



HAL
open science

On backward smoothing algorithms

Hai-Dang Dau, Nicolas Chopin

► **To cite this version:**

Hai-Dang Dau, Nicolas Chopin. On backward smoothing algorithms. *Annals of Statistics*, 2023, 51 (5), 10.1214/23-AOS2324 . hal-04417980

HAL Id: hal-04417980

<https://hal.science/hal-04417980v1>

Submitted on 6 Jan 2025

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.

ON BACKWARD SMOOTHING ALGORITHMS

HAI-DANG DAU & NICOLAS CHOPIN

ABSTRACT. In the context of state-space models, skeleton-based smoothing algorithms rely on a backward sampling step which by default has a $\mathcal{O}(N^2)$ complexity (where N is the number of particles). Existing improvements in the literature are unsatisfactory: a popular rejection sampling-based approach, as we shall show, might lead to badly behaved execution time; another rejection sampler with stopping lacks complexity analysis; yet another MCMC-inspired algorithm comes with no stability guarantee. We provide several results that close these gaps. In particular, we prove a novel non-asymptotic stability theorem, thus enabling smoothing with truly linear complexity and adequate theoretical justification. We propose a general framework which unites most skeleton-based smoothing algorithms in the literature and allows to simultaneously prove their convergence and stability, both in online and offline contexts. Furthermore, we derive, as a special case of that framework, a new coupling-based smoothing algorithm applicable to models with intractable transition densities. We elaborate practical recommendations and confirm those with numerical experiments.

1. INTRODUCTION

1.1. Background. A state-space model is composed of an unobserved Markov process X_0, \dots, X_T and observed data Y_0, \dots, Y_T . Given X_0, \dots, X_T , the data Y_0, \dots, Y_T are independent and generated through some specified emission distribution $Y_t|X_t \sim \mathbf{f}_t(dy_t|x_t)$. These models have wide-ranging applications (e.g. in biology, economics and engineering). Two important inference tasks related to state-space models are filtering (computing the distribution of X_t given Y_0, \dots, Y_t) and smoothing (computing the distribution of the whole trajectory (X_0, \dots, X_t) , again given all data until time t). Filtering is usually carried out through a particle filter, that is, a sequential Monte Carlo algorithm that propagates N weighted particles (realisations) through Markov and importance sampling steps; see [Chopin and Papaspiliopoulos \(2020\)](#) for a general introduction to state-space models (Chapter 2) and particle filters (Chapter 10).

This paper is concerned with skeleton-based smoothing algorithms, i.e. algorithms that approximate the smoothing distributions with empirical distributions based on the output of a particle filter (i.e. the locations and weights of the N particles at each time step). A simple example is genealogy tracking (empirically proposed in [Kitagawa, 1996](#) and theoretically analysed in [Del Moral and Miclo, 2001](#)) which keeps track of the ancestry (past states) of each particles. This smoother suffers from degeneracy: for t large enough, all the particles have the same ancestor at time 0.

The forward filtering backward smoothing (FFBS) algorithm ([Godsill et al., 2004](#)) has been proposed as a solution to this problem. Starting from the filtering approximation at time t , the algorithm samples successively particles at times $t-1$, $t-2$, etc. using backward kernels. Its theoretical properties, in particular the stability as $t \rightarrow \infty$, have been studied by [Del Moral et al. \(2010\)](#); [Douc et al. \(2011\)](#).

In many applications, one is mainly interested in approximating smoothing expectations of additive functions of the form

$$\mathbb{E}[\psi_0(X_0) + \psi_1(X_0, X_1) + \dots + \psi_t(X_{t-1}, X_t) | Y_0, \dots, Y_t]$$

for some functions ψ_0, \dots, ψ_t . Such expectations can be approximated in an online fashion by a procedure described in [Del Moral et al. \(2010\)](#). Inspired by this, the particle-based, rapid incremental smoother (PaRIS) algorithm of [Olsson and Westerborn \(2017\)](#) replaces some of the calculations with an additional layer of Monte Carlo approximation.

The backward sampling operation is central to both the FFBS and the PaRIS algorithms. The naive implementation has an $\mathcal{O}(N^2)$ cost. There have been numerous attempts at alleviating this problem in the literature, but, to our knowledge, they all lack formal support in terms of either computation complexity or stability as $T \rightarrow \infty$.

In the following five paragraphs, we elaborate on this limitation for each of the three major contenders, and we point out two related challenges with current backward sampling algorithms that we try to resolve in this article.

1.2. State of the art. First, [Douc et al. \(2011\)](#) proposed to use rejection sampling for the generation of backward indices in FFBS, and [Olsson and Westerborn \(2017\)](#) extended this technique to PaRIS. If the model has upper-and-lower bounded transition densities, this sampler has an $\mathcal{O}(N)$ expected execution time ([Douc et al. 2011](#), Proposition 2 and [Olsson and Westerborn 2017](#), Theorem 10). Unfortunately, most practical state space models (including linear Gaussian ones) violate this assumption, and the behaviour of the algorithm in this case is unclear. Empirically, it has been observed ([Taghavi et al. 2013](#); [Bunch and Godsill 2013](#); [Olsson and Westerborn 2017](#), Section 4.3) that in real examples, FFBS-reject and PaRIS-reject frequently suffer from low acceptance rates, in contrary to what users would expect from an algorithm with linear complexity. To cite [Bunch and Godsill \(2013\)](#), “[a]lthough theoretically elegant, the [...] algorithm has been found to suffer from such high rejection rates as to render it consistently slower than direct sampling implementation on problems with more than one state dimension”. To the best of our knowledge, no theoretical result has been put forward to formalise or quantify this bad behaviour.

Second, [Taghavi et al. \(2013\)](#) and [Olsson and Westerborn \(2017, Section 4.3\)](#) suggest putting a threshold on the number of rejection sampling trials to get more stable execution times. The thresholds are either chosen adaptively using a Kalman filter in [Taghavi et al. \(2013\)](#) or fixed at \sqrt{N} in [Olsson and Westerborn \(2017, Section 4.3\)](#). Although improvements are empirically observed, to the best of our knowledge, no theoretical analysis of the complexity of the resulting algorithm or formal justification of the proposed threshold is available.

Third, [Bunch and Godsill \(2013\)](#) use MCMC moves starting from the filtering ancestor instead of a full backward sampling step. Empirically, this algorithm seems to prevent the degeneracy associated with the genealogy tracking smoother using a very small number of MCMC steps (e.g. less than five). Unfortunately, as far as we know, this stability property is not proved anywhere in the literature, which deters the adoption of the algorithm. Using MCMC moves provides a procedure with truly linear and deterministic run time, and a stability result is the only missing piece of the puzzle to completely resolve the $\mathcal{O}(N^2)$ problem. We believe one reason for the current state of affair is that the stability proof techniques employed by [Douc et al. \(2011\)](#) and [Olsson and Westerborn \(2017\)](#) are difficult to extend to the MCMC case.

Fourth, and this is related to the third point, the stability of the PaRIS algorithm has only been proved in the asymptotic regime. More specifically, [Olsson and Westerborn \(2017\)](#) established a central limit theorem as $N \rightarrow \infty$ in Corollary 5, then showed that the corresponding asymptotic error remains controlled as $T \rightarrow \infty$ in Theorem 8 and Theorem 9. While non-asymptotic stability bounds for the FFBS algorithm are already available in [Del Moral et al. \(2010\)](#); [Douc et al. \(2011\)](#); [Dubarry and Le Corff \(2013\)](#), we do not think that they can be extended straightforwardly to PaRIS and we are not aware of any such attempt in the literature.

Fifth, all backward samplers mentioned thus far require access to the transition density. Many models have dynamics that can be simulated from but transition densities that are not explicitly calculable. Enabling backward sampling in this scenario is challenging and will certainly require some kind of problem-specific knowledge to extract information from the transition densities, despite not being able to evaluate them exactly.

1.3. Structure and contribution. Section 2 presents a general framework which admits as particular cases a wide variety of existing algorithms (e.g. FFBS, forward-additive smoothing, PaRIS) as well as the novel ones considered later in the paper. It allows to simultaneously prove the consistency as $N \rightarrow \infty$ and the stability as $T \rightarrow \infty$ for all of them. The main ingredient is the discrete backward kernels, which are essentially random $N \times N$ matrices employed differently in the offline and the online settings. On the technical side, the stability result is proved using a new technique, yielding a non-asymptotic bound that addresses the fourth point in subsection 1.2.

Next, we closely look at the use of rejection sampling and realise that in many models, the resulting execution time may be significantly heavy-tailed; see Section 3. For instance, the run time of PaRIS may have infinite expectation, whereas the run time of FFBS may have infinite variance. (Since it is

technically more involved, the material for the FFBS algorithm is delegated to Supplement B.) These results address the first point in subsection 1.2 and we discuss their severe practical implications.

We then derive and analyse hybrid rejection sampling schemes (i.e. schemes that use rejection sampling only up to a certain number of attempts, and then switch to the standard method). We show that they lead to a nearly $\mathcal{O}(N)$ algorithm (up to some log factor) in Gaussian models; again see Section 3. This stems from the subtle interaction between the tail of Gaussian densities and the finite Feynman-Kac approximation. Outside this class of model, the hybrid algorithm can still escape the $\mathcal{O}(N^2)$ complexity, although it might not reach the ideal linear run time target. These results shed some light on the second issue mentioned in subsection 1.2.

In Section 4, we look at backward kernels that are more efficient to simulate than the FFBS and the PaRIS ones. Section 4.1 describes backward kernels based on MCMC (Markov chain Monte Carlo) following Bunch and Godsill (2013) and extends them to the online scenario. We cast this family of algorithms as a particular case of the general framework developed in Section 2, which allows convergence and stability to follow immediately. This solves the long-standing problem described in the third point of subsection 1.2.

MCMC methods require evaluation of the likelihood and thus cannot be applied to models with intractable transition densities. In Section 4.2, we show how the use of forward coupling can replace the role of backward MCMC steps in these scenarios. This makes it possible to obtain stable performance in both on-line and off-line scenarios (with intractable transition densities) and provides a possible solution to the fifth challenge describe in subsection 1.2.

Section 5 illustrates the aforementioned algorithms in both on-line and off-line uses. We highlight how hybrid and MCMC samplers lead to a more user-friendly (i.e. smaller, less random and less model-dependent) execution time than the pure rejection sampler. We also apply our smoother for intractable densities to a continuous-time diffusion process with discretization. We observe that our procedure can indeed prevent degeneracy as $T \rightarrow \infty$, provided that some care is taken to build couplings with good performance. Section 6 concludes the paper with final practical recommendations and further research directions. Table 1 gives an overview of existing and novel algorithms as well as our contributions for each.

1.4. Related work. Proposition 1 of Douc et al. (2011) states that under certain assumptions, the FFBS-reject algorithm has an *asymptotic* $\mathcal{O}_{\mathbb{P}}(N)$ complexity. This does not contradict our results, which point out the undesirable properties of the *non-asymptotic* execution time. Clearly, non-asymptotic behaviours are what users really observe in practice. From a technical point of view, the proof of Douc et al. (2011, Prop. 1) is a simple application of Theorem 5 of the same paper. In contrast, non-asymptotic results such as Theorem B.1 and Theorem B.2 require more delicate finite sample analyses.

Figure 1 of Olsson and Westerborn (2017) and the accompanying discussion provide an excellent intuition on the stability of smoothing algorithms based on the support size of the backward kernels. We formalise this support size condition for the first time by the inequality (13) and construct a novel non-asymptotic stability result based on it. In contrast, Olsson and Westerborn (2017) depart from their initial intuition and use an entirely different technique to establish stability. Their result is asymptotic in nature.

Gloaguen et al. (2022) briefly mention the use of MCMC in PaRIS algorithm, but their algorithm is fundamentally different to and less efficient than Bunch and Godsill (2013). Indeed, they do not start the MCMC chains at the ancestors previously obtained during the filtering step. They are thus obliged to perform a large number of MCMC iterations for decorrelation, whereas the algorithms described in our Proposition 4, built upon the ideas of Bunch and Godsill (2013), only require a single MCMC step to guarantee stability. However, we would like to stress again that Bunch and Godsill (2013) did not prove this important fact.

Another way to reduce the computation time is to perform the expensive backward sampling steps at certain times t only. For other values of t , the naive genealogy tracking smoother is used instead. This idea has been recently proposed by Mastrototaro et al. (2021), who also provided a practical recipe for deciding at which values of t the backward sampling should take place and derived corresponding theoretical results.

Mode \ Kernel	FFBS kernel	PaRIS kernel	MCMC kernels	Intract.
Offline	(*) FFBS (+) Thm. B.1, Thm. B.2 (+) Thm. B.3, Cor. 2		(*) FFBS-MCMC (+) Prop. 4	(**)
Online	(*) Forward-additive	(*) PaRIS (+) Thm. 2.2, Prop. 2 (+) Thm. 3.1, Thm. 3.2	(**)	(**)

TABLE 1. Summary of smoothing algorithms considered in this paper (classified by the backward kernel and the execution mode) and our contributions. (*) means an existing algorithm, (+) means a novel theoretical result and (**) means a novel algorithm

Smoothing in models with intractable transition densities is very challenging. If these densities can be estimated accurately, the algorithms proposed by [Gloaguen et al. \(2022\)](#) permit to attack this problem. A case of particular interest is diffusion models, where unbiased transition density estimators are provided in [Beskos et al. \(2006\)](#); [Fearnhead et al. \(2008\)](#). More recently, [Yonekura and Beskos \(2022\)](#) use a special bridge path-space construction to overcome the unavailability of transition densities when the diffusion (possibly with jumps) must be discretised.

Our smoother for intractable models are based on a general coupling principle that is not specific to diffusions. We only require users to be able to simulate their dynamics (e.g. using discretisation in the case of diffusions) and to manipulate random numbers in their simulations so that dynamics starting from two different points can meet with some probability. Our method does not directly provide an estimator for the gradient of the transition density with respect to model parameters and thus cannot be used in its current form to perform maximum likelihood estimation (MLE) in intractable models; whereas the aforementioned work have been able to do so in the case of diffusions. However, the main advantage of our approach lies in its generality beyond the diffusion case. Furthermore, modifications allowing to perform MLE are possible and might be explored in further work specifically dedicated to the parameter estimation problem.

The idea of coupling has been incorporated in the smoothing problem in a different manner by [Jacob et al. \(2019\)](#). There, the goal is to provide offline unbiased estimates of the expectation under the smoothing distribution. Coupling and more generally ideas based on correlated random numbers are also useful in the context of partially observed diffusions via the multilevel approach ([Jasra et al., 2017](#)).

In this work, we consider smoothing algorithms that are based on a unique pass of the particle filter. Offline smoothing can be done using repeated iterations of the conditional particle filter ([Andrieu et al., 2010](#)). Full trajectories can also be constructed in an online manner if one is willing to accept some lag approximations ([Duffield and Singh, 2022](#)). Another approach to smoothing consists of using an additional information filter ([Fearnhead et al., 2010](#)), but it is limited to functions depending on one state only. Each of these algorithmic families has their own advantages and disadvantages, of which a detailed discussion is out of the scope of this article (see however [Nordh and Antonsson, 2015](#)).

2. GENERAL STRUCTURE OF SMOOTHING ALGORITHMS

In this section, we decompose each smoothing algorithm into two separate parts: the backward kernel (which determines its theoretical properties such as the convergence and the stability) and the execution mode (which is either online or offline and determines its implementation). This has two advantages: first, it induces an easy-to-navigate categorization of algorithms (see Table 1); and second, it allows to prove the convergence and the stability for each of them by verifying sufficient conditions on the backward kernel component only.

2.1. Notations. *Measure-kernel-function notations.* Let \mathcal{X} and \mathcal{Y} be two measurable spaces with respective σ -algebras $\mathcal{B}(\mathcal{X})$ and $\mathcal{B}(\mathcal{Y})$. The following definitions involve integrals and only make sense when they are well-defined. For a measure μ on \mathcal{X} and a function $f : \mathcal{X} \rightarrow \mathbb{R}$, the notations μf and $\mu(f)$ refer to $\int f(x)\mu(dx)$. A kernel (resp. Markov kernel) K is a mapping from $\mathcal{X} \times \mathcal{B}(\mathcal{Y})$ to \mathbb{R} (resp. $[0, 1]$) such that, for $B \in \mathcal{B}(\mathcal{Y})$ fixed, $x \mapsto K(x, B)$ is a measurable function on \mathcal{X} ; and for x fixed, $B \mapsto K(x, B)$ is a measure (resp. probability measure) on \mathcal{Y} . For a real-valued function g defined on \mathcal{Y} , let $Kg : \mathcal{X} \rightarrow \mathbb{R}$ be

the function $Kg(x) := \int g(y)K(x, dy)$. We sometimes write $K(x, g)$ for the same expression. The product of the measure μ on \mathcal{X} and the kernel K is a measure on \mathcal{Y} , defined by $\mu K(B) := \int K(x, B)\mu(dx)$. *Other notations.* • The notation $X_{0:t}$ is a shorthand for (X_0, \dots, X_t) • We denote by $\mathcal{M}(W^{1:N})$ the multinomial distribution supported on $\{1, 2, \dots, N\}$. The respective probabilities are W_1, \dots, W_N . If they do not sum to 1, we implicitly refer to the normalised version obtained by multiplication of the weights with the appropriate constant • The symbol $\xrightarrow{\mathbb{P}}$ means convergence in probability and \Rightarrow means convergence in distribution • The geometric distribution with parameter λ is supported on $\mathbb{Z}_{\geq 1}$, has probability mass function $f(n) = \lambda(1-\lambda)^{n-1}$ and is noted by $\text{Geo}(\lambda)$ • Let \mathcal{X} and \mathcal{Y} be two measurable spaces. Let μ and ν be two probability measures on \mathcal{X} and \mathcal{Y} respectively. The o-times product measure $\mu \otimes \nu$ is defined via $(\mu \otimes \nu)(h) := \int \int h(x, y)\mu(dx)\nu(dy)$ for bounded functions $h : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$. If $X \sim \mu$ and $Y \sim \nu$, we sometimes note $\mu \otimes \nu$ by $X \otimes Y$.

2.2. Feynman-Kac formalism and the bootstrap particle filter. Let $\mathcal{X}_{0:T}$ be a sequence of measurable spaces and $M_{1:T}$ be a sequence of Markov kernels such that M_t is a kernel from \mathcal{X}_{t-1} to \mathcal{X}_t . Let $X_{0:T}$ be an unobserved inhomogeneous Markov chain with starting distribution $X_0 \sim \mathbb{M}_0(dx_0)$ and Markov kernels $M_{1:T}$; i.e. $X_t|X_{t-1} \sim M_t(X_{t-1}, dx_t)$ for $t \geq 1$. We aim to study the distribution of $X_{0:T}$ given observed data $Y_{0:T}$. Conditioned on $X_{0:T}$, the data Y_0, \dots, Y_T are independent and $Y_t|X_{0:T} \equiv Y_t|X_t \sim \mathbf{f}_t(\cdot|X_t)$ for a certain emission distribution $\mathbf{f}_t(dy_t|x_t)$. Assume that there exists dominating measures $\tilde{\lambda}_t$ not depending on x_t such that

$$\mathbf{f}_t(dy_t|x_t) = f_t(y_t|x_t)\tilde{\lambda}_t(dy_t).$$

The distribution of $X_{0:t}|Y_{0:t}$ is then given by

$$(1) \quad \mathbb{Q}_t(dx_{0:t}) = \frac{1}{L_t} \mathbb{M}_0(dx_0) \prod_{s=1}^t M_s(x_{s-1}, dx_s) G_s(x_s)$$

where $G_s(x_s) := f(y_s|x_s)$ and $L_t > 0$ is the normalising constant. Moreover, $\mathbb{Q}_{-1} := \mathbb{M}_0$ and $L_{-1} := 1$ by convention. Equation (1) defines a Feynman-Kac model (Del Moral, 2004). It does not require M_t to admit a transition density, although herein we only consider models where this assumption holds. Let λ_t be a dominating measure on \mathcal{X}_t in the sense that there exists a function m_t (not necessarily tractable) such that

$$(2) \quad M_t(x_{t-1}, dx_t) = m_t(x_{t-1}, x_t)\lambda_t(dx_t).$$

A special case of the current framework are linear Gaussian state space models. They will serve as a running example for the article, and some of the results will be specifically demonstrated for models of this class. The rationale is that many real-world dynamics are partly, or close to, Gaussian. The notations for linear Gaussian models are given in Supplement A.1 and we will refer to them whenever this model class is discussed.

Particle filters are algorithms that sample from $\mathbb{Q}_t(dx_t)$ in an on-line manner. In this article, we only consider the bootstrap particle filter (Gordon et al., 1993) and we detail its notations in Algorithm 1. Many results in the following do apply to the auxiliary filter (Pitt and Shephard, 1999) as well, and we shall as a rule indicate explicitly when it is *not* the case.

We end this subsection with the definition of two sigma-algebras that will be referred to throughout the paper. Using the notations of Algorithm 1, let

$$(3) \quad \begin{aligned} \mathcal{F}_t &:= \sigma(X_{0:t}^{1:N}, A_{1:t}^{1:N}), \\ \mathcal{F}_t^- &:= \sigma(X_{0:t}^{1:N}). \end{aligned}$$

2.3. Backward kernels and off-line smoothing. In this subsection, we first describe three examples of backward kernels, in which we emphasise both the random measure and the random matrix viewpoints. We then formalise their use by stating a generic off-line smoothing algorithm.

Example 1 (FFBS algorithm, Godsill et al., 2004). Once Algorithm 1 has been run, the FFBS procedure generates a trajectory approximating the smoothing distribution in a backward manner. More precisely, it starts by simulating index $\mathcal{I}_T \sim \mathcal{M}(W_T^{1:N})$ at time T . Then, recursively for $t = T, \dots, 1$, given indices

Algorithm 1: Bootstrap particle filter**Input:** Feynman-Kac model (1)Simulate $X_0^{1:N} \stackrel{\text{i.i.d.}}{\sim} \mathbb{M}_0$ Set $\omega_0^n \leftarrow G_0(X_0^n)$ for $n = 1, \dots, N$ Set $\ell_0^N \leftarrow \sum_{n=1}^N \omega_0^n / N$ Set $W_0^n \leftarrow \omega_0^n / N \ell_0^N$ for $n = 1, \dots, N$ **for** $t \leftarrow 1$ **to** T **do** Resample. Simulate $A_t^{1:N} \stackrel{\text{i.i.d.}}{\sim} \mathcal{M}(W_{t-1}^{1:N})$ Move. Simulate $X_t^n \sim M_t(X_{t-1}^n, dx_t)$ for $n = 1, \dots, N$ Reweight. Set $\omega_t^n \leftarrow G_t(X_t^n)$ for $n = 1, 2, \dots, N$ Set $\ell_t^N \leftarrow \sum_{n=1}^N \omega_t^n / N$ Set $W_t^n \leftarrow \omega_t^n / N \ell_t^N$ for $n = 1, 2, \dots, N$ **Output:** For all $t \geq 0$ and function $\varphi : \mathcal{X}_t \rightarrow \mathbb{R}$, the quantity $\sum_{n=1}^N W_t^n \varphi(X_t^n)$ approximates $\int Q_t(dx_{0:t}) \varphi(x_t)$ and the quantity ℓ_t^N approximates L_t / L_{t-1}

$\mathcal{I}_{t:T}$, it generates $\mathcal{I}_{t-1} \in \{1, \dots, N\}$ with probability proportional to $W_{t-1}^n m_t(X_{t-1}^n, X_t^{\mathcal{I}_t})$. The smoothing trajectory is returned as $(X_0^{\mathcal{I}_0}, \dots, X_T^{\mathcal{I}_T})$. Formally, given \mathcal{F}_T , the indices $\mathcal{I}_{0:T}$ are generated according to the distribution

$$\mathcal{M}(W_t^{1:N})(di_T) \left[B_T^{N,\text{FFBS}}(i_T, di_{T-1}) B_{T-1}^{N,\text{FFBS}}(i_{T-1}, di_{T-2}) \dots B_1^{N,\text{FFBS}}(i_1, di_0) \right]$$

where the (random) backward kernels $B_t^{N,\text{FFBS}}$ are defined by

$$(4) \quad B_t^{N,\text{FFBS}}(i_t, di_{t-1}) := \sum_{n=1}^N \frac{W_{t-1}^n m_t(X_{t-1}^n, X_t^{i_t})}{\sum_{k=1}^N W_{t-1}^k m_t(X_{t-1}^k, X_t^{i_t})} \delta_n(di_{t-1}).$$

More simply, we can also look at these random kernels as random $N \times N$ matrices of which entries are given by

$$(5) \quad \hat{B}_t^{N,\text{FFBS}}[i_t, i_{t-1}] := \frac{W_{t-1}^{i_{t-1}} m_t(X_{t-1}^{i_{t-1}}, X_t^{i_t})}{\sum_{k=1}^N W_{t-1}^k m_t(X_{t-1}^k, X_t^{i_t})}.$$

We will need both the kernel viewpoint (4) and the matrix viewpoint (5) in this paper as the better choice depends on the context.

Example 2 (Genealogy tracking, Kitagawa, 1996; Del Moral and Miclo, 2001). It is well known that Algorithm 1 already gives as a by-product an approximation of the smoothing distribution. This information can be extracted from the genealogy, by first simulating index $\mathcal{I}_T \sim \mathcal{M}(W_T^{1:N})$ at time T , then successively appending ancestors until time 0 (i.e. setting sequentially $\mathcal{I}_{t-1} \leftarrow A_t^{\mathcal{I}_t}$). The smoothed trajectory is returned as $(X_0^{\mathcal{I}_0}, \dots, X_T^{\mathcal{I}_T})$. More formally, conditioned on \mathcal{F}_T , we simulate the indices $\mathcal{I}_{0:T}$ according to

$$\mathcal{M}(W_t^{1:N})(di_T) \left[B_T^{N,\text{GT}}(i_T, di_{T-1}) B_{T-1}^{N,\text{GT}}(i_{T-1}, di_{T-2}) \dots B_1^{N,\text{GT}}(i_1, di_0) \right]$$

where GT stands for ‘‘genealogy tracking’’ and the kernels $B_t^{N,\text{GT}}$ are simply

$$(6) \quad B_t^{N,\text{GT}}(i_t, di_{t-1}) := \delta_{A_t^{i_t}}(di_{t-1}).$$

Again, it may be more intuitive to view this random kernel as a random $N \times N$ matrix, the elements of which are given by

$$\hat{B}_t^{N,\text{GT}}[i_t, i_{t-1}] := \mathbb{1} \{i_{t-1} = A_t^{i_t}\}.$$

Example 3 (MCMC backward samplers, Bunch and Godsill, 2013). In Example 2, the backward variable \mathcal{I}_{t-1} is simply set to $A_t^{\mathcal{I}_t}$. On the contrary, in Example 1, we need to launch a simulator for the discrete measure $W_{t-1}^n m_t(X_{t-1}^n, X_t^{\mathcal{I}_t})$. Interestingly, the current value of $A_t^{\mathcal{I}_t}$ is not taken into account in that simulator. Therefore, a natural idea to combine the two previous examples is to apply one (or

several) MCMC steps to $A_t^{\mathcal{I}_t}$ and assign the result to \mathcal{I}_{t-1} . The MCMC algorithm operates on the space $\{1, 2, \dots, N\}$ and targets the invariant measure $W_{t-1}^n m_t(X_{t-1}^n, X_t^{\mathcal{I}_t})$. If only one independent Metropolis-Hastings (MH) step is used and the proposal is $\mathcal{M}(W_{t-1}^{1:N})$, the corresponding random matrix $\hat{B}_t^{N, \text{IMH}}$ has values

$$\hat{B}_t^{N, \text{IMH}}[i_t, i_{t-1}] = W_{t-1}^{i_{t-1}} \min \left(1, m_t(X_{t-1}^{i_{t-1}}, X_t^{i_t}) / m_t(X_{t-1}^{A_t^{i_t}}, X_t^{i_t}) \right)$$

if $i_{t-1} \neq A_t^{i_t}$, and

$$\hat{B}_t^{N, \text{IMH}}[i_t, A_t^{i_t}] = 1 - \sum_{n \neq A_t^{i_t}} \hat{B}_t^{N, \text{IMH}}[i_t, n].$$

This third example shows that some elements of the matrix $\hat{B}_t^{N, \text{IMH}}$ might be expensive to calculate. If several MCMC steps are performed, all elements of $\hat{B}_t^{N, \text{IMH}}$ will have non-trivial expressions. Still, simulating from $B_t^{N, \text{IMH}}(i_t, di_{t-1})$ is easy as it amounts to running a standard MCMC algorithm. MCMC backward samples are studied in more details in Section 4.1.

We formalise how off-line smoothing can be done given random matrices $\hat{B}_{1:T}^N$; see Algorithm 2. Note that in the above examples, our matrices \hat{B}_t^N are \mathcal{F}_t -measurable (i.e. they depend on particles and indices up to time t), but this is not necessarily the case in general (i.e. they may also depend on additional random variables, see Section 2.5). Furthermore, Algorithm 2 describes how to perform smoothing using the matrices $\hat{B}_{1:T}^N$, but does not say where they come from. At this point, it is useful to keep in mind the above three examples. In Section 2.4, we will give a general recipe for constructing valid matrices \hat{B}_t^N (i.e. those that give a consistent algorithm).

