h \in \mathbf{F}(\mathcal{X}_{m+1})$, using A 3.1,

$$|Q_m\langle z \rangle h(x) - Q_m h(x)| \leq \frac{1}{N} \|g_m\|_\infty \|h\|_\infty \leq \frac{1}{N} \bar{\tau}_m \|h\|_\infty,$$

applying [Gloaguen et al., 2022, Theorem 4.3] yields the following.

Proposition 6. Assume A 3.1. Then there exists $c > 0$ such that for every $t \in \mathbb{N}$, $N \in \mathbb{N}$, and $z_{0:t} \in \mathbf{X}_{0:t}$,

$$|\eta_{0:t}\langle z_{0:t} \rangle h_t - \eta_{0:t} h_t| \leq c N^{-1} \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty.$$

Note that assuming, in addition, that $\sup_{t \in \mathbb{N}} \|\tilde{h}_t\|_\infty < \infty$ yields an $\mathcal{O}(n/N)$ bound in Proposition 6. Finally, by combining these ingredients we are now ready to present a proof of Theorem 4.

Proof of Theorem 4. Write, using the tower property,

$$\mathbb{E}_\xi [\mu(\beta_t[\ell])(\text{id})] = \mathbb{E}_\xi [\mathbb{E}_{\zeta_{0:t}[\ell]} [\mu(\beta_t[0])(\text{id})]] = \int \xi K_t^\ell \mathbb{C}_t \mathbb{S}_t(d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}).$$

Thus, by the unbiasedness property in Theorem 5,

$$\begin{aligned} |\mathbb{E}_\xi [\mu(\beta_t[\ell])(\text{id})] - \eta_{0:t} h_t| &= \left| \int \xi K_t^\ell \mathbb{C}_t \mathbb{S}_t(d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) - \int \eta_{0:t} \mathbb{C}_t \mathbb{S}_t(d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) \right| \\ &\leq \|\xi K_t^\ell - \eta_{0:t}\|_{\text{TV}} \text{osc} \left(\int \mathbb{C}_t \mathbb{S}_t(\cdot, d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) \right), \end{aligned}$$

where, by Theorem 6, $\|\xi K_t^\ell - \eta_{0:t}\|_{\text{TV}} \leq \kappa_{N,t}^\ell$. Moreover, to derive an upper bound on the oscillation, we consider the decomposition

$$\text{osc} \left(\int \mathbb{C}_t \mathbb{S}_t(\cdot, d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) \right) \leq 2 \left(\left\| \int \mathbb{C}_t \mathbb{S}_t(\cdot, d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) - \eta_{0:t}(\cdot) h_t \right\|_\infty + \|\eta_{0:t}(\cdot) h_t - \eta_{0:t} h_t\|_\infty \right),$$

where the two terms on the right-hand side can be bounded using Proposition 6 and Proposition 5, respectively. This completes the proof of (A.44). We now consider the proof of (A.45). Writing

$$\mathbb{E}_\xi \left[(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t)^2 \right] = \int \xi K_t^\ell(dz_{0:t}) \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) (\mu(\mathbf{b}_t)(\text{id}) - \eta_{0:t} h_t)^2,$$

we may establish (A.45) using Proposition 4 and Proposition 6. We finally consider (A.46). Using the Markov property we obtain

$$\begin{aligned} \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t) (\mu(\beta_t[\ell+s])(\text{id}) - \eta_{0:t} h_t)] \\ = \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t) (\mathbb{E}_{\mathbb{C}_{0:t}[\ell]}[\mu(\beta_t[s])(\text{id})] - \eta_{0:t} h_t)], \end{aligned}$$

from which (A.46) follows by (A.44) and (A.45). \square

We are finally equipped to prove Theorem 1.

Proof of Theorem 1. We first consider the bias, which can be bounded according to

$$\begin{aligned} |\mathbb{E}_\xi[\Pi_{(k_0,k),N}(f)] - \eta_{0:t} h_t| &\leq (k - k_0)^{-1} \sum_{\ell=k_0+1}^k |\mathbb{E}_\xi \mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t| \\ &\leq (k - k_0)^{-1} N^{-1} \mathbf{c}_t^{\text{bias}} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right) \sum_{\ell=k_0+1}^k \kappa_{N,t}^\ell, \end{aligned}$$

from which the bound (3.15) follows immediately.

We turn to the MSE. Using the decomposition

$$\begin{aligned} \mathbb{E}_\xi [(\Pi_{(k_0,k),N}(f) - \eta_{0:t} h_t)^2] &\leq (k - k_0)^{-2} \left\{ \sum_{\ell=k_0+1}^k \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t)^2] \right. \\ &\quad \left. + 2 \sum_{\ell=k_0+1}^k \sum_{j=\ell+1}^k \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t) (\mu(\beta_t[j])(\text{id}) - \eta_{0:t} h_t)] \right\}, \end{aligned}$$

the MSE bound in Theorem 4 implies that

$$\sum_{\ell=k_0+1}^k \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t)^2] \leq \mathbf{c}_t^{\text{mse}} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2 N^{-1} (k - k_0).$$

Moreover, using the covariance bound in Theorem 4, we deduce that

$$\sum_{\ell=k_0+1}^k \sum_{j=\ell+1}^k \mathbb{E}_\xi [(\mu(\beta_t[\ell])(\text{id}) - \eta_{0:t} h_t) (\mu(\beta_t[j])(\text{id}) - \eta_{0:t} h_t)] \leq \mathbf{c}_t^{\text{cov}} \left(\sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty \right)^2 N^{-3/2} \left(\sum_{\ell=k_0+1}^k \sum_{j=\ell+1}^k \kappa_{N,t}^{(j-\ell)} \right).$$

Thus, the proof is concluded by noting that $\sum_{\ell=k_0+1}^k \sum_{j=\ell+1}^k \kappa_{N,t}^{(j-\ell)} \leq (k - k_0)/(1 - \kappa_{N,t})$. \square

A.6 Proofs of intermediate results

A.6.1 Proof of Proposition 1

Using the identity

$$\eta_0 Q_0 \cdots Q_{t-1} \mathbb{1}_{\mathbf{X}_t} = \prod_{m=0}^{t-1} \eta_m Q_m \mathbb{1}_{\mathbf{X}_{m+1}}$$

and the fact that each kernel Q_m has a transition density, write, for $h \in \mathbf{F}(\mathcal{X}_{0:t})$,

$$\begin{aligned} \eta_{0:t} h &= \int \cdots \int h(x_{0:t}) \eta_0(dx_0) \prod_{m=0}^{t-1} \left(\frac{\eta_m[q_m(\cdot, x_{m+1})] \lambda_{m+1}(dx_{m+1})}{\eta_m Q_m \mathbb{1}_{\mathbf{X}_{m+1}}} \right) \left(\frac{q_m(x_m, x_{m+1})}{\eta_m[q_m(\cdot, x_{m+1})]} \right) \\ &= \int \cdots \int h(x_{0:t}) \eta_t(dx_t) \prod_{m=0}^{t-1} \frac{\eta_m(dx_m) q_m(x_m, x_{m+1})}{\eta_m[q_m(\cdot, x_{m+1})]} \\ &= \left(\overleftarrow{Q}_{0, \eta_0} \otimes \cdots \otimes \overleftarrow{Q}_{n-1, \eta_{t-1}} \otimes \eta_t \right) h, \end{aligned} \quad (\text{A.47})$$

which was to be established.

A.6.2 Proof of Theorem 3

Lemma 1. For all $t \in \mathbb{N}$, $\mathbf{x}_t \in \mathbf{X}_t$, and $h \in \mathbf{F}(\mathcal{X}_{t+1} \otimes \mathcal{X}_{t+1})$,

$$\iint h(\mathbf{x}_{t+1}, z_{t+1}) \mathbf{Q}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}) \mu(\mathbf{x}_{t+1})(dz_{t+1}) = \iint h(\mathbf{x}_{t+1}, z_{t+1}) \mu(\mathbf{x}_t) Q_t(dz_{t+1}) \mathbf{M}_t \langle z_{t+1} \rangle(\mathbf{x}_t, d\mathbf{x}_{t+1}). \quad (\text{A.48})$$

In addition, for all $h \in \mathbf{F}(\mathcal{X}_0 \otimes \mathcal{X}_0)$,

$$\iint h(\mathbf{x}_0, z_0) \boldsymbol{\eta}_0(d\mathbf{x}_0) \mu(\mathbf{x}_0)(dz_0) = \iint h(\mathbf{x}_0, z_0) \boldsymbol{\eta}_0 \langle z_0 \rangle(d\mathbf{x}_0) \eta_0(dz_0). \quad (\text{A.49})$$

Proof. Since $\mu(\mathbf{x}_t) Q_t(dz_{t+1}) = \mathbf{g}_t(\mathbf{x}_t) \Phi_t(\mu(\mathbf{x}_t))(dz_{t+1})$, we may rewrite the right-hand side of (A.48) according to

$$\begin{aligned} &\iint h(\mathbf{x}_{t+1}, z_{t+1}) \mu(\mathbf{x}_t) Q_t(dz_{t+1}) \mathbf{M}_t \langle z_{t+1} \rangle(\mathbf{x}_t, d\mathbf{x}_{t+1}) \\ &= \mathbf{g}_t(\mathbf{x}_t) \frac{1}{N} \sum_{i=0}^{N-1} \iint h(\mathbf{x}_{t+1}, z_{t+1}) \Phi_t(\mu(\mathbf{x}_t))(dz_{t+1}) \\ &\quad \times \left(\Phi_t(\mu(\mathbf{x}_t))^{\otimes i} \otimes \delta_{z_{t+1}} \otimes \Phi_t(\mu(\mathbf{x}_t))^{\otimes (N-i-1)} \right) (d\mathbf{x}_{t+1}) \\ &= \mathbf{g}_t(\mathbf{x}_t) \frac{1}{N} \sum_{i=1}^N \int \cdots \int h((x_{t+1}^1, \dots, x_{t+1}^{i-1}, z_{t+1}, x_{t+1}^{i+1}, \dots, x_{t+1}^N), z_{t+1}) \\ &\quad \times \Phi_t(\mu(\mathbf{x}_t))(dz_{t+1}) \prod_{\ell \neq i} \Phi_t(\mu(\mathbf{x}_t))(dx_{t+1}^\ell) \\ &= \mathbf{g}_t(\mathbf{x}_t) \frac{1}{N} \sum_{i=1}^N \int h(\mathbf{x}_{t+1}, x_{t+1}^i) \mathbf{M}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}). \end{aligned}$$

On the other hand, note that the left-hand side of (A.48) can be expressed as

$$\iint h(\mathbf{x}_{t+1}, z_{t+1}) \mathbf{Q}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}) \mu(\mathbf{x}_{t+1})(dz_{t+1}) = \mathbf{g}_t(\mathbf{x}_t) \frac{1}{N} \sum_{i=1}^N \int h(\mathbf{x}_{t+1}, x_{t+1}^i) \mathbf{M}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}),$$

which establishes the identity. The identity (A.49) is established along similar lines. \square

We establish Theorem 3 by induction; thus, assume that the claim holds true for n and show that for all $h \in \mathbf{F}(\mathcal{X}_{0:t+1} \otimes \mathcal{X}_{0:t+1})$,

$$\begin{aligned} & \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \gamma_{0:t+1}(d\mathbf{x}_{0:t+1}) \mathbb{B}_{t+1}(\mathbf{x}_{0:t+1}, dz_{0:t+1}) \\ &= \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \gamma_{0:t+1}(dz_{0:t+1}) \mathbb{C}_{t+1}(z_{0:t+1}, d\mathbf{x}_{0:t+1}). \end{aligned} \quad (\text{A.50})$$

To prove this, we proceed, using definition (C.85), the left-hand side of (A.50) according to

$$\begin{aligned} & \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \gamma_{0:t+1}(d\mathbf{x}_{0:t+1}) \mathbb{B}_{t+1}(\mathbf{x}_{0:t+1}, dz_{0:t+1}) \\ &= \iint \gamma_{0:t}(d\mathbf{x}_{0:t}) \mathbb{B}_t(\mathbf{x}_{0:t}, dz_{0:t}) \\ & \quad \times \iint \bar{h}(\mathbf{x}_{0:t+1}, z_{0:t+1}) \mathbf{Q}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}) \mu(\mathbf{x}_{t+1})(dz_{t+1}), \end{aligned} \quad (\text{A.51})$$

where we have defined the function

$$\bar{h}(\mathbf{x}_{0:t+1}, z_{0:t+1}) := \frac{q_t(z_t, z_{t+1}) h(\mathbf{x}_{0:t+1}, z_{0:t+1})}{\mu(\mathbf{x}_t)[q_t(\cdot, z_{t+1})]}.$$

Now, applying Lemma 1 to the inner integral and using that

$$\mu(\mathbf{x}_t) \mathbf{Q}_t(dz_{t+1}) = \mu(\mathbf{x}_t)[q_t(\cdot, z_{t+1})] \lambda_{t+1}(dz_{t+1})$$

yields, for every $\mathbf{x}_{0:t}$ and $z_{0:t}$,

$$\begin{aligned} & \iint \bar{h}(\mathbf{x}_{0:t+1}, z_{0:t+1}) \mathbf{Q}_t(\mathbf{x}_t, d\mathbf{x}_{t+1}) \mu(\mathbf{x}_{t+1})(dz_{t+1}) \\ &= \iint \bar{h}(\mathbf{x}_{0:t+1}, z_{0:t+1}) \mu(\mathbf{x}_t) \mathbf{Q}_t(dz_{t+1}) \mathbf{M}_t\langle z_{t+1} \rangle(\mathbf{x}_t, d\mathbf{x}_{t+1}) \\ &= \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \mathbf{Q}_t(z_t, dz_{t+1}) \mathbf{M}_t\langle z_{t+1} \rangle(\mathbf{x}_t, d\mathbf{x}_{t+1}). \end{aligned}$$

Inserting the previous identity into (A.51) and using the induction hypothesis provides

$$\begin{aligned} & \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \gamma_{0:t+1}(d\mathbf{x}_{0:t+1}) \mathbb{B}_{t+1}(\mathbf{x}_{0:t+1}, dz_{0:t+1}) \\ &= \iint \gamma_{0:t}(dz_{0:t}) \mathbb{C}_t(z_{0:t}, d\mathbf{x}_{0:t}) \\ & \quad \times \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \mathbf{Q}_t(z_t, dz_{t+1}) \mathbf{M}_t\langle z_{t+1} \rangle(\mathbf{x}_t, d\mathbf{x}_{t+1}) \\ &= \iint h(\mathbf{x}_{0:t+1}, z_{0:t+1}) \gamma_{0:t+1}(dz_{0:t+1}) \mathbb{C}_{t+1}(z_{0:t+1}, d\mathbf{x}_{0:t+1}), \end{aligned}$$

which establishes (A.50).

A.6.3 Proof of Theorem 5

First, define, for $m \in \mathbb{N}$,

$$\mathbf{P}_m : \mathbf{Y}_m \times \mathcal{Y}_{m+1} \ni (\mathbf{y}_m, A) \mapsto \int \mathbf{M}_m(\mathbf{x}_{m|m}, d\mathbf{x}_{m+1}) \mathbf{S}_m(\mathbf{y}_m, \mathbf{x}_{m+1}, A). \quad (\text{A.52})$$

For any given initial distribution $\psi_0 \in \mathbf{M}_1(\mathcal{Y}_0)$, let $\mathbb{P}_{\psi_0}^{\mathbf{P}}$ be the distribution of the canonical Markov chain induced by the Markov kernels $\{\mathbf{P}_m\}_{m \in \mathbb{N}}$ and the initial distribution ψ_0 . By abuse of notation we write, for $\eta_0 \in \mathbf{M}_1(\mathcal{X}_0)$, $\mathbb{P}_{\eta_0}^{\mathbf{P}}$ instead of $\mathbb{P}_{\psi_0[\eta_0]}^{\mathbf{P}}$, where we have defined the extension $\psi_0[\eta_0](A) = \int \mathbb{1}_A(\mathbf{J}\mathbf{x}_0) \eta_0(d\mathbf{x}_0)$, $A \in \mathcal{Y}_0$. We preface the proof of Theorem 5 by some technical lemmas and a proposition.

Lemma 2. *For all $t \in \mathbb{N}$ and $(f_{t+1}, \tilde{f}_{t+1}) \in \mathbf{F}(\mathcal{X}_{t+1})^2$,*

$$\gamma_{t+1}(f_{t+1}B_{t+1}h_{t+1} + \tilde{f}_{t+1}) = \gamma_t\{Q_t f_{t+1}B_t h_t + Q_t(\tilde{h}_t f_{t+1} + \tilde{f}_{t+1})\}.$$

Proof. Pick arbitrarily $\varphi \in \mathbf{F}(\mathcal{X}_{t:t+1})$ and write, using definition (A.27) and the fact that Q_t has a transition density,

$$\begin{aligned} & \iint \varphi(x_{t:t+1}) \gamma_t(dx_t) Q_t(x_t, dx_{t+1}) \\ &= \iint \varphi(x_{t:t+1}) \gamma_t[q_t(\cdot, x_{t+1})] \lambda_{t+1}(dx_{t+1}) \frac{\gamma_t(dx_t) q_t(x_t, x_{t+1})}{\gamma_t[q_t(\cdot, x_{t+1})]} \\ &= \iint \varphi(x_{t:t+1}) \gamma_{t+1}(dx_{t+1}) \overleftarrow{Q}_{n, \eta_t}(x_{t+1}, dx_t). \end{aligned} \tag{A.53}$$

Now, by (A.34) it holds that

$$B_{t+1}h_{t+1}(x_{t+1}) = \int \overleftarrow{Q}_{n, \eta_t}(x_{t+1}, dx_t) \left(\tilde{h}_t(x_{t:t+1}) + \int h_t(x_{0:t}) B_t(x_t, dx_{0:t-1}) \right);$$

therefore, by applying (A.53) with

$$\varphi(x_{t:t+1}) := f_{t+1}(x_{t+1}) \left(\tilde{h}_t(x_{t:t+1}) + \int h_t(x_{0:t}) B_t(x_t, dx_{0:t-1}) \right)$$

we obtain that

$$\begin{aligned} \gamma_{t+1}(f_{t+1}B_{t+1}h_{t+1}) &= \iint \varphi(x_{t:t+1}) \gamma_{t+1}(dx_{t+1}) \overleftarrow{Q}_{n, \eta_t}(x_{t+1}, dx_t) \\ &= \iint \varphi(x_{t:t+1}) \gamma_t(dx_t) Q_t(x_t, dx_{t+1}) \\ &= \gamma_t(Q_t f_{t+1} B_t h_t + Q_t \tilde{h}_t f_{t+1}). \end{aligned}$$

Now the proof is concluded by noting that since $\gamma_{t+1} = \gamma_t Q_t$, $\gamma_{t+1} \tilde{f}_{t+1} = \gamma_t Q_t \tilde{f}_{t+1}$. \square

Lemma 3. *For every $t \in \mathbb{N}^*$, $h_t \in \mathbf{F}(\mathcal{Y}_t)$, and $\eta_0 \in \mathbf{M}_1(\mathcal{X}_0)$ it holds that*

$$\mathbb{E}_{\eta_0}^{\mathbf{P}}[h_t(\mathbf{v}_t) \mid \boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t}] = \mathbb{S}_t h_t(\boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t}), \quad \mathbb{P}_{\eta_0}^{\mathbf{P}}\text{-a.s.}$$