Algorithm 2: Generic off-line smoother

Input: Filtering results $X_{0:T}^{1:N}$, $W_{0:T}^{1:N}$, and $A_{1:T}^{1:N}$ from Algorithm 1; random matrices $\hat{B}_{1:T}^N$ (see Section 2.3 for two examples of such matrices and Section 2.4 for a general recipe to construct them)

for $n \leftarrow 1$ **to** N **do**

Simulate $\mathcal{I}_T^n \sim \mathcal{M}(W_T^{1:N})$

for $t \leftarrow T$ **to** 1 **do**

Simulate $\mathcal{I}_{t-1}^n \sim B_t^N(\mathcal{I}_t^n, di_{t-1})$ (the kernel $B_t^N(\mathcal{I}_t, \cdot)$ is defined by the \mathcal{I}_t -th row of the input matrix \hat{B}_t^N)

Output: The N smoothed trajectories $(X_0^{\mathcal{I}_0^n}, \dots, X_T^{\mathcal{I}_T^n})$ for $n = 1, \dots, N$

Algorithm 2 simulates, given \mathcal{F}_T and $\hat{B}_{1:T}^N$, N i.i.d. index sequences $\mathcal{I}_{0:T}^n$, each distributed according to

$$\mathcal{M}(W_T^{1:N})(di_T) \prod_{t=T}^1 B_t^N(i_t, di_{t-1}).$$

Once the indices $\mathcal{I}_{0:T}^n$ are simulated, the N smoothed trajectories are returned as $(X_0^{\mathcal{I}_0^n}, \dots, X_T^{\mathcal{I}_T^n})$. Given \mathcal{F}_T and $\hat{B}_{1:T}^N$, they are thus conditionally i.i.d. and their conditional distribution is described by the $x_{0:T}$ component of the joint distribution

$$(7) \quad \bar{\mathbb{Q}}_T^N(dx_{0:T}, di_{0:T}) := \mathcal{M}(W_T^{1:N})(di_T) \left[\prod_{t=T}^1 B_t^N(i_t, di_{t-1}) \right] \left[\prod_{t=T}^0 \delta_{X_t^{i_t}}(dx_t) \right].$$

Throughout the paper, the symbol $\bar{\mathbb{Q}}_T^N$ will refer to this joint distribution, while the symbol \mathbb{Q}_T^N will refer to the $x_{0:T}$ -marginal of $\bar{\mathbb{Q}}_T^N$ only. This allows the notation $\mathbb{Q}_T^N \varphi$ to make sense, where $\varphi = \varphi(x_0, \dots, x_T)$ is a real-valued function defined on the hidden states.




FIGURE 1. Relation between variables described in Theorem 2.1.

2.4. Validity and convergence. The kernels $B_t^{N,\text{FFBS}}$ and $B_t^{N,\text{GT}}$ are both valid backward kernels to generate convergent approximation of the smoothing distribution (Del Moral, 2004; Douc et al., 2011). This subsection shows that they are not the only ones and gives a sufficient condition for a backward kernel to be valid. It will prove a necessary tool to build more efficient B_t^N later in the paper.

Recall that Algorithm 1 outputs particles $X_{0:T}^{1:N}$, weights $W_{0:T}^{1:N}$ and ancestor variables $A_{1:T}^{1:N}$. Imagine that the $A_{1:T}^{1:N}$ were discarded after filtering has been done and we wish to simulate them back. We note that, since the $X_{0:T}^{1:N}$ are given, the $T \times N$ variables $A_{1:T}^{1:N}$ are conditionally i.i.d. We can thus simulate them back from

$$p(a_t^n | x_{0:T}^{1:N}) = p(a_t^n | x_{t-1}^{1:N}, x_t^n) \propto w_{t-1}^{a_t^n} m_t(x_{t-1}^{a_t^n}, x_t^n).$$

This is precisely the distribution of $B_t^{N,\text{FFBS}}(n, \cdot)$. It turns out that any other invariant kernel that can be used for simulating back the discarded $A_{1:T}^{1:N}$ will lead to a convergent algorithm as well. For instance, $B_t^{N,\text{GT}}(n, \cdot)$ (Example 2) simply returns back the old A_t^n , unlike $B_t^{N,\text{FFBS}}(n, \cdot)$ which creates a new version. The kernel $B_t^{N,\text{IMH}}(n, \cdot)$ (Example 3) is somewhat an intermediate between the two. We formalise these intuitions in the following theorem. It is stated for the bootstrap particle filter, but as a matter of fact, the proof can be extended straightforwardly to auxiliary particle filters as well.

Assumption 1. For all $0 \leq t \leq T$, $G_t(x_t) > 0$ and $\|G_t\|_\infty < \infty$.

Theorem 2.1. *We use the same notations as in Algorithms 1 and 2 (in particular, \hat{B}_t^N denotes the transition matrix that corresponds to the considered kernel B_t^N). Assume that for any $1 \leq t \leq T$, the random matrix \hat{B}_t^N satisfies the following conditions:*

- given \mathcal{F}_{t-1} and $\hat{B}_{1:t-1}^N$, the variables $(X_t^n, A_t^n, \hat{B}_t^N(n, \cdot))$ for $n = 1, \dots, N$ are i.i.d. and their distribution only depends on $X_{t-1}^{1:N}$, where $\hat{B}_t^N(n, \cdot)$ is the n -th row of matrix \hat{B}_t^N ;
- if J_t^n is a random variable such that

$$J_t^n | X_{t-1}^{1:N}, X_t^n, \hat{B}_t^N(n, \cdot) \sim B_t^N(n, \cdot)$$

, then (J_t^n, X_t^n) has the same distribution as (A_t^n, X_t^n) given $X_{t-1}^{1:N}$.

Then under Assumption 1, there exists constants $C_T > 0$ and $S_T < \infty$ such that, for any $\delta > 0$ and function $\varphi = \varphi(x_0, \dots, x_T)$:

$$(8) \quad \mathbb{P} \left(\left| \mathbb{Q}_T^N \varphi - \mathbb{Q}_T \varphi \right| \geq \frac{\sqrt{-2 \log(\delta/2C_T)} S_T \|\varphi\|_\infty}{\sqrt{N}} \right) \leq \delta$$

where \mathbb{Q}_T^N is defined by (7).

A typical relation between variables defined in the statement of the theorem is illustrated by a graphical model in Figure 1. (See Bishop 2006, Chapter 8 for the formal definition of graphical models and how to use them.) By “typical”, we mean that Theorem 2.1 technically allows for more complicated relations, but the aforementioned figure captures the most essential cases.

Theorem 2.1 is a generalisation of Douc et al. (2011, Theorem 5). Its proof thus follows the same lines (Supplement E.1). However, in our case the measure $\mathbb{Q}_T^N(dx_{0:T})$ is no longer Markovian. This is because the backward kernel $B_t^N(i_t, di_{t-1})$ does not depend on $X_t^{i_t}$ alone, but also possibly on its ancestor and

extra random variables. This small difference has a big consequence: compared to Douc et al. (2011, Theorem 5), Theorem 2.1 has a much broader applicability and encompasses, for instance, the MCMC-based algorithms presented in Section 4.1 and novel kernels presented in Section 4.2 for intractable densities.

As we have seen in (7), \mathbb{Q}_T^N is fundamentally a discrete measure of which the support contains N^{T+1} elements. As such, $\mathbb{Q}_T^N \varphi$ cannot be computed exactly in general and must be approximated using N trajectories $(X_0^{T_n}, \dots, X_T^{T_n})$ simulated via Algorithm 2. Theorem 2.1 is thus completed by the following corollary, which is an immediate consequence of Hoeffding inequality (Supplement E.13).

Corollary 1. *Under the same setting as Theorem 2.1, we have*

$$\mathbb{P} \left(\left| \frac{1}{N} \sum_n \varphi(X_0^{T_n}, \dots, X_T^{T_n}) - \mathbb{Q}_T \varphi \right| \geq \frac{\sqrt{-2 \log \left(\frac{\delta}{2(C_T+1)} \right)} (S_T + 1) \|\varphi\|_\infty}{\sqrt{N}} \right) \leq \delta.$$

2.5. Generic on-line smoother. As we have seen in Section 2.3 and Section 2.4, in general, the expectation $\mathbb{Q}_T^N \varphi$, for a real-valued function $\varphi = \varphi(x_0, \dots, x_T)$ of the hidden states, cannot be computed exactly due to the large support (N^{T+1} elements) of \mathbb{Q}_T^N . Moreover, in certain settings we are interested in the quantities $\mathbb{Q}_t^N \varphi_t$ for different functions φ_t . They cannot be approximated in an on-line manner without more assumptions on the connection between φ_{t-1} and φ_t . If the family (φ_t) is additive, i.e. there exists functions ψ_t such that

$$(9) \quad \varphi_t(x_{0:t}) := \psi_0(x_0) + \psi_1(x_0, x_1) + \dots + \psi_t(x_{t-1}, x_t)$$

then we can calculate $\mathbb{Q}_t^N \varphi_t$ both exactly *and* on-line. The procedure was first described in Del Moral et al. (2010) for the kernel $\mathbb{Q}_t^{N, \text{FFBS}}$ (i.e. the measure defined by (7) and the random kernels $B_t^{N, \text{FFBS}}$), but we will use the idea for other kernels as well. In this subsection, we first explain the principle of the method, then discuss its computational complexity and the link to the PaRIS algorithm (Olsson and Westerborn, 2017).

Principle. For simplicity, we start with the special case $\varphi_t(x_{0:t}) = \psi_0(x_0)$. Equation (7) and the matrix viewpoint of Markov kernels then give

$$\mathbb{Q}_t^N \varphi_t = [W_t^1 \dots W_t^N] \hat{B}_t^N \hat{B}_{t-1}^N \dots \hat{B}_1^N \begin{bmatrix} \psi_0(X_0^1) \\ \vdots \\ \psi_0(X_0^N) \end{bmatrix}.$$

This naturally suggests the following recursion formula to compute $\mathbb{Q}_t^N \varphi_t$:

$$\mathbb{Q}_t^N \varphi_t = [W_t^1 \dots W_t^N] \hat{S}_t^N$$

with $\hat{S}_0^N = [\psi_0(X_0^1) \dots \psi_0(X_0^N)]^\top$ and

$$(10) \quad \hat{S}_t^N := \hat{B}_t^N \hat{S}_{t-1}^N.$$

In the general case where functions φ_t are given by (9), simple calculations (Supplement E.2) show that (10) is replaced by

$$(11) \quad \hat{S}_t^N := \hat{B}_t^N \hat{S}_{t-1}^N + \text{diag}(\hat{B}_t^N \hat{\psi}_t^N)$$

where the $N \times N$ matrix $\hat{\psi}_t^N$ is defined by

$$\hat{\psi}_t^N[i_{t-1}, i_t] := \psi_t(X_{t-1}^{i_{t-1}}, X_t^{i_t})$$

and the operator $\text{diag} : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^N$ extracts the diagonal of a matrix. This is exactly what is done in Algorithm 3.

Algorithm 3: Generic on-line smoother for additive functions (one step)

Input: Particles $X_{t-1}^{1:N}$ and weights $W_{t-1}^{1:N}$ at time $t-1$; the $N \times 1$ vector \hat{S}_{t-1}^N (see text); additive function (9)

Generate $X_t^{1:N}$ and $W_t^{1:N}$ according to the particle filter (Algorithm 1)

Calculate the random matrix \hat{B}_t^N (see Section 2.3 and Section 2.4)

Create the $N \times 1$ vector \hat{S}_t^N according to (11). More precisely:

for $i_t \leftarrow 1$ **to** N **do**

$$\left[\hat{S}_t^N[i_t] \leftarrow \sum_{i_{t-1}} \hat{B}_t^N[i_t, i_{t-1}] \left(\hat{S}_{t-1}^N[i_{t-1}] + \psi_t(X_{t-1}^{i_{t-1}}, X_t^{i_t}) \right) \right]$$

Output: Quantity $\sum_n W_t^n \hat{S}_t^N[n]$ which is equal to $\mathbb{Q}_t^N \varphi_t$ and is an estimate of $\mathbb{Q}_t(\varphi_t)$; particles $X_t^{1:N}$, weights $W_t^{1:N}$ and vector S_t^N for the next step

Computational complexity and the PaRIS algorithm. Equations (10) and (11) involve a matrix-vector multiplication and thus require, in general, $\mathcal{O}(N^2)$ operations to be evaluated. When $\hat{B}_t^N \equiv \hat{B}_t^{N, \text{FFBS}}$, Algorithm 3 becomes the $\mathcal{O}(N^2)$ on-line smoothing algorithm of Del Moral et al. (2010). The $\mathcal{O}(N^2)$ complexity can however be lowered to $\mathcal{O}(N)$ if the matrices \hat{B}_t^N are *sparse*. This is the idea behind the PaRIS algorithm (Olsson and Westerborn, 2017), where the full matrix $\hat{B}_t^{N, \text{FFBS}}$ is unbiasedly estimated by a sparse matrix $\hat{B}_t^{N, \text{PaRIS}}$. More specifically, for any integer $\tilde{N} > 1$, for any $n \in 1, \dots, N$, let $J_t^{n,1}, \dots, J_t^{n,\tilde{N}}$ be conditionally i.i.d. random variables simulated from $B_t^{N, \text{FFBS}}(n, \cdot)$. The random matrix $\hat{B}_t^{N, \text{PaRIS}}$ is then defined as

$$\hat{B}_t^{N, \text{PaRIS}}[n, m] := \frac{1}{\tilde{N}} \sum_{\tilde{n}=1}^{\tilde{N}} \mathbb{1} \left\{ J_t^{n, \tilde{n}} = m \right\}$$

and the corresponding random kernel is

$$(12) \quad B_t^{N, \text{PaRIS}}(n, dm) = \frac{1}{\tilde{N}} \sum_{\tilde{n}=1}^{\tilde{N}} \delta_{J_t^{n, \tilde{n}}} (dm).$$

The following straightforward proposition establishes the validity of the $B_t^{N, \text{PaRIS}}$ kernel. Together with Theorem 2.1, it can be thought of as a reformulation of the consistency of the PaRIS algorithm (Olsson and Westerborn, 2017, Corollary 2) in the language of our framework.

Proposition 1. *The matrix $\hat{B}_t^{N, \text{PaRIS}}$ has only $\mathcal{O}(N\tilde{N})$ non-zero elements out of N^2 . It is an unbiased estimate of $\hat{B}_t^{N, \text{FFBS}}$ in the sense that*

$$\mathbb{E} \left[\hat{B}_t^{N, \text{PaRIS}} \middle| \mathcal{F}_t \right] = \hat{B}_t^{N, \text{FFBS}}.$$

Moreover, the sequence of matrices $B_{1:T}^{N, \text{PaRIS}}$ satisfies the two conditions of Theorem 2.1.

The proposition also justifies the $\mathcal{O}(N)$ complexity of (10) and (11), as long as \tilde{N} is fixed as $N \rightarrow \infty$. But it is important to remark that the preceding $\mathcal{O}(N)$ complexity does not include the cost of generating the matrices $\hat{B}_t^{N, \text{PaRIS}}$ themselves, i.e., the operations required to simulate the indices $J_t^{n, \tilde{n}}$. In Olsson and Westerborn (2017) it is argued that such simulations have an $\mathcal{O}(N)$ cost using the rejection sampling method whenever the transition density is both upper and lower bounded. Section 3 investigates the claim when this hypothesis is violated.

2.6. Stability. When $\hat{B}_t^N \equiv \hat{B}_t^{N, \text{GT}}$, Algorithms 2 and 3 reduce to the genealogy tracking smoother (Kitagawa, 1996). The matrix $\hat{B}_t^{N, \text{GT}}$ is indeed sparse, leading to the well-known $\mathcal{O}(N)$ complexity of this on-line procedure. As per Theorem 2.1, smoothing via genealogy tracking is convergent at rate $\mathcal{O}(N^{-1/2})$ if T is fixed. When $T \rightarrow \infty$ however, all particles will eventually share the same ancestor at time 0 (or any fixed time t). Mathematically, this phenomenon is manifested in two ways: (a) for fixed t and function $\phi_t : \mathcal{X}_t \rightarrow \mathbb{R}$, the error of estimating $\mathbb{E}[\phi_t(X_t) | Y_{0:T}]$ grows linearly with T ; and (b) the error of estimating $\mathbb{E} \left[\sum_{t=0}^T \psi_t(x_{t-1}, x_t) \middle| Y_{0:T} \right]$ grows quadratically with T . These correspond respectively to the degeneracy for the fixed marginal smoothing and the additive smoothing problems;

see also the introductory section of [Olsson and Westerborn \(2017\)](#) for a discussion. The random matrices $\hat{B}_t^{N,\text{GT}}$ are therefore said to be *unstable* as $T \rightarrow \infty$, which is not the case for $\hat{B}_t^{N,\text{FFBS}}$ or $\hat{B}_t^{N,\text{PaRIS}}$. This subsection gives sufficient conditions to ensure the stability of a general \hat{B}_t^N .

The essential point behind smoothing stability is simple: the support of $B_t^{N,\text{FFBS}}(n, \cdot)$ or $B_t^{N,\text{PaRIS}}(n, \cdot)$ for $\tilde{N} \geq 2$ contains more than one element, contrary to that of $B_t^{N,\text{GT}}(n, \cdot)$. This property is formalised by (13). To explain the intuitions, we use the notations of Algorithm 2 and consider the estimate

$$N^{-1} \left(\psi_0(X_0^{\mathcal{I}_0^1}) + \dots + \psi_0(X_0^{\mathcal{I}_0^N}) \right)$$

of $\mathbb{E}[\psi_0(X_0) | Y_{0:T}]$ when $T \rightarrow \infty$. The variance of the quantity above is a sum of $\text{Cov}(\psi_0(X_0^{\mathcal{I}_0^i}), \psi_0(X_0^{\mathcal{I}_0^j}))$ terms. It can therefore be understood by looking at a pair of trajectories simulated using Algorithm 2.

At final time $t = T$, \mathcal{I}_T^1 and \mathcal{I}_T^2 both follow the $\mathcal{M}(W_T^{1:N})$ distribution. Under regularity conditions (e.g. no extreme weights), they are likely to be different, i.e., $\mathbb{P}(\mathcal{I}_T^1 = \mathcal{I}_T^2) = \mathcal{O}(1/N)$. This property can be propagated backward: as long as $\mathcal{I}_t^1 \neq \mathcal{I}_t^2$, the two variables \mathcal{I}_{t-1}^1 and \mathcal{I}_{t-1}^2 are also likely to be different, with however a small $\mathcal{O}(1/N)$ chance of being equal. Moreover, as long as the two trajectories have not met, they can be simulated independently given \mathcal{F}_T^- (the sigma algebra defined in (3)). In mathematical terms, under the two hypotheses of Theorem 2.1, given \mathcal{F}_T^- and $\mathcal{I}_{t:T}^{1,2}$, it can be proved that the two variables \mathcal{I}_{t-1}^1 and \mathcal{I}_{t-1}^2 are independent if $\mathcal{I}_t^1 \neq \mathcal{I}_t^2$ (Lemma E.1, Supplement E.3).

Since there is an $\mathcal{O}(1/N)$ chance of meeting at each time step, if $T \gg N$, it is likely that the two paths will meet at some point $t \gg 0$. When $\mathcal{I}_t^1 = \mathcal{I}_t^2$, the two indices \mathcal{I}_{t-1} and \mathcal{I}_{t-2} are both simulated according to $B_t^N(\mathcal{I}_t^1, \cdot)$. In the genealogy tracking algorithm, $B_t^{N,\text{GT}}(i, \cdot)$ is a Dirac measure, leading to $\mathcal{I}_{t-1}^1 = \mathcal{I}_{t-1}^2$ almost surely. This spreads until time 0, so $\text{Corr}(\psi_0(X_0^{\mathcal{I}_0^1}), \psi_0(X_0^{\mathcal{I}_0^2}))$ is almost 1 if $T \gg N$.

Other kernels like $B_t^{N,\text{FFBS}}$ or $B_t^{N,\text{PaRIS}}$ do not suffer from the same problem. For these, the support size of $B_t^N(\mathcal{I}_t^1, \cdot)$ is greater than one and thus there is some real chance that $\mathcal{I}_{t-1}^1 \neq \mathcal{I}_{t-1}^2$. If that does happen, we are again back to the regime where the next states of the two paths can be simulated independently. Note also that the support of $B_t^N(\mathcal{I}_t^1, \cdot)$ does not need to be large and can contain as few as 2 elements. Even if \mathcal{I}_{t-1}^1 might still be equal to \mathcal{I}_{t-1}^2 with some probability, the two paths will have new chances to diverge at times $t-2$, $t-3$ and so on. Overall, this makes $\text{Corr}(\psi_0(X_0^{\mathcal{I}_0^1}), \psi_0(X_0^{\mathcal{I}_0^2}))$ quite small (Lemma E.3, Supplement E.3).

We formalise these arguments in the following theorem, whose proof (Supplement E.3) follows them very closely. The price for proof intuitiveness is that the theorem is specific to the bootstrap filter, although numerical evidence (Section 5) suggests that other filters are stable as well.

Assumption 2. The transition densities m_t are upper and lower bounded:

$$\bar{M}_\ell \leq m_t(x_{t-1}, x_t) \leq \bar{M}_h$$

for constants $0 < \bar{M}_\ell < \bar{M}_h < \infty$.

Assumption 3. The potential functions G_t are upper and lower bounded:

$$\bar{G}_\ell \leq G_t(x_t) \leq \bar{G}_h$$

for constants $0 < \bar{G}_\ell < \bar{G}_h < \infty$.

Remark. Since Assumption 2 implies that the \mathcal{X}_t 's are compact, Assumption 1 automatically implies Assumption 3 as soon as the G_t 's' are continuous functions.

Theorem 2.2. *We use the notations of Algorithms 1 and 2. Suppose that Assumptions 2 and 3 hold and the random kernels $B_{1:T}^N$ satisfy the conditions of Theorem 2.1. If, in addition, for the pair of random variables $(J_t^{n,1}, J_t^{n,2})$ whose distribution given $X_{t-1}^{1:N}$, X_t^n and $\hat{B}_t^N(n, \cdot)$ is defined by $B_t^N(n, \cdot) \otimes B_t^N(n, \cdot)$, we have*

$$(13) \quad \mathbb{P} \left(J_t^{n,1} \neq J_t^{n,2} \mid X_{t-1}^{1:N}, X_t^n \right) \geq \varepsilon_S$$

for some $\varepsilon_S > 0$ and all t, n ; then there exists a constant C not depending on T such that:

- *fixed marginal smoothing is stable, i.e. for $s \in \{0, \dots, T\}$ and a real-valued function $\phi_s : \mathcal{X}_s \rightarrow \mathbb{R}$ of the hidden state X_s , we have*

$$(14) \quad \mathbb{E} \left[\left(\int \mathbb{Q}_T^N(dx_s) \phi_s(x_s) - \mathbb{E}[\phi_s(X_s) | Y_{0:T}] \right)^2 \right] \leq \frac{C \|\phi_s\|_\infty^2}{N};$$

- *additive smoothing is stable, i.e. for $T \geq 2$ and the function φ_T defined in (9), we have*

$$(15) \quad \mathbb{E} \left[(\mathbb{Q}_T^N(\varphi_T) - \mathbb{Q}_T(\varphi_T))^2 \right] \leq \frac{C \sum_{t=0}^T \|\psi_t\|_\infty^2}{N} \left(1 + \sqrt{\frac{T}{N}} \right)^2.$$

In particular, when B_t^N is the PaRIS kernel with $\tilde{N} \geq 2$, Theorem 2.2 implies a novel non-asymptotic bound for the PaRIS algorithm. Olsson and Westerborn (2017) first established a central limit theorem as $N \rightarrow \infty$ and T fixed, then showed that the asymptotic variance is controlled as $T \rightarrow \infty$. In contrast, we follow an original approach (whose intuition is explained at the beginning of this subsection) in order to derive a finite sample size bound.

The main technical difficulty is to prove the fast mixing of the Markov kernel product $B_t^N B_{t-1}^N \dots B_{t'}^N$ in terms of $t - t'$. For the original FFBS kernel, the stability proof by Douc et al. (2011) relies on the uniform Doeblin property of each of the term $B_s^{N, \text{FFBS}}$ (page 2136, towards the end of their proof of Lemma 10) and from there, deduces the exponentially fast mixing of the product. When $B_s^{N, \text{FFBS}}$ is approximated by a sparse matrix B_s^N (which is the case for PaRIS, but also for certain MCMC-based and coupling-based smoothers that we shall see later), the aforementioned property no longer holds for each individual term B_s^N . Interestingly however, the good mixing of $B_t^{N, \text{FFBS}} \dots B_{t'}^{N, \text{FFBS}}$ is still conserved in the product $B_t^N \dots B_{t'}^N$. In Lemma E.3, we show that two trajectories generated via the latter kernel have such a small correlation that they are virtually indistinguishable from two independent trajectories generated via the former one.

Theorem 2.2 is stated under strong assumptions (similar to those used in Chopin and Papaspiliopoulos 2020, Chapter 11.4, and slightly stronger than Douc et al. 2011, Assumption 4). On the other hand, it applies to a large class of backward kernels (rather than only FFBS), including the new ones introduced in the forthcoming sections.

In the proof of this theorem, we proceed in two steps: first, we apply existing bounds (Dubarry and Le Corff, 2013, Theorem 3.1 and Del Moral, 2013, Chapter 17) for the error between the $B_t^{N, \text{FFBS}}$ -induced distribution and the true target; and second, we use our own techniques to control the error when $B_t^{N, \text{FFBS}}$ is replaced by any other kernel B_t^n satisfying (13). The $(1 + \sqrt{T/N})^2$ term in (15) comes from the first part and we do not know whether it can be dropped. However, it does not affect the scaling of the algorithm. Indeed, with or without it, the inequality implies that in order to have a constant error in the additive smoothing problem, one only has to take $N = \mathcal{O}(T)$ (instead of $N = \mathcal{O}(T^2)$ without backward sampling). Moreover, from an asymptotic point of view, we always have $\sigma^2(T) = \mathcal{O}(T)$ regardless of the presence of the $(1 + \sqrt{T/N})^2$ term, where $\sigma^2(T) := \lim_{N \rightarrow \infty} N \mathbb{E} \left[(\mathbb{Q}_T^N(\varphi_T) - \mathbb{Q}_T(\varphi_T))^2 \right]$.

3. SAMPLING FROM THE FFBS BACKWARD KERNELS

Sampling from the FFBS backward kernel lies at the heart of both the FFBS algorithm (Example 1) and the PaRIS one (Section 2.5). Indeed, at time t , they require generating random variables distributed according to $B_t^{N, \text{FFBS}}(i_t, di_{t-1})$ for i_t running from 1 to N . Since sampling from a discrete measure on N elements requires $\mathcal{O}(N)$ operations (e.g. via CDF inversion), the total computational cost becomes $\mathcal{O}(N^2)$. To reduce this, we start by considering the subclass of models satisfying the following assumption, which is much weaker than Assumption 2.

Assumption 4. The transition density $m_t(x_{t-1}, x_t)$ is strictly positive and upper bounded, i.e. there exists $\bar{M}_h > 0$ such that $0 < m_t(x_{t-1}, x_t) \leq \bar{M}_h, \forall (x_{t-1}, x_t)$.

The motivation for the first condition $0 < m_t(x_{t-1}, x_t)$ will be clear after Assumption 5 is defined. For now, we see that it is possible to sample from $B_t^{N, \text{FFBS}}(i_t, di_{t-1})$ using rejection sampling via the proposal

distribution $\mathcal{M}(W_{t-1}^{1:N})$. After an $\mathcal{O}(N)$ -cost initialisation, new draws can be simulated from the proposal in amortised $\mathcal{O}(1)$ time; see [Chopin and Papaspiliopoulos \(2020, Python Corner, Chapter 9\)](#), see also [Douc et al. \(2011, Appendix B.1\)](#) for an alternative algorithm with an $\mathcal{O}(\log N)$ cost per draw. The resulting procedure is summarised in [Algorithm 4](#). Compared to traditional FFBS or PaRIS implementations, these rejection-based variants have a random execution time that is more difficult to analyse. Under [Assumption 2](#), [Douc et al. \(2011\)](#) and [Olsson and Westerborn \(2017\)](#) derive an $\mathcal{O}(N\bar{M}_h/\bar{M}_\ell)$ expected complexity. However, the general picture, where the state space is not compact and only [Assumption 4](#) holds, is less clear.

Algorithm 4: Pure rejection sampler for simulating from $B_t^{N,\text{FFBS}}(i_t, di_{t-1})$

Input: Particles $X_{t-1}^{1:N}$ and weights $W_{t-1}^{1:N}$ at time $t-1$; particle $X_t^{i_t}$ at time t ; constant \bar{M}_h ; pre-initialised $\mathcal{O}(1)$ sampler for $\mathcal{M}(W_{t-1}^{1:N})$

repeat

$\mathcal{I}_{t-1} \sim \mathcal{M}(W_{t-1}^{1:N})$ using the pre-initialised $\mathcal{O}(1)$ sampler
 $U \sim \text{Unif}[0, 1]$

until $U \leq m_t(X_{t-1}^{\mathcal{I}_{t-1}}, X_t^{i_t})/\bar{M}_h$

Output: \mathcal{I}_{t-1} , which is distributed according to $B_t^{N,\text{FFBS}}(i_t, di_{t-1})$.

The present subsection intends to fill this gap. Our main focus is the PaRIS algorithm of which the presentation is simpler. Results for the FFBS algorithm can be found in [Supplement B](#). We restrict ourselves to the case where $\mathcal{X}_t = \mathbb{R}^{d_t}$, although extensions to other non compact state spaces are possible. Only the bootstrap particle filter is considered, and results from this section do *not* extend trivially to other filtering algorithms. In [Section 5](#), we shall employ different types of particle filters and see that the performance could change from one type to another, which is an additional weak point of rejection-based algorithms.

Assumption 5. The hidden state X_t is defined on the space $\mathcal{X}_t = \mathbb{R}^{d_t}$. The measure $\lambda_t(dx_t)$ with respect to which the transition density $m_t(x_{t-1}, x_t)$ is defined (cf. [\(2\)](#)) is the Lebesgue measure on \mathbb{R}^{d_t} .

This assumption together with the condition $m_t(x_{t-1}, x_t) > 0$ of [Assumption 4](#) ensures that the state space model is “truly non-compact”. Indeed, if $m_t(x_{t-1}, x_t)$ is zero whenever $x_{t-1} \notin \mathcal{C}_{t-1}$ or $x_t \notin \mathcal{C}_t$, where \mathcal{C}_{t-1} and \mathcal{C}_t are respectively two compact subsets of $\mathbb{R}^{d_{t-1}}$ and \mathbb{R}^{d_t} , then we are basically reduced to a state space model where $\mathcal{X}_{t-1} = \mathcal{C}_{t-1}$ and $\mathcal{X}_t = \mathcal{C}_t$.

3.1. Complexity of PaRIS algorithm with pure rejection sampling. We consider the PaRIS algorithm (i.e. [Algorithm 3](#) using the $B_t^{N,\text{PaRIS}}$ kernels). [Algorithm 5](#) provides a concrete description of the resulting procedure, using the bootstrap particle filter. At each time t , let $\tau_t^{n,\text{PaRIS}}$ be the number of rejection trials required to sample from $B_t^{N,\text{FFBS}}(n, dm)$. We then have

$$(16) \quad \tau_t^{n,\text{PaRIS}} \mid \mathcal{F}_{t-1}, X_t^n \sim \text{Geo} \left(\frac{\sum_i W_{t-1}^i m_t(X_{t-1}^i, X_t^n)}{\bar{M}_h} \right)$$

with \bar{M}_h defined in [Assumption 4](#).

By exchangeability of particles, the expected cost of the PaRIS algorithm at step t is proportional to $N\tilde{N}\mathbb{E}[\tau_t^{1,\text{PaRIS}}]$, where \tilde{N} is a fixed user-chosen parameter. Occasionally, X_t^1 falls into an unlikely region of \mathbb{R}^d and the acceptance rate becomes low. In other words, $\tau_t^{1,\text{PaRIS}}$ is a mixture of geometric distribution, some components of which might have a large expectation. Unfortunately, these inefficiencies add up and produce an unbounded execution time in expectation, as shown in the following proposition.

Proposition 2. *Under [Assumptions 4](#) and [5](#), the version of [Algorithm 5](#) using the pure rejection sampler satisfies $\mathbb{E}[\tau_t^{1,\text{PaRIS}}] = \infty$, where $\tau_t^{1,\text{PaRIS}}$ is defined in [\(16\)](#).*

Algorithm 5: Concrete implementation of PaRIS algorithm (i.e. Algorithm 3 with the $B_t^{N,\text{PaRIS}}$ backward kernel) using the bootstrap particle filter

Input: Particles $X_{t-1}^{1:N}$; weights $W_{t-1}^{1:N}$; vector S_{t-1}^N in \mathbb{R}^N ; pre-initialised sampler for $\mathcal{M}(W_{t-1}^{1:N})$; function ψ_t (cf. (9)); user-specified parameter \tilde{N}

for $n \leftarrow 1$ **to** N **do**

$A_t^n \sim \mathcal{M}(W_{t-1}^{1:N})$ (\star)
 $X_t^n \sim M_t(X_{t-1}^{A_t^n}, dx_t)$
 Simulate $J_t^{n,1:\tilde{N}}$ i.i.d. $B_t^{N,\text{FFBS}}(n, dn')$ using either the pure rejection sampler (Algorithm 4) or the hybrid rejection sampler (Algorithm 6)
 $S_t^N[n] \leftarrow \tilde{N}^{-1} \sum_{\tilde{n}=1}^{\tilde{N}} \left\{ S_{t-1}^N[J_t^{n,\tilde{n}}] + \psi_t(X_{t-1}^{J_t^{n,\tilde{n}}}, X_t^n) \right\}$

for $n \leftarrow 1$ **to** N **do**

$W_t^n \leftarrow G_t(X_t^n) / \sum_i G_t(X_t^i)$

$\mu_t^N \leftarrow \sum_{n=1}^N W_t^n S_t^N(n)$

Initialise a sampler for $\mathcal{M}(W_t^{1:N})$

Output: Estimate μ_t^N of $\mathbb{E}[\varphi(X_{0:t}) | Y_{0:t}]$; particles $X_t^{1:N}$; weights $W_t^{1:N}$; vector S_t^N in \mathbb{R}^N and pre-initialised sampler $\mathcal{M}(W_t^{1:N})$ for the next iteration

Proof. We have

$$\begin{aligned}
 \mathbb{E}[\tau_t^{1,\text{PaRIS}}] &= \bar{M}_h \mathbb{E} \left[\frac{1}{\sum_n m_t(X_{t-1}^n, X_t^1) W_{t-1}^n} \right] \quad \text{via (16)} \\
 &= \bar{M}_h \mathbb{E} \left[\mathbb{E} \left[\frac{1}{\sum_n m_t(X_{t-1}^n, X_t^1) W_{t-1}^n} \middle| \mathcal{F}_{t-1} \right] \right] \\
 &= \bar{M}_h \mathbb{E} \left[\int_{\mathcal{X}_t} \frac{1}{\sum_n m_t(X_{t-1}^n, x) W_{t-1}^n} \left(\sum m_t(X_{t-1}^n, x) W_{t-1}^n \right) \lambda_t(dx) \right] \\
 &= \bar{M}_h \mathbb{E} \left[\int_{\mathcal{X}_t} 1 \times \lambda_t(dx) \right] = \infty \quad \text{by Assumption 5.}
 \end{aligned}$$

□

In highly parallel computing architectures, each processor only handles one or a small number of particles. As such, the heavy-tailed nature of the execution time means that a few machines might prevent the whole system from moving forward. In *all* computing architectures, an execution time without expectation is essentially unpredictable. A common practice to estimate execution time is to run a certain algorithm with a small number N of particles, then “extrapolate” to the N_{final} of the definitive run. However, as $\mathbb{E}[\tau_t^{1,\text{PaRIS}}]$ is infinite for any N , it is unclear what kind of information we might get from preliminary runs. In Supplement B, besides studying the execution time of rejection-based implementations of the FFBS algorithm, we will delve deeper into the difference between the non-parallel and parallel settings.

From the proof of Proposition 2, it is clear that the quantity $\sum_n W_{t-1}^n m_t(X_{t-1}^n, x_t)$ will play a key role in the upcoming developments. We thus define it formally.

Definition 1. The true predictive density function r_t and its approximation r_t^N are defined as

$$\begin{aligned}
 r_t(x_t) &:= \frac{(\mathbb{Q}_{t-1} M_t)(dx_t)}{\lambda_t(dx_t)} \\
 r_t^N(x_t) &:= \sum W_{t-1}^n m_t(X_{t-1}^n, x_t)
 \end{aligned}$$

where the first equation is understood in the sense of the Radon-Nikodym derivative and the density $m_{t-1}(x_{t-1}, x_t)$ is defined with respect to the dominating measure $\lambda_t(dx_t)$ on \mathcal{X}_t (cf. (2)).

3.2. Hybrid rejection sampling. To solve the aforementioned issues of the pure rejection sampling procedure, we propose a hybrid rejection sampling scheme. The basic observation is that, for a single m , direct simulation (e.g. via CDF inversion) of $B_t^{N,\text{FFBS}}(i_t, di_{t-1})$ costs $\mathcal{O}(N)$. Thus, once $K = \mathcal{O}(N)$

rejection sampling trials have been attempted, one should instead switch to a direct simulation method. In other words, it does not make sense (at least asymptotically) to switch to direct sampling after K trials if $K \ll \mathcal{O}(N)$ or $K \gg \mathcal{O}(N)$. The validity of this method is established in the following proposition, where we actually allow K to depend on trials drawn so far. The proof, which is *not* an immediate consequence of the validity of ordinary rejection sampling, is given in Supplement E.4.

Proposition 3. *Let $\mu_0(x)$ and $\mu_1(x)$ be two probability densities defined on some measurable space \mathcal{X} with respect to a dominating measure $\lambda(dx)$. Suppose that there exists $C > 0$ such that $\mu_1(x) \leq C\mu_0(x)$. Let $(X_1, U_1), (X_2, U_2), \dots$ be a sequence of i.i.d. random variables distributed according to $\mu_0 \otimes \text{Unif}[0, 1]$ and let $X^* \sim \mu_1$ be independent of that sequence. Put*

$$K^* := \inf \left\{ n \in \mathbb{Z}_{\geq 1} \text{ such that } U_n \leq \frac{\mu_1(X_n)}{C\mu_0(X_n)} \right\}$$

and let K be any stopping time with respect to the natural filtration associated with the sequence $\{(X_n, U_n)\}_{n=1}^\infty$. Let Z be defined as X_{K^} if $K^* \leq K$ and X^* otherwise. Then Z is μ_1 -distributed.*

Proposition 3 thus allows users to pick $K = \alpha N$, where $\alpha > 0$ might be chosen somehow adaptively from earlier trials. In the following, we only consider the simple rule $K = N$, which does not induce any loss of generality in terms of the asymptotic behaviour and is easy to implement. The resulting iteration is described in Algorithm 6.

Algorithm 6: Hybrid rejection sampler for simulating from $B_t^{N, \text{FFBS}}(i_t, di_{t-1})$

Input: Particles $X_{t-1}^{1:N}$ and weights $W_{t-1}^{1:N}$ at time $t-1$; particle $X_t^{i_t}$ at time t ; constant \bar{M}_h ;
pre-initialised $\mathcal{O}(1)$ sampler for $\mathcal{M}(W_{t-1}^{1:N})$

accepted \leftarrow False

for $i \leftarrow 1$ **to** N **do**

$\mathcal{I}_{t-1} \sim \mathcal{M}(W_{t-1}^{1:N})$ using the pre-initialised $\mathcal{O}(1)$ sampler
 $U \sim \text{Unif}[0, 1]$
if $U \leq m_t(X_{t-1}^{\mathcal{I}_{t-1}}, X_t^{i_t}) / \bar{M}_h$ **then**
 $\left| \begin{array}{l} \textit{accepted} \leftarrow \text{True} \\ \textbf{break} \end{array} \right.$

if not *accepted* **then**

$\mathcal{I}_{t-1} \sim \mathcal{M}(W_{t-1}^n m(X_{t-1}^n, X_t^{i_t}))$

Output: \mathcal{I}_{t-1} , which is distributed according to $B_t^{N, \text{FFBS}}(i_t, di_{t-1})$.

When applied in the context of Algorithm 5, Algorithm 6 gives a smoother of expected complexity proportional to

$$N\tilde{N}\mathbb{E}[\min(\tau_t^{1, \text{PaRIS}}, N)]$$

at time t , where $\tau_t^{1, \text{PaRIS}}$ is defined in (16)). This quantity is no longer infinite, but its growth when $N \rightarrow \infty$ might depend on the model. Still, in all cases, it remains strictly larger than $\mathcal{O}(N)$ and strictly smaller than $\mathcal{O}(N^2)$. Perhaps more surprisingly, in linear Gaussian models (see Supplement A.1 for detailed notations), the smoother is of near-linear complexity (up to log factors). The following two theorems formalise these claims.

Assumption 6. The predictive density r_t of X_t given $Y_{0:t-1}$ and the potential function G_t are continuous functions on \mathbb{R}^{d_t} . The transition density $m_t(x_{t-1}, x_t)$ is a continuous function on $\mathbb{R}^{d_{t-1}} \times \mathbb{R}^{d_t}$.

Theorem 3.1. *Under Assumptions 1, 4, 5 and 6, the version of Algorithm 5 using the hybrid rejection sampler (Algorithm 6) satisfies $\lim_{N \rightarrow \infty} \mathbb{E}[\min(\tau_t^{1, \text{PaRIS}}, N)] = \infty$ and $\lim_{N \rightarrow \infty} \mathbb{E}[\min(\tau_t^{1, \text{PaRIS}}, N)]/N = 0$, where $\tau_t^{1, \text{PaRIS}}$ is defined in (16).*

Theorem 3.2. *We assume the same setting as Theorem 3.1. In linear Gaussian state space models (Supplement A.1), we have $\mathbb{E}[\min(\tau_t^{1,\text{PaRIS}}, N)] = \mathcal{O}((\log N)^{d_t/2})$.*

While Proposition 2 shows that $\tau_t^{1,\text{PaRIS}}$ has infinite expectation, Theorem 3.2 implies that its N -thresholded version only displays a slowly increasing mean. To give a very rough intuition on the phenomenon, consider $X \sim \mathcal{N}(0, 1)$. Then

$$\mathbb{E} \left[e^{X^2/2} \right] = \int_{\mathbb{R}} e^{x^2/2} \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx = +\infty$$

whereas

$$\begin{aligned} \mathbb{E} \left[\min(e^{X^2/2}, N) \right] &= \int_{\mathbb{R}} \min(e^{x^2/2}, N) \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\ (17) \quad &= \int_{|x| \leq \sqrt{2 \log N}} \frac{1}{\sqrt{2\pi}} dx + N \int_{|x| > \sqrt{2 \log N}} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\ &\leq \sqrt{\frac{4 \log N}{\pi}} + \frac{1}{\sqrt{\pi \log N}} \end{aligned}$$

using the bound $\mathbb{P}(X > x) \leq \frac{e^{-x^2/2}}{x\sqrt{2\pi}}$ for $x > 0$. The main technical difficulty of the proof of Theorem 3.2 (see Supplement E.6) is to perform this kind of argument under the error induced by the finite sample size particle approximation. In the language of this oversimplified example, we want (17) to hold when X does not follow $\mathcal{N}(0, 1)$ any more, but only an N -dependent approximation of it.

4. EFFICIENT BACKWARD KERNELS

4.1. MCMC Backward Kernels. This subsection analyses and extends the MCMC backward kernel defined in Example 3. As we remarked there, the matrix $\hat{B}_t^{N,\text{IMH}}$ is not sparse and even has some expensive-to-evaluate entries. We thus reserve it for use in the off-line smoother (Algorithm 2) whereas in the on-line scenario (Algorithm 3), we use its PaRIS-like counterpart

$$(18) \quad \hat{B}_t^{N,\text{IMHP}}[i_t, i_{t-1}] := \frac{1}{\tilde{N}} \sum_{\tilde{n}=1}^{\tilde{N}} \mathbb{1} \left\{ i_{t-1} = \tilde{J}_t^{i_t, \tilde{n}} \right\}$$

where $\tilde{J}_t^{i_t, 1:\tilde{N}}$ is an independent Metropolis-Hastings chain started at $J_t^{i_t, 1} := A_t^{i_t}$, targeting the measure $B_t^{N,\text{FFBS}}(i_t, di_{t-1})$ and using the proposal distribution $\mathcal{M}(W_{t-1}^{1:\tilde{N}})$. Thus, the parameter \tilde{N} signifies that $\tilde{N} - 1$ MCMC steps are applied to $A_t^{i_t}$, and we shall use the same convention for the kernel $B_t^{N,\text{IMH}}$. In both cases, the complexity of the corresponding algorithms are $\mathcal{O}((\tilde{N} - 1)N)$ which is equivalent to $\mathcal{O}(N)$ as long as \tilde{N} remains fixed when $N \rightarrow \infty$.

The validity and the stability of $\hat{B}_t^{N,\text{IMH}}$ and $\hat{B}_t^{N,\text{IMHP}}$ are established in the following proposition (proved in Supplement E.9). For simplicity, only the case $\tilde{N} = 2$ is examined, but as a matter of fact, the proposition remains true for $\tilde{N} \geq 2$.

Proposition 4. *The kernels $B_t^{N,\text{IMH}}$ and $B_t^{N,\text{IMHP}}$ with $\tilde{N} = 2$ satisfy the hypotheses of Theorem 2.1 and, under Assumptions 2 and 3, those of Theorem 2.2. Hence, their respective uses in Algorithms 2 and 3 guarantee a convergent and stable smoother.*

From a theoretical viewpoint, Proposition 4 is the first result establishing the stability for the use of MCMC moves inside backward sampling. It relies on technical innovations that we have explained in Section 2.6, in particular after the statement of Theorem 2.2.

From a practical viewpoint, the advantages of independent Metropolis-Hastings MCMC kernels compared to the rejection samplers of Section 3 are the dispensability of specifying an explicit \bar{M}_h and the deterministic $\mathcal{O}(N)$ nature of the execution time. In practice, we observe that the MCMC smoothers are usually 10-20 times faster than the rejection sampling-based counterparts (see e.g. Figure 4) while producing essentially the same sample quality. Finally, it is not hard to imagine situations where some proposal smarter than $\mathcal{M}(W_{t-1}^{1:\tilde{N}})$ would be beneficial. However, we only consider that one here, mainly because it already performs satisfactorily in our numerical examples.

4.2. Dealing with intractable transition densities.

4.2.1. *Intuition and formulation.* The purpose of backward sampling is to re-generate, for each particle, a new ancestor that is different from that of the filtering step. However, backward sampling is infeasible if the transition density $m_t(x_{t-1}, x_t)$ cannot be calculated. To get around this, we modify the particle filter so that each particle might, in some sense, have two ancestors right from the *forward* pass.

Consider the standard PF (Algorithm 1). Among the N resampled particles $X_{t-1}^{1:N}$, let us track two of them, say x_{t-1} and x'_{t-1} for simplicity. The move step of Algorithm 1 will push them through M_t using independent noises, resulting in x_t and x'_t (that is, given x_{t-1} and x'_{t-1} , we have $x_t \sim M_t(x_{t-1}, \cdot)$ and $x'_t \sim M_t(x'_{t-1}, \cdot)$ such that x_t and x'_t are independent). Thus, for e.g. linear Gaussian models, we have $\mathbb{P}(x_t = x'_t) = 0$. However, if the two simulations $x_t \sim M_t(x_{t-1}, \cdot)$ and $x'_t \sim M_t(x'_{t-1}, \cdot)$ are done with specifically correlated noises, it can happen that $\mathbb{P}(x_t = x'_t) > 0$. The joint distribution (x_t, x'_t) given (x_{t-1}, x'_{t-1}) is called a *coupling* of $M_t(x_{t-1}, \cdot)$ and $M_t(x'_{t-1}, \cdot)$; the event $x_t = x'_t$ is called the *meeting* event and we say that the coupling is *successful* when it occurs. In that case, the particle x_t automatically has two ancestors x_{t-1} and x'_{t-1} at time $t-1$ without needing any backward sampling.

The precise formulation of the modified forward pass is detailed in Algorithm 7. It consists of building in an on-line manner the backward kernels $B_t^{N, \text{ITR}}$ (where ITR stands for “intractable”). The main interest of this algorithm lies in the fact that while the function m_t may prove impossible to evaluate, it is usually possible to make x_t and x'_t meet by correlating somehow the random numbers used in their simulations. One typical example which this article focuses on is the coupling of continuous-time processes, but it is useful to keep in mind that Algorithm 7 is conceptually more general than that.

Algorithm 7: Modified forward pass for smoothing of intractable models (one time step)

Input: Feynman-Kac model (1), particles $X_{t-1}^{1:N}$ and weights $W_{t-1}^{1:N}$ that approximate the filtering distribution at time $t-1$

for $n \leftarrow 1$ **to** N **do**

Resample. Simulate $(A_t^{n,1}, A_t^{n,2})$ such that marginally each component is distributed according to $\mathcal{M}(W_{t-1}^{1:N})$

Move. Simulate $(X_t^{n,1}, X_t^{n,2})$ such that marginally the two components are distributed respectively according to $M_t(X_{t-1}^{A_t^{n,1}}, dx_t)$ and $M_t(X_{t-1}^{A_t^{n,2}}, dx_t)$

Choose $L \sim \text{Uniform}(\{1, 2\})$

Set $X_t^n \leftarrow X_t^{n,L}$

Calculate backward kernel.

if $X_t^{n,1} = X_t^{n,2}$ **then**

$B_t^{N, \text{ITR}}(n, di_{t-1}) \leftarrow (\delta \{A_t^{n,1}\} + \delta \{A_t^{n,2}\}) / 2$

else

$B_t^{N, \text{ITR}}(n, di_{t-1}) \leftarrow \delta \{A_t^{n,L}\}$

Reweight. Set $\omega_t^n \leftarrow G_t(X_t^n)$ for $n = 1, 2, \dots, N$

Set $\ell_t^N \leftarrow \sum_{n=1}^N \omega_t^n / N$

Set $W_t^n \leftarrow \omega_t^n / N \ell_t^N$ for $n = 1, 2, \dots, N$

Output: Particles $X_t^{1:N}$ and weights $W_t^{1:N}$ that approximate the filtering distribution at time t ;
backward kernel $B_t^{N, \text{ITR}}$ that can be used in either Algorithm 2 or 3

4.2.2. *Validity and stability.* The consistency of Algorithm 7 follows straightforwardly from Theorem 2.1. To produce a stable routine however, some conditions must be imposed on the couplings $(A_t^{n,1}, A_t^{n,2})$ and $(X_t^{n,1}, X_t^{n,2})$. We want $A_t^{n,1}$ to be different from $A_t^{n,2}$ as frequently as possible. On the contrary, we aim for a coupling of the two distributions $M_t(X_{t-1}^{A_t^{n,1}}, \cdot)$ and $M_t(X_{t-1}^{A_t^{n,2}}, \cdot)$ with high success rate so as to maximise the probability that $X_t^{n,1} = X_t^{n,2}$.

Assumption 7. There exists an $\varepsilon_A > 0$ such that $\mathbb{P}(A_t^{n,1} \neq A_t^{n,2} | X_{t-1}^{1:N}) \geq \varepsilon_A$.

Assumption 8. There exists an $\varepsilon_D > 0$ such that

$$\mathbb{P}(X_t^{n,2} = X_t^{n,1} | X_{t-1}^{1:N}, A_t^{n,1}, A_t^{n,2}, X_t^{n,1}) \geq \varepsilon_D \left(1 \wedge \frac{m_t(X_{t-1}^{A_t^{n,2}}, X_t^{n,1})}{m_t(X_{t-1}^{A_t^{n,1}}, X_t^{n,1})} \right).$$

The letters A and D in ε_A and ε_D stand for “ancestors” and “dynamics”. Assumption 8 means that the user-chosen coupling of $M_t(X_{t-1}^{A_t^{n,1}}, \cdot)$ and $M_t(X_{t-1}^{A_t^{n,2}}, \cdot)$ must be at least as ε_D times as efficient as their maximal couplings. For details on this interpretation, see Proposition 10 in the Supplement. In Lemma E.12, we also show that in spite of its appearance, Assumption 8 is actually symmetric with regards to $X_t^{n,1}$ and $X_t^{n,2}$.

We are now ready to state the main theorem of this subsection (see Supplement E.11 for a proof).

Theorem 4.1. *The kernels $B_t^{N,\text{ITR}}$ generated by Algorithm 7 satisfy the hypotheses of Theorem 2.1. Thus, under Assumption 1, Algorithm 7 provides a consistent smoothing estimate. If, in addition, the Feynman-Kac model (1) satisfies Assumptions 2 and 3 and the user-chosen couplings satisfy Assumptions 7 and 8, the kernels $B_t^{N,\text{ITR}}$ also fulfil (13) and the smoothing estimates generated by Algorithm 7 are stable.*

4.2.3. Good ancestor couplings. It is notable that Assumption 7 only considers the event $A_t^{n,1} \neq A_t^{n,2}$, which is a pure index condition that does not take into account the underlying particles $X_{t-1}^{A_t^{n,1}}$ and $X_{t-1}^{A_t^{n,2}}$. Indeed, if smoothing algorithms prevent degeneracy by creating multiple ancestors for a particle, we would expect that their separation (i.e. that they are far away in the state space \mathcal{X}_{t-1} , e.g. \mathbb{R}^d) is critical to the performance. Surprisingly, it is unnecessary: two very close particles (in \mathbb{R}^d) at time $t-1$ may have ancestors far away at time $t-2$ thanks to the mixing of the model.

We advise choosing an ancestor coupling $(A_t^{n,1}, A_t^{n,2})$ such that the distance between $X_{t-1}^{A_t^{n,1}}$ and $X_{t-1}^{A_t^{n,2}}$ is small. It will then be easier to design a dynamic coupling of $M_t(X_{t-1}^{A_t^{n,1}}, \cdot)$ and $M_t(X_{t-1}^{A_t^{n,2}}, \cdot)$ with a high success rate. Furthermore, simulating the dynamic coupling with two close rather than far away starting points can also take less time when, for instance, the dynamic involves multiple intermediate steps, but the two processes couple early. One way to achieve an ancestor coupling with the aforementioned property is to first simulate $A_t^{n,1} \sim \mathcal{M}(W_{t-1}^{1:N})$, then move $A_t^{n,1}$ through an MCMC algorithm which keeps invariant $\mathcal{M}(W_{t-1}^{1:N})$ and set the result to $A_t^{n,2}$. It suffices to use a proposal looking at indices whose underlying particles are close (in \mathbb{R}^d) to $X_{t-1}^{A_t^{n,1}}$. Finding nearby particles are efficient if they are first sorted using the Hilbert curve, hashed using locality-sensitive hashing or put in a KD-tree (see Samet, 2006, for a comprehensive review). In the context of particle filters, such techniques have been studied for different purposes in Gerber and Chopin (2015), Jacob et al. (2019) and Sen et al. (2018).

4.2.4. Conditionally-correlated version. In Algorithm 7, the ancestor pairs $(A_t^{n,1}, A_t^{n,2})_{n=1}^N$ are conditionally independent given \mathcal{F}_t^- and the same holds for the particles $(X_t^n)_{n=1}^N$. These conditional independences allow easier theoretical analysis, in particular, the casting of Algorithm 7 in the framework of Theorems 2.1 and 2.2. However, they are not optimal for performance in two important ways: (a) they do not allow keeping both $X_t^{n,1}$ and $X_t^{n,2}$ when the two are not equal, and (b) the set of ancestor variables $(A_t^{n,1})_{n=1}^N$ is multinomially resampled from $\{1, 2, \dots, N\}$ with weights $W_{t-1}^{1:N}$. We know that multinomial resampling is not the ideal scheme, see Supplement C.1 for discussion.

Consequently, in practice, we shall allow ourselves to break free from conditional independence. The resulting procedure is described in Algorithm 9 (Supplement C). Despite a lack of rigorous theoretical support, this is the algorithm that we will use in Section 5 since it enjoys better performance and it constitutes a fair comparison with practical implementations of the standard particle filter, which are mostly based on alternative resampling schemes.

5. NUMERICAL EXPERIMENTS

5.1. Linear Gaussian state-space models. Linear Gaussian models constitute a particular class of state space models. They are characterised by Markov dynamics that are Gaussian and observations that are projection of hidden states plus some Gaussian noises. Supplement A.1 defines, for different components of these models, the notations that we shall use here. In this section, we consider an instance described in Guarniero et al. (2017), where the matrix F_X satisfies $F_X[i, j] = \alpha^{1+|i-j|}$ for some α . We consider the problem with $\dim_X = \dim_Y = 2$ and the observations are noisy versions of the hidden states with C_Y being σ_Y^2 times the identity matrix of size 2. Unless otherwise specified, we take $\alpha = 0.4$ and $\sigma_Y^2 = 0.5$.

In this section, we focus on the performance of different online smoothers based on either genealogy tracking, pure/hybrid rejection sampling or MCMC. Rejection-based online smoothing amounts to the PaRIS algorithm, for which we use $\tilde{N} = 2$ for the $B_t^{N, \text{PaRIS}}$ kernel. We take $T = 3000$ and simulate the data from the model. The benchmark additive function is simply $\varphi_t(x_{0:t}) = \sum_{s=0}^t x_s(0)$ where $x_s(0)$ is the first coordinate of the \mathbb{R}^2 vector $x_s = [x_s(0), x_s(1)]$. For a study of offline smoothers (including FFBS), see Supplement D.1. In all programs here and there, we choose $N = 1000$ and use systematic resampling for the forward particle filters (see section C.1). Regarding MCMC smoothers, we employ the kernels $B_t^{N, \text{IMH}}$ or $B_t^{N, \text{IMHP}}$ consisting of only one MCMC step. All results are based on 150 independent runs.

Although our theoretical results are only proved for the bootstrap filter, we stress throughout that some of them extend to other filters as well. Therefore, we will also consider guided particle filters in the simulations. An introduction to this topic can be found in Chopin and Papaspiliopoulos (2020, Chapter 10.3.2), where the expression for the optimal proposal is also provided. In linear Gaussian models, this proposal is fully tractable and is the one we use.

To present efficiently the combination of two different filters (bootstrap and guided) and four different algorithms (naive genealogy tracking, pure/hybrid rejection and MCMC) we use the following abbreviations: “B” for bootstrap, “G” for guided, “N” for naive genealogy tracking, “P” for pure rejection, “H” for hybrid rejection and “M” for MCMC. For instance, the algorithm referred to as “BM” uses the bootstrap filter for the forward pass and the MCMC backward kernels to perform smoothing. Furthermore, the letter “R” will refer to the rejection kernel whenever the distinction between pure rejection and hybrid rejection is not necessary. (Recall that the two rejection methods produce estimators with the same distribution.)

Figure 2 shows the squared interquartile range for the online smoothing estimates $Q_t(\varphi_t)$ with respect to t . It verifies the rates of Theorem 2.2, although linear Gaussian models are not strongly mixing in the sense of Assumptions 2 and 3: the grid lines hint at a variance growth rate of $\mathcal{O}(T)$ for the MCMC and reject-based smoothers and of $\mathcal{O}(T^2)$ for the genealogy tracking ones. Unsurprisingly guided filters have better performance than bootstrap.

Figure 3 show box-plots of the execution time (divided by NT) for different algorithms over 150 runs. By execution time, we mean the number of Markov kernel transition density evaluations. We see that the bootstrap particle filter coupled with pure rejection sampling has a very heavy-tailed execution time. This behaviour is expected as per Proposition 2. Using the guided particle filter seems to fare better, but Figure 4 (for the same model but with $\sigma_Y^2 = 2$) makes it clear that this cannot be relied on either. Overall, these results highlight two fundamental problems with pure rejection sampling: the computational time has heavy tails and depends on the type of forward particle filter being used.

On the other hand, hybrid rejection sampling, despite having random execution time in principle, displays a very consistent number of transition density evaluations over different independent runs. Thus it is safe to say that the algorithm has a virtually deterministic execution time. The catch is that the average computational load (which is around 16 in Figure 3) cannot be easily calculated beforehand. In any case, it is much larger than the value 1 of MCMC smoothers (since only 1 MCMC step is performed in the kernel $B_t^{N, \text{IMHP}}$); whereas the performance (Figure 2) is comparable.

The bottom line is that MCMC smoothers should be the default option, and one MCMC step seems to be enough. If for some reason one would like to use rejection-based methods, using hybrid rejection is a must.




FIGURE 2. Squared interquartile range of the estimators $Q_t(\varphi_t)$ with respect to t , for different online smoothing algorithms. The model is linear Gaussian with parameters specified in section 5.1. See text for full explanation of the legend. For readability, the curves are down-sampled to 50 points before being drawn.




FIGURE 3. Box plots (based on 150 runs) of averaged execution times (numbers of transition density evaluations divided by NT) for different algorithms on the linear Gaussian model of section 5.1. Left: original figure, right: zoomed-in version.