Proof. Pick arbitrarily $v_t \in \mathbf{F}(\mathcal{X}_{0:t})$. We show that

$$\mathbb{E}_{\eta_0}^{\mathbf{P}}[v_t(\boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t}) h_t(\mathbf{v}_t)] = \mathbb{E}_{\eta_0}^{\mathbf{P}}[v_t(\boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t}) \mathbb{S}_t h_t(\boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t})], \tag{A.54}$$

from which the claim follows. Using the definition (A.52), the left-hand side of the previous identity

may be rewritten as

$$\begin{aligned}
& \int \cdots \int \psi_0[\boldsymbol{\eta}_0](d\mathbf{y}_0) \prod_{m=0}^{t-1} \mathbf{P}_m(\mathbf{y}_m, d\mathbf{y}_{m+1}) h_t(\mathbf{y}_t) v_t(\mathbf{x}_{0|0}, \dots, \mathbf{x}_{t|t}) \\
&= \int \cdots \int \boldsymbol{\eta}_0(d\mathbf{x}_{0|0}) \prod_{m=0}^{t-1} \mathbf{M}_m(\mathbf{x}_{m|m}, d\mathbf{x}_{m+1}) \mathbf{S}_0(\mathbf{J}\mathbf{x}_{0|0}, \mathbf{x}_1, d\mathbf{y}_1) \\
&\quad \times \prod_{m=0}^{t-1} \mathbf{S}_m(\mathbf{y}_m, \mathbf{x}_{m+1}, d\mathbf{y}_{m+1}) h_t(\mathbf{y}_t) v_t(\mathbf{x}_{0|0}, \dots, \mathbf{x}_{t|t}) \\
&= \int \cdots \int \boldsymbol{\eta}_0(d\mathbf{x}_0) \prod_{m=0}^{t-1} \mathbf{M}_m(\mathbf{x}_m, d\mathbf{x}_{m+1}) \mathbf{S}_0(\mathbf{J}\mathbf{x}_0, \mathbf{x}_1, d\mathbf{y}_1) \\
&\quad \times \prod_{m=0}^{t-1} \mathbf{S}_m(\mathbf{y}_m, \mathbf{x}_{m+1}, d\mathbf{y}_{m+1}) h_t(\mathbf{y}_t) v_t(\mathbf{x}_0, \dots, \mathbf{x}_t).
\end{aligned}$$

Thus, we may conclude the proof by using the definition (A.40) of \mathbb{S}_t together with Fubini's theorem. \square

Lemma 4. For every $t \in \mathbb{N}^*$ and $h_t \in \mathbf{F}(\mathcal{Y}_t)$,

$$\mathbb{E}_{\boldsymbol{\eta}_0} \left[\left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) h_t(\mathbf{v}_t) \right] = \int \gamma_{0:t} \mathbb{S}_t(d\mathbf{y}_t) h_t(\mathbf{y}_t).$$

Proof. The claim of the lemma is a direct implication of Lemma 3; indeed, by applying the tower property and the latter we obtain

$$\begin{aligned}
& \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathbf{P}} \left[\left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) h_t(\mathbf{v}_t) \right] \\
&= \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathbf{P}} \left[\left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \mathbb{S}_t h_t(\boldsymbol{\xi}_{0|0}, \dots, \boldsymbol{\xi}_{t|t}) \right] \\
&= \int \cdots \int \boldsymbol{\eta}_0(d\mathbf{x}_0) \prod_{m=0}^{t-1} \mathbf{g}_m(\mathbf{x}_m) \mathbf{M}_m(\mathbf{x}_m, d\mathbf{x}_{m+1}) \mathbb{S}_t h_t(\mathbf{x}_{0:t}) \\
&= \int \gamma_{0:t} \mathbb{S}_t(d\mathbf{y}_t) h_t(\mathbf{y}_t).
\end{aligned}$$

\square

Proposition 7. For all $t \in \mathbb{N}^*$, $(N, M) \in (\mathbb{N}^*)^2$, and $(f_t, \tilde{f}_t) \in \mathbf{F}(\mathcal{X}_t)^2$,

$$\int \gamma_{0:t} \mathbb{S}_t(d\mathbf{y}_t) \left(\frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{t|t}^i) + \tilde{f}_t(x_{t|t}^i)\} \right) = \gamma_t(f_t B_t h_t + \tilde{f}_t).$$

Proof. Applying Lemma 4 yields

$$\int \gamma_{0:t} \mathbb{S}_t(d\mathbf{y}_t) \left(\frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{t|t}^i) + \tilde{f}_t(x_{t|t}^i)\} \right) = \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathbf{P}} \left[\left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{i=1}^N \{\beta_t^i f_t(\boldsymbol{\xi}_{t|t}^i) + \tilde{f}_t(\boldsymbol{\xi}_{t|t}^i)\} \right]. \tag{A.55}$$

In the following we will use repeatedly the following filtrations. Let $\tilde{\mathcal{F}}_t := \sigma(\{\mathbf{v}_m\}_{m=0}^t)$ be the σ -field generated by the output of the PARIS (Algorithm 1) during the first t iterations. In addition, let $\mathcal{F}_t := \tilde{\mathcal{F}}_{t-1} \vee \sigma(\boldsymbol{\xi}_{t|t})$.

We proceed by induction. Thus, assume that the statement of the proposition holds true for a given $t \in \mathbb{N}^*$ and consider, for arbitrarily chosen $(f_{t+1}, \tilde{f}_{t+1}) \in \mathbf{F}(\mathcal{X}_{t+1})^2$,

$$\begin{aligned} \mathbb{E}_{\eta_0}^{\mathbf{P}} \left[\left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{i=1}^N \{ \beta_{t+1}^i f_{t+1}(\xi_{t+1|t+1}^i) + \tilde{f}_{t+1}(\xi_{t+1|t+1}^i) \} \mid \tilde{\mathcal{F}}_t \right] \\ = \left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \mathbb{E}_{\eta_0}^{\mathbf{P}} [\beta_{t+1}^1 f_{t+1}(\xi_{t+1|t+1}^1) + \tilde{f}_{t+1}(\xi_{t+1|t+1}^1) \mid \tilde{\mathcal{F}}_t], \end{aligned}$$

where we used that the variables $\{ \beta_{t+1}^i f_{t+1}(\xi_{t+1|t+1}^i) + \tilde{f}_{t+1}(\xi_{t+1|t+1}^i) \}_{i=1}^N$ are conditionally i.i.d. given $\tilde{\mathcal{F}}_t$. Note that, by symmetry,

$$\begin{aligned} \mathbb{E}_{\eta_0}^{\mathbf{P}} [\beta_{t+1}^1 \mid \mathcal{F}_{t+1}] &= \int \mathbf{S}_t(\mathbf{v}_t, \boldsymbol{\xi}_{t+1|t+1}, d\mathbf{y}_{t+1}) b_{t+1}^1 \\ &= \int \cdots \int \left(\prod_{j=1}^M \sum_{\ell=1}^N \frac{q_t(\xi_{t|t}^\ell, \xi_{t+1|t+1}^1)}{\sum_{\ell'=1}^N q_t(\xi_{t|t}^{\ell'}, \xi_{t+1|t+1}^1)} \delta_{(\xi_{t|t}^\ell, \beta_t^\ell)}(d\tilde{x}_t^{1,j}, d\tilde{b}_t^{1,j}) \right) \\ &\quad \times \frac{1}{M} \sum_{j=1}^M \left(\tilde{b}_t^{1,j} + \tilde{h}_t(\tilde{x}_t^{1,j}, \xi_{t+1|t+1}^1) \right) \\ &= \sum_{\ell=1}^N \frac{q_t(\xi_{t|t}^\ell, \xi_{t+1|t+1}^1)}{\sum_{\ell'=1}^N q_t(\xi_{t|t}^{\ell'}, \xi_{t+1|t+1}^1)} \left(\beta_t^\ell + \tilde{h}_t(\xi_{t|t}^\ell, \xi_{t+1|t+1}^1) \right). \end{aligned} \tag{A.56}$$

Thus, using the tower property,

$$\begin{aligned} \mathbb{E}_{\eta_0}^{\mathbf{P}} \left[\beta_{t+1}^1 f_{t+1}(\xi_{t+1|t+1}^1) \mid \tilde{\mathcal{F}}_t \right] \\ = \int \Phi_t(\mu(\boldsymbol{\xi}_{t|t})) (dx_{t+1}) f_{t+1}(x_{t+1}) \sum_{\ell=1}^N \frac{q_t(\xi_{t|t}^\ell, x_{t+1})}{\sum_{\ell'=1}^N q_t(\xi_{t|t}^{\ell'}, x_{t+1})} \left(\beta_t^\ell + \tilde{h}_t(\xi_{t|t}^\ell, x_{t+1}) \right), \end{aligned}$$

and consequently, using definition (A.25),

$$\begin{aligned} \left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \mathbb{E}_{\eta_0}^{\mathbf{P}} \left[\beta_{t+1}^1 f_{t+1}(\xi_{t+1|t+1}^1) \mid \tilde{\mathcal{F}}_t \right] \\ = \left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \int \frac{1}{N} \sum_{i=1}^N q_t(\xi_{t|t}^i, x_{t+1}) \\ \times f_{t+1}(x_{t+1}) \sum_{\ell=1}^N \frac{q_t(\xi_{t|t}^\ell, x_{t+1})}{\sum_{\ell'=1}^N q_t(\xi_{t|t}^{\ell'}, x_{t+1})} \left(\beta_t^\ell + \tilde{h}_t(\xi_{t|t}^\ell, x_{t+1}) \right) \lambda_{t+1}(dx_{t+1}) \\ = \left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{\ell=1}^N \left(\beta_t^\ell Q_t f_{t+1}(\xi_{t|t}^\ell) + Q_t(\tilde{h}_t f_{t+1})(\xi_{t|t}^\ell) \right). \end{aligned}$$

Thus, applying the induction hypothesis,

$$\begin{aligned}
& \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathcal{P}} \left[\left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{i=1}^N \beta_{t+1}^i f_{t+1}(\xi_{t+1|t+1}^i) \right] \\
&= \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathcal{P}} \left[\left(\prod_{m=0}^{t-1} \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{\ell=1}^N \left(\beta_t^\ell Q_t f_{t+1}(\xi_{t|t}^\ell) + Q_t(\tilde{h}_t f_{t+1})(\xi_{t|t}^\ell) \right) \right] \\
&= \gamma_t \left(Q_t f_{t+1} B_t h_t + Q_t(\tilde{h}_t f_{t+1}) \right). \tag{A.57}
\end{aligned}$$

In the same manner, it can be shown that

$$\mathbb{E}_{\boldsymbol{\eta}_0}^{\mathcal{P}} \left[\left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{i=1}^N \tilde{f}_{t+1}(\xi_{t+1|t+1}^i) \right] = \gamma_t Q_t \tilde{f}_{t+1}. \tag{A.58}$$

Now, by (A.57–A.58) and Lemma 2,

$$\begin{aligned}
& \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathcal{P}} \left[\left(\prod_{m=0}^t \mathbf{g}_m(\boldsymbol{\xi}_{m|m}) \right) \frac{1}{N} \sum_{i=1}^N \{ \beta_{t+1}^i f_{t+1}(\xi_{t+1|t+1}^i) + \tilde{f}_{t+1}(\xi_{t+1|t+1}^i) \} \right] \\
&= \gamma_t \left(Q_t f_{t+1} B_t h_t + Q_t(\tilde{h}_t f_{t+1}) + Q_t \tilde{f}_{t+1} \right) \\
&= \gamma_{t+1} (f_{t+1} B_{t+1} h_{t+1} + \tilde{f}_{t+1}),
\end{aligned}$$

which shows that the claim of the proposition holds at time $n + 1$.

It remains to check the base case $n = 0$, which holds trivially true as $\beta_0 = \mathbf{0}$, $B_0 h_0 = 0$ by convention, and the initial particles $\boldsymbol{\xi}_{0|0}$ are drawn from $\boldsymbol{\eta}_0$. This completes the proof. \square

Proof of Theorem 5. The identity $\int \boldsymbol{\eta}_{0:t}(\mathbf{d}\mathbf{x}_{0:t}) \mathbb{S}_t(\mathbf{x}_{0:t}, \mathbf{d}\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) = \eta_{0:t} h_t$ follows immediately by letting $f_t \equiv 1$ and $\tilde{f}_t \equiv 0$ in Proposition 7 and using that $\gamma_{0:t}(\mathbf{X}_{0:t}) = \gamma_{0:t}(\mathbf{X}_{0:t})$. Moreover, applying Theorem 3 yields

$$\begin{aligned}
\int \eta_{0:t} \mathbb{C}_t \mathbb{S}_t(\mathbf{d}\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) &= \iint \eta_{0:t}(\mathbf{d}z_{0:t}) \mathbb{C}_t(z_{0:t}, \mathbf{d}\mathbf{x}_{0:t}) \int \mathbb{S}_t(\mathbf{x}_{0:t}, \mathbf{d}\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) \\
&= \iint \boldsymbol{\eta}_{0:t}(\mathbf{d}\mathbf{x}_{0:t}) \mathbb{B}_t(\mathbf{x}_{0:t}, \mathbf{d}z_{0:t}) \int \mathbb{S}_t(\mathbf{x}_{0:t}, \mathbf{d}\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}) \\
&= \int \boldsymbol{\eta}_{0:t} \mathbb{S}_t(\mathbf{d}\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id}).
\end{aligned}$$

Finally, the first identity holds true since K_t leaves $\eta_{0:t}$ invariant. \square

A.6.4 Proof of Proposition 2

First, note that, by definitions (A.39) and (A.40),

$$\begin{aligned}
H_t(\mathbf{x}_{0:t}) &:= \int \mathbb{S}_t(\mathbf{x}_{0:t}, \mathbf{d}\mathbf{y}_t) \mu(\mathbf{x}_{[0:n|n]}) h \\
&= \int \cdots \int \left(\frac{1}{N} \sum_{j_t=1}^N h(x_{0:t-1|t}^{j_t}, x_t^{j_t}) \right) \\
&\quad \times \prod_{m=0}^{t-1} \prod_{i_{m+1}=1}^N \int \sum_{j_m=1}^N \frac{q_m(x_m^{j_m}, x_{m+1}^{i_{m+1}})}{\sum_{j'_m=1}^N q_m(x_m^{j'_m}, x_{m+1}^{i_{m+1}})} \delta_{x_{0:m|m}^{j_m}}(\mathbf{d}x_{0:m|m+1}^{i_{m+1}}),
\end{aligned}$$

where $x_{0:-1|0}^i = \emptyset$ for all $i \in \llbracket 1, N \rrbracket$ by convention. We will show that for every $k \in \llbracket 0, t \rrbracket$, $H_{k,t} \equiv H_t$, where

$$H_{k,n}(\mathbf{x}_{0:t}) := \frac{1}{N} \sum_{j_t=1}^N \cdots \sum_{j_k=1}^N \prod_{\ell=k}^{t-1} \frac{q_\ell(x_\ell^{j_\ell}, x_{\ell+1}^{j_{\ell+1}})}{\sum_{j'_\ell=1}^N q_\ell(x_\ell^{j'_\ell}, x_{\ell+1}^{j'_{\ell+1}})} a_{k,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}, x_k^{j_k}, \dots, x_t^{j_t})$$

with

$$\begin{aligned} & a_{k,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}, x_k^{j_k}, \dots, x_t^{j_t}) \\ &= \int \prod_{m=0}^{k-1} \prod_{i_{m+1}=1}^N \sum_{j_m=1}^N \frac{q_m(x_m^{j_m}, x_{m+1}^{i_{m+1}})}{\sum_{j'_m=1}^N q_m(x_m^{j'_m}, x_{m+1}^{i_{m+1}})} \delta_{x_{0:m|m}^{j_m}} (dx_{0:m|m+1}^{i_{m+1}}) h(x_{0:k-1|k}^{j_k}, x_k^{j_k}, \dots, x_t^{j_t}). \end{aligned}$$

Since, by convention, $\prod_{\ell=n}^{t-1} \dots = 1$, $H_{n,n}(\mathbf{x}_{0:t}) = N^{-1} \sum_{j_t=1}^N a_{n,n}(\mathbf{x}_0, \dots, \mathbf{x}_{[n-1]}, x_t^{j_t})$, and we note that $H_t \equiv H_{n,n}$. We now show that $H_{k,n} \equiv H_{k-1,n}$ for every $k \in \llbracket 1, t \rrbracket$; for this purpose, note that

$$\begin{aligned} & a_{k,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}, x_k^{j_k}, \dots, x_t^{j_t}) \\ &= \int \prod_{m=0}^{k-2} \prod_{i_{m+1}=1}^N \sum_{j_m=1}^N \frac{q_m(x_m^{j_m}, x_{m+1}^{i_{m+1}})}{\sum_{j'_m=1}^N q_m(x_m^{j'_m}, x_{m+1}^{i_{m+1}})} \delta_{x_{0:m|m}^{j_m}} (dx_{0:m|m+1}^{i_{m+1}}) \\ &\quad \times \int \prod_{i_k=1}^N \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{i_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{i_k})} \delta_{x_{0:k-1|k-1}^{j_{k-1}}} (dx_{0:k-1|k}^{i_k}) h(x_{0:k-1|k}^{j_k}, x_k^{j_k}, \dots, x_t^{j_t}), \end{aligned}$$

and since $x_{0:k-1|k-1}^{j_{k-1}} = (x_{0:k-2|k-1}^{j_{k-1}}, x_{k-1}^{j_{k-1}})$, it holds that

$$\begin{aligned} & \int \prod_{i_k=1}^N \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{i_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{i_k})} \delta_{x_{0:k-1|k-1}^{j_{k-1}}} (dx_{0:k-1|k}^{i_k}) h(x_{0:k-1|k}^{j_k}, x_k^{j_k}, \dots, x_t^{j_t}) \\ &= \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{j_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{j_k})} h(x_{0:k-2|k-1}^{j_{k-1}}, x_{k-1}^{j_{k-1}}, x_k^{j_k}, \dots, x_t^{j_t}). \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} & a_{k,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}, x_k^{j_k}, \dots, x_t^{j_t}) \\ &= \int \prod_{m=0}^{k-2} \prod_{i_{m+1}=1}^N \sum_{j_m=1}^N \frac{q_m(x_m^{j_m}, x_{m+1}^{i_{m+1}})}{\sum_{j'_m=1}^N q_m(x_m^{j'_m}, x_{m+1}^{i_{m+1}})} \delta_{x_{0:m|m}^{j_m}} (dx_{0:m|m+1}^{i_{m+1}}) \\ &\quad \times \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{j_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{j_k})} h(x_{0:k-2|k-1}^{j_{k-1}}, x_{k-1}^{j_{k-1}}, x_k^{j_k}, \dots, x_t^{j_t}). \end{aligned}$$

Now, changing the order of summation with respect to j_{k-1} and integration on the right hand side of the previous display yields

$$\begin{aligned} & a_{k,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}, x_k^{j_k}, \dots, x_t^{j_t}) \\ &= \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{j_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{j_k})} a_{k-1,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-2}, x_{k-1}^{j_{k-1}}, \dots, x_t^{j_t}). \end{aligned}$$

Thus,

$$\begin{aligned}
& H_{k,n}(\mathbf{x}_{0:t}) \\
&= \frac{1}{N} \sum_{j_t=1}^N \cdots \sum_{j_k=1}^N \prod_{\ell=k}^{t-1} \frac{q_\ell(x_\ell^{j_\ell}, x_{\ell+1}^{j_{\ell+1}})}{\sum_{j'_\ell=1}^N q_\ell(x_\ell^{j'_\ell}, x_{\ell+1}^{j'_{\ell+1}})} \\
&\quad \times \sum_{j_{k-1}=1}^N \frac{q_{k-1}(x_{k-1}^{j_{k-1}}, x_k^{j_k})}{\sum_{j'_{k-1}=1}^N q_{k-1}(x_{k-1}^{j'_{k-1}}, x_k^{j_k})} a_{k-1,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-2}, x_{k-1}^{j_{k-1}}, \dots, x_t^{j_t}) \\
&= \frac{1}{N} \sum_{j_t=1}^N \cdots \sum_{j_{k-1}=1}^N \prod_{\ell=k-1}^{t-1} \frac{q_\ell(x_\ell^{j_\ell}, x_{\ell+1}^{j_{\ell+1}})}{\sum_{j'_\ell=1}^N q_\ell(x_\ell^{j'_\ell}, x_{\ell+1}^{j'_{\ell+1}})} a_{k-1,n}(\mathbf{x}_0, \dots, \mathbf{x}_{k-2}, x_{k-1}^{j_{k-1}}, \dots, x_t^{j_t}) \\
&= H_{k-1,n}(\mathbf{x}_{0:t}),
\end{aligned}$$

which establishes the recursion. Therefore, $H_t \equiv H_{0,n}$ and we may now conclude the proof by noting that $\mathbb{B}_t h \equiv H_{0,n}$.