FIGURE 4. Same as Figure 3, but for the modified model where $\sigma_Y^2 = 2$.

5.2. Lotka-Volterra SDE. Lotka-Volterra models (originated in [Lotka, 1925](#) and [Volterra, 1928](#)) describe the population fluctuation of species due to natural birth and death as well as the consumption of one species by others. The emblematic case of two species is also known as the predator-prey model. In this subsection, we study the stochastic differential equation (SDE) version that appears in [Hening and Nguyen \(2018\)](#). Let $X_t = (X_t(0), X_t(1))$ represent respectively the populations of the prey and the predator at time t and let us consider the dynamics

$$(19) \quad \begin{cases} dX_t(0) = [\beta_0 X_t(0) - \frac{1}{2}\tau_0 [X_t(0)]^2 - \tau_1 X_t(0)X_t(1)]dt + X_t(0)dE_t(0) \\ dX_t(1) = [-\beta_1 X_t(1) + \tau_1 X_t(0)X_t(1)]dt + X_t(1)dE_t(1) \end{cases}$$

where $E_t = \Gamma W_t$ with W_t being the standard Brownian motion in \mathbb{R}^2 and Γ being some 2×2 matrix. The parameters β_0 and β_1 are the natural birth rate of the prey and death rate of the predator. The predator interacts with (eats) the prey at rate τ_1 . The quantity τ_0 encodes intra-species competition in the prey population. The $\frac{1}{2}$ in its parametrisation is to line up with the Lotka Volterra jump process in \mathbb{Z}^2 where the population sizes are integers and the interaction term becomes $\tau_0 X_t(0)[X_t(0) - 1]/2$.

The state space model is comprised of the process X_t and its noisy observations Y_t recorded at integer times. The Markov dynamics cannot be simulated exactly, but can be approximated through (Euler) discretisation. Nevertheless, the Euler transition density $m_t^E(x_{t-1}, x_t)$ remains intractable (unless the step size is exactly 1). Thus, the algorithms presented in Subsection 4.2 are useful. The missing bit is a method to efficiently couple $m_t^E(x_{t-1}, \cdot)$ and $m_t^E(x'_{t-1}, \cdot)$, which we carefully describe in Supplement D.2.1.

We consider the model with $\tau_0 = 1/800$, $\tau_1 = 1/400$, $\beta_0 = 0.3125$ and $\beta_1 = 0.25$. The matrix Γ is such that the covariance matrix of E_1 is $\begin{bmatrix} 1/100 & 1/200 \\ 1/200 & 1/100 \end{bmatrix}$. The observations are recorded on the log scale with Gaussian error of covariance matrix $\begin{bmatrix} 0.04 & 0.02 \\ 0.02 & 0.04 \end{bmatrix}$. The distribution of X_0 is two-dimensional normal with mean $[100, 100]$ and covariance matrix $\begin{bmatrix} 100 & 50 \\ 50 & 100 \end{bmatrix}$. This choice is motivated by the fact that the preceding parameters give the stationary population vector $[100, 100]$. According to [Hening and Nguyen \(2018\)](#), they also guarantee that neither animal goes extinct almost surely as $t \rightarrow \infty$.

By discretising (19) with time step $\delta = 1$, one can get some very rough intuition on the dynamics. For instance, per second there are about 31 preys born. Approximately the same number die (to maintain equilibrium), of which 6 die due to internal competition and 25 are eaten by the predator. The duration between two recorded observations corresponds more or less to one-third generation of the prey and one-fourth generation of the predator. The standard deviation of the variation due to environmental noise is about 10 individuals per observation period, for each animal.

Again, these intuitions are highly approximate. For readers wishing to get more familiar with the model, Supplement D.2.2 contains real plots of the states and the observations; as well as data on the performance of different smoothing algorithms for moderate values of T . We now showcase the results obtained in a large scale problem where $T = 3000$ and the data is simulated from the model.

We consider the additive function $\varphi_t(x_{0:t}) := \sum_{s=0}^t [x_s(0) - 100]$. Figure 5 represents using box plots the distributions of the estimators for $\mathbb{Q}_T(\varphi_T)$ using either the genealogy tracking smoother (with systematic resampling; see Supplement C.1) or Algorithm 9. Our proposed smoother greatly reduces the variance, at a computational cost which is empirically 1.5 to 2 times greater than the naive method. Since we used Hilbert curve to design good ancestor couplings (see Section 4.2.3), coupling of the dynamics succeeds 80% of the time. As discussed in the aforementioned section, starting two diffusion dynamics from nearby points make them couple earlier, which reduces the computational load afterwards.

Figure 6 plots with respect to t the squared interquartile range of the two methods for the estimation of $\mathbb{Q}_t(\varphi_t)$. Grid lines hint at a quadratic growth for the genealogy tracking smoother (as analysed in [Olsson and Westerborn, 2017](#), Sect. 1) and a linear growth for the kernel $B_t^{N, \text{ITRC}}$ (as described in Theorem 2.2).

Finally, Figure 20 (Supplement D.2.2) shows properties of the effective sample size (ESS) ratio for this model. In a nutshell, while being globally stable (between 40% and 70%), it has a tendency to drift towards near 0 from time to time due to unusual data points. At these moments, resampling kills most




FIGURE 5. Box plot of estimators (over 50 independent runs with $N = 1000$ particles) for $\mathbb{Q}_T(\varphi_T)$ in the Lotka-Volterra SDE model with $T = 3000$. They are calculated using either the naive genealogy tracking smoother or our smoother developed for intractable models (Algorithm 9).




FIGURE 6. Squared interquartile range for the genealogy tracking smoother and our proposed one. Same context as in Figure 5

of the particles and aggravates the degeneracy problem for the naive smoother. As we have seen in the above figures, systematic resampling is not enough to mitigate this in the long run.

6. CONCLUSION

6.1. Practical recommendations. Our first recommendation does not concern the smoothing algorithm per se. It is of paramount importance that the particle filter used in the preliminary filtering step performs reasonably well, since its output defines the support of the approximations generated by the subsequent smoothing algorithm. (Standard recommendations to obtain good performance from a particle filter are to increase N , or to use better proposal distributions, or both.)

When the transition density is tractable, we recommend the MCMC smoother by default (rather than even the standard, $\mathcal{O}(N^2)$ approach). It has a deterministic, $\mathcal{O}(N)$ complexity, it does not require the transition density to be bounded, and it seems to perform well even with one or two MCMC steps. If one still wants to use the rejection smoother instead, it is safe to say that there is no reason not to use the hybrid method.

Although the assumptions under which we prove the stability of the smoothing estimates are strong, the general message still holds. The Markov kernel and the potential functions must make the model forget

its past in some ways. Otherwise, we get an unstable model for which no smoothing methods can work. The rejection sampling – based smoothing algorithms can therefore serve as the ultimate test. Since they simulate exactly independent trajectories given the skeleton, there is no hope to perform better, unless one switches to another family of smoothing algorithms.

For intractable models, the key issue is to design couplings with high meeting probability. Fortunately, the inherent chaos of the model makes it possible to choose two very close starting points for the dynamics and thus easy to obtain a reasonable meeting probability. If further difficulties persist, there is a practical (and very heuristic) recipe to test whether one coupling of $M_t(x, \cdot)$ and $M_t(x', \cdot)$ is close to optimal. It consists in approximating $M_t(x, \cdot)$ and $M_t(x', \cdot)$ by Gaussian distributions and deduce the optimal coupling rate from their total variation distance. There is no closed formula for the total variation distance between two Gaussian distributions in high dimensions. However, it can be reliably estimated using the geometric interpretation of the total variation distance being one minus the area of the intersection created by the corresponding density graphs. In this way, one can get a very rough idea of to what extent a certain coupling realises the meeting potential that the two distributions have. If the coupling seems good and the trajectories still look degenerate, it can very well be that the model is unstable.

6.2. Further directions. The major limitation of our work is the exclusive theoretical analysis under the bootstrap particle filter. Moreover, we require that the N new particles generated at step t are conditionally independent given previous particles at time $t - 1$. This excludes practical optimisations like systematic resampling and Algorithm 9. Finally, the backward sampling step is also used in other algorithms (in particular Particle Markov Chain Monte Carlo, see [Andrieu et al., 2010](#)) and it would be interesting to see to what extent our techniques can be applied there.

6.3. Data and code. The code used to run numerical experiments is available at <https://github.com/hai-dang-dau/backward-samplers-code>. Some of the algorithms are already available in an experimental branch of the `particles` Python package at <https://github.com/nchopin/particles>.

ACKNOWLEDGEMENTS

The first author acknowledges a CREST PhD scholarship via AMX funding, and wish to thank the members of his PhD defense jury (Stéphanie Allasonnière, Randal Douc, Arnaud Doucet, Anthony Lee, Pierre del Moral, Christian Robert) for helpful comments on the corresponding chapter in his thesis.

We also thank Adrien Corenflos, Samuel Duffield, the associate editor, and the referees for their comments on a preliminary version of the paper.

REFERENCES

- Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):269–342.
- Beskos, A., Papaspiliopoulos, O., Roberts, G. O., and Fearnhead, P. (2006). Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(3):333–382. With discussions and a reply by the authors.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Information Science and Statistics. Springer, New York.
- Bou-Rabee, N., Eberle, A., and Zimmer, R. (2020). Coupling and convergence for Hamiltonian Monte Carlo. *The Annals of Applied Probability*, 30(3):1209 – 1250.
- Bunch, P. and Godsill, S. (2013). Improved particle approximations to the joint smoothing distribution using Markov chain Monte Carlo. *IEEE Transactions on Signal Processing*, 61(4):956–963.
- Carpenter, J., Clifford, P., and Fearnhead, P. (1999). Improved particle filter for nonlinear problems. *IEE Proc. Radar, Sonar Navigation*, 146(1):2–7.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Statist.*, 32(6):2385–2411.
- Chopin, N. and Papaspiliopoulos, O. (2020). *An Introduction to Sequential Monte Carlo*. Springer Series in Statistics. Springer.
- Corenflos, A. and Särkkä, S. (2022). The Coupled Rejection Sampler. *arXiv preprint arXiv:2201.09585*.
- Del Moral, P. (2004). *Feynman-Kac formulae. Genealogical and interacting particle systems with applications*. Probability and its Applications. Springer Verlag, New York.

- Del Moral, P. (2013). *Mean field simulation for Monte Carlo integration*, volume 126 of *Monographs on Statistics and Applied Probability*. CRC Press, Boca Raton, FL.
- Del Moral, P., Doucet, A., and Singh, S. S. (2010). A backward particle interpretation of Feynman-Kac formulae. *M2AN Math. Model. Numer. Anal.*, 44(5):947–975.
- Del Moral, P. and Miclo, L. (2001). Genealogies and increasing propagation of chaos for Feynman-Kac and genetic models. *Ann. Appl. Probab.*, 11(4):1166–1198.
- Douc, R., Cappé, O., and Moulines, E. (2005). Comparison of resampling schemes for particle filtering. In *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*, pages 64–69. IEEE.
- 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):2109–2145.
- Douc, R., Moulines, E., Priouret, P., and Soulier, P. (2018). *Markov chains*. Springer Series in Operations Research and Financial Engineering. Springer, Cham.
- 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.
- Duffield, S. and Singh, S. S. (2022). Online particle smoothing with application to map-matching. *IEEE Trans. Signal Process.*, 70:497–508.
- Fearnhead, P., Papaspiliopoulos, O., and Roberts, G. O. (2008). Particle filters for partially observed diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 70(4):755–777.
- Fearnhead, P., Wyncoll, D., and Tawn, J. (2010). A sequential smoothing algorithm with linear computational cost. *Biometrika*, 97(2):447–464.
- Gerber, M. and Chopin, N. (2015). Sequential quasi Monte Carlo. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 77(3):509–579.
- Gerber, M., Chopin, N., and Whiteley, N. (2019). Negative association, ordering and convergence of resampling methods. *Ann. Statist.*, 47(4):2236–2260.
- Gloaguen, P., Le Corff, S., and Olsson, J. (2022). A pseudo-marginal sequential Monte Carlo online smoothing algorithm. *Bernoulli*, 28(4):2606–2633.
- Godsill, S. J., Doucet, A., and West, M. (2004). Monte Carlo smoothing for nonlinear times series. *J. Amer. Statist. Assoc.*, 99(465):156–168.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. F, Comm., Radar, Signal Proc.*, 140(2):107–113.
- Guarniero, P., Johansen, A. M., and Lee, A. (2017). The iterated auxiliary particle filter. *J. Amer. Statist. Assoc.*, 112(520):1636–1647.
- Hening, A. and Nguyen, D. H. (2018). Stochastic Lotka-Volterra food chains. *J. Math. Biol.*, 77(1):135–163.
- Jacob, P. E., Lindsten, F., and Schön, T. B. (2019). Smoothing with couplings of conditional particle filters. *Journal of the American Statistical Association*.
- Jacob, P. E., O’Leary, J., and Atchadé, Y. F. (2020). Unbiased Markov chain Monte Carlo methods with couplings. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 82(3):543–600.
- Janson, S. (2011). Probability asymptotics: notes on notation. *arXiv preprint arXiv:1108.3924*.
- Jasra, A., Kamatani, K., Law, K. J., and Zhou, Y. (2017). Multilevel particle filters. *SIAM Journal on Numerical Analysis*, 55(6):3068–3096.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D):35–45.
- Kalman, R. E. and Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Trans. ASME Ser. D. J. Basic Engrg.*, 83:95–108.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Statist.*, 5(1):1–25.
- Lévy, P. (1940). Sur certains processus stochastiques homogènes. *Compositio mathematica*, 7:283–339.
- Lindvall, T. and Rogers, L. C. G. (1986). Coupling of multidimensional diffusions by reflection. *Ann. Probab.*, 14(3):860–872.
- Lotka, A. J. (1925). *Elements of physical biology*. Williams & Wilkins.
- Mastrototaro, A., Olsson, J., and Alenlöv, J. (2021). Fast and numerically stable particle-based online additive smoothing: the adasmooth algorithm. *arXiv preprint arXiv:2108.00432*.
- Mörters, P. and Peres, Y. (2010). *Brownian motion*, volume 30. Cambridge University Press.

- Nordh, J. and Antonsson, J. (2015). A Quantitative Evaluation of Monte Carlo Smoothers. *Technical report*.
- Olsson, J. and Westerborn, J. (2017). Efficient particle-based online smoothing in general hidden Markov models: the PaRIS algorithm. *Bernoulli*, 23(3):1951–1996.
- Olver, F. W. J., Lozier, D. W., Boisvert, R. F., and Clark, C. W., editors (2010). *NIST handbook of mathematical functions*. U.S. Department of Commerce, National Institute of Standards and Technology, Washington, DC; Cambridge University Press, Cambridge. With 1 CD-ROM (Windows, Macintosh and UNIX).
- Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: auxiliary particle filters. *J. Amer. Statist. Assoc.*, 94(446):590–599.
- Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods, 2nd ed.* Springer-Verlag, New York.
- Roberts, G. O. and Rosenthal, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probab. Surv.*, 1:20–71.
- Samet, H. (2006). *Foundations of multidimensional and metric data structures*. Morgan Kaufmann.
- Sen, D., Thiery, A. H., and Jasra, A. (2018). On coupling particle filter trajectories. *Statistics and Computing*, 28(2):461–475.
- Taghavi, E., Lindsten, F., Svensson, L., and Schön, T. B. (2013). Adaptive stopping for fast particle smoothing. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6293–6297.
- Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Volterra, V. (1928). Variations and fluctuations of the number of individuals in animal species living together. *ICES Journal of Marine Science*, 3(1):3–51.
- Wang, G., O’Leary, J., and Jacob, P. (2021). Maximal couplings of the Metropolis-Hastings algorithm. In *International Conference on Artificial Intelligence and Statistics*, pages 1225–1233. PMLR.
- Yonekura, S. and Beskos, A. (2022). Online smoothing for diffusion processes observed with noise. *Journal of Computational and Graphical Statistics*, 0(0):1–17.

APPENDIX A. ADDITIONAL NOTATIONS

This section defines new notations that do not appear in the main text (except notations for linear Gaussian models) but are used in the Supplement.

A.1. Linear Gaussian models. Let \dim_X and \dim_Y be two strictly positive integers and F_X and F_Y be two full-rank matrices of sizes $\dim_X \times \dim_X$ and $\dim_Y \times \dim_X$ respectively. Let C_X and C_Y be two symmetric positive definite matrices of respective sizes $\dim_X \times \dim_X$ and $\dim_Y \times \dim_Y$. A linear Gaussian state space model has the underlying Markov process defined by

$$X_t | X_{0:t-1} \sim \mathcal{N}(F_X X_{t-1}, C_X),$$

where X_0 also follows a Gaussian distribution; and admits the observation process

$$Y_t | X_t \sim \mathcal{N}(F_Y X_t, C_Y).$$

The predictive (X_t given $Y_{0:t-1}$), filtering (X_t given $Y_{0:t}$) and smoothing (X_t given $Y_{0:T}$) distributions are all Gaussian and their parameters can be explicitly calculated via recurrence formulas (Kalman, 1960; Kalman and Bucy, 1961). We shall denote their respective mean vectors and covariance matrices by $(\mu_t^{\text{pred}}, \Sigma_t^{\text{pred}})$, $(\mu_t^{\text{filt}}, \Sigma_t^{\text{filt}})$ and $(\mu_t^{\text{smth}}, \Sigma_t^{\text{smth}})$. In particular, the starting distribution X_0 is $\mathcal{N}(\mu_0^{\text{pred}}, \Sigma_0^{\text{pred}})$.

A.2. Total variation distance. Let μ and ν be two probability measures on \mathcal{X} . The total variation distance between μ and ν , sometimes also denoted $\text{TV}(\mu, \nu)$, is defined as $\|\mu - \nu\|_{\text{TV}} := \sup_{f: \mathcal{X} \rightarrow [0,1]} |\mu(f) - \nu(f)|$. The definition remains valid if f is restricted to the class of indicator functions on measurable subsets of \mathcal{X} . It implies in particular that $|\mu(f) - \nu(f)| \leq \|f\|_{\text{osc}} \text{TV}(\mu, \nu)$.

Next, we state a lemma summarising basic properties of the total variation distance and defining coupling-related notions (see, e.g. Proposition 3 and formula (13) of Roberts and Rosenthal (2004)). While the last property (covariance bound) is not in the aforementioned reference and does not seem popular in the literature, its proof is straightforward and therefore omitted.

Lemma A.1. *The total variation distance has the following properties:*

- (Alternative expressions.) If μ and ν admit densities $f(x)$ and $g(x)$ respectively with reference to a dominating measure λ , we have

$$\text{TV}(\mu, \nu) = \frac{1}{2} \int |f(x) - g(x)| \lambda(dx) = 1 - \int \min(f(x), g(x)) \lambda(dx).$$

- (Coupling inequality & maximal coupling.) For any pair of random variables (M, N) such that $M \sim \mu$ and $N \sim \nu$, we have

$$\mathbb{P}(M \neq N) \geq \text{TV}(\mu, \nu).$$

There exist pairs (M^*, N^*) for which equality holds. They are called maximal couplings of μ and ν .

- (Contraction property.) Let (X_n) be a Markov chain with invariant measure μ^* . Then

$$\text{TV}(X_n, \mu^*) \geq \text{TV}(X_{n+1}, \mu^*).$$

- (Covariance bound.) For any pair of random variables (M, N) such that $M \sim \mu$ and $N \sim \nu$ and real-valued functions h_1 and h_2 , we have

$$|\text{Cov}(h_1(M), h_2(N))| \leq 2 \|h_1\|_{\infty} \|h_2\|_{\infty} \text{TV}((M, N), \mu \otimes \nu).$$

A.3. Cost-to-go function. In the context of the Feynman-Kac model (1), define the associated cost-to-go function $H_{t:T}$ as (see e.g. Chopin and Papaspiliopoulos (2020, Chapter 5))

$$(20) \quad H_{t:T}(x_t) := \prod_{s=t+1}^T M_{s-1}(x_{s-1}, dx_s) G_s(x_s).$$

This function bridges $\mathbb{Q}_t(dx_t)$ and $\mathbb{Q}_T(dx_t)$, since $\mathbb{Q}_T(dx_t) \propto \mathbb{Q}_t(dx_t) H_{t:T}(x_t)$.

A.4. The projection kernel. Let \mathcal{X} and \mathcal{Y} be two measurable spaces. The projection kernel $\Pi_{\mathcal{X}}^{(\mathcal{X}, \mathcal{Y})}$ is defined by

$$\Pi_{\mathcal{X}}^{(\mathcal{X}, \mathcal{Y})}((x, y), dx^*) := \delta_x(dx^*).$$

In particular, for any function $g : \mathcal{X} \rightarrow \mathbb{R}$ and measure $\mu(dx, dy)$ defined on $\mathcal{X} \times \mathcal{Y}$, we have

$$\begin{aligned} (\Pi_{\mathcal{X}}^{(\mathcal{X}, \mathcal{Y})} g)(x, y) &= g(x) \\ (\mu \Pi_{\mathcal{X}}^{(\mathcal{X}, \mathcal{Y})})(g) &= \iint g(x) \mu(dx, dy) = \int g(x) \mu(dx) \end{aligned}$$

where the second identity shows the marginalising action of $\Pi_{\mathcal{X}}^{(\mathcal{X}, \mathcal{Y})}$ on μ . In the context of state space models, we define the shorthand

$$\Pi_t^{0:T} := \Pi_{\mathcal{X}_t}^{(\mathcal{X}_0, \dots, \mathcal{X}_T)}.$$

A.5. Other notations. For a real number x , let $\lfloor x \rfloor$ be the largest integer not exceeding x . The mapping $x \mapsto \lfloor x \rfloor$ is called the floor function • The Gamma function $\Gamma(a)$ is defined for $a > 0$ and is given by $\Gamma(a) := \int_{\mathbb{R}_+} e^{-x} x^{a-1} dx$ • Let \mathcal{X} and \mathcal{Y} be two measurable spaces. Let $K(x, dy)$ be a (not necessarily probability) kernel from \mathcal{X} to \mathcal{Y} . The norm of K is defined by $\|K\|_{\infty} := \sup_{f: \mathcal{X} \rightarrow \mathcal{Y}, f \neq 0} \|Kf\|_{\infty} / \|f\|_{\infty}$. In particular, for any function $f : \mathcal{X} \rightarrow \mathcal{Y}$, we have $\|Kf\|_{\infty} \leq \|K\|_{\infty} \|f\|_{\infty}$ • Let X_n be a sequence of random variables. We say that $X_n = \mathcal{O}_{\mathbb{P}}(1)$ if for any $\varepsilon > 0$, there exists $M > 0$ and N_0 , both depending on ε , such that $\mathbb{P}(|X_n| \geq M) \leq \varepsilon$ for all $n \geq N_0$. For a strictly positive deterministic sequence a_n , we say that $X_n = \mathcal{O}_{\mathbb{P}}(a_n)$ if $X_n/a_n = \mathcal{O}_{\mathbb{P}}(1)$. See [Janson \(2011\)](#) for discussions • We use the notation $\mathcal{N}(x|\mu, \Sigma)$ to refer to the value at x of the density function of the normal distribution $\mathcal{N}(\mu, \Sigma)$ • Let $f : U \rightarrow V$ be a function from some space U to another space V . Let S be a subset of U . The restriction of f to S , written $f|_S$, is the function from S to V defined by $f|_S(x) = f(x)$, $\forall x \in S$.

APPENDIX B. FFBS COMPLEXITY FOR DIFFERENT REJECTION SCHEMES

B.1. Framework and notations. The FFBS algorithm is a particular instance of Algorithm 2 where $B_t^{N, \text{FFBS}}$ kernels are used. If backward simulation is done using pure rejection sampling (Algorithm 4), the computational cost to simulate the $t - 1$ -th index of the n -th trajectory has conditional distribution

$$(21) \quad \tau_t^{n, \text{FFBS}} | \mathcal{F}_T, \mathcal{I}_{t:T}^n \sim \text{Geo} \left(\frac{\sum_i W_{t-1}^i m_t(X_{t-1}^i, X_t^{\mathcal{I}_t^n})}{M_h} \right).$$

At this point, it would be useful to compare this formula with (16) of the PaRIS algorithm. The difference is subtle but will drive interesting changes to the way rejection-based FFBS behaves.

If hybrid rejection sampling (Algorithm 6) is to be used instead, we are interested in the distribution of $\min(\tau_t^{n, \text{FFBS}}, N)$, for reasons discussed in Subsection 3.2. In a highly parallel setting, it is preferable that the distribution of *individual* execution times, i.e. $\tau_t^{n, \text{FFBS}}$ or $\min(\tau_t^{n, \text{FFBS}}, N)$, are not heavy-tailed. In contrast, for non-parallel hardware, only *cumulative* execution times, i.e. $\sum_{n=1}^N \tau_t^{n, \text{FFBS}}$ or $\sum_{n=1}^N \min(\tau_t^{n, \text{FFBS}}, N)$, matter. Even though the individual times might behave badly, the cumulative times could be much more regular thanks to effect of the central limit theorem, whenever applicable. Nevertheless, studying the finiteness of the k -th order moment of $\tau_t^{1, \text{FFBS}}$ is still a good way to get information about both types of execution times, since it automatically implies k -th order moment (in)finiteness for both of them.

B.2. Execution time for pure rejection sampling. We show that under certain circumstances, the execution time of the pure rejection procedure has infinite expectation. Proposition 1 in [Douc et al. \(2011\)](#) hints that the cost per trajectory for FFBS-reject might tend to infinity when $N \rightarrow \infty$. In contrast, we show that infinite expectation might very well happen for *finite* sample sizes. We first give the statement for general state space models, then focus on their implications for Gaussian ones. In particular, while infinite expectations occur only under certain configurations, infinite higher moments happen in *all* linear Gaussian models with non-degenerate dynamics.

Theorem B.1. *Using the setting and notations of Supplement B.1, under Assumptions 1 and 4, we have $\mathbb{E}[\tau_t^{1,\text{FFBS}}] = \infty$ whenever*

$$\int_{\mathcal{X}_t} G_t(x_t) H_{t:T}(x_t) \lambda_t(dx_t) = \infty$$

where the cost-to-go function $H_{t:T}$ is defined in (20) and the measure λ_t is defined in (2).

Theorem B.2. *Using the setting and notations of Supplement B.1, we consider linear Gaussian models and their notations defined in Supplement A.1. Then we have $\mathbb{E}[(\tau_t^{1,\text{FFBS}})^k] = \infty$ whenever k is greater than a certain k_0 being the smallest eigenvalue of the matrix $\text{Id} + C_X^{1/2} \left((\Sigma_t^{\text{smth}})^{-1} - (\Sigma_t^{\text{pred}})^{-1} \right) C_X^{1/2}$.*

The proofs of the two assertions are given in Supplement E.7. We now look at how they are manifested in concrete examples. The first remark is that for technical reasons, Theorem B.2 gives no information on the finiteness of $E[(\tau_t^{1,\text{FFBS}})^k]$ for $k = 1$ (since k_0 is already greater than or equal to 1 by definition). To study the finiteness of $\mathbb{E}[\tau_t^{1,\text{FFBS}}]$, we thus turn to Theorem B.1.

Example 4. In linear Gaussian models, the integral of Theorem B.1 is equal to

$$\int \mathcal{N}(y_t | F_Y x_t, C_Y) \prod_{s=t+1}^T \mathcal{N}(x_s | F_X x_{s-1}, C_X) \mathcal{N}(y_s | F_Y x_s, C_Y) dx_{t:T}$$

where the notation $\mathcal{N}(\mu, \Sigma)$ refers to the density of the normal distribution. The integrand is proportional to $\exp[-0.5(Q(x_{t:T}) - R(x_{t:T}))]$ for some quadratic form $Q(x_{t:T})$ and linear form $R(x_{t:T})$. The integral is finite if and only if Q is positive definite. In our case, this means that there is no non-trivial root for the equation $Q(x_{t:T}) = 0$, which is equivalent to

$$\begin{cases} F_Y x_s & = 0, \forall s = t, \dots, T \\ F_X x_{s-1} & = x_s, \forall s = t+1, \dots, T. \end{cases}$$

Put another way, $\mathbb{E}[\tau_t^{1,\text{FFBS}}]$ is infinite whenever the intersection

$$\bigcap_{k=0}^{T-t} \text{Ker}(F_Y F_X^k) = \bigcap_{k=0}^{T-t} F_X^{-k}(\text{Ker}(F_Y))$$

contains other things than the zero vector. A common and particularly troublesome situation is when $F_X = c \text{Id}$ for some $c > 0$ (but C_X can be arbitrary) and the dimension of the states (\dim_X) is greater than that of the observations (\dim_Y). Then the above intersection remains non-trivial no matter how big $T - t$ is. Thus, $\mathbb{E}[\tau_t^{1,\text{FFBS}}]$ has no expectation for any t . In general, the problem is less severe as successive intersections will shrink the space quickly to $\{0\}$. Consequently, Theorem B.1 only points out infiniteness of $\mathbb{E}[\tau_t^{1,\text{FFBS}}]$ for t close to T . The bad news however will come from higher moments, as seen in the below example.

We will now focus on a simple but particularly striking example. Our purpose here is to illustrate the concepts as well as to show that their implications are relevant even in small, familiar settings. More advanced scenarios are presented in Section 5 devoted to numerical experiments.

Example 5. We consider two one-dimensional Gaussian state-space models: they both have $F_X = 0.5$, $C_X = 1$, $X_0 \sim \mathcal{N}(0, C_X^2 / (1 - F_X^2))$ and $T = 3$. The only difference between them is that one has $\sigma_y^2 := C_Y = 0.5^2$ and another has $\sigma_y^2 = 3^2$. We are interested in the execution times $\tau_1^{n,\text{FFBS}}$ at time $t = 1$ (i.e. the rejection-based simulation of indices \mathcal{I}_0^n at time $t = 0$). Theorem B.2 then gives $k_0 \approx 1.14$ for $\sigma_y = 3$ and $k_0 \approx 5$ for $\sigma_y = 0.5$. The first implication is that in both cases, $\tau_1^{n,\text{FFBS}}$ is a heavy-tailed random variable and therefore FFBS-reject is not a viable option in a highly parallel setting. But an interesting phenomenon happens in the sequential hardware scenario where one is rather interested in the cumulative execution time, i.e. $\sum_{n=1}^N \tau_1^{n,\text{FFBS}}$, or equivalently, the mean number of trials per particle. In the $\sigma_y = 3$ case, non-existence of second moment prevents the cumulative regularisation effect of the central limit theorem. This is not the case for $\sigma_y = 0.5$, in which the cumulative execution time actually behaves nicely (Figures 7 and 8). However, the most valuable message from this example is perhaps that the performance of FFBS-reject depends in a non-trivial (hard to predict) way on the model parameters.