A.6.5 Proof of Theorem 7

In order to establish Theorem 7 we will prove the following more general result, of which Theorem 7 is a direct consequence.

Proposition 8. *For every $t \in \mathbb{N}$ and $M \in \mathbb{N}^*$ there exist $c_t > 0$ and $d_t > 0$ such that for every $N \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, $(f_t, \tilde{f}_t) \in \mathbf{F}(\mathcal{X}_t)^2$, and $\varepsilon > 0$,*

$$\int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) \mathbb{1} \left\{ \left| \frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{t|t}^i) + \tilde{f}_t(x_{t|t}^i)\} - \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \right| \geq \varepsilon \right\} \leq c_t \exp \left(-\frac{d_t N \varepsilon^2}{2 \kappa_t^2} \right),$$

where

$$\kappa_t := \|f_t\|_\infty \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty + \|\tilde{f}_t\|_\infty. \tag{A.59}$$

To prove Proposition 8 we need the following technical lemma.

Lemma 5. *For every $t \in \mathbb{N}$, $(f_{t+1}, \tilde{f}_{t+1}) \in \mathbf{F}(\mathcal{X}_{t+1})^2$, $z_{0:t+1} \in \mathbf{X}_{0:t+1}$, and $N \in \mathbb{N}^*$,*

$$\begin{aligned}
& \gamma_{t+1} \langle z_{0:t+1} \rangle (f_{t+1} B_{t+1} \langle z_{0:t} \rangle h_{t+1} + \tilde{f}_{t+1}) \\
&= \left(1 - \frac{1}{N}\right) \gamma_t \langle z_{0:t} \rangle \{Q_t f_{t+1} B_t \langle z_{0:t-1} \rangle h_t + Q_t (\tilde{h}_t f_{t+1} + \tilde{f}_{t+1})\} \\
&\quad + \frac{1}{N} \gamma_t \langle z_{0:t} \rangle g_t \left(f_{t+1}(z_{t+1}) B_{t+1} \langle z_{0:t} \rangle h_{t+1}(z_{t+1}) + \tilde{f}_{t+1}(z_{t+1}) \right).
\end{aligned}$$

Proof. Since Lemma 2 holds also for the Feynman–Kac model with a frozen path, we obtain

$$\gamma_{t+1} \langle z_{0:t+1} \rangle (f_{t+1} B_{t+1} \langle z_{0:t} \rangle h_{t+1} + \tilde{f}_{t+1}) = \gamma_t \langle z_{0:t} \rangle \{Q_t \langle z_{t+1} \rangle f_{t+1} B_t \langle z_{0:t} \rangle h_t + Q_t \langle z_{t+1} \rangle (\tilde{h}_t f_{t+1} + \tilde{f}_{t+1})\}.$$

Thus, the proof is concluded by noting that for every $x_t \in \mathbf{X}_t$ and $h \in \mathbf{F}(\mathcal{X}_{t+1})$,

$$Q_t \langle z_{t+1} \rangle h(x_t) = \left(1 - \frac{1}{N}\right) Q_t h(x_t) + \frac{1}{N} g(x_t) h(x_t, z_{t+1}).$$

□

Finally, before proceeding to the proof of Proposition 8, we introduce the law of the PARIS evolving conditionally on a frozen path $z = \{z_m\}_{m \in \mathbb{N}}$. Define, for $m \in \mathbb{N}$ and $z_{m+1} \in \mathbf{X}_{m+1}$,

$$\mathbf{P}_m \langle z_{m+1} \rangle : \mathbf{Y}_m \times \mathbf{Y}_{m+1} \ni (\mathbf{y}_m, A) \mapsto \int \mathbf{M}_m \langle z_{m+1} \rangle (\mathbf{x}_{m|m}, d\mathbf{x}_{m+1}) \mathbf{S}_m (\mathbf{y}_m, \mathbf{x}_{m+1}, A).$$

For any given initial distribution $\psi_0 \in \mathbf{M}_1(\mathbf{Y}_0)$, let $\mathbb{P}_{\psi_0}^{\mathbf{P},z}$ be the distribution of the canonical Markov chain induced by the Markov kernels $\{\mathbf{P}_m \langle z_{m+1} \rangle\}_{m \in \mathbb{N}}$ and the initial distribution ψ_0 . By abuse of notation we write $\mathbb{P}_{\eta_0}^{\mathbf{P},z}$ instead of $\mathbb{P}_{\psi_0[\eta_0]}^{\mathbf{P},z}$, where the extension $\psi_0[\eta_0]$ is defined in Appendix A.6.3.

Proof of Proposition 8. We proceed by forward induction over t . Let the σ -fields $\tilde{\mathcal{F}}_t$ and \mathcal{F}_t be defined as in the proof of Theorem 5, but for the conditional PARIS dual process. Then, under the law $\mathbb{P}_{\eta_0}^{\mathbf{P},z}$, reusing (A.56),

$$\begin{aligned} & \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\beta_t^1 f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right] \\ &= \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\beta_t^1 \mid \mathcal{F}_t \right] f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right] \\ &= \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[f_t(\xi_t^1) \sum_{\ell=1}^N \frac{q_{t-1}(\xi_{t-1}^\ell, \xi_t^1)}{\sum_{\ell'=1}^N q_{t-1}(\xi_{t-1}^{\ell'}, \xi_t^1)} \left(\beta_{t-1}^\ell + \tilde{h}_{t-1}(\xi_{t-1}^\ell, \xi_t^1) \right) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right]. \end{aligned}$$

Using (A.30), we get

$$\begin{aligned} & \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\beta_t^1 f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right] \\ &= \left(1 - \frac{1}{N} \right) \frac{\sum_{\ell=1}^N \{ \beta_{t-1}^\ell Q_{t-1} f_t(\xi_{t-1}^\ell) + Q_{t-1} (\tilde{h}_{t-1} f_t + \tilde{f}_t)(\xi_{t-1}^\ell) \}}{\sum_{\ell'=1}^N g_{t-1}(\xi_{t-1}^{\ell'})} \\ & \quad + \frac{1}{N} \left(f_t(z_t) \sum_{\ell=1}^N \frac{q_{t-1}(\xi_{t-1}^\ell, z_t)}{\sum_{\ell'=1}^N q_{t-1}(\xi_{t-1}^{\ell'}, z_t)} \left(\beta_{t-1}^\ell + \tilde{h}_t(\xi_{t-1}^\ell, z_t) \right) + \tilde{f}_t(z_t) \right). \end{aligned} \quad (\text{A.60})$$

In order to apply the induction hypothesis to each term on the right-hand side of the previous identity, note that

$$B_t \langle z_{0:t-1} \rangle h_t(z_t) = \frac{\eta_{t-1} \langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t) \{ B_{t-1} \langle z_{0:t-2} \rangle h_{t-1}(\cdot) + \tilde{h}_{t-1}(\cdot, z_t) \}]}{\eta_{t-1} \langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t)]}.$$

Therefore, using Lemma 5 and noting that $\gamma_t \langle z_{0:t} \rangle \mathbb{1}_{X_t} / \gamma_{t-1} \langle z_{0:t} \rangle \mathbb{1}_{X_{t-1}} = \eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1}$ yields

$$\begin{aligned} \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) &= \frac{1}{N} \left(f_t(z_t) B_t \langle z_{0:t-1} \rangle h_t(z_t) + \tilde{f}_t(z_t) \right) \\ & \quad + \left(1 - \frac{1}{N} \right) \frac{\eta_{t-1} \langle z_{0:t-1} \rangle \{ Q_{t-1} f_t B_{t-1} \langle z_{0:t-2} \rangle h_t + Q_{t-1} (\tilde{h}_{t-1} f_t + \tilde{f}_t) \}}{\eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1}}. \end{aligned} \quad (\text{A.61})$$

By combining (A.60) with (A.61), we decompose the error according to

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \{ \beta_t^i f_t(\xi_{t|t}^i) + \tilde{f}_t(\xi_{t|t}^i) \} - \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \\ &= \frac{1}{N} \sum_{i=1}^N \{ \beta_t^i f_t(\xi_{t|t}^i) + \tilde{f}_t(\xi_{t|t}^i) \} - \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\beta_t^1 f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right] \\ & \quad + \mathbb{E}_{\eta_0}^{\mathbf{P},z} \left[\beta_t^1 f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right] - \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \\ &= \mathbf{I}_N^{(1)} + \left(1 - \frac{1}{N} \right) \mathbf{I}_N^{(2)} + \frac{1}{N} \mathbf{I}_N^{(3)}, \end{aligned} \quad (\text{A.62})$$

where

$$\begin{aligned}
\mathbb{I}_N^{(1)} &:= \frac{1}{N} \sum_{i=1}^N \{\beta_t^i f_t(\xi_t^i) + \tilde{f}_t(\xi_t^i)\} - \mathbb{E}_{\boldsymbol{\eta}_0}^{\mathbf{P},z} \left[\beta_t^1 f_t(\xi_t^1) + \tilde{f}_t(\xi_t^1) \mid \tilde{\mathcal{F}}_{t-1} \right], \\
\mathbb{I}_N^{(2)} &:= \frac{\sum_{\ell=1}^N \{\beta_{t-1}^\ell Q_{t-1} f_t(\xi_{t-1}^\ell) + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)(\xi_{t-1}^\ell)\}}{\sum_{\ell'=1}^N g_{t-1}(\xi_{t-1}^{\ell'})} \\
&\quad - \frac{\eta_{t-1} \langle z_{0:t-1} \rangle \{Q_{t-1} f_t B_t \langle z_{0:t-1} \rangle h_t + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)\}}{\eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1}}, \tag{A.63}
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{I}_N^{(3)} &:= f_t(z_t) \sum_{\ell=1}^N \frac{q_{t-1}(\xi_{t-1}^\ell, z_t)}{\sum_{\ell'=1}^N q_{t-1}(\xi_{t-1}^{\ell'}, z_t)} \left(\beta_{t-1}^\ell + \tilde{h}_{t-1}(\xi_{t-1}^\ell, z_t) \right) \\
&\quad - f_t(z_t) \frac{\eta_{t-1} \langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t) \{B_{t-1} \langle z_{0:t-2} \rangle h_{t-1}(\cdot) + \tilde{h}_{t-1}(\cdot, z_t)\}]}{\eta_{t-1} \langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t)]}. \tag{A.64}
\end{aligned}$$

The proof is now completed by treating the terms $\mathbb{I}_N^{(1)}$, $\mathbb{I}_N^{(2)}$, and $\mathbb{I}_N^{(3)}$ separately, using Hoeffding's inequality and its generalisation in [Douc et al., 2011, Lemma 4]. Choose $\varepsilon > 0$; then, by Hoeffding's inequality,

$$\mathbb{P}_{\boldsymbol{\eta}_0}^{\mathbf{P},z} \left(|\mathbb{I}_N^{(1)}| \geq \varepsilon \right) \leq 2 \exp \left(-\frac{1}{2} \frac{\varepsilon^2}{\kappa_t^2} N \right). \tag{A.65}$$

To treat $\mathbb{I}_N^{(2)}$, we apply the induction hypothesis to the numerator and denominator, each normalised by $1/N$, yielding, since $\|Q_{t-1} h\|_\infty \leq \bar{\tau}_{t-1} \|h\|_\infty$ for all $h \in \mathbb{F}(\mathcal{X}_{t-1} \otimes \mathcal{X}_t)$,

$$\begin{aligned}
\mathbb{P}_{\boldsymbol{\eta}_0}^{\mathbf{P},z} \left(\left| \frac{1}{N} \sum_{\ell=1}^N \{\beta_{t-1}^\ell Q_{t-1} f_t(\xi_{t-1}^\ell) + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)(\xi_{t-1}^\ell)\} \right. \right. \\
\left. \left. - \eta_{t-1} \langle z_{0:t-1} \rangle \{Q_{t-1} f_t B_t \langle z_{0:t-1} \rangle h_t + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)\} \right| \geq \varepsilon \right) \\
\leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \frac{\varepsilon^2}{\bar{\tau}_{t-1}^2 \kappa_t^2} N \right)
\end{aligned}$$

and

$$\mathbb{P}_{\boldsymbol{\eta}_0}^{\mathbf{P},z} \left(\left| \frac{1}{N} \sum_{\ell=1}^N g_{t-1}(\xi_{t-1}^\ell) - \eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1} \right| \geq \varepsilon \right) \leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \frac{\varepsilon^2}{\bar{\tau}_{t-1}^2} N \right).$$

Combining the previous two bounds with the generalised Hoeffding inequality in [Douc et al., 2011, Lemma 4] yields, using also the bounds

$$\frac{\sum_{\ell=1}^N \{\beta_{t-1}^\ell Q_{t-1} f_t(\xi_{t-1}^\ell) + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)(\xi_{t-1}^\ell)\}}{\sum_{\ell'=1}^N g_{t-1}(\xi_{t-1}^{\ell'})} \leq \kappa_t$$

and $\eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1} \geq \tau_{t-1}$, the inequality

$$\mathbb{P}_{\boldsymbol{\eta}_0}^{\mathbf{P},z} \left(|\mathbb{I}_N^{(2)}| \geq \varepsilon \right) \leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \frac{\tau_{t-1}^2 \varepsilon^2}{\bar{\tau}_{t-1}^2 \kappa_t^2} N \right). \tag{A.66}$$

The last term $I_N^{(3)}$ is treated along similar lines; indeed, by the induction hypothesis, since $\|q_{t-1}\|_\infty \leq \bar{\tau}_{t-1}\bar{\sigma}_{t-1}$,

$$\begin{aligned} \mathbb{P}_{\eta_0}^{P,z} \left(\left| \frac{1}{N} \sum_{\ell=1}^N q_{t-1}(\xi_{t-1}^\ell, z_t) \left(\beta_{t-1}^\ell + \tilde{h}_{t-1}(\xi_{t-1}^\ell, z_t) \right) \right. \right. \\ \left. \left. - \eta_{t-1}\langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t) \{B_{t-1}\langle z_{0:t-1} \rangle h_{t-1}(\cdot) + \tilde{h}_{t-1}(\cdot, z_t)\}] \right| \geq \varepsilon \right) \\ \leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \left(\frac{\varepsilon}{\bar{\tau}_{t-1}\bar{\sigma}_{t-1} \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty} \right)^2 N \right) \end{aligned}$$

and

$$\mathbb{P}_{\eta_0}^{P,z} \left(\left| \frac{1}{N} \sum_{\ell=1}^N q_{t-1}(\xi_{t-1}^\ell, z_t) - \eta_{t-1}\langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t)] \right| \geq \varepsilon \right) \leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \left(\frac{\varepsilon}{\bar{\tau}_{t-1}\bar{\sigma}_{t-1}} \right)^2 N \right).$$

Thus, since

$$\sum_{\ell=1}^N \frac{q_{t-1}(\xi_{t-1}^\ell, z_t)}{\sum_{\ell'=1}^N q_{t-1}(\xi_{t-1}^{\ell'}, z_t)} \left(\beta_{t-1}^\ell + \tilde{h}_{t-1}(\xi_{t-1}^\ell, z_t) \right) \leq \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty$$

and $\eta_{t-1}\langle z_{0:t-1} \rangle [q_{t-1}(\cdot, z_t)] \geq \mathcal{I}_{t-1}$, the generalised Hoeffding inequality provides

$$\mathbb{P}_{\eta_0}^{P,z} \left(|I_N^{(3)}| \geq \varepsilon \right) \leq c_{t-1} \exp \left(-\mathbf{d}_{t-1} \left(\frac{\mathcal{I}_{t-1}\varepsilon}{2\bar{\tau}_{t-1}\bar{\sigma}_{t-1}\|f_t\|_\infty \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty} \right)^2 N \right). \quad (\text{A.67})$$

Finally, combining the bounds (A.65–A.67) completes the proof. \square

A.6.6 Proof of Proposition 4

The statement of Proposition 4 is implied by the following more general result, which we will prove below.

Proposition 9. *For every $t \in \mathbb{N}$, $M \in \mathbb{N}^*$, $N \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, $(f_t, \tilde{f}_t) \in \mathbf{F}(\mathcal{X}_t)^2$, and $p \geq 2$, it holds that*

$$\int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, \mathbf{d}\mathbf{b}_t) \left| \frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{t|t}^i) + \tilde{f}_t(x_{t|t}^i)\} - \eta_t\langle z_{0:t} \rangle (f_t B_t\langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \right|^p \leq c_t (p/\mathbf{d}_t)^{p/2} N^{-p/2} \kappa_t^p,$$

where $c_t > 0$, $\mathbf{d}_t > 0$ and κ_t are defined in Proposition 8 and (A.59), respectively.

Before proving Proposition 9, we establish the following result.

Lemma 6. *Let X be an \mathbb{R}^d -valued random variable, defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$, satisfying $\mathbb{P}(|X| \geq t) \leq c \exp(-t^2/(2\sigma^2))$ for every $t \geq 0$ and some $c > 0$ and $\sigma > 0$. Then for every $p \geq 2$ it holds that $\mathbb{E}[|X|^p] \leq cp^{p/2}\sigma^p$.*

Proof. Using Fubini's theorem and the change of variable formula,

$$\mathbb{E}[|X|^p] = \int_0^\infty pt^{p-1} \mathbb{P}(|X| \geq t) dt = cp2^{p/2-1}\sigma^p \Gamma(p/2),$$

where Γ is the Gamma function. It remains to apply the bound $\Gamma(p/2) \leq (p/2)^{p/2-1}$ (see [Anderson and Qiu, 1997]), which holds for $p \geq 2$ by [2, Theorem 1.5]. \square

Proof of Proposition 9. By combining Proposition 8 and Lemma 6 we obtain

$$\begin{aligned} N \int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) & \left| \frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{i|t}^i) + \tilde{f}_t(x_{i|t}^i)\} - \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \right|^2 \\ & \leq c_t (p/d_t)^{p/2} N^{-p/2} \left(\|f_t\|_\infty \sum_{m=0}^{t-1} \|\tilde{h}_m\|_\infty + \|\tilde{f}_t\|_\infty \right)^p, \end{aligned}$$

which was to be established. \square

A.6.7 Proof of Proposition 5

Like previously, we establish Proposition 5 via a more general result, namely the following.

Proposition 10. *For every $t \in \mathbb{N}$, there exists $\bar{c}_t^{bias} < \infty$ such that for every $M \in \mathbb{N}^*$, $N \in \mathbb{N}^*$, $z_{0:t} \in \mathbf{X}_{0:t}$, and $(f_t, \tilde{f}_t) \in \mathbf{F}(\mathcal{X}_t)^2$,*

$$\left| \int \mathbb{C}_t \mathbb{S}_t(z_{0:t}, d\mathbf{b}_t) \frac{1}{N} \sum_{i=1}^N \{b_t^i f_t(x_{i|t}^i) + \tilde{f}_t(x_{i|t}^i)\} - \eta_t \langle z_{0:t} \rangle (f_t B_t \langle z_{0:t-1} \rangle h_t + \tilde{f}_t) \right| \leq \bar{c}_t^{bias} \kappa_t N^{-1},$$

where κ_t is defined in (A.59).

We preface the proof of Proposition 10 by a technical lemma providing a bound on the bias of ratios of random variables.