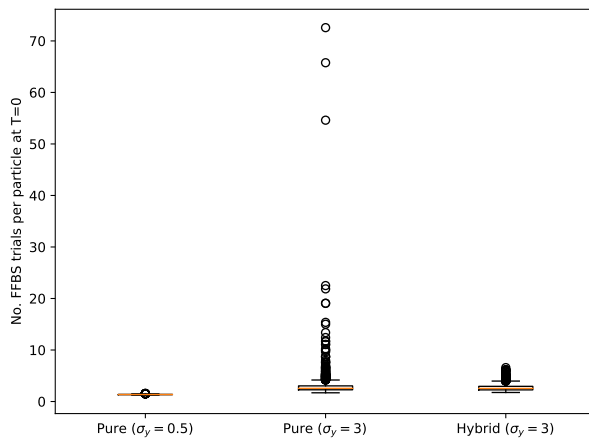


FIGURE 7. Box plots for the mean number of trials per particles to simulate indices at time 0, for models described in Example 5 and for FFBS algorithms based on pure and hybrid rejection sampling. The figure is obtained by running bootstrap particle filters with $N = 500$ over 1500 independent executions.

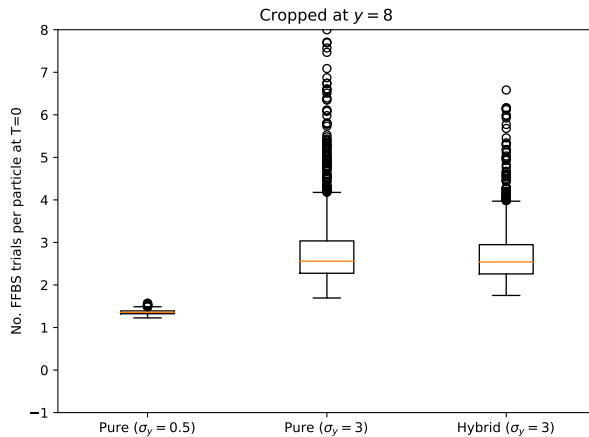


FIGURE 8. Zoom of Figure 7 to $0 \leq y \leq 8$

B.3. Execution time for hybrid rejection sampling. Formula (21) suggests defining the limit distribution $\tau_t^{\infty, \text{FFBS}}$ as

$$\tau_t^{\infty, \text{FFBS}} \mid X_t^{\infty, \text{FFBS}} \sim \text{Geo} \left(\frac{r_t(X_t^{\infty, \text{FFBS}})}{\bar{M}_h} \right)$$

where $X_t^{\infty, \text{FFBS}} \sim Q_T(dx_t)$ and r_t given in Definition 1. These quantities provide the following characterisation of the cumulative execution time for the hybrid FFBS algorithm (proved in Section E.8).

Theorem B.3. *Under Assumptions 1 and 4 and the setting of Section B.1, we have*

$$\frac{\sum_{n=1}^N \min(\tau_t^{n, \text{FFBS}}, N)}{N} = \mathcal{O}_{\mathbb{P}} \left(\mathbb{E}[\min(\tau_t^{\infty, \text{FFBS}}, N)] \right)$$

where the notation $\mathcal{O}_{\mathbb{P}}$ is defined in Supplement A.

This theorem admits the following corollary for linear Gaussian models (also proved in Section E.8).

Corollary 2. *For linear Gaussian models (Supplement A.1), if smoothing is performed using the hybrid rejection version of the FFBS algorithm, the mean execution time per particle at time step t is $\mathcal{O}_{\mathbb{P}}(\log^{d_t/2} N)$ where d_t is the dimension of \mathcal{X}_t .*

The bound $\mathcal{O}_{\mathbb{P}}(\log^{d_t/2} N)$ is actually quite conservative. For instance, with either $\sigma_y = 0.5$ or $\sigma_y = 3$, the model considered in Example 5 admits $\mathbb{E}[\tau_t^{\infty, \text{FFBS}}] < \infty$. (Gaussian dynamics can be handled using exact analytic calculations and enables to verify the claim straightforwardly.) Theorem B.3 then gives an execution time per particle of order $\mathcal{O}_{\mathbb{P}}(1)$ for hybrid FFBS, which is better than the $\mathcal{O}_{\mathbb{P}}(\sqrt{\log N})$ predicted by Corollary 2. Yet another unsatisfactory point of the result is its failure to make sense of the spectacular improvement brought by hybrid rejection sampling over the ordinary procedure in the $\sigma_y = 3$ case (see Figure 7). As explained in Example 5, this is connected to the variance of $\mathbb{E}[\tau_t^{1, \text{FFBS}}]$ and not merely the expectation; so a study of second order properties of $N^{-1} \sum_n \min(\tau_t^{n, \text{FFBS}}, N)$ would be desirable.

APPENDIX C. CONDITIONALLY-CORRELATED VERSIONS OF PARTICLE ALGORITHMS

C.1. Alternative resampling schemes. In Algorithm 1, the indices $A_t^{1:N}$ are drawn conditionally i.i.d. from the multinomial distribution $\mathcal{M}(W_{t-1}^{1:N})$. They satisfy

$$\mathbb{E} \left[\sum_{j=1}^N \mathbb{1}_{A_t^j=i} \middle| \mathcal{F}_{t-1} \right] = N W_{t-1}^i$$

for any $i = 1, \dots, N$. There are other ways to generate $A_t^{1:N}$ from $W_{t-1}^{1:N}$ that still verify this identity. We call them *unbiased resampling schemes*, and the natural one used in Algorithm 1 *multinomial resampling*.

The main motivation for alternative resampling schemes is performance. We refer to Chopin (2004); Douc et al. (2005); Gerber et al. (2019) for more details, but would like to mention that the theoretical studies of particle algorithms using other resampling schemes are more complicated since $X_t^{1:N}$ are no longer i.i.d. given \mathcal{F}_{t-1} . We use systematic resampling (Carpenter et al., 1999) in our experiments. See Algorithm 8 for a succinct description and Chopin and Papaspiliopoulos (2020, Chapter 9) for efficient implementations in $\mathcal{O}(N)$ running time.

Algorithm 8: Systematic resampling

Input: Weights $W_{t-1}^{1:N}$ summing to 1

Generate $U \sim \text{Uniform}[0, 1]$

for $n \leftarrow 1$ **to** N **do**

 Set A_t^n to the unique index k satisfying

$$W_1 + \dots + W_{k-1} \leq \frac{n-1+u}{N} < W_1 + \dots + W_k$$

Output: Resampled indices $A_t^{1:N}$

C.2. Conditionally-correlated version of Algorithm 7. In this part, we present an alternative version of Algorithm 7 that does not create conditionally i.i.d. particles at each time step. The procedure is detailed in Algorithm 9. It creates on the fly backward kernels $B_t^{N, \text{ITRC}}$ (for “intractable, conditionally correlated”). It involves a resampling step which can be done in principle using any unbiased resampling scheme. Following the intuitions of Subsection 4.2.3 and the notations of Algorithm 9, we want a scheme such that in most cases, $A_t^{2k-1} \neq A_t^{2k}$ but the Euclidean distance between $X_{t-1}^{A_t^{2k-1}}$ and $X_{t-1}^{A_t^{2k}}$ is small. Algorithm 10 proposes such a method (which we name the Adjacent Resampler). It can run in $\mathcal{O}(N)$ time using a suitably implemented linked list.

APPENDIX D. ADDITIONAL INFORMATION ON NUMERICAL EXPERIMENTS

D.1. Offline smoothing in linear Gaussian models. In this section, we study offline smoothing for the linear Gaussian model specified in Section 5.1. Since offline processing requires storing particles at all times t in the memory, we use $T = 500$ here instead of $T = 3000$. Apart from that, the algorithmic and benchmark settings remain the same.

Algorithm 9: Conditionally-correlated version of Algorithm 7

Input: Feynman-Kac model (1), particles $X_{t-1}^{1:N}$ and weights $W_{t-1}^{1:N}$ that approximate $\mathbb{Q}_{t-1}(dx_{t-1})$

Resample $A_t^{1:N}$ from $\{1, 2, \dots, N\}$ with weights $W_{t-1}^{1:N}$ using any resampling scheme (such as the Adjacent Resampler in Algorithm 10)

for $k \leftarrow 1$ **to** $N/2$ **do**

Move. Simulate X_t^{2k-1} and X_t^{2k} such that marginally, $X_t^{2k-1} \sim M_t(X_{t-1}^{A_t^{2k-1}}, \cdot)$ and

$X_t^{2k} \sim M_t(X_{t-1}^{A_t^{2k}}, \cdot)$

Calculate backward kernel.

if $X_t^{2k-1} = X_t^{2k}$ **then**

 Set $B_t^{N, \text{ITRC}}(2k-1, \cdot) \leftarrow (\delta \{A_t^{2k-1}\} + \delta \{A_t^{2k}\}) / 2$

 Set $B_t^{N, \text{ITRC}}(2k, \cdot) \leftarrow (\delta \{A_t^{2k-1}\} + \delta \{A_t^{2k}\}) / 2$

else

 Set $B_t^{N, \text{ITRC}}(2k-1, \cdot) \leftarrow \delta \{A_t^{2k-1}\}$

 Set $B_t^{N, \text{ITRC}}(2k, \cdot) \leftarrow \delta \{A_t^{2k}\}$

Reweight. Set $\omega_t^n \leftarrow G_t(X_t^n)$ for $n = 1, 2, \dots, N$

Set $\ell_t^N \leftarrow \sum_{n=1}^N \omega_t^n / N$

Set $W_t^n \leftarrow \omega_t^n / N \ell_t^N$ for $n = 1, 2, \dots, N$

Output: Particles $X_t^{1:N}$ and weights $W_t^{1:N}$ that approximate $\mathbb{Q}_t(dx_t)$; backward kernel $B_t^{N, \text{ITRC}}$ for use in Algorithms 2 and 3

Algorithm 10: The Adjacent Resampler

Input: Particles $X_{t-1}^{1:N}$, weights $W_{t-1}^{1:N}$

Sort the particles $X_{t-1}^{1:N}$ using the Hilbert curve. Let $s \leftarrow [s_1 \dots s_N]$ be the corresponding *indices*

Resample from $\{1, \dots, N\}$ with weights $W_{t-1}^{1:N}$ using systematic resampling (Carpenter et al., 1999; Gerber et al., 2019), then let $f : \{1, \dots, N\} \rightarrow \mathbb{Z}$ be the function defined by $f(i)$ being the number of times the index s_i was resampled. Obviously $\sum_{i=1}^N f(i) = N$

Initialise $i \leftarrow 1$

for $n \leftarrow 1$ **to** N **do**

 Set $A_t^n \leftarrow s_i$

 Update $f(i) \leftarrow f(i) - 1$

 Let Ω_1 be the set $\{\min \{\ell > i \mid f_\ell > 0\}\}$ (which has one element if the minimum is well-defined and zero element otherwise)

 Let Ω_2 be the set $\{\max \{\ell < i \mid f_\ell > 0\}\}$ (which has one element if the maximum is well-defined and zero element otherwise)

 If $\Omega_1 \cup \Omega_2$ is not empty, update $i \leftarrow \operatorname{argmax} f|_{\Omega_1 \cup \Omega_2}$ (see section A.5 for the restriction notation).

 If there is more than one argmax, pick one randomly

Output: Resampled indices $A_t^{1:N}$

Figure 9 plots the squared interquartile range of the estimators $\mathbb{Q}_T(\varphi_t)$ with respect to t , for different algorithms. For small t , the function φ_t only looks at states close to time 0, whereas for bigger t , recent states less affected by degeneracy are also taken into account. In all cases though, we see that MCMC and rejection-based smoothers have superior performance.

Figure 10 shows box plots of the averaged execution times (per particle N per time t) based on 150 runs. The observations are comparable to those in Section 5.1. We see a performance difference between the rejection-based smoothers using the bootstrap and the guided filters. Both have an execution time that is much more variable than hybrid rejection algorithms. The latter still need around 10 times more CPU load than MCMC smoothers, for essentially the same precision.

We now take a closer look at the reason behind the performance difference between the bootstrap filter and the guided one when pure rejection sampling is used. Figure 11 shows the effective sample size




FIGURE 9. Squared interquartile range of the estimators of $\mathbb{Q}_T(\varphi_t)$ with respect to t , for different algorithms applied to the model of Section D.1. See Section 5.1 for the meaning of the acronyms in the legend.




FIGURE 10. Box plots of the number of transition density evaluations divided by NT for different algorithms in the offline linear Gaussian model of Section D.1.

(ESS) of both filters as a function of time. We can see that there is an outlier in the data around time $t = 40$. Figure 12 box-plots the execution times divided by N at $t = 40$ for the pure rejection sampling algorithm, whereas Figures 13 and 14 do the same for $t = 38$ and $t = 42$. The root of the problem is now clear: at most times t there is very few difference between the execution times of the bootstrap and the guided filters. However, if an outlier is present in the data, the guided filter suddenly requires a very high number of transition density evaluation in the rejection sampler. This gives yet another reason to avoid using pure rejection sampling.

D.2. Lotka-Volterra SDE.

D.2.1. *Coupling of Euler discretisations.* Consider the SDE

$$(22) \quad dX_t = b(X_t)dt + \sigma(X_t)dW_t$$

and two starting points X_0^A and X_0^B in \mathbb{R}^d . We wish to simulate X_1^A and X_1^B such that the transitions from X_0^A to X_1^A and X_0^B to X_1^B both follow the Euler-discretised version of the equation, but X_1^A and X_1^B are correlated in a way that increases, as much as we can, the probability that they are equal. Algorithm 11 makes it clear that it all boils down to the coupling of two Gaussian distributions.




FIGURE 11. Evolution of the ESS for the linear Gaussian model of Section D.1.




FIGURE 12. Box plots of the average execution time per particle for the pure rejection algorithm at time $t = 40$. Figure produced based on 150 independent runs of the model described in Section D.1.





FIGURE 13. Same as Figure 12, but for $t = 38$.

FIGURE 14. Same as Figure 12, but for $t = 42$.**Algorithm 11:** Coupling of two Euler discretisations

Input: Functions $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, two starting points X_0^A and X_0^B at time 0, number of discretisation step N_{dist}

Initialise $X^A \leftarrow X_0^A$

Initialise $X^B \leftarrow X_0^B$

Set $\delta \leftarrow 1/N_{\text{dist}}$

for $i \leftarrow 1$ **to** N_{dist} **do**

 Simulate $(\tilde{X}^A, \tilde{X}^B)$ from a coupling of

$$\mathcal{N}(X^A + \delta b(X^A), \delta \sigma(X^A) \sigma(X^A)^\top)$$

 and

$$\mathcal{N}(X^B + \delta b(X^B), \delta \sigma(X^B) \sigma(X^B)^\top),$$

 such as Algorithm 14

 Update $(X^A, X^B) \leftarrow (\tilde{X}^A, \tilde{X}^B)$

Set $(X_1^A, X_1^B) \leftarrow (X^A, X^B)$

Output: Two endpoints X_1^A and X_1^B at time 1, obtained by passing X_0^A and X_0^B in a correlated manner through a discretised version of (22)

Lindvall and Rogers (1986) propose the following construction: if two diffusions X_t^A and X_t^B both follow the dynamics of (22), that is,

$$dX_t^A = b(X_t^A)dt + \sigma(X_t^A)dW_t^A$$

$$dX_t^B = b(X_t^B)dt + \sigma(X_t^B)dW_t^B$$

and the two Brownian motions are correlated via

$$(23) \quad dW_t^B = [\text{Id} - 2u(X^A, X^B)u(X^A, X^B)^\top]dW_t^A$$

where Id is the identity matrix and the vector u is defined by

$$u(x, x') = \frac{\sigma(x')^{-1}(x - x')}{\|\sigma(x')^{-1}(x - x')\|_2},$$

then under some regularity conditions, the two diffusions meet almost surely. (Note two special features of (23): it is valid because the term in the square bracket is an orthogonal matrix; and it ceases to be well-defined once the two trajectories have met.) Simulating the meeting time τ turns out to be very challenging. The Euler discretisation (Algorithm 11 + Algorithm 12) has a fixed step size δ , and there is zero probability that τ is of the form $k\delta$ for some integer k . Since the coupling transform is deterministic, the two Euler-simulated trajectories will *never* meet. Figure 15 depicts this difficulty in the special case of two Brownian motions in dimension 1 (i.e. $b(x) \equiv 0$ and $\sigma \equiv 1$). Under this setting, (23) means that

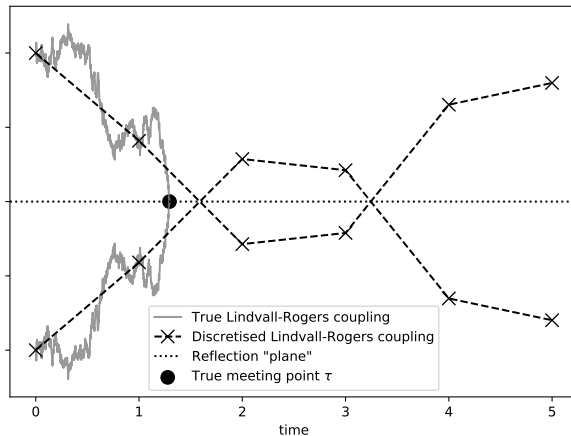


FIGURE 15. Coupling of two Brownian motions in \mathbb{R} starting from 0 and 4 respectively. The true Lindvall-Rogers coupling (23) is represented by the continuous grey lines. The discretised simulation (Algorithm 11 + Algorithm 12) is shown by the dashed lines. The discretised trajectories not only miss the true meeting point τ but also never meet afterwards (see text).

the two Brownian increments are symmetric with respect to the midpoint of the segment connecting their initial states. Note that the two dashed lines do cross at two points, but using them as meeting points is invalid: since they are not part of the discretisation but the result of some heuristic “linear interpolation”, it would change the distribution of the trajectories.

Algorithm 12: Lindvall-Rogers coupling of two Gaussian distributions

Input: Two vectors μ^A, μ^B in \mathbb{R}^d and two $d \times d$ matrices σ^A and σ^B

Calculate $u \leftarrow (\sigma^B)^{-1}(\mu^A - \mu^B)$

Normalise $u \leftarrow u / \|u\|_2$

Simulate $W^A \sim \mathcal{N}(0, \text{Id})$

Set $W^B \leftarrow (\text{Id} - 2uu^\top)W^A$

Set $X^A \leftarrow \mu^A + \sigma^A W^A$

Set $X^B \leftarrow \mu^B + \sigma^B W^B$

Output: Two correlated points X^A and X^B marginally distributed according to $\mathcal{N}(\mu^A, \sigma^A(\sigma^A)^\top)$ and $\mathcal{N}(\mu^B, \sigma^B(\sigma^B)^\top)$ respectively

We therefore need some coupling that has a non-zero meeting probability at each δ -step. This can be achieved by the rejection maximal coupling (Algorithm 13, see also, e.g. Roberts and Rosenthal, 2004) as well as the recently proposed coupled rejection sampler (Corenfos and Särkkä, 2022). However, they all make use of rejection sampling in one way or another, which renders the execution time random. We wish to avoid this if possible. The reflection-maximal coupling (Bou-Rabee et al., 2020; Jacob et al., 2020) has deterministic cost and optimal meeting probability, but is only applicable for two Gaussian distributions of the same covariance matrix, which is not our case.

As suggested by Figure 15, the discretised Lindvall-Rogers coupling (Algorithm 12) is actually great for bringing together two faraway trajectories. Only when they start getting closer that it misses out. At that moment, the two distributions corresponding to the next δ -step have non-negligible overlap and would preferably be coupled in the style of Algorithm 13. We propose a modified coupling scheme that acts like Algorithm 12 when the two trajectories are at a large distance and behaves as Algorithm 13 otherwise.

Algorithm 13: Rejection maximal coupler for two distributions**Input:** Two probability distributions f^A and f^B Simulate $X^A \sim f^A$ Simulate $U^A \sim \text{Uniform}[0, f^A(X^A)]$ **if** $U^A \leq f^B(X^A)$ **then** Set $X^B \leftarrow X^A$ **else** **repeat** Simulate $X^B \sim f^B$ Simulate $U^B \sim \text{Uniform}[0, f^B(X^B)]$ **until** $U^B > f^A(X^B)$ **Output:** Two maximally-coupled realisations X^A and X^B , marginally f^A -distributed and f^B -distributed respectively

The idea is to preliminarily generate a uniform draw in the “overlapping zone” of the two distributions (if they are close enough to make that easy). Next, we perform Algorithm 12 and then, any of the two simulations belonging to the overlapping zone will be replaced by the aforementioned preliminary draw (if it is available). The precise mathematical formulation is given in Algorithm 14 and the proof in Supplement E.12.

Algorithm 14: Modified Lindvall-Rogers (MLR) coupler of two Gaussian distributions**Input:** Two vectors μ^A and μ^B in \mathbb{R}^d , two $d \times d$ matrices σ^A and σ^B Let f^A and f^B be respectively the probability densities of $\mathcal{N}(\mu^A, \sigma^A(\sigma^A)^\top)$ and $\mathcal{N}(\mu^B, \sigma^B(\sigma^B)^\top)$ Simulate X^A and X^B from Algorithm 12Simulate $U \sim \text{Uniform}[0, 1]$ Set $U^A \leftarrow U f^A(X^A)$ and $U^B \leftarrow U f^B(X^B)$ Simulate $Y \sim f^A$ and $V \sim \text{Uniform}[0, f^A(Y)]$ **if** $V \leq f^B(Y)$ **then** **if** $U^A \leq f^B(X^A)$ **then** update $(X^A, U^A) \leftarrow (Y, V)$ **if** $U^B \leq f^A(X^B)$ **then** update $(X^B, U^B) \leftarrow (Y, V)$ **Output:** Two correlated random vectors X^A and X^B , distributed marginally according to $\mathcal{N}(\mu^A, \sigma^A(\sigma^A)^\top)$ and $\mathcal{N}(\mu^B, \sigma^B(\sigma^B)^\top)$

Algorithm 14 has a deterministic execution time, but it does not attain the optimal coupling rate. Yet, as $\delta \rightarrow 0$, we see empirically that it still recovers the oracle coupling time defined by (23) (although we did not try to prove this formally). In Figure 16, we couple two standard Brownian motions starting from $a = 0$ and $b = 1.5$ using Algorithm 14 with different values of δ . It is known, by a simple application of the reflection principle (Lévy, 1940; see also Chapter 2.2 of Mörters and Peres, 2010), that the reflection coupling (23) succeeds after a $\text{Levy}(0, (b-a)^2/4)$ -distributed time. We therefore have to deal with a heavy-tailed distribution and restrict ourselves to the interval $[0, 5]$. We see that the law of the meeting time is stable and convergent as $\delta \rightarrow 0$. Thus, at least empirically, Algorithm 14 does not suffer from the instability problem as $\delta \rightarrow 0$, contrary to a naive path space augmentation approach (see Yonekura and Beskos, 2022 for a discussion).

D.2.2. *Supplementary figures.* Figure 17 plots a realisation of the states and data with parameters given in Subsection 5.2, for a relatively small scale dataset ($T = 50$). While the periodic trait seen in classical deterministic Lotka-Volterra equations is still visible (with a period of around 20), it is clear that here random perturbations have added considerable chaos to the system. Figures 18 and 19 show respectively the performances of the naive genealogy tracking smoother and ours (Algorithm 9) on the dataset of Figure 17. Our smoother has successfully prevented the degeneracy phenomenon, particularly for times close to 0. Figure 20 shows, in two different ways, the properties of effective sample sizes (ESS) in the $T = 3000$ scenario (see Section 5.2).




FIGURE 16. Densities of the meeting times restricted to $[0, 5]$ for two Brownian motions started from 0 and 1.5. The curves are drawn using 20 000 simulations from either a Levy distribution (for the “True distribution” curve) or Algorithm 14 (for the MLR ones). The boundary effect of kernel density estimators causes spills beyond 0 and 5.




FIGURE 17. A realisation of the Lotka-Volterra SDE with parameters described in Section 5.2. The stationary point of the system is $[100, 100]$.

APPENDIX E. PROOFS

E.1. Proof of Theorem 2.1 (general convergence theorem). In line with (7), we define the distribution $\mathbb{Q}_t^N(dx_{0:t})$ for $t < T$ as the $x_{0:t}$ marginal of the joint distribution

$$(24) \quad \bar{\mathbb{Q}}_t^N(dx_{0:t}, di_{0:t}) := \mathcal{M}(W_t^{1:N})(di_t) \left[\prod_{s=t}^1 B_s^N(i_s, di_{s-1}) \right] \left[\prod_{s=t}^0 \delta_{X_s^{i_s}}(dx_s) \right].$$

The proof builds up on an inductive argument which links \mathbb{Q}_t^N with \mathbb{Q}_{t-1}^N through new innovations at time t . More precisely, we have the following fundamental proposition, where \mathcal{F}_t^+ is defined as the smallest σ -algebra containing \mathcal{F}_t and $\hat{B}_{1:t}^N$.

Proposition 5. \mathbb{Q}_t^N is a mixture distribution that admits the representation

$$(25) \quad \mathbb{Q}_t^N(dx_{0:t}) = (\ell_t^N)^{-1} N^{-1} \sum_n G_t(x_t) K_t^N(n, dx_{0:t})$$

where ℓ_t^N is defined in Algorithm 1 and $K_t^N(n, dx_{0:t})$ is a certain probability measure satisfying

$$(26) \quad \mathbb{E} [K_t^N(n, dx_{0:t}) | \mathcal{F}_{t-1}^+] = \mathbb{Q}_{t-1}^N(dx_{0:t-1}) M_t(x_{t-1}, dx_t).$$




FIGURE 18. Smoothing trajectories for the dataset of Figure 17 using the naive genealogy tracking smoother ($B_t^{N,GT}$ kernels) with systematic resampling (see Section C.1). We took $N = 100$ and randomly plotted 30 smoothing trajectories.




FIGURE 19. Same as Figure 18, but smoothing was done using Algorithm 9 instead.

In other words, for any (possibly random) function $\varphi_t^N : \mathcal{X}_0 \times \dots \times \mathcal{X}_t \rightarrow \mathbb{R}$ such that $\varphi_t^N(x_{0:t})$ is \mathcal{F}_{t-1}^+ -measurable, we have

$$\mathbb{E} \left[\int K_t^N(n, dx_{0:t}) \varphi_t^N(x_{0:t}) \middle| \mathcal{F}_{t-1}^+ \right] = \int \mathbb{Q}_{t-1}^N(dx_{0:t-1}) M_t(x_{t-1}, dx_t) \varphi_t^N(x_{0:t}).$$

Moreover, $\int K_t^N(n, dx_{0:t}) \varphi_t^N(x_{0:t})$, for $n = 1, \dots, N$ are i.i.d. given \mathcal{F}_{t-1}^+ .

The proof is postponed until the end of this subsection. This proposition gives the expression (25) for \mathbb{Q}_t^N , which is easier to manipulate than (24) and which highlights, through (26), its connection to \mathbb{Q}_{t-1}^N . To further simplify the notations, let us define, following Douc et al. (2011), the kernel $L_{t_1:t_2}$, for $t_1 \leq t_2$, as

$$(27) \quad L_{t_1:t_2}(x_{0:t_1}^*, dx_{0:t_2}) := \delta_{x_{0:t_1}^*}(dx_{0:t_1}) \prod_{s=t_1+1}^{t_2} M_s(x_{s-1}, dx_s) G_s(x_s).$$

In other words, for real-valued functions $\varphi_{t_2} = \varphi_{t_2}(x_0, \dots, x_{t_2})$, we have

$$L_{t_1:t_2}(x_{0:t_1}^*, \varphi_{t_2}) = \int \varphi_{t_2}(x_0^*, \dots, x_{t_1}^*, x_{t_1+1}, \dots, x_{t_2}) \prod_{s=t_1+1}^{t_2} M_s(x_{s-1}, dx_s) G_s(x_s).$$




FIGURE 20. Effective sample size (ESS) for the Lotka-Volterra SDE model with $T = 3000$ (Section 5.2). The left pane draws the box plot of the collection of all estimated ESS for $t = 0, \dots, 3000$. The right pane plots the evolution of ESS with time. The quantity changes so chaotically that the curve only plots one value every 20 time steps for readability.

The usefulness of these kernels will come from the simple remark $\mathbb{Q}_{t_2} \propto \mathbb{Q}_{t_1} L_{t_1:t_2}$. We also see that

$$\|L_{t_1:t_2} \varphi_{t_2}\|_\infty \leq \|\varphi_{t_2}\|_\infty \prod_{s=t_1+1}^{t_2} \|G_s\|_\infty,$$

which gives $\|L_{t_1:t_2}\|_\infty < \infty$, where the norm of a kernel is defined in Subsection A. We are now in a position to state an importance sampling-like representation of \mathbb{Q}_t^N .