Lemma 7. *Let α and β be (possibly dependent) random variables defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and such that $\mathbb{E}[\alpha^2] < \infty$ and $\mathbb{E}[\beta^2] < \infty$. Moreover, assume that there exist $c > 0$ and $d > 0$ such that $|\alpha/\beta| \leq c$, \mathbb{P} -a.s., $|a/b| \leq c$, $\mathbb{E}[(\alpha - a)^2] \leq c^2 d^2$, and $\mathbb{E}[(\beta - b)^2] \leq d^2$. Then*

$$|\mathbb{E}[\alpha/\beta] - a/b| \leq 2c(d/b)^2 + c|\mathbb{E}[\beta - b]|/|b| + |\mathbb{E}[\alpha - a]|/|b|. \quad (\text{A.68})$$

Proof. Using the identity

$$\mathbb{E}[\alpha/\beta] - a/b = \mathbb{E}[(\alpha/\beta)(b - \beta)^2]/b^2 + \mathbb{E}[(\alpha - a)(b - \beta)]/b^2 + a\mathbb{E}[b - \beta]/b^2 + \mathbb{E}[\alpha - a]/b,$$

the claim is established by applying the Cauchy–Schwarz inequality and the assumptions of the lemma according to

$$\begin{aligned} |\mathbb{E}[\alpha/\beta] - a/b| & \leq c\mathbb{E}[(\beta - b)^2]/b^2 + \{\mathbb{E}[(\alpha - a)^2]\mathbb{E}[(\beta - b)^2]\}^{1/2}/b^2 + |a|\mathbb{E}[b - \beta]/b^2 + |\mathbb{E}[\alpha - a]|/b^2 \\ & \leq 2c(d/b)^2 + c|\mathbb{E}[\beta - b]|/|b| + |\mathbb{E}[\alpha - a]|/|b|. \end{aligned}$$

\square

Proof of Proposition 5. We proceed by induction and assume that the claim holds true for $n - 1$. Reusing the error decomposition (A.62), it is enough to bound the expectations of the terms $\mathbb{I}_N^{(2)}$ and $\mathbb{I}_N^{(3)}$ given in (A.63) and (A.64), respectively (since $\mathbb{E}_{\boldsymbol{\eta}_0^{\mathbf{P},z}}[\mathbb{I}_N^{(1)}] = 0$). This will be done using the induction hypothesis, Lemma 7, and Proposition 9. More precisely, to bound the expectation of $\mathbb{I}_N^{(2)}$, we use Lemma 7 with $\alpha \leftarrow \alpha_t$, $\beta \leftarrow \beta_t$, $a \leftarrow a_t$, and $b \leftarrow b_t$, where

$$\begin{aligned} \alpha_t & := \frac{1}{N} \sum_{\ell=1}^N \{\beta_{t-1}^\ell Q_{t-1} f_t(\xi_{t-1}^\ell) + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)(\xi_{t-1}^\ell)\}, & \beta_t & := \frac{1}{N} \sum_{\ell=1}^N g_{t-1}(\xi_{t-1}^\ell), \\ a_t & := \eta_{t-1} \langle z_{0:t-1} \rangle \{Q_{t-1} f_t B_t \langle z_{0:t-1} \rangle h_t + Q_{t-1}(\tilde{h}_{t-1} f_t + \tilde{f}_t)\}, & b_t & := \eta_{t-1} \langle z_{0:t-1} \rangle g_{t-1}. \end{aligned}$$

For this purpose, note that $|\alpha_t/\beta_t| \leq \kappa_t$ and $|a_t/b_t| \leq \kappa_t$, where κ_t is defined in (A.59). On the other hand, using Proposition 9 (applied with $p = 2$), we obtain

$$\mathbb{E}_{\eta_0}^{\mathcal{P},z}[(\alpha_t - a_t)^2] \leq d_t^2 \kappa_t^2 \quad \text{and} \quad \mathbb{E}_{\eta_0}^{\mathcal{P},z}[(\beta_t - b_t)^2] \leq d_t^2,$$

where $d_t^2 := c_t \bar{\tau}_{t-1}^2 / (d_t N)$. Using the induction assumption, we get

$$|\mathbb{E}_{\eta_0}^{\mathcal{P},z}[\alpha_t] - a_t| \leq \bar{c}_{t-1}^{bias} N^{-1} \bar{\tau}_{t-1} \kappa_t \quad \text{and} \quad |\mathbb{E}_{\eta_0}^{\mathcal{P},z}[\beta_t] - b_t| \leq \bar{c}_{t-1}^{bias} N^{-1} \bar{\tau}_{t-1}.$$

Hence, the conditions of Lemma 7 are satisfied and we deduce that

$$|\mathbb{E}_{\eta_0}^{\mathcal{P},z}[\mathbb{I}_N^{(2)}]| = |\mathbb{E}_{\eta_0}^{\mathcal{P},z}[\alpha_t/\beta_t] - a_t/b_t| \leq 2\kappa_t \frac{c_t}{d_t N} \frac{\bar{\tau}_{t-1}^2}{\bar{\tau}_{t-1}} + 2\bar{c}_{t-1}^{bias} \kappa_t \frac{\bar{\tau}_{t-1}}{\bar{\tau}_{t-1} N}.$$

The bound on $|\mathbb{E}_{\eta_0}^{\mathcal{P},z}[\mathbb{I}_N^{(2)}]|$ is obtained along the same lines. \square

B. Learning with PPG

This section is divided into three subsections. Appendix B.1 establishes, following closely [Karimi et al., 2019], a non-asymptotic bound for stochastic approximation schemes under general assumptions. Appendix B.2 shows how assumptions A 4.1 and A 3.1 imply the assumptions provided in Appendix B.1 and therefore allow to establish Theorem 2. Finally, Appendix B.3 provides sufficient assumptions on the model ensuring that A 4.1 holds.

B.1 Non-asymptotic bound

We follow closely [Karimi et al., 2019]. Consider the recursion

$$\theta_{n+1} = \theta_n - \gamma_{n+1} H_{\theta_n}(X_{n+1}), \quad n \in \mathbb{N},$$

where $\theta_n \in \Theta \subset \mathbb{R}^d$ for some $d \in \mathbb{N}^*$ and $\{X_n\}_{n \in \mathbb{N}}$ is a *state-dependent* Markov chain on some measurable space $(\mathsf{X}, \mathcal{X})$ in the sense that $X_{n+1} \sim \mathbb{P}_{\theta_n}(X_n, \cdot)$ with \mathbb{P}_{θ} being some Markov kernel on $(\mathsf{X}, \mathcal{X})$. Let $h(\theta) = \int H_{\theta}(x) \pi_{\theta}(dx)$, where π_{θ} is the invariant measure of \mathbb{P}_{θ} and $e_{n+1} := H_{\theta_n}(X_{n+1}) - h(\theta_n)$. As all norms are equivalent in finite dimensional vector spaces, we use $\|\cdot\|$ to denote a generic norm. We denote by $\{\mathcal{F}_n\}_{n \in \mathbb{N}}$ the natural filtration of the Markov chain $\{X_n\}_{n \in \mathbb{N}}$.

A B.1. There exists a Borel measurable function $V : \Theta \rightarrow \mathbb{R}$ such that for every $\theta \in \Theta$, $\nabla V(\theta) = h(\theta)$.

A B.2. There exists $L^V \in \mathbb{R}_{\geq 0}$ such that for every $(\theta, \theta') \in \Theta^2$,

$$\|\nabla V(\theta) - \nabla V(\theta')\| \leq L^V \|\theta - \theta'\|.$$

A B.3. There exists a Borel measurable function $\hat{H} : \Theta \times \mathsf{X} \rightarrow \Theta$ such that for every $\theta \in \Theta$ and $x \in \mathsf{X}$,

$$\hat{H}_{\theta}(x) - \mathbb{P}_{\theta} \hat{H}_{\theta}(x) = H_{\theta}(x) - h(\theta).$$

A B.4. There exists $L^{\mathbb{P}\hat{H}} \in \mathbb{R}_{\geq 0}$ such that for every $(\theta_0, \theta_1) \in \Theta^2$,

$$\sup_{x \in \mathsf{X}} \|\mathbb{P}_{\theta_0} \hat{H}_{\theta_0}(x) - \mathbb{P}_{\theta_1} \hat{H}_{\theta_1}(x)\| \leq L^{\mathbb{P}\hat{H}} \|\theta_0 - \theta_1\|.$$

A B.5. There exists $L_0^{\mathbb{P}\hat{H}} \in \mathbb{R}_{\geq 0}$ such that

$$\sup_{\theta \in \Theta} \|\mathbb{P}_{\theta} \hat{H}_{\theta}\| \leq L_0^{\mathbb{P}\hat{H}}.$$

A B.6. There exists $\sigma_{mse} \in \mathbb{R}_{\geq 0}$ such that for every $x \in \mathsf{X}$ and $\theta \in \Theta$,

$$\int \|H_\theta(x') - h(\theta)\|^2 \mathbb{P}_\theta(x, dx') \leq \sigma_{mse}^2 .$$

A B.7. There exists $L^{\hat{H}} \in \mathbb{R}_{\geq 0}$ such that for every $x \in \mathsf{X}$,

$$\sup_{\theta \in \Theta} \int \|\hat{H}_\theta\| \mathbb{P}_\theta(x, dx') \leq L^{\hat{H}} .$$

Theorem 8. Assume that **A B.1**–**A B.7** hold. In addition, assume that there exist $a > 0$ and $a' > 0$ such that for all $n \in \mathbb{N}$,

$$\gamma_{n+1} \leq \gamma_n \leq a\gamma_{n+1}, \quad \gamma_n - \gamma_{n+1} \leq a' \gamma_n^2, \quad \gamma_1 \leq (L^V + C_h)^{-1}/2 .$$

Moreover, for any $n \in \mathbb{N}^*$, let ϖ be a $\llbracket 0, n \rrbracket$ -valued random variable, independent of $\{\mathcal{F}_\ell\}_{\ell \geq 0}$ and such that $\mathbb{P}(\varpi = k) = \gamma_{k+1} / \sum_{\ell=0}^n \gamma_{\ell+1}$ for $k \in \llbracket 0, n \rrbracket$. Then,

$$\mathbb{E} [\|h(\theta_\varpi)\|^2] \leq 2 \frac{V_{0,n} + C_{0,n} + (\sigma_{mse}^2 L^V + C_\gamma) \sum_{k=0}^n \gamma_{k+1}^2}{\sum_{k=0}^n \gamma_{k+1}} ,$$

where $V_{0,n} := \mathbb{E} [V(\theta) - V(\theta_n)]$ and

$$C_{0,n} := \gamma_1 h(\theta_0) L^{\hat{H}} + L_0^{\mathbb{P}^{\hat{H}}} (\gamma_1 - \gamma_{n+1} + 1) , \quad (\text{B.69})$$

$$C_\gamma := \sigma_{mse} L^{\mathbb{P}^{\hat{H}}} + (1 + \sigma_{mse}) L^V L_0^{\mathbb{P}^{\hat{H}}} , \quad (\text{B.70})$$

$$C_h := L^{\mathbb{P}^{\hat{H}}} ((a + 1)/2 + a\sigma_{mse}) + (L^V + a' + 1) L_0^{\mathbb{P}^{\hat{H}}} . \quad (\text{B.71})$$

Proof. We follow closely the proof of [Karimi et al., 2019, Theorem 2] and adapt it to our setting. First, note that by **A B.1**, assumptions **A1** and **A2** of [Karimi et al., 2019, Theorem 2] hold with $c_0 = d_0 = 0$ and $c_1 = d_1 = 1$. In addition, the claim in [Karimi et al., 2019, Lemma 1] holds true since by **A B.2**, **A3** holds. Moreover, [Karimi et al., 2019, Equation 17] can also be established under **A B.6**, as we may rewrite it as

$$\sum_{\ell=0}^n \gamma_{\ell+1}^2 \mathbb{E} [\|e_{\ell+1}\|^2] = \sum_{\ell=0}^n \gamma_{\ell+1}^2 \mathbb{E} [\mathbb{E} [\|e_{\ell+1}\|^2 \mid \mathcal{F}_\ell]] \leq \sigma_{mse}^2 \sum_{\ell=0}^n \gamma_{\ell+1}^2 .$$

Following the proof of [Karimi et al., 2019, Lemma 2], consider the decomposition

$$\mathbb{E} \left[- \sum_{\ell=0}^n \gamma_{\ell+1} \langle \nabla V(\theta_\ell), e_{\ell+1} \rangle \right] = \mathbb{E} [A_1 + A_2 + A_3 + A_4 + A_5] ,$$

where

$$\begin{aligned} A_1 &:= - \sum_{\ell=1}^n \gamma_{\ell+1} \left\langle \nabla V(\theta_\ell), \hat{H}_{\theta_\ell}(X_{\ell+1}) - \mathbb{P}_{\theta_\ell} \hat{H}_{\theta_\ell}(X_\ell) \right\rangle, \\ A_2 &:= - \sum_{\ell=1}^n \gamma_{\ell+1} \left\langle \nabla V(\theta_\ell), \mathbb{P}_{\theta_\ell} \hat{H}_{\theta_\ell}(X_\ell) - \mathbb{P}_{\theta_{\ell-1}} \hat{H}_{\theta_{\ell-1}}(X_\ell) \right\rangle, \\ A_3 &:= - \sum_{\ell=1}^n \gamma_{\ell+1} \left\langle \nabla V(\theta_\ell) - \nabla V(\theta_{\ell-1}), \mathbb{P}_{\theta_{\ell-1}} \hat{H}_{\theta_{\ell-1}}(X_\ell) \right\rangle, \\ A_4 &:= - \sum_{\ell=1}^n (\gamma_{\ell+1} - \gamma_\ell) \left\langle \nabla V(\theta_{\ell-1}), \mathbb{P}_{\theta_{\ell-1}} \hat{H}_{\theta_{\ell-1}}(X_\ell) \right\rangle, \\ A_5 &:= -\gamma_1 \left\langle \nabla V(\theta_0), \hat{H}_{\theta_0}(X_1) \right\rangle + \gamma_{n+1} \left\langle \nabla V(\theta_n), \mathbb{P}_{\theta_n} \hat{H}_{\theta_n}(X_{n+1}) \right\rangle. \end{aligned}$$

As $\widehat{H}_{\theta_\ell}(X_{\ell+1}) - \mathbb{P}_{\theta_\ell} \widehat{H}_{\theta_\ell}(X_\ell)$ is a martingale difference, it holds that $\mathbb{E}[A_1] = 0$. The upper bounds on the expectations of A_2 , A_3 and A_4 are obtained similarly as in [Karimi et al., 2019]. Using **A B.4**,

$$A_2 \leq L^{\mathbb{P}^{\widehat{H}}} \left(\sigma_{mse} \sum_{k=1}^n \gamma_k^2 + \frac{1}{2} (1 + 2a\sigma_{mse} + a) \sum_{k=0}^n \gamma_{k+1}^2 \|h(\theta_k)\|^2 \right).$$

By **A B.2** and **B.5**,

$$A_3 \leq L^V L_0^{\mathbb{P}^{\widehat{H}}} \left((1 + \sigma_{mse}) \sum_{k=1}^n \gamma_k^2 + \sum_{k=1}^n \gamma_k^2 \|h(\theta_k)\|^2 \right).$$

On the other hand,

$$A_4 \leq L_0^{\mathbb{P}^{\widehat{H}}} \left(\gamma_1 - \gamma_{n+1} + a' \sum_{k=1}^n \gamma_k^2 \|h(\theta_{k-1})\|^2 \right).$$

We now focus on A_5 . As in the proof of [Karimi et al., 2019, Lemma 2], the expectation of the first term can be straightforwardly bounded by $\gamma_1 \|h(\theta_0)\| L^{\widehat{H}}$ using the Cauchy–Schwarz inequality and **A B.7**. The second term can, using **A B.5** and $\gamma_{n+1} \|h(\theta_n)\| \leq 1 + \gamma_{n+1}^2 \|h(\theta_n)\|^2$, be bounded in the same way according to

$$\begin{aligned} \gamma_{n+1} \left\langle \nabla V(\theta_n), \mathbb{P}_{\theta_n} \widehat{H}_{\theta_n}(X_{n+1}) \right\rangle &\leq L_0^{\mathbb{P}^{\widehat{H}}} \gamma_{n+1} \|h(\theta_n)\| \leq L_0^{\mathbb{P}^{\widehat{H}}} (1 + \gamma_{n+1}^2 \|h(\theta_n)\|^2) \\ &\leq L_0^{\mathbb{P}^{\widehat{H}}} \left(1 + \sum_{\ell=0}^n \gamma_{\ell+1}^2 \|h(\theta_\ell)\|^2 \right). \end{aligned}$$

The rest of the proof follows that of [Karimi et al., 2019, Theorem 2]. \square

B.2 Application to Theorem 2

The goal of this section is to establish that the assumptions of Theorem 2 ensure all the assumptions in appendix B.1, which in turn allows Theorem 8 to be applied. First, we start by explicitly defining the kernel \mathbb{P}_θ and the function h in terms of the kernels presented in appendix A. We write $\mathbb{P}_{\theta,t}$ instead of \mathbb{P}_θ to explicit the dependence of the kernel on the *fixed* number of observations t .

B.2.1 Verification of the assumptions of Theorem 8

For $(k_0, k) \in (\mathbb{N}^*)^2$ such that $k_0 < k$, define

$$\mathbb{P}_{\theta,t} : \mathbf{E}_t^{k-k_0} \times \mathcal{E}_t^{\otimes(k-k_0)} \ni (\mathbf{y}_t[k_0 : k], z_{0:t}[k_0 : k], A) \mapsto \mathbb{K}_{\theta,t}^{k_0} \otimes \mathbb{K}_{\theta,t}^{\otimes(k-k_0)}(z_{0:t}[k], A), \quad (\text{B.72})$$

where $\mathbb{K}_{\theta,t}$ is the PPG kernel defined in (A.41). Note that $\mathbb{P}_{\theta,t}$ depends only on the last frozen path, namely $z_{0:t}[k]$. Note also that, since $\mathbb{K}_{\theta,t}$ depends only on the paths, there is no dependence between $\mathbf{y}_{t,\ell}[k_0 : k]$ and $\mathbf{y}_{t,\ell+1}[k_0 : k]$. The score ascent algorithm (Algorithm 4) can be formulated as follows.

1. Sample $(z_{0:t,\ell}[k_0 : k], \mathbf{y}_{t,\ell}[k_0 : k]) \sim \mathbb{P}_{\theta,t}((z_{0:t,\ell-1}[k_0 : k], \mathbf{y}_{t,\ell-1}[k_0 : k]), \cdot)$.
2. Update the parameter according to $\eta_{\ell+1} = \eta_\ell + \gamma_{\ell+1} H(z_{0:t,\ell}[k_0 : k], \mathbf{y}_{t,\ell}[k_0 : k])$, where

$$H(z_{0:t,\ell}[k_0 : k], \mathbf{y}_{t,\ell}[k_0 : k]) = \frac{1}{k - k_0 + 1} \sum_{i=k_0}^k \mu(\boldsymbol{\beta}_{t,\ell}[i])(\text{id}) = \Pi_{(k_0-1,k),N}(h_t),$$

where $\Pi_{(k_0-1,k),N}(h_t)$ is defined in (3.12). We denote by $\pi_{\theta,t}$ the invariant distribution of $\mathbb{P}_{\theta,t}$, which, by Proposition 3, is given by $\pi_{\theta,t} = (\eta_{0:t} \otimes \mathbb{C}_t \mathcal{S}_t)^{\otimes(k-k_0)}$.

We also require the strong mixing assumption to hold uniformly in θ .

A B.8 (Strong mixing uniformly in θ). For every $s \in \mathbb{N}$ there exist $\underline{\tau}_s, \bar{\tau}_s, \underline{\sigma}_s$, and $\bar{\sigma}_s$ in \mathbb{R}_+^* such that for all $\theta \in \Theta$,

- (i) $\underline{\tau}_s \leq g_{s,\theta}(x_s) \leq \bar{\tau}_s$ for every $x_s \in \mathbf{X}_s$,
- (ii) $\underline{\sigma}_s \leq m_{s,\theta}(x_s, x_{s+1}) \leq \bar{\sigma}_s$ for every $(x_s, x_{s+1}) \in \mathbf{X}_{s:s+1}$.

Note that the assumption above implies that $\kappa_{N,t}$ is also uniform in θ .

Proof that A B.1 holds.

Proposition 11. For all $\theta \in \Theta$, $h(\theta) = \nabla V(\theta)$, where $V(\theta) = \log \gamma_{0:t,\theta}(\mathbf{X}_{0:t})$ is the log-likelihood function.

Proof. By Theorem 5,

$$\begin{aligned} h(\theta) &= \int H(\tilde{\mathbf{y}}_t[k_0 : k], \tilde{x}_{0:t}[k_0 : k]) \pi_{\theta,t}(\mathrm{d}(\tilde{\mathbf{y}}_t[k_0 : k], \tilde{x}_{0:t}[k_0 : k])) \\ &= \frac{1}{k - k_0 + 1} \sum_{i=k_0}^k \int [\eta_{0:t,\theta} \otimes \mathbb{C}_{t,\theta} \mathbb{S}_{t,\theta}] (\mathrm{d}(\tilde{\mathbf{y}}_t[i], \tilde{x}_{0:t}[i])) \mu(\tilde{\beta}_{t,\ell}[i])(\mathrm{id}) \\ &= \eta_{0:t,\theta}(s_{0:t,\theta}) = \nabla V(\theta). \end{aligned}$$

□

Proof that A B.2 holds. A B.2 is trivially implied by A 4.1(i).

Proof that A B.3 and B.5 hold. Let \hat{H}_θ be given by

$$\hat{H}_\theta : \mathbf{E}_t^{k-k_0} \ni (\mathbf{y}_t[k_0 : k], z_{0:t}[k_0 : k]) \mapsto \sum_{r=0}^{\infty} \{\mathbb{P}_{\theta,t}^r H(\mathbf{y}_t[k_0 : k], z_{0:t}[k_0 : k]) - h(\theta)\}. \quad (\text{B.73})$$

Then the following holds true.

Lemma 8. Assume A B.8. Then for all $\theta \in \Theta$ and $t \in \mathbb{N}^*$,

$$\|\mathbb{P}_{\theta,t} \hat{H}_\theta\|_\infty \leq \sigma_{bias} (1 - \kappa_{N,t}^k)^{-1}.$$

Proof. By Theorem 1, we have for any $r > 0$

$$|\mathbb{P}_{\theta,t}^r H(\mathbf{y}_t[k_0 : k], z_{0:t}[k_0 : k]) - h(\theta)| \leq \sigma_{bias} \kappa_{N,t}^{(r-1)k}$$

and thus

$$\|\mathbb{P}_{\theta,t} \hat{H}_\theta\|_\infty \leq \sum_{r=1}^{\infty} \|\mathbb{P}_{\theta,t}^r H - h(\theta)\|_\infty \leq \sigma_{bias} \sum_{r=0}^{\infty} \kappa_{N,t}^{rk} \leq \sigma_{bias} (1 - \kappa_{N,t}^k)^{-1},$$

where $\kappa_{N,t} \in (0, 1)$. □

Lemma 8 proves A B.3 and B.5 with $L_0^{\mathbb{P}\hat{H}} := \sigma_{bias} (1 - \kappa_{N,t}^k)^{-1}$.

Proof that A B.4 holds.

Theorem 9. *Assume A B.8 and A 4.1. Then for every $t \in \mathbb{N}$, $\theta \in \Theta$ and $N \in \mathbb{N}^*$ such that $N > 1 + 5\rho_t^2/2$,*

$$\left\| \mathbb{P}_{\theta_1, t} \widehat{H}_{\theta_1} - \mathbb{P}_{\theta_2, t} \widehat{H}_{\theta_2} \right\|_{\infty} \leq L^{\mathbb{P} \widehat{H}} \|\theta_1 - \theta_2\| ,$$

where

$$L^{\mathbb{P} \widehat{H}} := \|L_2^P\|_{\infty} [1 + \kappa_{N, t}^k (1 - \kappa_{N, t}^k)] + L^V + \sigma_{bias} (1 - \kappa_{N, t})^{-1} (1 - \kappa_{N, t}^k)^{-1} [\|L_1^P\|_{\infty} (1 - \kappa_{N, t}^k)^{-1} + L^{\eta} \kappa_{N, t}^k] . \quad (\text{B.74})$$

Proof. We establish the claim by adapting the proof of [Karimi et al., 2019, Lemma 7]. First, recall that the kernel $K_{\theta, t}$ defined in (A.42) is the path marginalized version of $\mathbb{K}_{\theta, t}$ given in (A.41). Note that for every $x \in \mathbf{E}_t^{k-k_0}$,

$$\mathbb{P}_{\theta_1, t} \widehat{H}_{\theta_1}(x) = \sum_{n=0}^{\infty} \delta_x \mathbb{P}_{\theta_1, t} \{ \mathbb{P}_{\theta_1, t}^n H - h(\theta_1) \} = \sum_{n=0}^{\infty} \delta_x K_{\theta_1, t}^{kn} \{ \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_1} \mathbb{P}_{\theta_1, t} H \} ,$$

where we have used (i) the fact that the backward statistics output by $\mathbb{P}_{\theta, t}$ are independent of the input backward statistics and (ii) the penultimate line in the computation of $h(\theta)$ above. We follow the proof of [Fort et al., 2011, Lemma 4.2] and consider the following decomposition: for $n \in \mathbb{N}^*$,

$$\begin{aligned} & \delta_x K_{\theta_1, t}^{kn} (\mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_1} \mathbb{P}_{\theta_1, t} H) - \delta_x K_{\theta_2, t}^{kn} (\mathbb{P}_{\theta_2, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_2, t} H) \\ &= \sum_{j=0}^{n-1} \left(\delta_x K_{\theta_1, t}^{kj} - \eta_{0:t, \theta_1} \right) \left(K_{\theta_1, t}^{kj} - K_{\theta_2, t}^{kj} \right) \left(K_{\theta_2, t}^{k(n-j-1)} \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_1, t} H \right) \\ & \quad - \left(\delta_x K_{\theta_2, t}^{kn} \mathbb{P}_{\theta_2, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_2, t} H \right) + \left(\delta_x K_{\theta_2, t}^{kn} \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_1, t} H \right) \\ & \quad - \eta_{0:t, \theta_1} \left(K_{\theta_2, t}^{kn} \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_1, t} H \right) . \end{aligned} \quad (\text{B.75})$$

Applying Theorem 6 with $\mu = \delta_x$ and $\nu = \eta_{0:t, \theta}$ and using the fact that $\eta_{0:t, \theta} K_{\theta, t}^{\ell} = \eta_{0:t, \theta}$ for all $\ell \in \mathbb{N}$, we obtain that for all $\ell \in \mathbb{N}$ and all $\theta \in \Theta$, $\left\| \delta_x K_{\theta, t}^{\ell} - \eta_{0:t, \theta} \right\|_{\text{TV}} \leq \kappa_{N, t}^{\ell}$. Note that by A 4.1(iii), $K_{\theta, t}$ is Lipschitz; therefore, for all $r \in \mathbb{N}^*$, by Lemma 18, $K_{\theta, t}^r$ is Lipschitz with constant $\|L_1^P\|_{\infty} (1 - \kappa_{N, t})^{-1}$. Combining all this together, we obtain

$$\begin{aligned} & \left| \left(\delta_x K_{\theta_1, t}^{kj} - \eta_{0:t, \theta_1} \right) \left(K_{\theta_1, t}^{kj} - K_{\theta_2, t}^{kj} \right) \left(K_{\theta_2, t}^{k(n-j-1)} \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_1, t} H \right) \right| \\ &= \left| \left(\delta_x K_{\theta_1, t}^{kj} - \eta_{0:t, \theta_1} \right) \left(K_{\theta_1, t}^{kj} - K_{\theta_2, t}^{kj} \right) \left\{ K_{\theta_2, t}^{k(n-j-1)} [\mathbb{P}_{\theta_1, t} H - h(\theta_1)] - \eta_{0:t, \theta_2} [\mathbb{P}_{\theta_1, t} H - h(\theta_1)] \right\} \right| \\ &\leq \|L_1^P\|_{\infty} (1 - \kappa_{N, t})^{-1} \kappa_{N, t}^{kj} \kappa_{N, t}^{k(n-j-1)} \|\mathbb{P}_{\theta_1, t} H - h(\theta_1)\|_{\infty} \|\theta_1 - \theta_2\| \\ &\leq \sigma_{bias} \|L_1^P\|_{\infty} (1 - \kappa_{N, t})^{-1} \kappa_{N, t}^{k(n-1)} \|\theta_1 - \theta_2\| , \end{aligned}$$

where the last inequality is due to Theorem 1. Therefore, the first term of the right side of (B.75) is upper bounded by $\sigma_{bias} \|L_1^P\|_{\infty} (1 - \kappa_{N, t})^{-1} n \kappa_{N, t}^{k(n-1)} \|\theta_1 - \theta_2\|$. The second term of (B.75) can be written

$$\begin{aligned} & - \left(\delta_x K_{\theta_2, t}^{kn} \mathbb{P}_{\theta_2, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_2, t} H \right) + \left(\delta_x K_{\theta_2, t}^{kn} \mathbb{P}_{\theta_1, t} H - \eta_{0:t, \theta_2} \mathbb{P}_{\theta_1, t} H \right) \\ &= \left(\delta_x K_{\theta_2, t}^{kn} - \eta_{0:t, \theta_2} \right) (\mathbb{P}_{\theta_1, t} H - \mathbb{P}_{\theta_2, t} H) , \end{aligned}$$

and using again the ergodicity of $K_{\theta, t}$ and the fact that $\theta \mapsto \mathbb{P}_{\theta, t} H$ is uniformly Lipschitz by A 4.1(iv), we may conclude that it is upper bounded by $\|L_2^P\|_{\infty} \kappa_{N, t}^{kn} \|\theta_1 - \theta_2\|$. Finally, for the last term, using

the facts that $K_{\theta,t}^k$ is $\eta_{0:t,\theta}$ -invariant and geometrically ergodic and that $\theta \mapsto \eta_{0:t,\theta}$ is Lipschitz by **A 4.1(iv)** yields

$$\begin{aligned} & \left| \eta_{0:t,\theta_1} \left(K_{\theta_2,t}^{kn} \mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_1,t} H \right) \right| \\ &= \left| (\eta_{0:t,\theta_1} - \eta_{0:t,\theta_2}) \left\{ K_{\theta_2,t}^{kn} [\mathbb{P}_{\theta_1,t} H - h(\theta_1)] - \eta_{0:t,\theta_2} [\mathbb{P}_{\theta_1,t} H - h(\theta_1)] \right\} \right| \\ &\leq L^\eta \kappa_{N,t}^{kn} \|\mathbb{P}_{\theta_1,t} H - h(\theta_1)\|_\infty \|\theta_1 - \theta_2\| \\ &\leq L^\eta \sigma_{bias} (1 - \kappa_{N,t})^{-1} \kappa_{N,t}^{kn} \|\theta_1 - \theta_2\|. \end{aligned}$$

Therefore, we have that

$$\begin{aligned} & \delta_x K_{\theta_1,t}^{kn} (\mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_1} \mathbb{P}_{\theta_1,t} H) - \delta_x K_{\theta_2,t}^{kn} (\mathbb{P}_{\theta_2,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_2,t} H) \\ &\leq \left\{ \sigma_{bias} \|L_1^P\|_\infty (1 - \kappa_{N,t})^{-1} n \kappa_{N,t}^{k(n-1)} + [\|L_2^P\|_\infty + L^\eta \sigma_{bias} (1 - \kappa_{N,t})^{-1}] \kappa_{N,t}^{kn} \right\} \|\theta_1 - \theta_2\|. \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} & \left| \mathbb{P}_{\theta_1,t} \widehat{H}_{\theta_1}(x) - \mathbb{P}_{\theta_2,t} \widehat{H}_{\theta_2}(x) \right| \\ &\leq |\delta_x \mathbb{P}_{\theta_1,t} H - \delta_x \mathbb{P}_{\theta_2,t} H| + |\eta_{0:t,\theta_1} \mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_2,t} H| \\ &\quad + \left| \sum_{n=1}^{\infty} \delta_x K_{\theta_1,t}^{kn} (\mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_1} \mathbb{P}_{\theta_1,t} H) - \delta_x K_{\theta_2,t}^{kn} (\mathbb{P}_{\theta_2,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_2,t} H) \right| \\ &\leq |\delta_x \mathbb{P}_{\theta_1,t} H - \delta_x \mathbb{P}_{\theta_2,t} H| + |\eta_{0:t,\theta_1} \mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_2,t} H| \\ &\quad + \left\{ \sigma_{bias} \|L_1^P\|_\infty (1 - \kappa_{N,t})^{-1} (1 - \kappa_{N,t}^k)^{-2} \right. \\ &\quad \left. + [\|L_2^P\|_\infty + L^\eta \sigma_{bias} (1 - \kappa_{N,t})^{-1}] \kappa_{N,t}^k (1 - \kappa_{N,t}^k)^{-1} \right\} \|\theta_1 - \theta_2\|. \end{aligned}$$

To conclude, note that by **A 4.1(iv)**, $\|\delta_x \mathbb{P}_{\theta_1,t} H - \delta_x \mathbb{P}_{\theta_2,t} H\| \leq \|L_2^P\|_\infty \|\theta_1 - \theta_2\|$. Furthermore, note that by **Theorem 5** we obtain that for all $\theta \in \Theta$, $\eta_{0:t,\theta} \mathbb{P}_{\theta,t} H = \eta_{0:t,\theta} s_{0:t,\theta} = \nabla V(\theta)$. Therefore, by **A 4.1(i)** we obtain that $\|\eta_{0:t,\theta_1} \mathbb{P}_{\theta_1,t} H - \eta_{0:t,\theta_2} \mathbb{P}_{\theta_2,t} H\| \leq L^V \|\theta_1 - \theta_2\|$, concluding the proof. \square

Proof that A B.6 holds. **A B.6** is simply a bound on the MSE of the roll-out PPG estimator, given by **Theorem 1**.

Proof that A B.7 holds.

Proposition 12. For all $\theta \in \Theta$ and all $\ell \in \llbracket 1, t-1 \rrbracket$

$$\mathbb{E} \left[\|\widehat{H}_\theta\| \mid \mathcal{F}_\ell \right] \leq 2 \|s_{0:t,\theta}\|_\infty + \sigma_{bias} (1 - \kappa_{N,t}^k)^{-1}.$$

Proof. Note that for all $x \in \mathbf{E}_t^{k-k_0}$ and all $\theta \in \Theta$,

$$\widehat{H}_\theta(x) = H(x) - h(\theta) + \mathbb{P}_{\theta,t} \widehat{H}_\theta(x). \quad (\text{B.76})$$

Lemma 8 shows that $\|\mathbb{P}_{\theta,t} \widehat{H}_\theta\|_\infty \leq \sigma_{bias} (1 - \kappa_{N,t}^k)^{-1}$. Note that $h(\theta) \leq \|s_{0:t,\theta}\|_\infty$. We write

$$\mathbb{E} [\|H\| \mid \mathcal{F}_\ell] \leq \frac{1}{(k - k_0 + 1)N} \sum_{i=k_0}^k \sum_{j=1}^N \mathbb{E} \left[\|\beta_{t,\ell}^j[i]\| \mid \mathcal{F}_\ell \right].$$

By **Proposition 14**, $\mathbb{E} \left[\|\beta_{t,\ell}^j[i]\| \mid \mathcal{F}_\ell \right] \leq \|s_{0:t,\theta}\|_\infty$, concluding the proof. \square

A B.7 follows directly by **Proposition 12** and by considering $\sup_{\theta \in \Theta} \|s_{0:t,\theta}\|_\infty$.

B.2.2 Proof of Theorem 2

We have shown in Appendix B.2.1 that under **A** 4.1 and B.8, it is possible to apply Theorem 8. To conclude the proof of Theorem 2 we just have to rearrange the constants. We start by rewriting the constant in Theorem 9

$$L^{\mathbb{P}\hat{H}} = C_1 + \sigma_{bias}(1 - \kappa_{N,t})^{-1}(1 - \kappa_{N,t}^k)^{-1}C_2,$$

with

$$\begin{aligned} C_1 &= \|L_2^P\|_\infty [1 + \kappa_{N,t}^k(1 - \kappa_{N,t}^k)^{-1}] + L^V \\ C_2 &= \|L_1^P\|_\infty (1 - \kappa_{N,t}^k)^{-1} + L^\eta \kappa_{N,t}^k. \end{aligned}$$

By (B.70) and Lemma 8,

$$\begin{aligned} C_\gamma &= \sigma_{mse} L^{\mathbb{P}\hat{H}} + (1 + \sigma_{mse}) L^V L_0^{\mathbb{P}\hat{H}} \\ &= \sigma_{mse} [C_1 + \sigma_{bias}(1 - \kappa_{N,t})^{-1}(1 - \kappa_{N,t}^k)^{-1}C_2] + (1 + \sigma_{mse}) L^V \sigma_{bias}(1 - \kappa_{N,t}^k)^{-1} \\ &= \sigma_{mse} C_1 + \sigma_{mse} \sigma_{bias}(1 - \kappa_{N,t}^k)^{-1} [L^V + (1 - \kappa_{N,t})^{-1}C_2] + \sigma_{bias} L^V (1 - \kappa_{N,t}^k)^{-1}. \end{aligned}$$

Therefore,

$$\begin{aligned} C_{0,\gamma} &:= \sigma_{mse}^2 L^V + C_\gamma \\ &= \sigma_{mse}^2 L^V + \sigma_{mse} C_1 + \sigma_{mse} \sigma_{bias}(1 - \kappa_{N,t}^k)^{-1} [L^V + (1 - \kappa_{N,t})^{-1}C_2] + \sigma_{bias} L^V (1 - \kappa_{N,t}^k)^{-1}. \end{aligned}$$

In the same way, we can rewrite (B.71) as

$$\begin{aligned} C_h &= L^{\mathbb{P}\hat{H}} [(a+1)/2 + a\sigma_{mse}] + (L^V + a' + 1) L_0^{\mathbb{P}\hat{H}} \\ &= [C_1 + \sigma_{bias}(1 - \kappa_{N,t})^{-1}(1 - \kappa_{N,t}^k)^{-1}C_2] [(a+1)/2 + a\sigma_{mse}] + (L^V + a' + 1) \sigma_{bias}(1 - \kappa_{N,t}^k)^{-1}. \end{aligned}$$

The constant C_0 from Theorem 2 is $L^{\hat{H}} = 2 \sup_{\theta \in \Theta} \|s_{0:t,\theta}\|_\infty + \sigma_{bias}(1 - \kappa_{N,t}^k)^{-1}$ which completes the proof.

B.3 Conditions on the model to verify A 4.1

In our specific application to score ascent, we work with the following assumptions.

- A B.9** (Lipschitz). (i) For all $t \in \mathbb{N}$, there exists $L_t^s \in \mathbf{M}(\mathbf{X}_{t:t+1})$ such that for all $(x_t, x_{t+1}) \in \mathbf{X}_{t:t+1}$, the function $\theta \mapsto s_{t,\theta}(x_t, x_{t+1})$ is $L_t^s(x_t, x_{t+1})$ -Lipschitz and $\mathbf{X}_{t:t+1} \ni (x_t, x_{t+1}) \mapsto s_{t,\theta}(x_t, x_{t+1})$ is bounded by $\|s_t(\theta)\|_\infty$ for all $\theta \in \Theta$. Furthermore, $\|L_k^s\|_\infty < \infty$.
(ii) For all $t \in \mathbb{N}$, there exists $L_t^q \in \mathbf{X}_{t:t+1}$ such that $\|L_t^q\|_\infty < \infty$ and that for all $(x_t, x_{t+1}) \in \mathbf{X}_{t:t+1}$, $\theta \mapsto q_{t,\theta}(x_t, x_{t+1})$ is $L_t^q(x_t, x_{t+1})$ -Lipschitz.