Corollary 3. *Let $\varphi_t^N : \mathcal{X}_0 \times \dots \times \mathcal{X}_t \rightarrow \mathbb{R}$ be a (possibly random) function such that $\varphi_t^N(x_{0:t})$ is \mathcal{F}_{t-1}^+ -measurable. Suppose that φ_t^N is either uniformly non-negative (i.e. $\varphi_t^N(x_{0:t}) \geq 0$ almost surely) or uniformly bounded (i.e. there exists a deterministic C such that $|\varphi_t^N(x_{0:t})| \leq C$ almost surely). Then*

$$\mathbb{Q}_t^N \varphi_t^N = \frac{N^{-1} \sum_n \tilde{K}_t^N(n, \varphi_t^N)}{N^{-1} \sum_n \tilde{K}_t^N(n, \mathbb{1})},$$

where $\tilde{K}_t^N(n, \cdot)$ is a certain random kernel such that

- $\mathbb{E} \left[\tilde{K}_t^N(n, \varphi_t^N) \middle| \mathcal{F}_{t-1}^+ \right] = (\mathbb{Q}_{t-1}^N L_{t-1:t}) \varphi_t^N$;
- $N^{-1} \sum_n \tilde{K}_t^N(n, \mathbb{1}) = \ell_t^N$;
- $(\tilde{K}_t^N(n, \varphi_t^N))_{n=1, \dots, N}$ are i.i.d. given \mathcal{F}_{t-1}^+ ;
- almost surely, $|\tilde{K}_t^N(n, \varphi_t^N)| \leq \|\varphi_t^N\|_\infty \|G_t\|_\infty$ if φ_t^N is uniformly bounded and $\tilde{K}_t^N(n, \varphi_t^N) \geq 0$ if φ_t^N is uniformly non-negative.

These statements are valid for $t = 0$ under the convention $\mathbb{Q}_{-1}^N L_{-1:0} = \mathbb{Q}_{-1} L_{-1:0} = \mathbb{M}_0$ and \mathcal{F}_{-1} being the trivial σ -algebra.

Proof. Put $\tilde{K}_t^N(n, \varphi_t^N) := \int G_t(x_t) K_t^N(n, dx_{0:t}) \varphi_t^N(x_{0:t})$ where K_t^N is defined in Proposition 5. Then

$$\begin{aligned} \mathbb{Q}_t^N(\varphi_t^N) &= \frac{N^{-1} \sum_n \int G_t(x_t) K_t^N(n, dx_{0:t}) \varphi_t^N(x_{0:t})}{\ell_t^N} \\ &= \frac{N^{-1} \sum_n \tilde{K}_t^N(n, \varphi_t^N)}{\ell_t^N}. \end{aligned}$$

Since \mathbb{Q}_t^N is a probability measure, applying this identity twice yields

$$\mathbb{Q}_t^N(\varphi_t^N) = \frac{\mathbb{Q}_t^N(\varphi_t^N)}{\mathbb{Q}_t^N(\mathbb{1})} = \frac{N^{-1} \sum_n \tilde{K}_t^N(n, \varphi_t^N)}{N^{-1} \sum_n \tilde{K}_t^N(n, \mathbb{1})}.$$

The remaining points are simple consequences of the definition of \tilde{K}_t^N and $L_{t-1:t}$. \square

The corollary hints at a natural induction proof for Theorem 2.1.

Proof of Theorem 2.1. The following calculations are valid for all $T \geq 0$, under the convention defined at the end of Corollary 3. They will prove (8) for $T = 0$ and, at the same time, prove it for any $T \geq 1$ under the hypothesis that it already holds true for $T - 1$. Let $\varphi_T = \varphi_T(x_0, \dots, x_T)$ be a real-valued function on $\mathcal{X}_0 \times \dots \times \mathcal{X}_T$. Write

$$(28) \quad \sqrt{N}(\mathbb{Q}_T^N \varphi_T - \mathbb{Q}_T \varphi_T) = \sqrt{N} \left(\frac{N^{-1} \sum_n \tilde{K}_T^N(n, \varphi_T)}{N^{-1} \sum_n \tilde{K}_T^N(n, \mathbb{1})} - \frac{\mathbb{Q}_{T-1} L_{T-1:T} \varphi_T}{\mathbb{Q}_{T-1} L_{T-1:T} \mathbb{1}} \right)$$

where the rewriting of $\mathbb{Q}_T \varphi_T$ is a consequence of $\mathbb{Q}_T \propto \mathbb{Q}_{T-1} L_{T-1:T}$. We will bound this difference by Hoeffding's inequalities for ratios (see Supplement E.13 for notations, including the definition of sub-Gaussian variables that we shall use below). We have

- that $\sqrt{N}(N^{-1} \sum \tilde{K}_T^N(n, \varphi_T) - \mathbb{Q}_{T-1}^N L_{T-1:T} \varphi_T)$ is $(1, \|\varphi_T\|_\infty \|G_T\|_\infty)$ -sub-Gaussian conditioned on \mathcal{F}_{t-1}^+ because of Theorem E.15 (and thus unconditionally, by the law of total expectation);
- and that $\sqrt{N}(N^{-1} \mathbb{Q}_{T-1}^N L_{T-1:T} \varphi_T - \mathbb{Q}_{T-1} L_{T-1:T} \varphi_T)$ is sub-Gaussian with parameters

$$(C_{T-1}, S_{T-1} \|L_{T-1:T}\|_\infty \|\varphi_T\|_\infty)$$

if $T \geq 1$ by induction hypothesis. The quantity is equal to 0 if $T = 0$.

This permits to apply Lemma E.16, which results in the sub-Gaussian properties of

- the quantity $\sqrt{N}(N^{-1} \sum \tilde{K}_T^N(n, \varphi_T) - \mathbb{Q}_{T-1} L_{T-1:T} \varphi_T)$, with parameters $(1 + C_{T-1}, S'_{T-1} \|\varphi_T\|_\infty)$, for a certain constant S'_{T-1} ;
- and the quantity $\sqrt{N}(N^{-1} \sum \tilde{K}_T^N(n, \mathbb{1}) - \mathbb{Q}_{T-1} L_{T-1:T} \mathbb{1})$, which is a special case of the former one, with parameters $(1 + C_{T-1}, S'_{T-1})$.

Finally, we invoke Proposition 11 and deduce the sub-Gaussian property of (28) with parameters

$$\left(2 + 2C_{T-1}, 2 \frac{S'_{T-1} \|\varphi_T\|_\infty}{\mathbb{Q}_{T-1} L_{T-1:T} \mathbb{1}} \right)$$

which finishes the proof. \square

Proof of Proposition 5. From (24), we have

$$\begin{aligned} \mathbb{Q}_t^N(dx_{0:t}) &= \sum_{i_t} \bar{\mathbb{Q}}_t^N(di_t) \bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t) \\ &= (\ell_t^N)^{-1} N^{-1} \sum_{i_t} G_t(X_t^{i_t}) \bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t) \\ &= (\ell_t^N)^{-1} N^{-1} \sum_{i_t} G_t(x_t) \bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t) \end{aligned}$$

since $\bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t)$ has a $\delta_{X_t^{i_t}}(dx_t)$ term. In fact, the identity

$$\bar{\mathbb{Q}}_t^N(dx_{0:t}, di_{t-1}|i_t) = \delta_{X_t^{i_t}}(dx_t) B_t^N(i_t, di_{t-1}) \bar{\mathbb{Q}}_{t-1}^N(dx_{0:t-1}|i_{t-1})$$

follows directly from the backward recursive nature of Algorithm 2, and thus

$$(29) \quad \bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t) = \delta_{X_t^{i_t}}(dx_t) \int_{i_{t-1}} B_t^N(i_t, di_{t-1}) \bar{\mathbb{Q}}_{t-1}^N(dx_{0:t-1}|i_{t-1}).$$

The $\mathbb{Q}_{t-1}^N(dx_{0:t-1}|i_{t-1})$ term is \mathcal{F}_{t-1}^+ -measurable. We shall calculate the expectation of $\delta_{X_t^{i_t}}(dx_t)B_t^N(i_t, di_{t-1})$ given \mathcal{F}_{t-1}^+ . The following arguments are necessary for formal verification, but the result (30) is natural in light of the ancestor regeneration intuition explained in Section 2.4.

Let $f_t^N : \{1, \dots, N\} \times \mathcal{X}_t \rightarrow \mathbb{R}$ be a (possibly random) function such that $f_t^N(i_{t-1}, x_t)$ is \mathcal{F}_{t-1}^+ -measurable. Let $J_t^{i_t}$ be a random variable such that given \mathcal{F}_{t-1}^+ , $X_t^{i_t}$ and $\hat{B}_t^N(i_t, \cdot)$, $J_t^{i_t}$ is $B_t^N(i_t, di_{t-1})$ -distributed. This automatically makes $J_t^{i_t}$ satisfy the second hypothesis of Theorem 2.1. Additionally, by virtue of its first hypothesis, the distribution of $(J_t^{i_t}, A_t^{i_t})$ is the same given either \mathcal{F}_{t-1}^+ or $X_{t-1}^{1:N}$ (see also Figure 1). We can now write

$$\begin{aligned} & \mathbb{E} \left[\int f_t^N(i_{t-1}, x_t) \delta_{X_t^{i_t}}(dx_t) B_t^N(i_t, di_{t-1}) \middle| \mathcal{F}_{t-1}^+ \right] \\ &= \mathbb{E} \left[\int f_t^N(i_{t-1}, X_t^{i_t}) B_t^N(i_t, di_{t-1}) \middle| \mathcal{F}_{t-1}^+ \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[f_t^N(J_t^{i_t}, X_t^{i_t}) \middle| \mathcal{F}_{t-1}^+, X_t^{i_t}, \hat{B}_t^N(i_t, \cdot) \right] \middle| \mathcal{F}_{t-1}^+ \right] \\ &= \mathbb{E} \left[f_t^N(J_t^{i_t}, X_t^{i_t}) \middle| \mathcal{F}_{t-1}^+ \right] \text{ by the law of total expectation} \\ &= \mathbb{E} \left[f_t^N(A_t^{i_t}, X_t^{i_t}) \middle| \mathcal{F}_{t-1}^+ \right] \text{ by the second hypothesis of Theorem 2.1} \\ &= \int f_t^N(i_{t-1}, x_t) \mathcal{M}(W_{t-1}^{1:N})(di_{t-1}) M_t(X_{t-1}^{i_{t-1}}, dx_t). \end{aligned}$$

This equality means that

$$(30) \quad \mathbb{E} \left[\delta_{X_t^{i_t}}(dx_t) B_t^N(i_t, di_{t-1}) \middle| \mathcal{F}_{t-1}^+ \right] = \mathcal{M}(W_{t-1}^{1:N})(di_{t-1}) M_t(X_{t-1}^{i_{t-1}}, dx_t),$$

Now, put

$$K^N(i_t, dx_{0:t}) := \bar{\mathbb{Q}}_t^N(dx_{0:t}|i_t).$$

From (29) and (30), we have

$$\begin{aligned} \mathbb{E} \left[K^N(i_t, dx_{0:t}) \middle| \mathcal{F}_{t-1}^+ \right] &= \int_{i_{t-1}} \mathcal{M}(W_{t-1}^{1:N})(di_{t-1}) M_t(X_{t-1}^{i_{t-1}}, dx_t) \bar{\mathbb{Q}}_{t-1}^N(dx_{0:t-1}|i_{t-1}) \\ &= M_t(x_{t-1}, dx_t) \int_{i_{t-1}} \mathcal{M}(W_{t-1}^{1:N})(di_{t-1}) \bar{\mathbb{Q}}_{t-1}^N(dx_{0:t-1}|i_{t-1}) \\ &\text{since } \bar{\mathbb{Q}}_{t-1}^N(dx_{0:t-1}|i_{t-1}) \text{ has a } \delta_{X_{t-1}^{i_{t-1}}(dx_{t-1})} \text{ term} \\ &= M_t(x_{t-1}, dx_t) \mathbb{Q}_{t-1}^N(dx_{0:t-1}) \end{aligned}$$

which finishes the proof. \square

E.2. Proof of Equation (11) (online smoothing recursion).

Proof. Using (7) and the matrix notations, the distribution $\bar{\mathbb{Q}}_t^N(di_s)$ can be represented by the $1 \times N$ vector

$$\hat{q}_{s|t}^N := [W_t^1 \dots W_t^N] \hat{B}_t^N \dots \hat{B}_{s+1}^N.$$

Defining the $N \times N$ matrix $\hat{\psi}_s^N$ as

$$\hat{\psi}_s^N[i_{s-1}, i_s] := \psi_s(X_{s-1}^{i_{s-1}}, X_s^{i_s}),$$

we have

$$\begin{aligned} \mathbb{E}_{\bar{\mathbb{Q}}_t^N}[\psi_s(X_{s-1}, X_s)] &= \sum_{i_s, i_{s-1}} \hat{q}_{s|t}^N[1, i_s] \hat{B}_s^N[i_s, i_{s-1}] \hat{\psi}_s^N[i_{s-1}, i_s] \\ &= \sum_{i_s} \hat{q}_{s|t}^N[1, i_s] (\hat{B}_s^N \hat{\psi}_s^N)[i_s, i_s] \\ &= \hat{q}_{s|t}^N \text{diag}(\hat{B}_s^N \hat{\psi}_s^N). \end{aligned}$$




FIGURE 21. Directed graph representing the relations between variables generated in Algorithm 2. Only those necessary for the proof of Lemma E.1 are included.

Therefore,

$$\mathbb{Q}_t^N \varphi_t = \sum_{s=0}^t [W_t^1 \dots W_t^N] \hat{B}_t^N \dots \hat{B}_{s+1}^N \text{diag}(\hat{B}_s^N \hat{\psi}_s^N)$$

from which follows the recursion

$$\begin{cases} \mathbb{Q}_t^N \varphi_t &= [W_t^1 \dots W_t^N] \hat{S}_t^N, \\ \hat{S}_t^N &:= \hat{B}_t^N \hat{S}_{t-1}^N + \text{diag}(\hat{B}_t^N \hat{\psi}_t^N). \end{cases}$$

This is exactly (11). □

E.3. Proof of Theorem 2.2 (general stability theorem). The following lemma describes the simultaneous backward construction of two trajectories $\mathcal{I}_{0:T}^1$ and $\mathcal{I}_{0:T}^2$ given \mathcal{F}_T^- .

Lemma E.1. *We use the same notations as in Algorithms 1 and 2. Suppose that the hypotheses of Theorem 2.1 are satisfied. Then, given $\mathcal{I}_{t:T}^1$, $\mathcal{I}_{t:T}^2$ and \mathcal{F}_T^- ,*

- if $\mathcal{I}_t^1 \neq \mathcal{I}_t^2$, the two variables \mathcal{I}_{t-1}^1 and \mathcal{I}_{t-1}^2 are conditionally independent and their marginal distributions are respectively $B_t^{N,\text{FFBS}}(\mathcal{I}_t^1, \cdot)$ and $B_t^{N,\text{FFBS}}(\mathcal{I}_t^2, \cdot)$;
- if $\mathcal{I}_t^1 = \mathcal{I}_t^2$, under the aforementioned conditioning, the two variables \mathcal{I}_{t-1}^1 and \mathcal{I}_{t-1}^2 are both marginally distributed according to $B_t^{N,\text{FFBS}}(\mathcal{I}_t^1, \cdot)$. Moreover, if (13) holds, we have

$$(31) \quad \mathbb{P} \left(\mathcal{I}_{t-1}^1 \neq \mathcal{I}_{t-1}^2 \mid \mathcal{I}_{t:T}^{1,2}, \mathcal{F}_T^- \right) \mathbb{1}_{\mathcal{I}_t^1 = \mathcal{I}_t^2} \geq \varepsilon_S \mathbb{1}_{\mathcal{I}_t^1 = \mathcal{I}_t^2}.$$

In particular, the sequence of variables $(\mathcal{I}_{T-s}^1, \mathcal{I}_{T-s}^2)_{s=0}^T$ is a Markov chain given \mathcal{F}_T^- .

Proof. To simplify the notations, let \tilde{b}_t^n denote the \mathbb{R}^n vector $\hat{B}_t^N(n, \cdot)$. The relation between variables generated by Algorithm 2 is depicted as a graphical model in Figure 21. We consider

$$\begin{aligned}
(32) \quad p(\tilde{b}_t^{1:N}, i_{t-1}^{1:2} | \mathcal{F}_T^-, i_{t:T}^{1:2}) &= p(\tilde{b}_t^{1:N} | \mathcal{F}_T^-, i_{t:T}^{1:2}) p(i_{t-1}^{1:2} | \tilde{b}_t^{1:N}, \mathcal{F}_T^-, i_{t:T}^{1:2}) \\
&= p(\tilde{b}_t^{1:N} | x_{t-1}^{1:N}, x_t^{1:N}) p(i_{t-1}^{1:2} | \tilde{b}_t^{1:N}, i_t^{1:2}) \\
&\quad \text{(by properties of graphical models, see Figure 21)} \\
&= \left[\prod_n p(\tilde{b}_t^n | x_{t-1}^{1:N}, x_t^n) \right] \tilde{b}_t^{i_1^1}(i_{t-1}^1) \tilde{b}_t^{i_2^2}(i_{t-1}^2).
\end{aligned}$$

The distribution of i_{t-1}^1 given \mathcal{F}_T^- and $i_{t:T}^{1:2}$ is thus the i_{t-1}^1 -marginal of

$$(33) \quad p(\tilde{b}_t^{i_1^1} | x_{t-1}^{1:N}, x_t^{i_1^1}) \tilde{b}_t^{i_1^1}(i_{t-1}^1),$$

which is exactly the distribution of $p(j_t^{i_1^1} | x_{t-1}^{1:N}, x_t^{i_1^1})$ where the J 's are defined in the statement of Theorem 2.1. By the second hypothesis of that theorem, the aforementioned distribution is equal to $p(a_t^{i_1^1} | x_{t-1}^{1:N}, x_t^{i_1^1})$, which is in turn no other than $B_t^{N, \text{FFBS}}(i_t^1, \cdot)$. Moreover, if $i_t^1 \neq i_t^2$, (32) straightforwardly implies the conditional independence of i_{t-1}^1 and i_{t-1}^2 . When $i_t^1 = i_t^2$, the distribution of $i_{t-1}^{1:2}$ given \mathcal{F}_T^- and $i_{t:T}^{1:2}$ is the $i_{t-1}^{1:2}$ -marginal of

$$p(\tilde{b}_t^{i_1^1} | x_{t-1}^{1:N}, x_t^{i_1^1}) \tilde{b}_t^{i_1^1}(i_{t-1}^1) \tilde{b}_t^{i_1^1}(i_{t-1}^2).$$

Thus, we can apply (13) for $n = i_t^1$, where $i_{t-1}^{1:2}$ here plays the role of $J_t^{1:2}$ there. Equation (31) is now proved. \square

As Lemma E.1 describes the distribution of two trajectories, it immediately gives the distribution of a single trajectory.

Corollary 4. *Under the same settings as in Lemma E.1, given \mathcal{F}_T^- , the distribution of $\mathcal{I}_{0:T}^1$ is*

$$\mathcal{M}(W_T^{1:N})(di_T) B_T^{N, \text{FFBS}}(i_T, di_{T-1}) \dots B_1^{N, \text{FFBS}}(i_1, di_0).$$

Note that the corollary applies even if the backward kernel used in Algorithm 2 is *not* the FFBS one. This is due to the conditioning on \mathcal{F}_T^- and the second hypothesis of Theorem 2.1.

Proof of Theorem 2.2. First of all, we remark that as per Algorithm 2, using index variables $\mathcal{I}_{0:T}^{1:N}$ adds a level of Monte Carlo approximation to $\mathbb{Q}_T^N(dx_{0:T})$. Therefore

$$\begin{aligned}
(34) \quad \mathbb{E} [(\mathbb{Q}_T^N(\varphi_T) - \mathbb{Q}_T(\varphi_T))^2] &= \mathbb{E} \left[\left(\frac{1}{N} \sum_{n=1}^N \varphi_T(X_0^{\mathcal{I}_0^n}, \dots, X_T^{\mathcal{I}_T^n}) - \mathbb{Q}_T(\varphi_T) \right)^2 \right] \\
&= \mathbb{E} \left[(\mathbb{Q}_T^{N, \text{FFBS}}(\varphi_T) - \mathbb{Q}_T(\varphi_T))^2 \right] + \\
&\quad + \mathbb{E} \left[\text{Var} \left(\frac{1}{N} \sum_{n=1}^N \varphi_T(X_0^{\mathcal{I}_0^n}, \dots, X_T^{\mathcal{I}_T^n}) \middle| \mathcal{F}_T^- \right) \right]
\end{aligned}$$

where the ultimate inequality is justified by the law of total expectation and Corollary 4. (Note that $(\mathcal{I}_{0:T}^n)_{n=1}^N$ are identically distributed but *not* necessarily independent given \mathcal{F}_T^- .) Using Lemma E.3

(stated and proved below) and putting $\rho := 1 - \bar{M}_\ell/\bar{M}_h$, we have

$$\begin{aligned}
& \text{Var} \left(\frac{1}{N} \sum_{n=1}^N \varphi_T(X_0^{I_n^0}, \dots, X_T^{I_n^T}) \middle| \mathcal{F}_T^- \right) \\
&= \text{Var} \left(\frac{1}{N} \sum_{n=1}^N \sum_{t=0}^T \psi_t(X_{t-1}^{I_n^t}, X_t^{I_n^t}) \middle| \mathcal{F}_T^- \right) \\
&= \frac{1}{N^2} \sum_{n,m \leq N} \sum_{s,t \leq T} \text{Cov} \left(\psi_t(X_{t-1}^{I_n^t}, X_t^{I_n^t}), \psi_s(X_{s-1}^{I_m^s}, X_s^{I_m^s}) \middle| \mathcal{F}_T^- \right) \\
&\leq \frac{2}{N^2} \sum_{\substack{n,m \leq N \\ n=m}} \sum_{s,t \leq T} \|\psi_t\|_\infty \|\psi_s\|_\infty \rho^{|t-s|-1} + \\
&\quad + \frac{4}{N^2} \sum_{\substack{n,m \leq N \\ n \neq m}} \sum_{s,t \leq T} \frac{\tilde{C}}{N} \|\psi_t\|_\infty \|\psi_s\|_\infty \rho^{|t-s|-1} \\
&= \left(\sum_{s,t \leq T} 2 \|\psi_t\|_\infty \|\psi_s\|_\infty \rho^{|t-s|-1} \right) \frac{(2\tilde{C} + 1)N - 2\tilde{C}}{N^2} \\
&\leq \left[\sum_{s,t \leq T} \left(\|\psi_t\|_\infty^2 + \|\psi_s\|_\infty^2 \right) \rho^{|t-s|-1} \right] \frac{2\tilde{C} + 1}{N} \leq \frac{4(2\tilde{C} + 1)}{N\rho(1-\rho)} \sum \|\psi_t\|_\infty^2.
\end{aligned} \tag{35}$$

We now look at the first term of (34). In the fixed marginal smoothing case, for any $s \in \mathbb{Z}_+$, $s \leq T$ and any function $\phi_s : \mathcal{X}_s \rightarrow \mathbb{R}$, Douc et al. (2011) proved that

$$\mathbb{P} \left(\left| \mathbb{Q}_T^{N, \text{FFBS}}(\varphi_T) - \mathbb{Q}_T(\varphi_T) \right| \geq \varepsilon \right) \leq B' e^{-C' N \varepsilon^2 / \|\phi_s\|_\infty^2}$$

for $\varphi_T(x_{0:T}) = \phi_s(x_s)$ and constants B' and C' not depending on T . Using $\mathbb{E}[\Delta^2] = \int_0^\infty \mathbb{P}(\Delta^2 \geq t) dt$, the inequality implies

$$\mathbb{E} \left[\left| \mathbb{Q}_T^{N, \text{FFBS}}(\varphi_T) - \mathbb{Q}_T(\varphi_T) \right|^2 \right] \leq \frac{B' \|\phi_s\|_\infty^2}{C' N} \tag{36}$$

for $\varphi_T(x_{0:T}) = \phi_s(x_s)$. In the additive smoothing case, Dubarry and Le Corff (2013) proved that, for $T \geq 2$,

$$\mathbb{E} \left[\left| \mathbb{Q}_T^{N, \text{FFBS}}(\varphi_T) - \mathbb{Q}_T(\varphi_T) \right|^2 \right] \leq \frac{C'}{N} \left(\sum_{t=0}^T \|\psi_t\|_\infty^2 \right) \left(1 + \sqrt{\frac{T}{N}} \right)^2. \tag{37}$$

Equations (36), (37), (35) and (34) conclude the proof. \square

The following lemma quantifies the backward mixing property induced by Assumption 2.

Lemma E.2. *Under the same setting as Theorem 2.2, we have*

$$\text{TV} \left(B_t^{N, \text{FFBS}}(m, \cdot), B_t^{N, \text{FFBS}}(n, \cdot) \right) \leq 1 - \frac{\bar{M}_\ell}{\bar{M}_h}$$

for all $m, n \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$.

Proof. We have

$$\begin{aligned}
& 1 - \text{TV} \left(B_t^{N, \text{FFBS}}(m, \cdot), B_t^{N, \text{FFBS}}(n, \cdot) \right) \\
&= \left[\sum_{i=1}^N \min \left(\frac{G_{t-1}(X_{t-1}^i) m_t(X_{t-1}^i, X_t^m)}{\sum_{j=1}^N G_{t-1}(X_{t-1}^j) m_t(X_{t-1}^j, X_t^m)}, \right. \right. \\
&\quad \left. \left. \frac{G_{t-1}(X_{t-1}^i) m_t(X_{t-1}^i, X_t^n)}{\sum_{j=1}^N G_{t-1}(X_{t-1}^j) m_t(X_{t-1}^j, X_t^n)} \right) \right] \text{ by Lemma A.1 (Supplement A.2)} \\
&\geq \left[\sum_{i=1}^N \frac{G_t(X_{t-1}^i) \bar{M}_\ell}{\sum_{j=1}^N G_t(X_{t-1}^j) \bar{M}_h} \right] \text{ by Assumption 2} \\
&= (\bar{M}_\ell / \bar{M}_h).
\end{aligned}$$

□

Lemma E.3. *Under the same settings as in Theorem 2.2, for any $m, n \in \{1, \dots, N\}$ and $s, s' \in \{0, \dots, T\}$, we have*

$$\begin{aligned}
(38) \quad & \text{Cov} \left(\psi_s(X_{s-1}^{\mathcal{I}_s^m}, X_s^{\mathcal{I}_s^m}), \psi_{s'}(X_{s'-1}^{\mathcal{I}_{s'}^n}, X_{s'}^{\mathcal{I}_{s'}^n}) \middle| \mathcal{F}_T^- \right) \\
& \leq 2 \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right)^{|s-s'|-1} \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \times \begin{cases} 2\tilde{C}/N & \text{if } m \neq n \\ 1 & \text{if } m = n \end{cases}
\end{aligned}$$

where $\tilde{C} = \tilde{C}(\bar{M}_\ell, \bar{M}_h, \bar{G}_\ell, \bar{G}_h, \varepsilon_S)$ is a constant that does not depend on T (and which arises in the formulation of Lemma E.4). If s or s' is equal to 0, we adopt the natural convention $\psi_0(x_{-1}, x_0) := \psi_0(x_0)$.

Proof. We first handle the case $m \neq n$. Without loss of generality, assume that $m = 1, n = 2$ and $s \geq s'$. The covariance bound of Lemma A.1 yields

$$\begin{aligned}
(39) \quad & \text{Cov} \left(\psi_s(X_{s-1}^{\mathcal{I}_s^1}, X_s^{\mathcal{I}_s^1}), \psi_{s'}(X_{s'-1}^{\mathcal{I}_{s'}^2}, X_{s'}^{\mathcal{I}_{s'}^2}) \middle| \mathcal{F}_T^- \right) \\
& \leq 2 \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \text{TV} \left((\mathcal{I}_{s-1:s}^1, \mathcal{I}_{s'-1:s'}^2) \middle| \mathcal{F}_T^-, (\mathcal{I}_{s-1:s}^1 \middle| \mathcal{F}_T^-) \otimes (\mathcal{I}_{s'-1:s'}^2 \middle| \mathcal{F}_T^-) \right).
\end{aligned}$$

We shall bound this total variation distance via the coupling inequality of Lemma A.1 (Supplement A.2). The idea is to construct, in addition to $\mathcal{I}_{0:T}^1$ and $\mathcal{I}_{0:T}^2$, two trajectories $\mathcal{I}_{0:T}^{*1}$ and $\mathcal{I}_{0:T}^{*2}$ i.i.d. given \mathcal{F}_T^- such that each of them is conditionally distributed according to $\mathcal{I}_{0:T}^1$ (cf. Corollary 4). To make the coupling inequality efficient, it is desirable to make $\mathcal{I}_{0:T}^1$ and $\mathcal{I}_{0:T}^{*1}$ as similar as possible (same thing for $\mathcal{I}_{0:T}^2$ and $\mathcal{I}_{0:T}^{*2}$).

The detailed construction of the four trajectories $\mathcal{I}_{0:T}^1, \mathcal{I}_{0:T}^2, \mathcal{I}_{0:T}^{*1}$ and $\mathcal{I}_{0:T}^{*2}$ given \mathcal{F}_T^- is described in Algorithm 15. In particular, we ensure that $\forall t \geq s-1$, we have $\mathcal{I}_t^1 = \mathcal{I}_t^{*1}$. For $t \leq s-1$, if $\mathcal{I}_t^2 = \mathcal{I}_t^{*2}$, it is guaranteed that $\mathcal{I}_t^2 = \mathcal{I}_t^{*2}$ holds $\forall \ell \leq t$. The rationale for different coupling behaviours between the times $t \geq s-1$ and $t \leq s-1$ will become clear in the proof: the former aim to control the correlation between two different trajectories $m = 1$ and $n = 2$ and result in the \tilde{C}/N term of (38); the latter are for bounding the correlation between two times s and s' and result in the $(1 - \bar{M}_\ell / \bar{M}_h)^{|s-s'|-1}$ term of the same equation.

The correctness of Algorithm 15 is asserted by Lemma E.1. Step (\star) is valid because that lemma states that the distribution of \mathcal{I}_{t-1}^k given $\mathcal{F}_T^-, \mathcal{I}_{t:T}^{1,2}$ and $\mathcal{I}_{t:T}^{*1,2}$ is $B_t^{N, \text{FFBS}}(\mathcal{I}_t^k, \cdot)$. Furthermore, we note that $(R_{T-t})_{t=0}^T$ where

$$R_t := (\mathcal{I}_t^1, \mathcal{I}_t^2, \mathcal{I}_t^{*1}, \mathcal{I}_t^{*2}),$$

is a Markov chain given \mathcal{F}_T^- .