Lemma 9 (**A** B.2(i) holds). *Assume **A** B.8 and **A** 4.1. There exists a constant L^V such that the Lyapunov function V satisfies, for all $(\theta_1, \theta_2) \in \Theta^2$,*

$$\|\nabla V(\theta_1) - \nabla V(\theta_2)\| \leq L^V \|\theta_1 - \theta_2\|.$$

Proof. For all θ_1, θ_2 ,

$$\begin{aligned} \|\nabla V(\theta_1) - \nabla V(\theta_2)\| &= \|\eta_{0:t,\theta_1}(s_{0:t,\theta_1}) - \eta_{0:t,\theta_2}(s_{0:t,\theta_2})\| \\ &\leq \|\eta_{0:t,\theta_1}(s_{0:t,\theta_1}) - \eta_{0:t,\theta_1}(s_{0:t,\theta_2})\| + \|\eta_{0:t,\theta_1}(s_{0:t,\theta_2}) - \eta_{0:t,\theta_2}(s_{0:t,\theta_2})\|. \end{aligned}$$

By (3.1) and by [Gloaguen et al., 2022, Theorem 4.10] there exists a constant c such that

$$\|\eta_{0:t,\theta_1}(s_{0:t,\theta_2}) - \eta_{0:t,\theta_2}(s_{0:t,\theta_2})\| \leq ct \|\theta_1 - \theta_2\| \sup_\theta \sup_k \|s_k(\theta)\|_\infty,$$

Using **A 3.1** and **A 4.1**[i], we can write:

$$\begin{aligned} \|\eta_{0:t,\theta_1}(s_{0:t,\theta_1}) - \eta_{0:t,\theta_1}(s_{0:t,\theta_2})\| &\leq \sum_{u=0}^{t-1} \eta_{0:t,\theta_1} [\|s_{u,\theta_1}(x_{u:u+1}) - s_{u,\theta_2}(x_{u:u+1})\|], \\ &\leq \sum_{u=0}^{t-1} \eta_{0:t,\theta_1} [L_u^s(x_{u:u+1})] \|\theta_1 - \theta_2\|, \\ &\leq \frac{\sigma_+}{\sigma_-} \sup_{u \in [0, t-1]} [L_u^s] \|\theta_1 - \theta_2\| t. \end{aligned}$$

□

Theorem 10 (Lipschitz continuity of Particle Gibbs with Backward Sampling). *Assume **A B.9**. For every $t \in \mathbb{N}$, $\theta \in \Theta$ and $N \in \mathbb{N}^*$*

$$\sup_{x_{0:t} \in \mathbb{X}_{0:t}} \|K_{\theta_1,t}(x_{0:t}, \cdot) - K_{\theta_2,t}(x_{0:t}, \cdot)\|_{\text{TV}} \leq L_{t,N}^K \|\theta_1 - \theta_2\|,$$

where

$$L_{t,N}^K := \sum_{\ell=0}^{t-1} \bar{\tau}_\ell^{-1} [\bar{\sigma}_\ell^{-1} + (N-1)] \|L_\ell^q\|_\infty. \quad (\text{B.77})$$

Proof. We know that $K_{\theta,t} = \mathbb{C}_{m,\theta} \mathbb{B}_{t,\theta}$. Therefore, by Lemmas 14, 16 and 19, we have that $K_{\theta,t}$ is Lipschitz with constant equals $L_t^{\mathbb{C}} + \sup_\theta \mathbb{C}_{t,\theta} L_t^{\mathbb{B}}$. □

Corollary 1 (**A 4.1**(iii) holds.). *Assume **A B.9**. For every $t \in \mathbb{N}$, $\theta \in \Theta$, $r \in \mathbb{N}^*$ and $N \in \mathbb{N}^*$ such that $N > 1 + 5\rho_t^2 t/2$*

$$\sup_{x_{0:t} \in \mathbb{X}_{0:t}} \|K_{\theta_1,t}^r(x_{0:t}, \cdot) - K_{\theta_2,t}^r(x_{0:t}, \cdot)\|_{\text{TV}} \leq L_{t,N}^P \|\theta_1 - \theta_2\|$$

where

$$L_{t,N}^P := (1 - \kappa_{t,N})^{-1} \|L_{t,N}^K\|_\infty \quad (\text{B.78})$$

where $L_{t,N}^K$ is defined in (B.77).

Proof. Under B.8, the Particle Gibbs with backward sampling is geometrically ergodic with contraction rate $\kappa_{t,N}$ and thus $L_{t,N}^K$ is bounded and the result follows from Lemma 18 □

Corollary 2 (**A 4.1**(i)). *Assume **A B.8** and **A B.9**. For all $t \in \mathbb{N}^*$, $(\theta_0, \theta_1) \in \Theta^2$,*

$$\|\eta_{0:t,\theta_0} - \eta_{0:t,\theta_1}\|_{\text{TV}} \leq L^{\eta} \|\theta_0 - \theta_1\|,$$

where

$$L^{\eta} := L_{t,N^*}^P, \quad (\text{B.79})$$

and $L_{t,N}^P$ is defined in (B.78) and $N^* = \lceil 1 + 5\rho_t^2/2 \rceil$.

Proof. Consider the following decomposition, valid for all $k \in \mathbb{N}^*$ and $N \geq 1 + 5\rho_t^2/2$, and all $x_{0:t} \in \mathbb{X}_{0:t}$,

$$\begin{aligned} \|\eta_{0:t,\theta_1} - \eta_{0:t,\theta_2}\|_{\text{TV}} &\leq \|\eta_{0:t,\theta_1} - K_{\theta_1,t}^k(x_{0:t}, \cdot)\|_{\text{TV}} + \|\eta_{0:t,\theta_2} - K_{\theta_2,t}^k(x_{0:t}, \cdot)\|_{\text{TV}} + \|K_{\theta_1,t}^k(x_{0:t}, \cdot) - K_{\theta_2,t}^k(x_{0:t}, \cdot)\|_{\text{TV}} \\ &\leq \|\eta_{0:t,\theta_1} - K_{\theta_1,t}^k(x_{0:t}, \cdot)\|_{\text{TV}} + \|\eta_{0:t,\theta_2} - K_{\theta_2,t}^k(x_{0:t}, \cdot)\|_{\text{TV}} + L_{t,N}^P \|\theta_1 - \theta_2\|, \end{aligned}$$

where we applied Corollary 1. Since the Lipschitz constant of $K_{\theta,t}$ is independent of k , and $K_{\theta,t}$ is geometrically ergodic for all θ , we obtain by taking the limit when k goes to infinity with N fixed,

$$\|\eta_{0:t,\theta_1} - \eta_{0:t,\theta_2}\|_{\text{TV}} \leq \frac{\|L_{t,N}^K\|_\infty}{1 - \kappa_{t,N}} \|\theta_1 - \theta_2\|,$$

for all $N \geq 1 + 5\rho_t^2/2$, where the dependence in N is hidden in $L_{t,N}^P$. The result follows by choosing $N = \lceil 1 + 5\rho_t^2/2 \rceil$. □

Remark 2. As noted by [Lindholm and Lindsten, 2018], the Lipschitz constant appearing in Corollary 1 possesses an unexpected dependence on $N - 1$. One would expect it not to be true, in that we know that $\mathbb{K}_{\theta,t}$ converges geometrically fast and uniformly to $\eta_{0:t}$ and this is faster as N gets bigger. Therefore, for large N the Lipschitz constant is expected to converge to that of $\eta_{0:t}$ whose Lipschitz constant is independent of N .

Proposition 13 (Lipschitz continuity of $\theta \mapsto \mathbb{K}_{\theta,t}\mu(\beta_t)(\text{id})$). *Assume A B.9. For every $t \in \mathbb{N}$, $\theta \in \Theta$ and $N \in \mathbb{N}^*$,*

$$\|\mathbb{K}_{\theta_1,t}\mu(\beta_t)(\text{id}) - \mathbb{K}_{\theta_2,t}\mu(\beta_t)(\text{id})\|_\infty \leq L_t^{\mathbb{K}}\|\theta_1 - \theta_2\| ,$$

where

$$L_t^{\mathbb{K}} := (N - 1) \sum_{\ell=0}^{t-1} \bar{\tau}_\ell \|L_\ell^q\|_\infty + \sum_{j=1}^m \|L_j^{\tilde{Q}}\|_\infty \left[\sum_{\ell=0}^{m-1} s_\ell^\infty \right] + \sum_{j=1}^m \|L_j^s\|_\infty . \quad (\text{B.80})$$

Proof. Consider $e = (x_{0:t}, \mathbf{y}_{0:t}) \in \mathbf{E}_t$ and $f_\theta(e) := \int \mathbb{S}_{m,\theta}(x_{0:t}, d\tilde{\mathbf{y}}_t)\mu(\mathbf{b}_t)(\text{id})$. Then $\mathbb{K}_{\theta,t}\mu(\mathbf{b}_t)(\text{id}) = \mathbb{C}_{m,\theta}f_\theta(x_{0:t})$ is a composition of a Markov kernel and a Lipschitz function, therefore Lipschitz. \square

Corollary 3 (A 4.1(iv) holds.). *Assume A B.9. For every $t \in \mathbb{N}$, $\theta \in \Theta$ and $N \in \mathbb{N}^*$*

$$\sup_{x_{0:t} \in \mathbf{X}_{0:t}} \|\mathbb{P}_{\theta_1,t}H - \mathbb{P}_{\theta_2,t}H\| \leq L_2^P \|\theta_1 - \theta_2\| ,$$

where

$$L_2^P = L_{t,N}^P + L_t^{\mathbb{K}} , \quad (\text{B.81})$$

with L^P and $L_t^{\mathbb{K}}$ are defined in (B.80) and (B.78).

Proof. Let $\tilde{f} : \mathbf{E}^{k-k_0} \ni (x_{0:t}[k_0 : k], \mathbf{x}_{0:t|t}[k_0 : k], \mathbf{b}_t[k_0 : k]) \mapsto (k - k_0)^{-1} \sum_{\ell=k_0+1}^k \mu(\mathbf{b}_t[\ell])(\text{id})$. As $\mathbb{K}_{\theta,t}$ depends only on the path, with a slight abuse of notation, we can define $f_\theta(x_{0:t}) := \mathbb{K}_{\theta,t}^{\otimes k-k_0}(\tilde{f})(x_{0:t})$. By proposition 13, we have that f_θ is Lipschitz with $L^f = L_t^{\mathbb{K}}$. Note that $\mathbb{P}_{\theta,t}H(x_{0:t}, \mathbf{y}_t) = K_{\theta,t}^{k_0}f_\theta(x_{0:t})$, therefore, by lemma 19 Lipschitz with constant $L^P + L_t^{\mathbb{K}}$. \square

C. Lipschitz properties

C.1 Lipschitz continuity of \mathbb{P}_θ ,

In this section we prove the following items:

- $\mathbb{C}_{m,\theta}(z_{0:m}, \cdot)$ is Lipschitz, see Appendix C.1.1
- $\mathbb{B}_{m,\theta}(\mathbf{x}_{0:m}, \cdot)$ is Lipschitz, see Appendix C.1.2
- $\int \mathbb{S}_{m,\theta}(\mathbf{x}_{0:m}, d\mathbf{b}_m)\mu(\mathbf{b}_m)(\text{Id})$ is Lipschitz, see Appendix C.1.3

The following technical lemma will be useful.

Lemma 10. *Let $\alpha \in]0, 1]$, $x \in \mathbb{R}_{\geq 0}$ and $\ell \in \mathbb{N}$. Then for all $\lambda_i \in \mathbb{R}_{\geq 0}$, $i \in \llbracket 0, \ell \rrbracket$, such that $\alpha \geq \prod_{i=0}^{\ell} (1 - \lambda_i x)$ it holds that $\alpha \geq 1 - x \sum_{i=0}^{\ell} \lambda_i$.*

Proof. Consider first the case where $x\lambda_i \leq 1$ for all $i \in \llbracket 0, \ell \rrbracket$. We prove the result by induction. The case $\ell = 0$ is straightforward. Assume now that the result holds for some $r \in \llbracket 0, \ell - 1 \rrbracket$. Then,

$$\begin{aligned} \prod_{i=0}^{r+1} (1 - \lambda_i x) &= (1 - \lambda_{r+1} x) \prod_{i=0}^r (1 - \lambda_i x) \geq (1 - \lambda_{r+1} x) (1 - x \sum_{i=0}^r \lambda_i) \\ &= 1 - x \sum_{i=0}^{r+1} \lambda_i + x^2 \sum_{i=0}^r \lambda_i \lambda_{r+1} \geq 1 - x \sum_{i=0}^{r+1} \lambda_i . \end{aligned}$$

Consider now the case where there is a index $j \in \llbracket 0, \ell \rrbracket$ such that $x\lambda_j \geq 1$. Then $\alpha \geq 0 \geq 1 - (\sum_{i=0}^{\ell} \lambda_i)x$. \square

We begin with some important definitions. Let P and Q be probability distributions on some common measurable space (X, \mathcal{X}) , and assume that these distributions admit densities p and q w.r.t some common reference measure λ . Let $\mathbb{M}[P, Q]$ denote a maximal coupling between P and Q . As in [Lindholm and Lindsten, 2018, Theorem 2], it is possible to explicitly construct one such maximal coupling by

$$\mathbb{M}[P, Q](d(x, y)) := \min\{p(x), q(x)\}\lambda(dx)\delta_x(dy) + \frac{[P(dx) - \min\{p(x), q(x)\}\lambda(dx)][Q(dy) - \min\{p(y), q(y)\}\lambda(dy)]}{1 - \lambda(\min\{p, q\})}. \quad (\text{C.82})$$

From this definition it follows that for continuous and discrete dominating measures λ ,

$$\int \mathbb{1}_{\{x=y\}} \mathbb{M}[P, Q] d(x, y) = \int \min\{p(x), q(x)\}\lambda(dx).$$

Moreover, for two Markov transition kernels K_1 and K_2 on (X, \mathcal{X}) , which are assumed to admit transition densities with respect to some common dominating measure, we let, for $(x_1, x_2) \in X^2$, $\mathbb{M}[K_1, K_2]((x_1, x_2), \cdot)$ denote the maximal coupling between the measures $K_1(x_1, \cdot)$ and $K_2(x_2, \cdot)$. Defined in this way, $\mathbb{M}[K_1, K_2]$ defines a Markov transition kernel on the product space $(X^2, \mathcal{X}^{\otimes 2})$

The following Lemma will be crucial in what follows.

Lemma 11. *(i) Let (μ_1, μ_2) be two probability measures admitting a density with respect to a common dominating measure and let (K_1, K_2) two Markov transition kernels also admitting transition densities with respect to some dominating measure. Then the probability measure*

$$\mathbb{M}[\mu_1, \mu_2] \mathbb{M}[K_1, K_2](d(x_1, x_2)) = \int \mathbb{M}[\mu_1, \mu_2](d(z_1, z_2)) \mathbb{M}[K_1, K_2]((z_1, z_2), d(x_1, x_2)),$$

is a coupling of $(\mu_1 K_1, \mu_2 K_2)$, and it holds that

$$\begin{aligned} & \int \mathbb{1}_{x_1=x_2} \mathbb{M}[\mu_1 K_1, \mu_2 K_2](d(x_1, x_2)) \\ & \geq \int \int \mathbb{1}_{z_1=z_2} \mathbb{1}_{x_1=x_2} \mathbb{M}[\mu_1, \mu_2](d(z_1, z_2)) \mathbb{M}[K_1, K_2]((z_1, z_2), d(x_1, x_2)). \end{aligned}$$

(ii) Let (μ_1, \dots, μ_n) and (ν_1, \dots, ν_n) be probability measures such that for all $i \in \llbracket 1, n \rrbracket$, μ_i and ν_i admit densities with respect to the same dominating measure. Then $\bigotimes_{i=1}^n \mathbb{M}[\mu_i, \nu_i]$ is a coupling of $\bigotimes_{i=1}^n \mu_i$ and $\bigotimes_{i=1}^n \nu_i$, and thus

$$\begin{aligned} & \int \prod_{i=1}^n \mathbb{1}_{x_i=y_i} \mathbb{M} \left[\bigotimes_{i=1}^n \mu_i, \bigotimes_{i=1}^n \nu_i \right] (d(x_1, \dots, x_n, y_1, \dots, y_n)) \\ & \geq \int \prod_{i=1}^n \mathbb{1}_{x_i=y_i} \bigotimes_{i=1}^n \mathbb{M}[\mu_i, \nu_i] (d(x_1, \dots, x_n, y_1, \dots, y_n)). \end{aligned}$$

Proof. It is enough to show that $\mathbb{M}[\mu_1, \mu_2] \mathbb{M}[K_1, K_2]$ admits $\mu_1 K_1$ and $\mu_2 K_2$ as marginal distributions. This follows immediately from the fact that $\mathbb{M}[\mu_1, \mu_1]$ and $\mathbb{M}[K_1, K_2]$ admit the right marginal

distributions; indeed,

$$\begin{aligned}
& \mathbb{M}[\mu_1, \mu_2] \mathbb{M}[K_1, K_2](\mathbf{X} \times A) \\
&= \int \mathbb{M}[\mu_1, \mu_2](dz_1, d_2) \mathbb{M}[K_1, K_2](z_1, z_2, d(x_1, x_2)) \mathbb{1}_{\mathbf{X} \times A}(x_1, x_2) \mathbb{1}_{\mathbf{X}^2}(z_1, z_2) \\
&= \int \mathbb{M}[\mu_1, \mu_2](dz_1, d_2) K_2(z_2, A) \\
&= \int \mu_2(dz_2) K_2(z_2, A) \\
&= \mu_2 K_2(A).
\end{aligned}$$

The derivation for the first marginal distribution follows similarly. For the second point, since $\mathbb{M}[\mu_1, \mu_2] \mathbb{M}[K_1, K_2]$ is a coupling of $(\mu_1 K_1, \mu_2 K_2)$ and $\mathbb{M}[\mu_1 K_1, \mu_2 K_2]$ is the maximal coupling, we have that

$$\begin{aligned}
& \int \mathbb{1}_{x_1=x_2} \mathbb{M}[\mu_1 K_1, \mu_2 K_2](d(x_1, x_2)) \\
&\geq \iint \mathbb{1}_{x_1=x_2} \mathbb{M}[\mu_1, \mu_2](d(z_1, z_2)) \mathbb{M}[K_1, K_2](z_1, z_2; d(x_1, x_2)) \\
&\geq \iint \mathbb{1}_{x_1=x_2} \mathbb{1}_{z_1=z_2} \mathbb{M}[\mu_1, \mu_2](d(z_1, z_2)) \mathbb{M}[K_1, K_2](z_1, z_2; d(x_1, x_2)).
\end{aligned}$$

The proof of the second item follows similarly. \square

C.1.1 $\theta \mapsto \mathbb{C}_{m,\theta}$ is Lipschitz.

We proceed by a coupling method that is inspired by [Lindholm and Lindsten, 2018, Theorem 2]. The coupling we consider is that where the *selection* and *mutation* steps of the particle filter are respectively coupled maximally.

Algorithm 6 Coupling $\mathbb{C}_{m,\theta}$

Data: $\theta_1, \theta_2, \zeta_{0:m}$

Result: $\mathbf{x}_{0:m,1}, \mathbf{x}_{0:m,2}$

23 draw $\mathbf{x}_{0,1}, \mathbf{x}_{0,2} \sim \mathbb{M}[\boldsymbol{\eta}_0 \langle \zeta_0 \rangle, \boldsymbol{\eta}_0 \langle \zeta_0 \rangle]$

24 **for** $s \leftarrow 1$ **to** t **do**

25 \lfloor draw $(\mathbf{x}_{s,1}, \mathbf{x}_{s,2}) \sim \mathbb{M}[\mathbf{M}_{s-1,\theta_1} \langle \zeta_s \rangle(\mathbf{x}_{s-1,1}, \cdot), \mathbf{M}_{s-1,\theta_2} \langle \zeta_s \rangle(\mathbf{x}_{s-1,2}, \cdot)]$

First, let us prove that the one step *selection–mutation* kernel is Lipschitz.