Algorithm 15: Sampler for the variables $\mathcal{I}_{0:T}^1, \mathcal{I}_{0:T}^2, \mathcal{I}_{0:T}^{*1}$ and $\mathcal{I}_{0:T}^{*2}$ (see proof of Lemma E.3)

Input: Feynman-Kac model (1), variables $X_{0:T}^{1:N}$ from the output of Algorithm 1, integer $s \geq 0$ (see statement of Lemma E.3)

Sample $\mathcal{I}_T^1, \mathcal{I}_T^2 \stackrel{i.i.d.}{\sim} \mathcal{M}(W_T^{1:N})$

Set $\mathcal{I}_T^{*1} \leftarrow \mathcal{I}_T^1$ and $\mathcal{I}_T^{*2} \leftarrow \mathcal{I}_T^2$

for $t \leftarrow T$ **to** 1 **do**

if $\mathcal{I}_t^1 \neq \mathcal{I}_t^2$ **then**

for $k \in \{1, 2\}$ **do**

 Sample $(\mathcal{I}_{t-1}^k, \mathcal{I}_{t-1}^{*k})$ from any maximal coupling of $B_t^{N, \text{FFBS}}(\mathcal{I}_t^k, \cdot)$ and $B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*k}, \cdot)$
 (cf. Lemma A.1)

else

 Sample the \mathbb{R}^N vector $\hat{B}_t^N(\mathcal{I}_t^1, \cdot)$ from $p(\hat{b}_t^N(i_t^1, \cdot) | x_{t-1}^{1:N}, x_t^{i_t^1})$

 Sample $\mathcal{I}_{t-1}^1, \mathcal{I}_{t-1}^{*1} \stackrel{i.i.d.}{\sim} \hat{B}_t^N(\mathcal{I}_t^1, \cdot)$

 Set $k \leftarrow 1, \ell \leftarrow 2$ if $t \geq s$ and $k \leftarrow 2, \ell \leftarrow 1$ otherwise

 Sample $\mathcal{I}_{t-1}^{*k} \sim B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*k}, \cdot)$ such that $(\mathcal{I}_{t-1}^{*k}, \mathcal{I}_{t-1}^k)$ is any maximal coupling of $B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*k}, \cdot)$ and $B_t^{N, \text{FFBS}}(\mathcal{I}_t^k, \cdot)$ given $\mathcal{I}_{t:T}^{1:2}, \mathcal{I}_{t:T}^{*1:2}$ and \mathcal{F}_T^- (\star) - see text for validity of this step)

 Sample $\mathcal{I}_{t-1}^{*\ell} \sim B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*\ell}, \cdot)$

Output: Four trajectories $\mathcal{I}_{0:T}^1, \mathcal{I}_{0:T}^2, \mathcal{I}_{0:T}^{*1}, \mathcal{I}_{0:T}^{*2}$ to be used in the proof of Lemma E.3

From (39), applying the coupling inequality of Lemma A.1 gives

$$\begin{aligned}
(40) \quad \text{Cov} \left(\psi_s(X_{s-1}^{\mathcal{I}_{s-1}^1}, X_s^{\mathcal{I}_s^1}), \psi_{s'}(X_{s'-1}^{\mathcal{I}_{s'-1}^2}, X_{s'}^{\mathcal{I}_{s'}^2}) \middle| \mathcal{F}_T^- \right) \\
\leq 2 \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \mathbb{P} \left((\mathcal{I}_{s-1:s}^1, \mathcal{I}_{s'-1:s'}^2) \neq (\mathcal{I}_{s-1:s}^{*1}, \mathcal{I}_{s'-1:s'}^{*2}) \middle| \mathcal{F}_T^- \right) \\
= 2 \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \mathbb{P} \left(\mathcal{I}_{s'-1:s'}^2 \neq \mathcal{I}_{s'-1:s'}^{*2} \middle| \mathcal{F}_T^- \right)
\end{aligned}$$

where the last equality results from the construction of Algorithm 15. The sub-case $s = s'$ following directly from Lemma E.4, we now focus on the sub-case $s \geq s' + 1$. For all $t \leq s - 1$,

$$\begin{aligned}
(41) \quad & \mathbb{P} \left(\mathcal{I}_{t-1}^2 \neq \mathcal{I}_{t-1}^{*2} \middle| \mathcal{F}_T^- \right) \\
&= \mathbb{P} \left(\mathcal{I}_{t-1}^2 \neq \mathcal{I}_{t-1}^{*2}, \mathcal{I}_t^2 \neq \mathcal{I}_t^{*2} \middle| \mathcal{F}_T^- \right) \\
&\quad \text{by construction of Algorithm 15} \\
&= \mathbb{E} \left[\mathbb{P} \left(\mathcal{I}_{t-1}^2 \neq \mathcal{I}_{t-1}^{*2}, \mathcal{I}_t^2 \neq \mathcal{I}_t^{*2} \middle| R_t, \mathcal{F}_T^- \right) \middle| \mathcal{F}_T^- \right] \\
&\quad \text{by the law of total expectation} \\
&= \mathbb{E} \left[\text{TV} \left(B_t^{N, \text{FFBS}}(\mathcal{I}_t^2, \cdot), B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*2}, \cdot) \right) \mathbb{1} \{ \mathcal{I}_t^2 \neq \mathcal{I}_t^{*2} \} \middle| \mathcal{F}_T^- \right] \\
&\leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right) \mathbb{P} \left(\mathcal{I}_t^2 \neq \mathcal{I}_t^{*2} \middle| \mathcal{F}_T^- \right) \text{ by Lemma E.2.}
\end{aligned}$$

Thus

$$\begin{aligned}
& \mathbb{P} \left(\mathcal{I}_{s'-1:s'}^2 \neq \mathcal{I}_{s'-1:s'}^{*2} \middle| \mathcal{F}_T^- \right) \\
&= \mathbb{P} \left(\mathcal{I}_{s'}^2 \neq \mathcal{I}_{s'}^{*2} \middle| \mathcal{F}_T^- \right) \text{ by construction of Algorithm 15} \\
&\leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right)^{s-s'-1} \mathbb{P} \left(\mathcal{I}_{s-1}^2 \neq \mathcal{I}_{s-1}^{*2} \middle| \mathcal{F}_T^- \right) \text{ by applying (41) recursively} \\
&\leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right)^{s-s'-1} \frac{\tilde{C}}{N} \text{ by Lemma E.4,}
\end{aligned}$$

which, combined with (40) finishes the proof for the current sub-case $s \geq s' + 1$. It remains to show (38) when $m = n$. The proof follows the same lines as in the case $m \neq n$, although we shall briefly outline

some arguments to show how the factor \tilde{C}/N disappeared. The case $s = s'$ being trivial, suppose that $s \geq s' + 1$ and without loss of generality that $m = n = 3$. To use the coupling tools of Lemma A.1, we construct trajectories $\mathcal{I}_{0:T}^3$, $\mathcal{I}_{0:T}^{*3}$ and $\mathcal{I}_{0:T}^{*4}$ via Algorithm 16 and write, in the spirit of (40):

Algorithm 16: Sampler for the variables $\mathcal{I}_{0:T}^3$, $\mathcal{I}_{0:T}^{*3}$ and $\mathcal{I}_{0:T}^{*4}$ (see proof of Lemma E.3)

Input: Feynman-Kac model (1), variables $X_{0:T}^{1:N}$ from the output of Algorithm 1, integer $s \geq 0$ (see statement of Lemma E.3)

Sample $\mathcal{I}_T^{*3}, \mathcal{I}_T^{*4} \stackrel{i.i.d.}{\sim} \mathcal{M}(W_T^{1:N})$

Set $\mathcal{I}_T^3 \leftarrow \mathcal{I}_T^{*3}$

for $t \leftarrow T$ **to** 1 **do**

if $t \geq s$ **then**

 Sample $\mathcal{I}_{t-1}^{*3} \sim B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*3}, \cdot)$ and $\mathcal{I}_{t-1}^{*4} \sim B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*4}, \cdot)$

 Set $\mathcal{I}_{t-1}^3 \leftarrow \mathcal{I}_{t-1}^{*3}$

else

 Sample $(\mathcal{I}_{t-1}^3, \mathcal{I}_{t-1}^{*4})$ from a maximal coupling of $B_t^{N, \text{FFBS}}(\mathcal{I}_t^3, \cdot)$ and $B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*4}, \cdot)$

 Sample $\mathcal{I}_{t-1}^{*3} \sim B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*3}, \cdot)$

Output: Three trajectories $\mathcal{I}_{0:T}^3$, $\mathcal{I}_{0:T}^{*3}$ and $\mathcal{I}_{0:T}^{*4}$ to be used in the proof of Lemma E.3

$$(42) \quad \text{Cov} \left(\psi_s(X_{s-1}^{\mathcal{I}_{s-1}^3}, X_s^{\mathcal{I}_s^3}), \psi_{s'}(X_{s'-1}^{\mathcal{I}_{s'-1}^3}, X_{s'}^{\mathcal{I}_{s'}^3}) \middle| \mathcal{F}_T^- \right) \\ \leq 2 \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \mathbb{P} \left((\mathcal{I}_{s-1:s}^3, \mathcal{I}_{s'-1:s'}^3) \neq (\mathcal{I}_{s-1:s}^{*3}, \mathcal{I}_{s'-1:s'}^{*4}) \middle| \mathcal{F}_T^- \right) \\ = 2 \|\psi_s\|_\infty \|\psi_{s'}\|_\infty \mathbb{P} \left(\mathcal{I}_{s'}^3 \neq \mathcal{I}_{s'}^{*4} \middle| \mathcal{F}_T^- \right)$$

where the last equality follows from the construction of Algorithm 16 and the hypothesis $s \geq s' + 1$. For all $t \leq s - 1$, the inequality

$$(43) \quad \mathbb{P} \left(\mathcal{I}_{t-1}^3 \neq \mathcal{I}_{t-1}^{*4} \middle| \mathcal{F}_T^- \right) \leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right) \mathbb{P} \left(\mathcal{I}_t^3 \neq \mathcal{I}_t^{*4} \middle| \mathcal{F}_T^- \right)$$

can be proved using the same techniques as those used to prove (41): applying Lemma E.2 given $(\mathcal{I}_t^3, \mathcal{I}_t^{*3}, \mathcal{I}_t^{*4})$ then invoking the law of total expectation. Repeatedly instantiating (43) gives

$$\mathbb{P} \left(\mathcal{I}_{s'}^3 \neq \mathcal{I}_{s'}^{*4} \middle| \mathcal{F}_T^- \right) \leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right)^{s-s'-1} \mathbb{P} \left(\mathcal{I}_{s-1}^3 \neq \mathcal{I}_{s-1}^{*4} \middle| \mathcal{F}_T^- \right) \\ \leq \left(1 - \frac{\bar{M}_\ell}{\bar{M}_h} \right)^{s-s'-1}$$

which, when plugged into (42), finishes the proof. \square

Lemma E.4. For \mathcal{I}_s^2 and \mathcal{I}_s^{*2} defined by the output of Algorithm 15, we have

$$\mathbb{P} \left(\mathcal{I}_s^2 \neq \mathcal{I}_s^{*2} \middle| \mathcal{F}_T^- \right) \leq \tilde{C}/N, \text{ and} \\ \mathbb{P} \left(\mathcal{I}_{s-1}^2 \neq \mathcal{I}_{s-1}^{*2} \middle| \mathcal{F}_T^- \right) \leq \tilde{C}/N, \text{ if } s \geq 1,$$

for some constant $\tilde{C} = \tilde{C}(\bar{M}_\ell, \bar{M}_h, \bar{G}_\ell, \bar{G}_h, \varepsilon_S)$.

Proof. Define $A_t := \mathbb{1} \{ \mathcal{I}_t^1 \neq \mathcal{I}_t^2 \}$, $B_t := \mathbb{1} \{ \mathcal{I}_t^2 = \mathcal{I}_t^{*2} \}$ and $\Gamma_t := A_t B_t$ and recall that $R_t := (\mathcal{I}_t^1, \mathcal{I}_t^2, \mathcal{I}_t^{*1}, \mathcal{I}_t^{*2})$. The sequence $(R_{T-\ell})_{\ell=0}^T$ is a Markov chain given \mathcal{F}_T^- , but this is not necessarily the case for the sequence $(\Gamma_{T-\ell})_{\ell=0}^T$ of Bernoulli random variables. Nevertheless, Lemma E.5 below shows that one can get bounds on two-step “transition probabilities” for $(\Gamma_{T-\ell})$, i.e. the probabilities under \mathcal{F}_T^- that $\Gamma_{t-2} = 1$ given Γ_t

and R_t . This motivates our following construction of actual Markov chains approximating the dynamic of Γ_t . Let Γ_T^* and Γ_{T-1}^* be two independent Bernoulli random variables given \mathcal{F}_T^- such that

$$(44) \quad \begin{aligned} \mathbb{P}(\Gamma_T^* = 1 | \mathcal{F}_T^-) &= \mathbb{P}(\Gamma_T = 1 | \mathcal{F}_T^-) \\ \mathbb{P}(\Gamma_{T-1}^* = 1 | \mathcal{F}_T^-) &= \mathbb{P}(\Gamma_{T-1} = 1 | \mathcal{F}_T^-). \end{aligned}$$

Let $\Gamma_T^*, \Gamma_{T-2}^*, \Gamma_{T-4}^*, \dots$ and $\Gamma_{T-1}^*, \Gamma_{T-3}^*, \dots$ be two homogeneous Markov chains given \mathcal{F}_T^- with the same transition kernel $\overleftarrow{\mathfrak{C}}^2$ defined by

$$(45) \quad \begin{aligned} \mathbb{P}_{\mathcal{F}_T^-}(\Gamma_{t-2}^* = 1 | \Gamma_t^* = 1) &= 1 - \frac{2}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 =: \overleftarrow{\mathfrak{C}}_{11}^2 \\ \mathbb{P}_{\mathcal{F}_T^-}(\Gamma_{t-2}^* = 1 | \Gamma_t^* = 0) &= \frac{\bar{M}_\ell \varepsilon_S}{2\bar{M}_h} =: \overleftarrow{\mathfrak{C}}_{01}^2 \end{aligned}$$

where for two events E_1, E_2 , the notation $\mathbb{P}_{\mathcal{F}_T^-}(E_1 | E_2)$ is the ratio between $\mathbb{P}(E_1, E_2 | \mathcal{F}_T^-)$ and $\mathbb{P}(E_2 | \mathcal{F}_T^-)$. We shall now prove by backward induction the following statement:

$$(46) \quad \mathbb{P}(\Gamma_t = 1 | \mathcal{F}_T^-) \geq \mathbb{P}(\Gamma_t^* = 1 | \mathcal{F}_T^-), \forall t \geq s - 1.$$

Firstly, (46) holds for $t = T$ and $t = T - 1$. Now suppose that it holds for some $t \geq s + 1$ and we wish to justify it for $t - 2$. By Lemma E.5,

$$\begin{aligned} \mathbb{P}(\Gamma_{t-2} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{\Gamma_t=1} &\geq \overleftarrow{\mathfrak{C}}_{11}^2 \mathbb{1}_{\Gamma_t=1} \\ \mathbb{P}(\Gamma_{t-2} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{\Gamma_t=0} &\geq \overleftarrow{\mathfrak{C}}_{01}^2 \mathbb{1}_{\Gamma_t=0}. \end{aligned}$$

Applying the law of total expectation gives

$$\begin{aligned} \mathbb{P}(\Gamma_{t-2} = 1 | \mathcal{F}_T^-) &\geq \overleftarrow{\mathfrak{C}}_{11}^2 \mathbb{P}(\Gamma_t = 1 | \mathcal{F}_T^-) + \overleftarrow{\mathfrak{C}}_{01}^2 \mathbb{P}(\Gamma_t = 0 | \mathcal{F}_T^-) \\ &= \left(\overleftarrow{\mathfrak{C}}_{11}^2 - \overleftarrow{\mathfrak{C}}_{01}^2 \right) \mathbb{P}(\Gamma_t = 1 | \mathcal{F}_T^-) + \overleftarrow{\mathfrak{C}}_{01}^2 \\ &\geq \left(\overleftarrow{\mathfrak{C}}_{11}^2 - \overleftarrow{\mathfrak{C}}_{01}^2 \right) \mathbb{P}(\Gamma_t^* = 1 | \mathcal{F}_T^-) + \overleftarrow{\mathfrak{C}}_{01}^2 \\ &\quad \text{if } N \text{ is large enough, by induction hypothesis} \\ &= \mathbb{P}(\Gamma_{t-2}^* = 1 | \mathcal{F}_T^-) \end{aligned}$$

and (46) is now proved. To finish the proof of the lemma, it is necessary to lower bound its right hand side. We start by controlling the distribution Γ_t^* for $t = T$ and $t = T - 1$. We have

$$(47) \quad \begin{aligned} \mathbb{P}(\Gamma_T^* = 1 | \mathcal{F}_T^-) &= \mathbb{P}(\Gamma_T = 1 | \mathcal{F}_T^-) \text{ by (44)} \\ &= 1 - \mathbb{P}(A_T = 0 | \mathcal{F}_T^-) \text{ as } B_T = 1 \text{ by Algorithm 15} \\ &= 1 - \sum_{i=1}^N \mathbb{P}(\mathcal{I}_T^1 = \mathcal{I}_T^2 = i | \mathcal{F}_T^-) \\ &= 1 - \sum_{i=1}^N \left(\frac{G(X_T^i)}{\sum_{j=1}^N G(X_T^j)} \right)^2 \\ &\geq 1 - \frac{1}{N} \left(\frac{\bar{G}_h}{\bar{G}_\ell} \right)^2 \text{ by Assumption 3} \end{aligned}$$

and

$$\begin{aligned}
(48) \quad \mathbb{P}(\Gamma_{T-1}^* = 1 | \mathcal{F}_T^-) &\geq \mathbb{P}(\Gamma_T = 1, \Gamma_{T-1} = 1 | \mathcal{F}_T^-) \\
&= \mathbb{E}[\mathbb{P}(\Gamma_T = 1, \Gamma_{T-1} = 1 | R_T, \mathcal{F}_T^-) | \mathcal{F}_T^-] \\
&\text{by the law of total expectation} \\
&= \mathbb{E}[\mathbb{P}(\Gamma_{T-1} = 1 | R_T, \mathcal{F}_T^-) \mathbb{1}_{\Gamma_T=1} | \mathcal{F}_T^-] \\
&\geq \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell}\right)^2\right] \mathbb{P}(\Gamma_T = 1 | \mathcal{F}_T^-) \text{ via (53)} \\
&\geq \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell}\right)^2\right] \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h}{\bar{G}_\ell}\right)^2\right].
\end{aligned}$$

The contraction property of Lemma A.1 makes it possible to relate the intermediate distributions $\Gamma_t^* | \mathcal{F}_T^-$ to the end point ones $\Gamma_{T-1}^* | \mathcal{F}_T^-$ and $\Gamma_T^* | \mathcal{F}_T^-$. More specifically, (45) and Lemma A.1 lead to

$$(49) \quad \text{TV}(\Gamma_t^* | \mathcal{F}_T^-, \mu^*) \leq \max(\text{TV}(\Gamma_T^* | \mathcal{F}_T^-, \mu^*), \text{TV}(\Gamma_{T-1}^* | \mathcal{F}_T^-, \mu^*))$$

where μ^* is the invariant distribution of a Markov chain with transition matrix $\overleftarrow{\mathfrak{E}}^2$, namely

$$(50) \quad \begin{cases} \mu^*(\{0\}) &= \frac{\overleftarrow{\mathfrak{E}}_{10}^2}{\overleftarrow{\mathfrak{E}}_{01}^2 + \overleftarrow{\mathfrak{E}}_{10}^2} \\ \mu^*(\{1\}) &= 1 - \mu^*(\{0\}). \end{cases}$$

Furthermore, an alternative expression of the total variation distance given in Lemma A.1 implies that the total variation distance between two Bernoulli distributions of parameters p and q is $|p - q|$. Combining this with (49), the triangle inequality and the rough estimate $\max(a, b) \leq a + b \forall a, b \geq 0$, we get

$$\mathbb{P}(\Gamma_t^* = 0 | \mathcal{F}_T^-) \leq 3\mu^*(\{0\}) + \mathbb{P}(\Gamma_T^* = 0 | \mathcal{F}_T^-) + \mathbb{P}(\Gamma_{T-1}^* = 0 | \mathcal{F}_T^-) \leq \tilde{C}/N$$

where $\tilde{C} = \tilde{C}(\bar{M}_\ell, \bar{M}_h, \bar{G}_\ell, \bar{G}_h, \varepsilon_S)$. The last inequality is straightforwardly derived by plugging respectively (50), (47) and (48) into the three terms of the preceding sum. This combined with (46) finishes the proof. \square

Lemma E.5. For s defined in the statement of Lemma E.3; A_t , B_t and R_t defined in the proof of Lemma E.4 and all $t \geq s + 1$, we have

$$\begin{aligned}
\mathbb{P}(A_{t-2}B_{t-2} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t B_t = 1} &\geq \left(1 - \frac{2}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell}\right)^2\right) \mathbb{1}_{A_t B_t = 1}; \\
\mathbb{P}(A_{t-2}B_{t-2} = 1 | R_t, \mathcal{F}_T^-) &\geq \frac{\bar{M}_\ell \varepsilon_S}{2\bar{M}_h}
\end{aligned}$$

where the inequalities hold for N large enough, i.e., $N \geq N_0 = N_0(\bar{M}_\ell, \bar{M}_h, \bar{G}_\ell, \bar{G}_h, \varepsilon_S)$.

Proof. We start by showing the following three inequalities for all $t \geq s$ and N sufficiently large:

$$(51) \quad \mathbb{P}(A_{t-1} = 1 | R_t, \mathcal{F}_T^-) \geq \varepsilon_S;$$

$$(52) \quad \mathbb{P}(A_{t-1}B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t=1} \geq (\bar{M}_\ell/2\bar{M}_h) \mathbb{1}_{A_t=1};$$

$$(53) \quad \mathbb{P}(A_{t-1}B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t B_t = 1} \geq \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell}\right)^2\right] \mathbb{1}_{A_t B_t = 1}.$$

For (51), we have

$$(54) \quad \mathbb{P}(A_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t \neq 1} = \mathbb{P}(\mathcal{I}_{t-1}^1 \neq \mathcal{I}_{t-1}^2 | R_t, \mathcal{F}_T^-) \mathbb{1}_{\mathcal{I}_t^1 = \mathcal{I}_t^2} \geq \varepsilon_S \mathbb{1}_{A_t \neq 1}$$

by Lemma E.1. Next,

$$\begin{aligned}
& \mathbb{P} (A_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t=1} \\
&= \mathbb{P} (\mathcal{I}_{t-1}^1 \neq \mathcal{I}_{t-1}^2 | R_t, \mathcal{F}_T^-) \mathbb{1}_{\mathcal{I}_t^1 \neq \mathcal{I}_t^2} \\
&= \left[1 - \sum_i \mathbb{P} (\mathcal{I}_{t-1}^1 = \mathcal{I}_{t-1}^2 = i | R_t, \mathcal{F}_T^-) \right] \mathbb{1}_{\mathcal{I}_t^1 \neq \mathcal{I}_t^2} \\
(55) \quad &= \left[1 - \sum_{i=1}^N \prod_{k=1}^2 \frac{G_{t-1}(X_{t-1}^i) m_t(X_{t-1}^i, X_t^{\mathcal{I}_t^k})}{\sum_{j=1}^N G_{t-1}(X_{t-1}^j) m_t(X_{t-1}^j, X_t^{\mathcal{I}_t^k})} \right] \mathbb{1}_{\mathcal{I}_t^1 \neq \mathcal{I}_t^2} \text{ by Lemma E.1} \\
&\geq \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 \right] \mathbb{1}_{A_t=1} \text{ by Assumptions 2 and 3.}
\end{aligned}$$

Combining (54) and (55) yields (51) for N large enough. To prove (52), we write

$$\begin{aligned}
(56) \quad & \mathbb{P} (A_{t-1} B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t=1} \\
&= [1 - \mathbb{P} (A_{t-1} B_{t-1} = 0 | R_t, \mathcal{F}_T^-)] \mathbb{1}_{A_t=1} \\
&\geq [1 - \mathbb{P} (A_{t-1} = 0 | R_t, \mathcal{F}_T^-) - \mathbb{P} (B_{t-1} = 0 | R_t, \mathcal{F}_T^-)] \mathbb{1}_{A_t=1} \\
&= [\mathbb{P} (A_{t-1} = 1 | R_t, \mathcal{F}_T^-) + \mathbb{P} (B_{t-1} = 1 | R_t, \mathcal{F}_T^-) - 1] \mathbb{1}_{A_t=1}.
\end{aligned}$$

We analyse the second term in the above expression. We have

$$\begin{aligned}
(57) \quad & \mathbb{P} (B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t=1} \\
&= \mathbb{P} (\mathcal{I}_{t-1}^2 = \mathcal{I}_{t-1}^{*2} | R_t, \mathcal{F}_T^-) \mathbb{1}_{\mathcal{I}_t^1 \neq \mathcal{I}_t^2} \\
&= \left[1 - \text{TV} \left(B_t^{N, \text{FFBS}}(\mathcal{I}_t^2, \cdot), B_t^{N, \text{FFBS}}(\mathcal{I}_t^{*2}, \cdot) \right) \right] \mathbb{1}_{A_t=1} \\
&\text{by construction of Algorithm 15} \\
&\geq (\bar{M}_\ell / \bar{M}_h) \mathbb{1}_{A_t=1} \text{ by Lemma E.2.}
\end{aligned}$$

Plugging (55) and (57) into (56) yields

$$\mathbb{P} (A_{t-1} B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t=1} \geq \left(-\frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 + \frac{\bar{M}_\ell}{\bar{M}_h} \right) \mathbb{1}_{A_t=1}$$

and thus (52) follows if N is large enough. The inequality (53) is justified by combining (55), the simple decomposition $\mathbb{1}_{A_t B_t=1} = \mathbb{1}_{A_t=1} \mathbb{1}_{B_t=1}$ and the fact that Algorithm 15 guarantees $B_{t-1} = 1$ if $A_t = B_t = 1$.

We can now deduce the two inequalities in the statement of the Lemma. The first one is a straightforward double application of (53):

$$\begin{aligned}
& \mathbb{P} (A_{t-2} B_{t-2} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t B_t=1} \\
&\geq \mathbb{P} (A_{t-2} B_{t-2} = 1, A_{t-1} B_{t-1} = 1 | R_t, \mathcal{F}_T^-) \mathbb{1}_{A_t B_t=1} \\
&= \mathbb{E} \left[\mathbb{P} (A_{t-2} B_{t-2} = 1, A_{t-1} B_{t-1} = 1 | R_{t-1}, R_t, \mathcal{F}_T^-) | R_t, \mathcal{F}_T^- \right] \mathbb{1}_{A_t B_t=1} \\
&\text{by the law of total expectation} \\
&= \mathbb{E} \left[\mathbb{P} (A_{t-2} B_{t-2} = 1 | R_{t-1}, \mathcal{F}_T^-) \mathbb{1}_{A_{t-1} B_{t-1}=1} | R_t, \mathcal{F}_T^- \right] \mathbb{1}_{A_t B_t=1} \\
&\text{since } (R_{T-\ell})_{\ell=0}^T \text{ is Markov given } \mathcal{F}_T^- \\
&\geq \mathbb{E} \left[\left(1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 \right) \mathbb{1}_{A_{t-1} B_{t-1}=1} \middle| R_t, \mathcal{F}_T^- \right] \mathbb{1}_{A_t B_t=1} \\
&\geq \left[1 - \frac{1}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 \right]^2 \mathbb{1}_{A_t B_t=1} \geq \left(1 - \frac{2}{N} \left(\frac{\bar{G}_h \bar{M}_h}{\bar{G}_\ell \bar{M}_\ell} \right)^2 \right) \mathbb{1}_{A_t B_t=1}.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
& \mathbb{P} (A_{t-2}B_{t-2} = 1 | R_t, \mathcal{F}_T^-) \\
& \geq \mathbb{P} (A_{t-2}B_{t-2} = 1, A_{t-1} = 1 | R_t, \mathcal{F}_T^-) \\
& = \mathbb{E} [\mathbb{P} (A_{t-2}B_{t-2} = 1 | R_{t-1}, \mathcal{F}_T^-) \mathbb{1}_{A_{t-1}=1} | R_t, \mathcal{F}_T^-] \\
& \quad \text{using law of total expectation and the Markov property as above} \\
& \geq \frac{\bar{M}_\ell}{2\bar{M}_h} \mathbb{P} (A_{t-1} = 1 | R_t, \mathcal{F}_T^-) \text{ by (52)} \\
& \geq \frac{\bar{M}_\ell}{2\bar{M}_h} \varepsilon_S \text{ by (51)}
\end{aligned}$$

and the second inequality is proved. \square

E.4. Proof of Proposition 3 (hybrid rejection validity).

Proof. Put $Z_n := (X_n, U_n C \mu_0(X_n))$. Then Z_n is uniformly distributed on

$$\mathcal{G}_0 := \{(x, y) \in \mathcal{X} \times \mathbb{R}_+, y \leq C \mu_0(x)\}.$$

The proof would be done if one could show that, given $K^* \leq K$, the variable Z_{K^*} is uniformly distributed on

$$\mathcal{G}_1 := \{(x, y) \in \mathcal{X} \times \mathbb{R}_+, y \leq \mu_1(x)\}.$$

Note that K^* is, by definition, the first time index where the sequence (Z_n) touches \mathcal{G}_1 . Let B be any subset of \mathcal{G}_1 . We have

$$\begin{aligned}
(58) \quad & \mathbb{P} (Z_{K^*} \in B | K^* \leq K) \propto \mathbb{P} (Z_{K^*} \in B, K^* \leq K) \\
& = \sum_{k^*=1}^{\infty} \mathbb{P} (Z_{k^*} \in B, K^* = k^*, K \geq k^*) \\
& = \sum_{k^*=1}^{\infty} \mathbb{P} (Z_{k^*} \in B, Z_{1:k^*-1} \notin \mathcal{G}_1, K > k^* - 1) \\
& = \sum_{k^*=1}^{\infty} \mathbb{P} (Z_{k^*} \in B) \mathbb{P} (Z_{1:k^*-1} \notin \mathcal{G}_1, K > k^* - 1) \text{ since } K \text{ stopping time} \\
& = \mathbb{P} (Z_1 \in B) \sum_{k^*=1}^{\infty} \mathbb{P} (Z_{1:k^*-1} \notin \mathcal{G}_1, K > k^* - 1) \\
& \propto \mathbb{P} (Z_1 \in B) \propto \mathbb{P} (Z_1 \in B | Z_1 \in \mathcal{G}_1).
\end{aligned}$$

By considering the special case $B = \mathcal{G}_1$, we see that the constant of proportionality between the first and the last terms of (58) must be 1, from which the proof follows. \square

E.5. Proof of Theorem 3.1 (hybrid algorithm's intermediate complexity). From (16), one may have the correct intuition that as $N \rightarrow \infty$, $\tau_t^{1, \text{PaRIS}}$ tends in distribution to that of the variable $\tau_t^{\infty, \text{PaRIS}}$ defined as

$$(59) \quad \tau_t^{\infty, \text{PaRIS}} | X_t^{\infty, \text{PaRIS}} \sim \text{Geo} \left(\frac{r_t(X_t^{\infty, \text{PaRIS}})}{\bar{M}_h} \right)$$

where $X_t^{\infty, \text{PaRIS}} \sim \mathbb{Q}_{t-1} M_t(dx_t)$ is distributed according to the predictive distribution of X_t given $Y_{0:t-1}$ and r_t is the density of $X_t^{\infty, \text{PaRIS}}$ with respect to the Lebesgue measure (cf. Definition 1). The following proposition formalises the connection between $\tau_t^{1, \text{PaRIS}}$ and $\tau_t^{\infty, \text{PaRIS}}$.