Lemma 12. For all $t \in \mathbb{N}$, $\mathbf{x}_{t-1} \in \mathbf{X}_{t-1}$ and $(\theta_1, \theta_2) \in \Theta^2$,

$$\int \mathbb{1}_{\{x_1=x_2\}} \mathbb{M}[\Phi_{t-1,\theta_1}(\mu(\mathbf{x}_{t-1})), \Phi_{t-1,\theta_2}(\mu(\mathbf{x}_{t-1}))](d(x_1, x_2)) \geq 1 - \frac{\sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot))}{N \bar{\tau}_n} \|\theta_1 - \theta_2\|. \tag{C.83}$$

Proof. By **A 3.1(i)** and **A 4.1(iii)**,

$$\begin{aligned}
& \int \mathbb{1}_{\{x_1=x_2\}} \mathbb{M} [\Phi_{t-1,\theta_1}(\mu(\mathbf{x}_{t-1})), \Phi_{t-1,\theta_2}(\mu(\mathbf{x}_{t-1}))] (d(x_1, x_2)) \\
&= \int \min \left(\sum_{i=1}^N \frac{q_{t-1,\theta_1}(x_{t-1}^i, x)}{\sum_{j=1}^N g_{t-1,\theta_1}(x_{t-1}^j)}, \sum_{i=1}^N \frac{q_{t-1,\theta_2}(x_{t-1}^i, x)}{\sum_{j=1}^N g_{t-1,\theta_2}(x_{t-1}^j)} \right) \lambda_t(dx) \\
&\geq \sum_{j=1}^N \int \min \left(\frac{q_{t-1,\theta_1}(x_{t-1}^i, x)}{\sum_{j=1}^N g_{t-1,\theta_1}(x_{t-1}^j)}, \frac{q_{t-1,\theta_2}(x_{t-1}^i, x)}{\sum_{j=1}^N g_{t-1,\theta_2}(x_{t-1}^j)} \right) \lambda_t(dx) \\
&\geq \frac{1}{\sum_{j=1}^N \max(g_{t-1,\theta_1}(x_{t-1}^j), g_{t-1,\theta_2}(x_{t-1}^j))} \sum_{j=1}^N \int \min(q_{t-1,\theta_1}(x_{t-1}^j, x), q_{t-1,\theta_2}(x_{t-1}^j, x)) \lambda_t(dx) \\
&\geq \frac{\sum_{j=1}^N \max(g_{t-1,\theta_1}(x_{t-1}^j), g_{t-1,\theta_2}(x_{t-1}^j)) - \sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot)) \|\theta_1 - \theta_2\|}{\sum_{j=1}^N \max(g_{t-1,\theta_1}(x_{t-1}^j), g_{t-1,\theta_2}(x_{t-1}^j))} \\
&\geq 1 - \frac{\sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot))}{N\bar{\tau}_n} \|\theta_1 - \theta_2\|,
\end{aligned}$$

where we have used that

$$\begin{aligned}
\int \max(q_{t-1,\theta_1}(x_{t-1}^i, x), q_{t-1,\theta_2}(x_{t-1}^i, x)) \lambda_t(dx) &\geq \max \left(\int q_{t-1,\theta_1}(x_{t-1}^i, x) \lambda_t(dx), \int q_{t-1,\theta_2}(x_{t-1}^i, x) \lambda_t(dx) \right) \\
&\geq \max(g_{t-1,\theta_1}(x_{t-1}^i), g_{t-1,\theta_2}(x_{t-1}^i)).
\end{aligned}$$

□

Lemma 13. For all $t \in \mathbb{N}$, $\mathbf{x}_{t-1} \in \mathbf{X}_{t-1}$, $z \in \mathbf{X}_t$ and $(\theta_1, \theta_2) \in \Theta^2$,

$$\|\mathbf{M}_{t-1,\theta_1}\langle z \rangle(\mathbf{x}_{t-1}, \cdot) - \mathbf{M}_{t-1,\theta_2}\langle z \rangle(\mathbf{x}_{t-1}, \cdot)\|_{\text{TV}} \leq L_{t-1}^{\mathbf{M}}(\mathbf{x}_{t-1}) \|\theta_1 - \theta_2\|$$

where $L_{t-1}^{\mathbf{M}}(\mathbf{x}_{t-1}) = (1 - N^{-1})\bar{\tau}_{t-1}^{-1} \sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot))$.

Proof. Let us denote by $\mathbb{U}[[1, n]]$ the uniform distribution on $[[1, n]]$. By definition of the kernel $\mathbf{M}_{t-1,\theta}\langle z \rangle$, we have that

$$\mathbf{M}_{t-1,\theta}\langle z \rangle(\mathbf{x}_{t-1}, d\mathbf{x}_t) = \int \mathbb{U}[[1, n]](dj) \{\Phi_{t-1}(\mu(\mathbf{x}_{t-1}))^{\otimes j} \otimes \delta_z \otimes \Phi_{t-1}(\mu(\mathbf{x}_{t-1}))^{\otimes (N-j-1)}\} (d\mathbf{x}_t)$$

and thus, applying the two items of Lemma 11 combined with the fact that $\mathbb{M}[\mu, \mu](d(x_1, x_2)) =$

$\mu(dx_1)\delta_{x_1}(dx_2)$ for any probability measure μ , we get that

$$\begin{aligned}
& \int \mathbb{1}_{\{x_{t,1}=x_{t,2}\}} \mathbb{M}[\mathbf{M}_{t-1,\theta_1}\langle z\rangle(\mathbf{x}_{t-1}, \cdot), \mathbf{M}_{t-1,\theta_2}\langle z\rangle(\mathbf{x}_{t-1}, \cdot)] d(\mathbf{x}_{t,1}, \mathbf{x}_{t,2}) \\
& \geq \int \mathbb{1}_{x_{t,1}=x_{t,2}, i_1=i_2} \mathbb{M}[\mathbb{U}[1, n], \mathbb{U}[1, n]](d(i_1, i_2)) \\
& \quad \times \mathbb{M}[\Phi_{t-1,\theta_1}(\mu(\mathbf{x}_{t-1})), \Phi_{t-1,\theta_2}(\mu(\mathbf{x}_{t-1}))]^{\otimes i_1} \otimes \mathbb{M}[\delta_z, \delta_z] \\
& \quad \otimes \mathbb{M}[\Phi_{t-1,\theta_1}(\mu(\mathbf{x}_{t-1})), \Phi_{t-1,\theta_2}(\mu(\mathbf{x}_{t-1}))]^{\otimes N-i_1-1} d(\mathbf{x}_{t,1}, \mathbf{x}_{t,2}) \\
& = \frac{1}{N} \sum_{i=1}^N \int \prod_{k=1, k \neq i}^n \mathbb{1}_{x_{t,1}^i=x_{t,2}^i} \mathbb{M}[\Phi_{t-1,\theta_1}(\mu(\mathbf{x}_{t-1})), \Phi_{t-1,\theta_2}(\mu(\mathbf{x}_{t-1}))](d(x_{t,1}^i, x_{t,2}^i)) \\
& \geq \left(1 - \frac{\sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot))}{N\bar{\tau}_{t-1}} \|\theta_1 - \theta_2\|\right)^{N-1} \\
& \geq 1 - \frac{N-1}{\bar{\tau}_{t-1}N} \sum_{i=1}^N \lambda_t(L_{t-1}^q(x_{t-1}^i, \cdot)) \|\theta_1 - \theta_2\|.
\end{aligned}$$

where we have applied Lemma 12 in the penultimate line and Lemma 10 in the last one. \square

Lemma 14. For every $t \in \mathbb{N}^*$, there exists $L_t^{\mathbb{C}} \in \mathcal{M}(\mathcal{X}_{0:t})$ such that

$$\|\mathbb{C}_{t,\theta_1}(z_{0:t}) - \mathbb{C}_{t,\theta_2}(z_{0:t})\|_{\text{TV}} \leq L_t^{\mathbb{C}}(z_{0:t}) \|\theta_1 - \theta_2\|, \quad (\text{C.84})$$

where $L_t^{\mathbb{C}}(z_{0:t}) = \sup_{\theta} \mathbb{C}_{t,\theta} \left[\sum_{i=0}^{t-1} L_i^{\mathbb{M}} \right] (z_{0:t})$. Under **AB.9(i)**, we obtain that $\|L_t^{\mathbb{C}}\|_{\infty} \leq (N-1) \sum_{\ell=0}^{t-1} \bar{\tau}_{\ell} \|L_{\ell}^q\|_{\infty}$.

Proof. This is a direct application of lemma 20. \square

C.1.2 $\theta \mapsto \mathbb{B}_{t,\theta}(\mathbf{x}_{0:t}, \cdot)$ is Lipschitz

We start by recalling the definition of \mathbb{B}_m

$$\mathbb{B}_{t,\theta} : \mathbf{X}_{0:t} \times \mathcal{X}_{0:t} \ni (\mathbf{x}_{0:t}, A) \mapsto \int \cdots \int \mathbb{1}_A(x_{0:t}) \left(\prod_{s=0}^{t-1} \overleftarrow{Q}_{s,\mu(\mathbf{x}_s)}(x_{s+1}, dx_s) \right) \mu(\mathbf{x}_t)(dx_t). \quad (\text{C.85})$$

Lemma 15. For all $s \in \llbracket 0, t \rrbracket$, $x_{t+1} \in \mathbf{X}_{t+1}$, $\mathbf{x}_t \in \mathbf{X}_t$ and $(\theta_1, \theta_2) \in \Theta^2$

$$\left\| \overleftarrow{Q}_{s,\mu(\mathbf{x}_s),\theta_1}(x_{s+1}, \cdot) - \overleftarrow{Q}_{s,\mu(\mathbf{x}_s),\theta_2}(x_{s+1}, \cdot) \right\|_{\text{TV}} \leq L_s^{\overleftarrow{Q}}(x_{s+1}, \mathbf{x}_s) \|\theta_1 - \theta_2\|. \quad (\text{C.86})$$

with $L_s^{\overleftarrow{Q}}(x_{s+1}, \mathbf{x}_s) = (N\bar{\tau}_s\bar{\sigma}_s)^{-1} \sum_{i=1}^N L_s^q(x_s^i, x_{s+1})$. Under **AB.9(i)**, we have $\|L_m^{\overleftarrow{Q}}\|_{\infty} = (\bar{\tau}_m\bar{\sigma}_m)^{-1} \|L_m^q\|_{\infty}$.

Proof. Note that $\overleftarrow{Q}_{t,\mu(\mathbf{x}_t)}(x_{t+1}, \cdot) = \sum_{\ell=1}^N \frac{q_t(x_t^{\ell}, x_{t+1})}{\sum_{\ell'=1}^N q_t(x_t^{\ell'}, x_{t+1})} \delta_{x_t^{\ell}}$. Therefore, similarly to the proof of Lemma 12,

$$\begin{aligned}
& \int \mathbb{1}_{\{x_{t,1}=x_{t,2}\}} \mathbb{M} \left[\overleftarrow{Q}_{t,\mu(\mathbf{x}_t),\theta_1}(x_{t+1}, \cdot), \overleftarrow{Q}_{t,\mu(\mathbf{x}_t),\theta_2}(x_{t+1}, \cdot) \right] d(x_{t,1}, x_{t,2}) \\
& \geq \frac{\sum_{\ell=1}^N \max(q_{t,\theta_1}(x_t^{\ell}, x_{t+1}), q_{t,\theta_2}(x_t^{\ell}, x_{t+1})) - L_t^q(x_t^{\ell}, x_{t+1}) \|\theta_1 - \theta_2\|}{\sum_{\ell=1}^N \max(q_{t,\theta_1}(x_t^{\ell}, x_{t+1}), q_{t,\theta_2}(x_t^{\ell}, x_{t+1}))} \\
& \geq 1 - \frac{\sum_{\ell=1}^N L_t^q(x_t^{\ell}, x_{t+1})}{N\bar{\tau}_t\bar{\sigma}_t} \|\theta_1 - \theta_2\|.
\end{aligned}$$

\square

Lemma 16. For all $t \in \mathbb{N}$, $\mathbf{x}_{0:t} \in \mathbf{X}_{0:t}$ and $(\theta_1, \theta_2) \in \Theta^2$

$$\|\mathbb{B}_{t,\theta_1}(\mathbf{x}_{0:t}, \cdot) - \mathbb{B}_{t,\theta_2}(\mathbf{x}_{0:t}, \cdot)\|_{\text{TV}} \leq L_t^{\mathbb{B}}(\mathbf{x}_{0:t}) \|\theta_1 - \theta_2\| \quad (\text{C.87})$$

where $L_t^{\mathbb{B}}(\mathbf{x}_{0:t}) = \sup_{\theta} \mathbb{B}_t \left[\sum_{i=0}^{t-1} L_i^{\overleftarrow{Q}} \right] (\mathbf{x}_{0:t})$. Under **A B.9(i)**, we have that $\|L_t^{\mathbb{B}}\|_{\infty} = \sum_{i=0}^{t-1} (\bar{\tau}_i \bar{\sigma}_i)^{-1} \|L_i^q\|_{\infty}$.

Proof. Apply lemma 19 and lemma 15. \square

C.1.3 $\theta \mapsto \int \mathbb{S}_{t,\theta}(\mathbf{x}_{0:t}, d\mathbf{b}_t) \mu(\mathbf{b}_t)(\text{id})$ is Lipschitz

Define the backward ancestors kernel

$$\mathcal{B}_{\theta,t} : \mathbf{X}_{t+1} \times \mathbf{X}_t \times \sigma(\llbracket 1, N \rrbracket) \mapsto \int \mathbb{1}_A(\tilde{j}) \left(\sum_{\ell=1}^N \frac{q_t(x_t^{\ell}, x_{t+1})}{\sum_{\ell'=1}^N q_t(x_t^{\ell'}, x_{t+1})} \delta_{\ell}(d\tilde{j}) \right).$$

Lemma 17. ($\mathcal{B}_{\theta,t}$ is Lipschitz) For every $m \in \llbracket 0, t \rrbracket$, there exists $L_m^{BK} \in \mathbf{M}(\mathcal{X}_{m:m+1})$ such that

$$\|\mathcal{B}_{\theta_1,m}(x_{m+1}, \mathbf{x}_m) - \mathcal{B}_{\theta_2,m}(x_{m+1}, \mathbf{x}_m)\|_{\text{TV}} \leq L_m^{\overleftarrow{Q}}(x_{m+1}, \mathbf{x}_m) \|\theta_1 - \theta_2\|, \quad (\text{C.88})$$

where $L_m^{\overleftarrow{Q}}$ is defined in Lemma 15

Proof. $\mathcal{B}_{\theta,s}$ is the index version of the kernel (C.85) and thus it is Lipschitz with the same constant. \square

Proposition 14. For every $m \in \llbracket 0, t \rrbracket$, we have that

$$\left| \int \mathbb{C}_m \mathbb{S}_{m,\theta}(z_{0:m}, d\mathbf{b}_m) \mu(\mathbf{b}_m)(\text{Id}) \right| \leq \sum_{\ell=0}^{m-1} s_{\ell}^{\infty} \quad (\text{C.89})$$

and

$$\left| \int \mathbb{S}_{m,\theta_1}(\mathbf{x}_{0:m}, d\mathbf{b}_m) \mu(\mathbf{b}_m)(\text{Id}) - \int \mathbb{S}_{m,\theta_2}(\mathbf{x}_{0:m}, d\mathbf{b}_m) \mu(\mathbf{b}_m)(\text{Id}) \right| \leq L_m^{\mathbb{S}\mu}(\mathbf{x}_{0:m}) \|\theta_1 - \theta_2\|. \quad (\text{C.90})$$

where $L_m^{\mathbb{S}\mu}(\mathbf{x}_{0:m}) = N^{-1} \sum_{i=1}^N L_m^B(x_m^k, \mathbf{x}_{0:m})$ and L_m^B is defined recursively as

$$L_{m+1}^B(x_{m+1}^k, \mathbf{x}_{0:m}) = L_m^{\overleftarrow{Q}}(x_{m+1}^k, \mathbf{x}_m) \sum_{\ell=0}^m s_{\ell}^{\infty} + \int \mathcal{B}_{\theta,m}(x_{m+1}^k, \mathbf{x}_m, d\mathbf{J}) \left\{ L_m^{\mathbb{S}}(x_m^{\mathbf{J}}, x_{m+1}^k) + L_m^B(x_m^{\mathbf{J}}, \mathbf{x}_{0:m-1}) \right\}. \quad (\text{C.91})$$

In particular, under **A B.9**, we have that $L_m^B \leq \sum_{j=1}^m \|L_j^{\overleftarrow{Q}}\|_{\infty} \left[\sum_{\ell=0}^{m-1} s_{\ell}^{\infty} \right] + \sum_{j=1}^m \|L_j^{\mathbb{S}}\|_{\infty}$.

Proof. Consider the following kernels,

$$\tilde{\mathbb{S}}_{m,\theta}(\mathbf{x}_{0:m+1}, d(\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M}) := \prod_{\ell=0}^m \prod_{k=1}^N \tilde{\mathbb{S}}_{\ell,\theta}(x_{\ell+1}^k, \mathbf{x}_{\ell}, d(\mathbf{J}_{\ell}^{k,j})_{j=1}^M), \quad (\text{C.92})$$

$$\tilde{\mathbb{S}}_{\ell,\theta}(x_{\ell+1}^k, \mathbf{x}_{\ell}, d(\mathbf{J}_{\ell}^{k,j})_{j=1}^M) := \prod_{j=1}^M \mathcal{B}_{\theta,\ell}(x_{\ell+1}^k, \mathbf{x}_{\ell}, d\mathbf{J}_{\ell}^{k,j}). \quad (\text{C.93})$$

Define for all $k \in [1 : N]$, $m \in \mathbb{N}_{>0}$,

$$B_{m+1,k} : \theta \mapsto \int \tilde{\mathbb{S}}_{m,\theta}(\mathbf{x}_{0:m+1}, d(\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M}) b_{m+1}^k(\mathbf{x}_{0:m+1}, (\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M}),$$

where $b_{m+1}^k(\mathbf{x}_{0:m+1}, (\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M})$ is defined recursively as

$$b_{m+1}^k(\mathbf{x}_{0:m+1}, (\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M}) = M^{-1} \sum_{\ell=1}^M b_m^{J_m^{k,\ell}}(\mathbf{x}_{0:m}, (\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_{m-1}^{i,j})_{i=1,j=1}^{N,M}) + s_{m,\theta}(x_m^{J_m^{k,\ell}}, x_{m+1}^k).$$

For notational convenience, we henceforth drop the arguments and simply write b_{m+1}^k .

We herebelow show that $B_{m+1,k}$ is Lipschitz with constant $L_m^B(x_{m+1}^k, \mathbf{x}_m)$ and bounded by $\sum_{\ell=0}^{m-1} s_\ell^\infty$. For $m > 2$ and $k \in [1 : N]$,

$$\begin{aligned} B_{m+1,k}(\theta) &= \int \tilde{\mathbb{S}}_{m,\theta}(\mathbf{x}_{0:m+1}, d(\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_m^{i,j})_{i=1,j=1}^{N,M}) b_{m+1}^k \\ &= \int \cdots \int \tilde{\mathbb{S}}_{m-1,\theta}(\mathbf{x}_{0:m}, d(\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_{m-1}^{i,j})_{i=1,j=1}^{N,M}) \tilde{\mathbb{S}}_{m,\theta}(x_{m+1}^k, \mathbf{x}_m, d(\mathbf{J}_m^{k,j})_{j=1}^M) \\ &\quad \times \left\{ M^{-1} \sum_{\ell=1}^M b_m^{J_m^{k,\ell}} + s_{m,\theta}(x_m^{J_m^{k,\ell}}, x_{m+1}^k) \right\} \\ &= \int \cdots \int \tilde{\mathbb{S}}_{m,\theta}(x_{m+1}^k, \mathbf{x}_m, d\{\mathbf{J}_m^{k,j}\}_{j=1}^M) \left[M^{-1} \sum_{\ell=1}^M \left\{ s_{m,\theta}(x_m^{J_m^{k,\ell}}, x_{m+1}^k) \right. \right. \\ &\quad \left. \left. + \int \tilde{\mathbb{S}}_{m-1,\theta}(\mathbf{x}_{0:m}, d(\mathbf{J}_0^{i,j}, \dots, \mathbf{J}_{m-1}^{i,j})_{i=1,j=1}^{N,M}) b_m^{J_m^{k,\ell}} \right\} \right] \\ &= \int \cdots \int \tilde{\mathbb{S}}_{m,\theta}(x_{m+1}^k, \mathbf{x}_m, d(\mathbf{J}_m^{k,j})_{j=1}^M) \left[M^{-1} \sum_{\ell=1}^M \left\{ s_{m,\theta}(x_m^{J_m^{k,\ell}}, x_{m+1}^k) + B_{m,J_m^{k,\ell}}(\theta) \right\} \right] \\ &= \int \mathcal{B}_{\theta,m}(x_{m+1}^k, \mathbf{x}_m, d\mathbf{J}) \left\{ s_{m,\theta}(x_m^J, x_{m+1}^k) + B_{m,J}(\theta) \right\} \end{aligned}$$

Applying the induction hypothesis conditionally on $\mathbf{J}_m^{k,\ell}$, $B_{m,J_m^{k,\ell}}$ is Lipschitz with constant $L_m^B(x_m^{J_m^{k,\ell}}, \mathbf{x}_{0:m-1})$ and thus the Lipschitz constant of $B_{m+1,k}$ is

$$L_{m+1}^B(x_{m+1}^k, \mathbf{x}_{0:m}) = L_m^{\tilde{Q}}(x_{m+1}^k, \mathbf{x}_m) \sum_{\ell=0}^m s_\ell^\infty + \int \mathcal{B}_{\theta,m}(x_{m+1}^k, \mathbf{x}_m, d\mathbf{J}) \left\{ L_m^s(x_m^J, x_{m+1}^k) + L_m^B(x_m^J, \mathbf{x}_{0:m-1}) \right\}. \quad (\text{C.94})$$

where we have used the fact that $\mathcal{B}_{\theta,m}$ and $s_{m,\theta}$ are also Lipschitz. Again by induction $B_{m+1,k}$ is bounded uniformly by $\sum_{\ell=0}^m s_\ell^\infty$. The induction is concluded by noting that for the base case $m = 0$, $\beta_m^k = 0$ for all $k \in \mathbb{N}$ and thus the result holds.

It now remains to check that for all $\theta \in \Theta$, $m \in [0, t]$ and $k \in [1 : N]$,

$$B_{m,k}(\theta) = \int \mathbb{S}_m(\mathbf{x}_{0:m}, d\mathbf{b}_m) b_m^k.$$

Again, we proceed by induction.

$$\begin{aligned}
& \int \mathbb{S}_m(\mathbf{x}_{0:m}, d\mathbf{b}_m) b_m^k \\
&= \int \cdots \int \mathbb{S}_{m-1}(\mathbf{x}_{0:m-1}, d\mathbf{b}_{m-1}) \mathbf{S}_m(\mathbf{b}_{m-1}, \mathbf{x}_{m-1:m}, d\mathbf{b}_m) b_m^k \\
&= \int \cdots \int \mathbb{S}_{m-1}(\mathbf{x}_{0:m-1}, d\mathbf{b}_{m-1}) \\
&\quad \times \prod_{j=1}^M \left(\sum_{p=1}^N \frac{q_{m-1}(x_{m-1}^p, x_m^k)}{\sum_{\ell=1}^N q_{m-1}(x_{m-1}^\ell, x_m^k)} \delta_{x_{m-1}^p, b_{m-1}^p} (d(\tilde{x}_{m-1}^{k,j}, \tilde{b}_{m-1}^{k,j})) \right) \\
&\quad \times \left[M^{-1} \sum_{n=1}^M \left\{ \tilde{b}_{m-1}^{k,n} + s_{m,\theta}(\tilde{x}_{m-1}^{k,n}, x_m^k) \right\} \right] \\
&= \int \cdots \int \mathbb{S}_{m-1}(\mathbf{x}_{0:m-1}, d\mathbf{b}_{m-1}) \\
&\quad \times \prod_{j=1}^M \left(\sum_{p=1}^N \frac{q_{m-1}(x_{m-1}^p, x_m^k)}{\sum_{\ell=1}^N q_{m-1}(x_{m-1}^\ell, x_m^k)} \delta_p(dJ_{m-1}^{k,j}) \right) \left[M^{-1} \sum_{n=1}^M \left\{ b_{m-1}^{j^{k,n}} + s_{m,\theta}(x_{m-1}^{j^{k,n}}, x_m^k) \right\} \right] \\
&= \int \cdots \int \tilde{\mathbf{S}}_{m,\theta}(x_{m-1}^k, \mathbf{x}_{\ell-1}, d(J_{\ell-1}^{k,j})_{j=1}^M) \\
&\quad \times \left[M^{-1} \sum_{\ell=1}^M \left\{ s_{m,\theta}(x_{m-1}^{j^{k,\ell}}, x_m^k) + \mathbb{S}_{m-1}(\mathbf{x}_{0:m-1}, d\mathbf{b}_{m-1}) b_{m-1}^{j^{k,\ell}} \right\} \right] \\
&= \int \cdots \int \tilde{\mathbf{S}}_{m,\theta}(x_{m-1}^k, \mathbf{x}_{\ell-1}, d(J_{\ell-1}^{k,j})_{j=1}^M) \\
&\quad \times \left[M^{-1} \sum_{\ell=1}^M \left\{ s_{m,\theta}(x_{m-1}^{j^{k,\ell}}, x_m^k) + \int \mathbb{S}_{m-1}(\mathbf{x}_{0:m-1}, d\mathbf{b}_{m-1}) b_{m-1}^{j^{k,\ell}} \right\} \right] \\
&= \int \cdots \int \tilde{\mathbf{S}}_{m,\theta}(x_{m-1}^k, \mathbf{x}_{\ell-1}, d(J_{\ell-1}^{k,j})_{j=1}^M) \left[M^{-1} \sum_{\ell=1}^M \left\{ s_{m,\theta}(x_{m-1}^{j^{k,\ell}}, x_m^k) + B_{m-1, J_{m-1}^{k,\ell}}(\theta) \right\} \right] \\
&= B_{m,k}(\theta)
\end{aligned}$$

The proof is finalized by noting that

$$\int \mathbb{S}_m(\mathbf{x}_{0:m}, d\mathbf{b}_m) \mu(\mathbf{b}_m)(\text{Id}) = N^{-1} \sum_{k=1}^N B_{m,k}(\theta)$$

and thus it is Lipschitz with constant $L_m^{\mathbb{S}\mu}(\mathbf{x}_{0:m}) = N^{-1} \sum_{i=1}^N L_m^B(x_m^i, \mathbf{x}_{m-1})$. \square

C.2 Lipschitz properties of Markov Kernels

Lemma 18 (Composition of ergodic Lipschitz kernels is Lipschitz). *Let P_θ be a Markov kernel over $X \times \mathcal{Y}$ that is uniformly π -geometrically ergodic for any θ with contraction constant ρ independent of θ and such that there exists $L_\rho > 0$ such that for every $x \in X$*

$$\|P_{\theta_0}(x, \cdot) - P_{\theta_1}(x, \cdot)\|_{\text{TV}} \leq L_\rho \|\theta_0 - \theta_1\|.$$

Then, for all $k > 0$

$$\|P_{\theta_0}^k(x, \cdot) - P_{\theta_1}^k(x, \cdot)\|_{\text{TV}} \leq \frac{L_\rho}{1-\rho} \|\theta_0 - \theta_1\|.$$

Proof. We use the following decomposition borrowed from [Fort et al., 2011]. For any $k \geq 1$,

$$P_{\theta_0}^k f - P_{\theta_1}^k f = \sum_{j=0}^{k-1} P_{\theta_0}^j (P_{\theta_0} - P_{\theta_1}) (P_{\theta_1}^{k-j-1} f - \pi f).$$

Then, for any f s.t. $\|f\|_\infty \leq 1$ and $x \in X$,

$$\begin{aligned} |P_{\theta_0}^k f(x) - P_{\theta_1}^k f(x)| &\leq \sum_{j=0}^{k-1} \left| \int P_{\theta_0}^j(x, dy) \sup_{z \in X} |P_{\theta_1}^{k-j-1} f(z) - \pi f| \right| L_P \|\theta_0 - \theta_1\| \\ &\leq L_P \left(\sum_{j=0}^{k-1} \rho^{k-j-1} \right) \|\theta_0 - \theta_1\| \\ &\leq \frac{L_P}{1-\rho} \|\theta_0 - \theta_1\|. \end{aligned}$$

□

Lemma 19 (Composition of Lipschitz kernels is lipschitz). *Let P_θ, Q_θ be two kernels defined over $X \times \mathcal{Y}$ and $Y \times \mathcal{Z}$ such that for ever $x \in X, y \in Y$ there are $L_p \in \mathbf{M}(X), L_q \in \mathbf{M}(Y)$ that satisfy*

$$\|P_{\theta_0}(x, \cdot) - P_{\theta_1}(x, \cdot)\|_{\text{TV}} \leq L_p(x) \|\theta_0 - \theta_1\|$$

and

$$\|Q_{\theta_0}(y, \cdot) - Q_{\theta_1}(y, \cdot)\|_{\text{TV}} \leq L_q(y) \|\theta_0 - \theta_1\|.$$

Then

$$\|P_{\theta_0} Q_{\theta_0}(x, \cdot) - P_{\theta_1} Q_{\theta_1}(x, \cdot)\|_{\text{TV}} \leq L_{pq} \|\theta_0 - \theta_1\|,$$

where $L_{pq} = (\sup_\theta P_\theta L_q(x) + L_p(x) \sup_y \sup_\theta Q_\theta(y, Z))$.

Proof. Let $f \in \mathbf{M}$ such that $\|f\|_\infty \leq 1$.

$$\begin{aligned} \|P_{\theta_1} Q_{\theta_1} f - P_{\theta_2} Q_{\theta_2} f\| &\leq \|P_{\theta_1} [Q_{\theta_1} f - Q_{\theta_2} f]\| + \|(P_{\theta_1} - P_{\theta_2}) Q_{\theta_2} f\| \\ &\leq (P_{\theta_1} L_q(x) + L_p(x) \|Q_{\theta_2} f\|_\infty) \|\theta_1 - \theta_2\|. \end{aligned}$$

□

Corollary 4. *Let P_θ, Q_θ be two **Markov** kernels defined over $X \times \mathcal{Y}$ and $Y \times \mathcal{Z}$ such that for ever $x \in X, y \in Y$ there are $L_p \in \mathbf{M}(X), L_q \in \mathbf{M}(Y)$ that satisfy*

$$\|P_{\theta_0}(x, \cdot) - P_{\theta_1}(x, \cdot)\|_{\text{TV}} \leq L_p(x) \|\theta_0 - \theta_1\|$$

and

$$\|Q_{\theta_0}(y, \cdot) - Q_{\theta_1}(y, \cdot)\|_{\text{TV}} \leq L_q(y) \|\theta_0 - \theta_1\|.$$

Then

$$\|P_{\theta_0} Q_{\theta_0}(x, \cdot) - P_{\theta_1} Q_{\theta_1}(x, \cdot)\|_{\text{TV}} \leq L_{pq} \|\theta_0 - \theta_1\|,$$

where $L_{pq} = (\sup_\theta P_\theta L_q(x) + L_p(x))$.

Lemma 20 (Product of Lipschitz kernels is lipschitz). *Let P_θ, Q_θ be two markov kernels that are uniformly Lipschitz with constants L_P, L_Q . Then $P_\theta \otimes Q_\theta$ is uniformly Lipschitz with constant $L_P + L_Q$.*

Proof. Let $h_\theta : y \mapsto \int Q_\theta(y, dz) f(y, z)$. Then $(P_{\theta_i} \otimes Q_{\theta_i})(f) = P_{\theta_i}(h_{\theta_i})$ and the proof is similar to that of the previous Lemma since h_θ is Lipschitz with constant L_Q and $\|h_\theta\|_\infty \leq 1$. □




Figure 3: Output of the PPG roll-out estimator for the LGSSM. The curves describe the evolution of the bias with increasing k for different particle sample sizes N . The left and right panels correspond to $k_0 = k - 1$ and $k_0 = \lfloor k/2 \rfloor$, respectively.

D. Additional numerical results

D.1 PPG

D.2 Learning

For both experiments, all the parameters were initialized by sampling from a centered multivariate gaussian distribution with covariance matrix of $0.01I$. We have used the ADAM optimizer [Kingma and Ba, 2014] with a learning rate decay of $1/\sqrt{\ell}$ where ℓ is the iteration index, with a starting learning rate of 0.2. We rescale the gradients by T .

LGSSM For LGSSM we evaluated for fixed number of particles ($N = 64$) and number of gibbs iterations ($k = 8$) the influence of the burn-in phase (k_0) over the final distance obtained to the MLE estimator. Table 3 indicates that configurations with smaller k_0 perform better. A possible interpretation of this phenomenon is that, since between two gradient ascent iterates the conditioning path is being passed on, this conditioning path from a moment on makes the estimates less biased, so the importance of having k_0 high to have less bias vanishes, but the effect of augmenting the variance with k_0 is still shown, since the fact of having a conditioning particle from the right marginal does not affect the variance of the estimator, only its bias.

References

- [Anderson and Qiu, 1997] Anderson, G. D. and Qiu, S.-L. (1997). A monotonicity property of the gamma function. *Proc. Amer. Math. Soc.*, 125(11):3355–3362.
- [Andrieu and Doucet, 2003] Andrieu, C. and Doucet, A. (2003). Online Expectation–Maximization type algorithms for parameter estimation in general state space models. In *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, volume 6, pages 69–72.

Table 3: Distance to θ_{MLE} for each configuration in the LGSSM case.

Algorithm	N	k_0	k	D_{mle}
PPG	64	0	8	0.205 ± 0.013
PPG	64	1	8	0.213 ± 0.016
PPG	64	2	8	0.201 ± 0.010
PPG	64	3	8	0.201 ± 0.010
PPG	64	4	8	0.207 ± 0.012
PPG	64	5	8	0.212 ± 0.015
PPG	64	6	8	0.210 ± 0.017
PPG	64	7	8	0.211 ± 0.018

- [Andrieu et al., 2010a] Andrieu, C., Doucet, A., and Holenstein, R. (2010a). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342.
- [Andrieu et al., 2010b] Andrieu, C., Doucet, A., and Holenstein, R. (2010b). Particle Markov chain Monte Carlo methods (with discussion). *J. Roy. Statist. Soc. B*, 72:269–342.
- [Andrieu et al., 2018] Andrieu, C., Lee, A., and Vihola, M. (2018). Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. *Bernoulli*, 24(2):842–872.
- [Cappé, 2001] Cappé, O. (2001). Recursive computation of smoothed functionals of hidden Markovian processes using a particle approximation. *Monte Carlo Methods Appl.*, 7(1–2):81–92.
- [Cappé, 2011] Cappé, O. (2011). Online EM algorithm for hidden Markov models. *J. Comput. Graph. Statist.*, 20(3):728–749.
- [Cappé et al., 2007] Cappé, O., Godsill, S. J., and Moulines, E. (2007). An overview of existing methods and recent advances in sequential Monte Carlo. *IEEE Proceedings*, 95(5):899–924.
- [Cappé and Moulines, 2005] Cappé, O. and Moulines, E. (2005). On the use of particle filtering for maximum likelihood parameter estimation. In *European Signal Processing Conference (EUSIPCO)*, Antalya, Turkey.
- [Cappé et al., 2005] Cappé, O., Moulines, E., and Rydén, T. (2005). *Inference in Hidden Markov Models*. Springer.
- [Cappé et al., 2009] Cappé, O., Moulines, E., and Rydén, T. (2009). Inference in hidden markov models. In *Proceedings of EUSFLAT conference*, pages 14–16.
- [Chopin and Papaspiliopoulos, 2020] Chopin, N. and Papaspiliopoulos, O. (2020). *An Introduction to Sequential Monte Carlo*. Springer.
- [Chopin and Singh, 2015a] Chopin, N. and Singh, S. S. (2015a). On particle Gibbs sampling. *Bernoulli*, 21(3):1855–1883.
- [Chopin and Singh, 2015b] Chopin, N. and Singh, S. S. (2015b). On particle gibbs sampling. *Bernoulli*, 21(3):1855–1883.
- [Del Moral, 2004] Del Moral, P. (2004). *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications*. Springer.
- [Del Moral, 2013] Del Moral, P. (2013). *Mean Field Simulation for Monte Carlo Integration*. CRC Press.
- [Del Moral et al., 2010] Del Moral, P., Doucet, A., and Singh, S. S. (2010). A backward interpretation of Feynman–Kac formulae. *ESAIM: Mathematical Modelling and Numerical Analysis*, 44:947–975.

- [Del Moral and Jasra, 2018] Del Moral, P. and Jasra, A. (2018). A sharp first order analysis of Feynman–Kac particle models, part II: Particle Gibbs samplers. *Stoch. Proc. Appl.*, 128(1):354–371.
- [Del Moral et al., 2016] Del Moral, P., Kohn, R., and Patras, F. (2016). On particle Gibbs samplers. *Ann. Inst. H. Poincaré Probab. Statist.*, 52(4):1687–1733.
- [Douc et al., 2011] Douc, R., Garivier, A., Moulines, E., and Olsson, J. (2011). Sequential Monte Carlo smoothing for general state space hidden Markov models. *Ann. Appl. Probab.*, 21(6):1201–2145.
- [Douc and Moulines, 2008] Douc, R. and Moulines, E. (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *Ann. Statist.*, 36(5):2344–2376.
- [Douc et al., 2018] Douc, R., Moulines, E., Priouret, P., and Soulier, P. (2018). *Markov Chains*. Springer.
- [Douc et al., 2014] Douc, R., Moulines, E., and Stoffer, D. (2014). *Nonlinear time series: Theory, methods and applications with R examples*. CRC press.
- [Dubarry and Le Corff, 2013] Dubarry, C. and Le Corff, S. (2013). Non-asymptotic deviation inequalities for smoothed additive functionals in nonlinear state-space models. *Bernoulli*, 19(5B):2222 – 2249.
- [Fort et al., 2011] Fort, G., Moulines, E., and Priouret, P. (2011). Convergence of adaptive and interacting markov chain monte carlo algorithms. *The Annals of Statistics*, 39(6).
- [Gloaguen et al., 2022] Gloaguen, P., Le Corff, S., and Olsson, J. (2022). A pseudo-marginal sequential Monte Carlo online smoothing algorithm. *Bernoulli*, 28(4):2606–2633.
- [Glynn and Rhee, 2014] Glynn, P. W. and Rhee, C.-H. (2014). Exact estimation for markov chain equilibrium expectations. *Journal of Applied Probability*, 51(A):377–389.
- [Godsill et al., 2004] Godsill, S. J., Doucet, A., and West, M. (2004). Monte Carlo smoothing for non-linear time series. *J. Am. Statist. Assoc.*, 50:438–449.
- [Jacob et al., 2020] Jacob, P. E., O’Leary, J., and Atchadé, Y. F. (2020). Unbiased markov chain monte carlo methods with couplings. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(3):543–600.
- [Karimi et al., 2019] Karimi, B., Miasojedow, B., Moulines, E., and Wai, H.-T. (2019). Non-asymptotic analysis of biased stochastic approximation scheme. In Beygelzimer, A. and Hsu, D., editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1944–1974. PMLR.
- [Kingma and Ba, 2014] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization.
- [Lee et al., 2020] Lee, A., Singh, S. S., and Vihola, M. (2020). Coupled conditional backward sampling particle filter. *The Annals of Statistics*, 48(5):3066–3089.
- [Lindholm and Lindsten, 2018] Lindholm, A. and Lindsten, F. (2018). Learning dynamical systems with particle stochastic approximation em.
- [Lindsten et al., 2014a] Lindsten, F., Jordan, M. I., and Schön, T. B. (2014a). Particle Gibbs with ancestor sampling. *J. Mach. Learn. Res.*, 15(1):2145–2184.
- [Lindsten et al., 2014b] Lindsten, F., Jordan, M. I., and Schön, T. B. (2014b). Particle gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15(63):2145–2184.

- [Naesseth et al., 2020] Naesseth, C. A., Lindsten, F., and Blei, D. (2020). Markovian score climbing: Variational inference with $\text{kl}(p \rightarrow q)$.
- [Olsson and Westerborn, 2017] Olsson, J. and Westerborn, J. (2017). Efficient particle-based online smoothing in general hidden Markov models: The PaRIS algorithm. *Bernoulli*, 23(3):1951–1996.
- [Poyiadjis et al., 2005] Poyiadjis, G., Doucet, A., and Singh, S. S. (2005). Particle methods for optimal filter derivative: application to parameter estimation. In *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, pages v/925–v/928.
- [Poyiadjis et al., 2011] Poyiadjis, G., Doucet, A., and Singh, S. S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98(1):65–80.
- [Särkkä, 2013] Särkkä, S. (2013). *Bayesian Filtering and Smoothing*. Cambridge University Press.
- [Singh et al., 2017] Singh, S. S., Lindsten, F., and Moulines, E. (2017). Blocking strategies and stability of particle gibbs samplers. *Biometrika*, 104(4):953–969.
- [Whiteley, 2010] Whiteley, N. (2010). Discussion on particle markov chain monte carlo methods. pages 306–307.
- [Zhao et al., 2021] Zhao, Y., Nassar, J., Jordan, I., Bugallo, M., and Park, I. M. (2021). Streaming variational monte carlo.