Proposition 6. *We have $\tau_t^{1, \text{PaRIS}} \Rightarrow \tau_t^{\infty, \text{PaRIS}}$ as $N \rightarrow \infty$.*

Proof. From (16) and Definition 1 one has

$$(60) \quad \tau_t^{1, \text{PaRIS}} \mid X_t^1, \mathcal{F}_{t-1} \sim \text{Geo} \left(\frac{r_t^N(X_t^1)}{\bar{M}_h} \right).$$

In light of (59), it suffices to establish that

$$(61) \quad \frac{r_t^N(X_t^1)}{\bar{M}_h} \Rightarrow \frac{r_t(X_t^{\infty, \text{PaRIS}})}{\bar{M}_h}.$$

Indeed, this would mean that for any continuous bounded function ψ , we have

$$\begin{aligned} \mathbb{E}[\psi(\tau_t^{1, \text{PaRIS}})] &= \mathbb{E} \left[(\text{Geo}^* \psi) \left(\frac{r_t^N(X_t^1)}{\bar{M}_h} \right) \right] \rightarrow \mathbb{E} \left[(\text{Geo}^* \psi) \left(\frac{r_t(X_t^{\infty, \text{PaRIS}})}{\bar{M}_h} \right) \right] \\ &= \mathbb{E}[\psi(\tau_t^{\infty, \text{PaRIS}})] \end{aligned}$$

where Geo^* is the geometric Markov kernel that sends each λ to the geometric distribution of parameter λ , i.e. $\text{Geo}^*(\lambda, dx) = \text{Geo}(\lambda)$. To this end, write

$$(62) \quad \begin{aligned} r_t^N(X_t^1) - r_t(X_t^1) &= \frac{\sum_n G_{t-1}(X_{t-1}^n) m_t(X_{t-1}^n, X_t^1)}{\sum_n G_{t-1}(X_{t-1}^n)} - r_t(X_t^1) \\ &= \frac{\sum_n N^{-1} G_{t-1}(X_{t-1}^n) [m_t(X_{t-1}^n, X_t^1) - r_t(X_t^1)]}{N^{-1} \sum_n G_{t-1}(X_{t-1}^n)}. \end{aligned}$$

We study the mean squared error of the numerator:

$$\begin{aligned} &\mathbb{E} \left\{ \frac{1}{N} \sum_n G_{t-1}(X_{t-1}^n) [m_t(X_{t-1}^n, X_t^1) - r_t(X_t^1)] \right\}^2 \\ &= \frac{1}{N} \mathbb{E} \left\{ G_{t-1}(X_{t-1}^1)^2 [m_t(X_{t-1}^1, X_t^1) - r_t(X_t^1)]^2 \right\} \\ &\quad + \frac{N(N-1)}{N^2} \mathbb{E} \left\{ G_{t-1}(X_{t-1}^1) G_{t-1}(X_{t-1}^2) [m_t(X_{t-1}^1, X_t^1) - r_t(X_t^1)] \right. \\ &\quad \left. \times [m_t(X_{t-1}^2, X_t^1) - r_t(X_t^1)] \right\} \end{aligned}$$

where we have again used the exchangeability induced by step (\star) of Algorithm 5. The first term obviously tends to 0 as $N \rightarrow \infty$ by Assumptions 4 and 1. The second term also vanishes asymptotically thanks to Lemma E.6 below and Assumption 6. Assumption 1 also implies that the denominator of (62) converges in probability to some constant, via the consistency of particle approximations, see e.g. Del Moral (2004) or Chopin and Papaspiliopoulos (2020). Thus, $r_t^N(X_t^1) - r_t(X_t^1) \Rightarrow 0$ by Slutsky's theorem. Moreover, $r_t(X_t^1) \Rightarrow r_t(X_t^{\infty, \text{PaRIS}})$ by the continuity of r_t and the consistency of particle approximations. Using again Slutsky's theorem yields (61). \square

The following lemma is needed to complete the proof of Proposition 6 and is related to the propagation of chaos property, see Del Moral (2004, Chapter 8).

Lemma E.6. *We have $(X_{t-1}^1, X_{t-1}^2, X_t^1) \Rightarrow \mathbb{Q}_{t-2} M_{t-1} \otimes \mathbb{Q}_{t-2} M_{t-1} \otimes \mathbb{Q}_{t-1} M_t$.*

Proof. For vectors u , v , and w , we have, by the symmetry of the distribution of particles:

$$\begin{aligned} &\mathbb{E} [\exp(iuX_{t-1}^1 + ivX_{t-1}^2 + iwX_t^1)] \\ &= \mathbb{E} \left[\left(\frac{1}{N} \sum e^{iuX_{t-1}^n} \right) \left(\frac{1}{N} \sum e^{ivX_{t-1}^n} \right) \left(\frac{1}{N} \sum e^{iwX_t^n} \right) \right] \\ &\quad - \frac{N}{N^2} \mathbb{E} \left[e^{iuX_{t-1}^1} e^{ivX_{t-1}^1} \left(\frac{1}{N} \sum e^{iwX_t^n} \right) \right] \\ &\quad + \frac{N}{N^2} \mathbb{E} \left[e^{iuX_{t-1}^1} e^{ivX_{t-1}^2} \left(\frac{1}{N} \sum e^{iwX_t^n} \right) \right]. \end{aligned}$$

Note that

$$\frac{1}{N} \sum e^{iuX_{t-1}^n} \xrightarrow{\text{a.s.}} \mathbb{Q}_{t-2} M_{t-1}(\exp(iu \bullet))$$

and

$$\frac{1}{N} \sum e^{iwX_t^n} \xrightarrow{\text{a.s.}} \mathbb{Q}_{t-1} M_t(\exp(iw\bullet)).$$

The dominated convergence theorem, applicable since $|e^{iu}| \leq 1$ for $u \in \mathbb{R}$, finishes the proof. \square

Proof of Theorem 3.1. First of all,

$$(63) \quad \mathbb{E}[\tau_t^{\infty, \text{PaRIS}}] = \mathbb{E} \left[\frac{\bar{M}_h}{r_t(X_t^{\infty, \text{PaRIS}})} \right] = \int_{\mathcal{X}_t} \frac{\bar{M}_h}{r_t(x_t)} r_t(x_t) dx_t = \infty$$

by Assumption 5. Next, for any $x \in \mathbb{R} \setminus \mathbb{Z}$ and $N > x$,

$$\mathbb{P} \left(\min(\tau_t^{1, \text{PaRIS}}, N) \leq x \right) = \mathbb{P} \left(\tau_t^{1, \text{PaRIS}} \leq x \right) \rightarrow \mathbb{P} \left(\tau_t^{\infty, \text{PaRIS}} \leq x \right)$$

by Proposition 6. Thus, by Portmanteau theorem,

$$(64) \quad \min(\tau_t^{1, \text{PaRIS}}, N) \Rightarrow \tau_t^{\infty, \text{PaRIS}}.$$

Altogether, we have

$$\begin{aligned} \liminf_{N \rightarrow \infty} \mathbb{E} \left[\min(\tau_t^{1, \text{PaRIS}}, N) \right] &= \liminf_{N \rightarrow \infty} \sum k \mathbb{P} \left(\min(\tau_t^{1, \text{PaRIS}}, N) = k \right) \\ &\geq \sum \liminf_{N \rightarrow \infty} k \mathbb{P} \left(\min(\tau_t^{1, \text{PaRIS}}, N) = k \right) \text{ by Fatou's lemma} \\ &= \sum k \mathbb{P} \left(\tau_t^{\infty, \text{PaRIS}} = k \right) \text{ by (64)} \\ &= \infty \text{ by (63)} \end{aligned}$$

and

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\min(\tau_t^{1, \text{PaRIS}}, N) \right] = \lim_{N \rightarrow \infty} \mathbb{E} \left[\min \left(\frac{\tau_t^{1, \text{PaRIS}}}{N}, 1 \right) \right] \rightarrow 0$$

since $\tau_t^{1, \text{PaRIS}} \Rightarrow \tau_t^{\infty, \text{PaRIS}}$ implies that the sequence of random variables

$$\min \left(\frac{\tau_t^{1, \text{PaRIS}}}{N}, 1 \right)$$

converges to 0 in distribution while being bounded between 0 and 1. \square

E.6. Proof of Theorem 3.2 (hybrid PaRIS near-linear complexity). The following proposition shows that the real execution time for the hybrid algorithm is asymptotically at most of the same order as the ‘‘oracle’’ hybrid execution time.

Proposition 7. *We have*

$$\limsup_{N \rightarrow \infty} \frac{\mathbb{E} \left[\min(\tau_t^{1, \text{PaRIS}}, N) \right]}{\mathbb{E} \left[\min(\tau_t^{\infty, \text{PaRIS}}, N) \right]} < \infty.$$

Proof. Put

$$(65) \quad z^N(\lambda) := \frac{1 - (1 - \lambda)^N}{\lambda} = \sum_{n=0}^{N-1} (1 - \lambda)^n.$$

One can quickly verify (using the memorylessness of the geometric distribution for example) that $z^N(\lambda) = \mathbb{E}[\min(G, N) | G \sim \text{Geo}(\lambda)]$. It will be useful to keep in mind the elementary estimate $z^N(\lambda) \leq \min(N, \lambda^{-1})$.

We can now write

$$\begin{aligned}
\mathbb{E} \left[\min(\tau_t^{1,\text{PaRIS}}, N) \right] &= \mathbb{E} \left[z^N \left(\frac{r_t^N(X_t^1)}{\bar{M}_h} \right) \right] \quad (\text{by (60)}) = \mathbb{E} \left[\mathbb{E} \left[z^N \left(\frac{r_t^N(X_t^1)}{\bar{M}_h} \right) \middle| \mathcal{F}_{t-1} \right] \right] \\
&= \mathbb{E} \left[\int_{\mathcal{X}_t} z^N \left(\frac{r_t^N(x_t)}{\bar{M}_h} \right) r_t^N(x_t) \lambda_t(dx_t) \right] \\
&\leq c_t \left(\int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right) r_t(x_t) \lambda_t(dx_t) + b_t \right) \quad \text{by Lemma E.7} \\
&= c_t \left(\mathbb{E}[\min(\tau_t^{\infty,\text{PaRIS}}, N)] + b_t \right)
\end{aligned}$$

from which the proposition is immediate. \square

Lemma E.7. *In addition to notations of Algorithm 1, let the function z^N be defined as in (65) and the functions r_t and r_t^N be defined as in Definition 1. Let $\phi_t : \mathcal{X}_t \rightarrow \mathbb{R}_{>0}$ be a bounded non-negative deterministic function. Then, under Assumptions 1 and 4, there exist constants b_t and c_t depending only on the model such that*

$$\mathbb{E} \left[\int_{\mathcal{X}_t} z^N \left(\frac{r_t^N(x_t)}{\bar{M}_h} \right) r_t^N \phi_t \right] \leq c_t \left(\int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right) r_t \phi_t + b_t \|\phi_t\|_\infty \right)$$

where for brevity, we shortened the integration notation (e.g. dropping $\lambda_t(dx_t)$, dropping x_t from $\phi(x_t)$, etc.) whenever there is no ambiguity.

Proof. We have

$$(66) \quad \mathbb{E} \left[\int_{\mathcal{X}_t} z^N \left(\frac{r_t^N(x_t)}{\bar{M}_h} \right) r_t^N \phi_t \right] \leq \int_{\mathcal{X}_t} z^N \left(\mathbb{E} \left[\frac{r_t^N(x_t)}{\bar{M}_h} \right] \right) \mathbb{E} [r_t^N(x_t)] \phi_t$$

using Fubini's theorem and the concavity of $\lambda \mapsto \lambda z^N(\lambda)$ on $[0, 1]$. By a well-known result on the bias of a particle filter (which is in fact the propagation of chaos in the special case of $q = 1$ particle), we have:

$$\begin{aligned}
|\mathbb{E} [r_t^N(x_t)] - r_t(x_t)| &= \left| \mathbb{E} \left[\sum W_{t-1}^n m_t(X_{t-1}^n, x_t) \right] - r_t(x_t) \right| \\
&= \left| \mathbb{E} \left[m_t \left(X_{t-1}^{A_t^1}, x_t \right) \right] - \mathbb{Q}_{t-1}(m_t(\bullet, x_t)) \right| \\
&\leq \frac{b_t \bar{M}_h}{N}
\end{aligned}$$

for some constant b_t . We next show that such a bias does not change the asymptotic behavior of z^N . More precisely,

$$\begin{aligned}
z^N \left(\mathbb{E} \left[\frac{r_t^N(x_t)}{\bar{M}_h} \right] \right) &\leq z^N \left(\frac{r_t(x_t)}{\bar{M}_h} - \frac{b_t}{N} \right) \\
&= \sum_{n=0}^{N-1} \left(\frac{1 - r_t(x_t)/\bar{M}_h + b_t/N}{1 - r_t(x_t)/\bar{M}_h} \right)^n \left(1 - \frac{r_t(x_t)}{\bar{M}_h} \right)^n \\
(67) \quad &\leq \sum_{n=0}^{N-1} \left(1 + \frac{b_t}{N(1 - r_t(x_t)/\bar{M}_h)} \right)^N \left(1 - \frac{r_t(x_t)}{\bar{M}_h} \right)^n \\
&\leq \exp \left(\frac{b_t}{1 - r_t(x_t)/\bar{M}_h} \right) z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right) \\
&\leq e^{2b_t} z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right)
\end{aligned}$$

if x_t is such that $r_t(x_t)/\bar{M}_h \leq 1/2$. In contrast, if $r_t(x_t)/\bar{M}_h \geq 1/2$, then provided that $N \geq 6b_t$, we have

$$(68) \quad z^N \left(\mathbb{E} \left[\frac{r_t^N(x_t)}{\bar{M}_h} \right] \right) \leq z^N \left(\frac{r_t(x_t)}{\bar{M}_h} - \frac{b_t}{N} \right) \leq z^N \left(\frac{1}{3} \right) \leq 3z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right).$$

Putting together (67) and (68), we have, for $N \geq 6b_t$,

$$z^N \left(\mathbb{E} \left[\frac{r_t^N(x_t)}{\bar{M}_h} \right] \right) \leq (e^{2b_t} + 3) z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right)$$

and so, by (66),

$$\begin{aligned} \mathbb{E} \left[\int_{\mathcal{X}_t} z^N \left(\frac{r_t^N(x_t)}{\bar{M}_h} \right) r_t^N \phi_t \right] &\leq (e^{2b_t} + 3) \int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right) \mathbb{E} [r_t^N(x_t)] \phi_t \\ &= (e^{2b_t} + 3) \mathbb{E} \left[z^N \left(\frac{r_t(X_t^1)}{\bar{M}_h} \right) \phi_t(X_t^1) \right]. \end{aligned}$$

Again, using the result on the bias of a particle filter,

$$\left| \mathbb{E} \left[z^N \left(\frac{r_t(X_t^1)}{\bar{M}_h} \right) \phi_t(X_t^1) \right] - \int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{\bar{M}_h} \right) r_t \phi_t \right| \leq \frac{b_t \|z^N\|_\infty \|\phi_t\|_\infty}{N} = b_t \|\phi_t\|_\infty$$

which, together with the previous inequality, implies the desired result. \square

Proposition 8. *In linear Gaussian state space models, we have*

$$\mathbb{E} \left[\min(\tau_t^{\infty, \text{PaRIS}}, N) \right] = \mathcal{O} \left((\log N)^{d_t/2} \right).$$

Proof. Let μ_t and Σ_t be such that $X_t^{\infty, \text{PaRIS}} \sim \mathcal{N}(\mu_t, \Sigma_t)$. Then

$$\log(r_t(X_t^{\infty, \text{PaRIS}})/\bar{M}_h) = b'_t - W_t$$

where b'_t is some constant and

$$W_t := \frac{(X_t^{\infty, \text{PaRIS}} - \mu_t)^\top \Sigma_t^{-1} (X_t^{\infty, \text{PaRIS}} - \mu_t)}{2} \sim \text{Gamma} \left(\frac{d_t}{2}, 1 \right).$$

We have

$$\begin{aligned} \mathbb{E} \left[\min(\tau_t^{\infty, \text{PaRIS}}, N) \right] &= \mathbb{E} \left[z^N \left(\frac{r_t(X_t^{\infty, \text{PaRIS}})}{\bar{M}_h} \right) \right] = \mathbb{E} \left[z^N (e^{b'_t - W_t}) \right] \\ &= \int_0^\infty z^N (e^{b'_t - w}) \frac{w^{d_t/2-1} e^{-w}}{\Gamma(d_t/2)} dw \\ &\leq \int_0^{\log N} e^{w - b'_t} \frac{w^{d_t/2-1} e^{-w}}{\Gamma(d_t/2)} dw + \int_{\log N}^\infty N \frac{w^{d_t/2-1} e^{-w}}{\Gamma(d_t/2)} dw \end{aligned}$$

using the bound $z^N(\lambda) \leq \min(N, 1/\lambda)$. The first term is of order $\mathcal{O}(\log^{d_t/2} N)$ by elementary calculus, while the second term is of order $\mathcal{O}(\log^{d_t/2-1} N)$ using asymptotic properties of the incomplete Gamma function, see [Olver et al. \(2010, Section 8.11\)](#). \square

Proof of Theorem 3.2. The theorem is a straightforward consequence of Proposition 7 and Proposition 8. \square

E.7. Proof of Theorems B.1 and B.2 (pure rejection FFBS complexity). We start with a useful remark linking the projection kernels Π and the cost-to-go functions defined in Supplement A with the L-kernels formulated in (27). The proof is simple and therefore omitted.

Lemma E.8. *We have $L_{t:T}(x_{0:t}, \mathbb{1}) = H_{t:T}(x_t)$ for all $x_{0:t}$. Moreover, for any function $\phi_t : \mathcal{X}_t \rightarrow \mathbb{R}$, we have*

$$L_{t:T} \Pi_t^{0:T} \phi_t = \Pi_t^{0:t} (\phi_t \times H_{t:T}).$$

Theorems B.1 and B.2 both rely on an induction argument wrapped up in the following proposition.

Proposition 9. *We use the notations of Algorithm 2. Let \mathbb{Q}_t^N be defined as in (24), where the B_s^N kernels can be $B_s^{N, \text{FFBS}}$ or any other kernels satisfying the hypotheses of Theorem 2.1. Suppose that Assumption 1 holds. Let $f_t^N : \mathcal{X}_t \rightarrow \mathbb{R}_{\geq 0}$ be a (possibly random) function such that $f_t^N(x_t)$ is \mathcal{F}_{t-1} -measurable. Then the following assertions are true:*

(a) Suppose that $\mathbb{E} \left[\int_{\mathcal{X}_t} \{r_t^N \times f_t^N \times G_t \times H_{t:T}\}(x_t) \lambda_t(dx_t) \right] = \infty$, where r_t^N and λ_t are defined in Definition 1. Then

$$\mathbb{E} \left[\int \mathbb{Q}_T^N(dx_t) f_t^N(x_t) \right] = \infty.$$

(b) Suppose that $\int_{\mathcal{X}_t} \{r_t^N \times f_t^N \times G_t \times H_{t:T}\}(x_t) \lambda_t(dx_t) \xrightarrow{\mathbb{P}} 0$. Then

$$\int \mathbb{Q}_T^N(dx_t) f_t^N(x_t) \xrightarrow{\mathbb{P}} 0.$$

Proof. Part (a). We shall prove by induction the statement

$$\mathbb{E} [\mathbb{Q}_s^N L_{s:T} \Pi_t^{0:T} f_t^N] = \infty, \forall t-1 \leq s \leq T.$$

For $s = t-1$, it follows from part (a)'s hypothesis and Lemma E.8. Indeed,

$$\begin{aligned} & \mathbb{Q}_{t-1}^N L_{t-1:T} \Pi_t^{0:T} f_t^N \\ &= \mathbb{Q}_{t-1}^N L_{t-1:t} L_{t:T} \Pi_t^{0:T} f_t^N = \mathbb{Q}_{t-1}^N L_{t-1:t} \Pi_t^{0:t} (f_t^N \times H_{t:T}) \\ &= \iint_{\mathcal{X}_{t-1} \times \mathcal{X}_t} \mathbb{Q}_{t-1}^N(dx_{t-1}) m_t(x_{t-1}, x_t) \lambda_t(dx_t) G_t(x_t) (f_t^N \times H_{t:T})(x_t) \\ &= \int_{\mathcal{X}_t} \{r_t^N \times f_t^N \times G_t \times H_{t:T}\}(x_t) \lambda_t(dx_t). \end{aligned}$$

For $s \geq t$, we have

$$\begin{aligned} \mathbb{E} [\mathbb{Q}_s^N L_{s:T} \Pi_t^{0:T} f_t^N] &= \mathbb{E} \left[\frac{N^{-1} \sum \tilde{K}_s^N(n, L_{s:T} \Pi_t^{0:T} f_t^N)}{\ell_s^N} \right] \text{ by Corollary 3} \\ &\geq \frac{1}{\|G_s\|_\infty} \mathbb{E} \left[N^{-1} \sum \tilde{K}_s^N(n, L_{s:T} \Pi_t^{0:T} f_t^N) \right] \\ &\text{by Assumption 1 and definition of } \ell_s^N \text{ (see Algorithm 1)} \\ &\geq \frac{1}{\|G_s\|_\infty} \mathbb{E} [\mathbb{Q}_{s-1}^N L_{s-1:s} L_{s:T} \Pi_t^{0:T} f_t^N] \\ &\text{by Corollary 3 and law of total expectation} \\ &= \mathbb{E} [\mathbb{Q}_{s-1}^N L_{s-1:T} \Pi_t^{0:T} f_t^N] = \infty \text{ (induction hypothesis)}. \end{aligned}$$

Part (b). Similar to part (a), we shall prove by induction the statement

$$\mathbb{Q}_s^N L_{s:T} \Pi_t^{0:T} f_t^N \xrightarrow{\mathbb{P}} 0, \forall t-1 \leq s \leq T.$$

Again, by Corollary 3, this quantity is equal to

$$\frac{N^{-1} \sum \tilde{K}_s^N(n, L_{s:T} \Pi_t^{0:T} f_t^N)}{\ell_s^N},$$

and the expectation of the numerator given \mathcal{F}_{s-1} is $\mathbb{Q}_{s-1}^N L_{s-1:T} \Pi_t^{0:T} f_t^N$, which tends to 0 in probability by induction hypothesis. Lemma E.11 (see below at the end of the section), the classical result $\ell_s^N \xrightarrow{\mathbb{P}} \ell_s := \mathbb{Q}_{s-1} M_s(G_s)$ and Stutsky's theorem concludes the proof. \square

Proof of Theorem B.1. By (21), we have $\mathbb{E}[\tau_t^{1, \text{FFBS}}] = \mathbb{E}[\int \mathbb{Q}_T^N(dx_t) f_t^N(x_t)]$ where

$$f_t^N(x_t) = \frac{\bar{M}_h}{\sum W_{t-1}^n m_t(X_{t-1}^n, x_t)} = \frac{\bar{M}_h}{r_t^N}$$

with r_t^N given in Definition 1. Proposition 9(a) gives a sufficient condition for $\mathbb{E}[\tau_t^{1, \text{FFBS}}] = \infty$ to hold, namely

$$\int_{\mathcal{X}_t} (r_t^N \times f_t^N \times G_t \times H_{t:T})(x_t) \lambda_t(dx_t) = \infty,$$

which is equivalent to the hypothesis of the theorem. \square

Proof of Theorem B.2. We use notations from Definition 1 and Supplement A.1. We note $\mathcal{N}(x|\mu, \Sigma)$ the density of the specified normal distribution at point x . Using Lemma E.9, Proposition 9 and (21), we have

$$\begin{aligned} \mathbb{E}[(\tau_t^{1, \text{FFBS}})] = \infty &\Leftrightarrow \mathbb{E} \left[\frac{1}{r_t^N (X_t^{T_t})^k} \right] = \infty \\ &\Leftrightarrow \mathbb{E} \left[\int \mathbb{Q}_T^N(dx_t) \frac{1}{r_t^N (x_t)^k} \right] = \infty \\ &\Leftrightarrow \mathbb{E} \left[\int_{\mathcal{X}_t} \frac{r_t^N G_t H_{t:T}}{(r_t^N)^k} (x_t) \lambda_t(dx_t) \right] = \infty \\ &\Leftrightarrow \int_{\mathcal{X}_t} \frac{r_t G_t H_{t:T}}{(r_t^N)^{k-1} r_t} (x_t) \lambda_t(dx_t) = \infty \text{ almost surely} \\ &\Leftrightarrow \int_{\mathcal{X}_t} \frac{\mathcal{N}(x_t | \mu_t^{\text{smth}}, \Sigma_t^{\text{smth}})}{r_t^N (x_t)^{k-1} \mathcal{N}(x_t | \mu_t^{\text{pred}}, \Sigma_t^{\text{pred}})} \lambda_t(dx_t) = \infty \text{ a.s.} \end{aligned}$$

The theorem then follows from elementary arguments, by noting that r_t^N is a mixture of N Gaussian distributions with covariance matrix C_X . \square

Lemma E.9. *Let L be a $]0, 1]$ -valued random variable. Suppose X is another random variable such that $X|L \sim \text{Geo}(L)$. Then for any real number $k > 0$,*

$$\mathbb{E}[X^k] = \infty \Leftrightarrow \mathbb{E}[L^{-k}] = \infty.$$

Proof. By the definition of X , we have

$$\mathbb{E}[X^k] = \mathbb{E} \left[\sum_{x=1}^{\infty} x^k (1-L)^{x-1} L \right].$$

A natural idea is then to approximate the sum by the integral $\int_0^{\infty} x^k (1-L)^{x-1} L dx$, from which one easily extracts the L^{-k} factor. This is however technically laborious, since the function $x \mapsto x^k (1-L)^{x-1} L$ is not monotone on the whole real line. It is only so starting from a certain x_0 which itself depends on L . We would therefore rather write

$$\begin{aligned} \mathbb{E}[X^k] &= \int_0^{\infty} \mathbb{P}(X^k \geq x) dx = \int_0^{\infty} \mathbb{P}(X \geq x^{1/k}) dx \\ &= \int_0^{\infty} \mathbb{E} \left[(1-L)^{\lfloor x^{1/k} \rfloor} \right] dx \end{aligned}$$

where the two integrands are equal Lebesgue-almost-everywhere

$$= \mathbb{E} \left[\int_0^{\infty} \exp \left(-|\log(1-L)| \lfloor x^{1/k} \rfloor \right) dx \right]$$

with the natural interpretation of expressions when $L = 1$. Using $u \sim v$ as a shorthand for “ u and v are either both finite or both infinite”, we have

$$\begin{aligned} \mathbb{E}[X^k] &\sim \mathbb{E} \left[\int_0^{\infty} \exp \left(-|\log(1-L)| x^{1/k} \right) dx \right] \text{ by Lemma E.10} \\ &= k \Gamma(k) \mathbb{E} \left[\frac{1}{|\log(1-L)|^k} \right] \sim \mathbb{E} \left[\frac{1}{L^k} \right] \text{ by Lemma E.10 again.} \end{aligned}$$

\square

The following lemma is elementary. Its proof is therefore omitted.

Lemma E.10. *Let L be a $]0, 1]$ -valued random variable and let f_1 and f_2 be two continuous functions from $]0, 1]$ to \mathbb{R} . Suppose that $\limsup_{\ell \rightarrow 0^+} f_1(\ell)/f_2(\ell)$ and $\limsup_{\ell \rightarrow 0^+} f_2(\ell)/f_1(\ell)$ are both finite. Then $\mathbb{E}[f_1(L)]$ is finite if and only if $\mathbb{E}[f_2(L)]$ is so.*

Lemma E.11. *Let Z_1, Z_2, \dots be non-negative random variables. Suppose that there exist σ -algebras $\mathcal{F}_1, \mathcal{F}_2, \dots$ such that $\mathbb{E}[Z_n | \mathcal{F}_n] \xrightarrow{\mathbb{P}} 0$. Then $Z_n \xrightarrow{\mathbb{P}} 0$.*

Proof. Fix $\varepsilon > 0$. By Markov's inequality, $\mathbb{P}(Z_n \geq \varepsilon | \mathcal{F}_n) \leq \varepsilon^{-1} \mathbb{E}[Z_n | \mathcal{F}_n]$. Therefore, the $[0, 1]$ -bounded random variable $\mathbb{P}(Z_n \geq \varepsilon | \mathcal{F}_n)$ tends to 0 in probability, hence also in expectation. The law of total expectation then gives $\mathbb{P}(Z_n \geq \varepsilon) \rightarrow 0$, which, by varying ε , establishes the convergence of Z_n to 0 in probability. \square

E.8. Proof of Theorem B.3 and Corollary 2 (hybrid FFBS complexity).

Proof of Theorem B.3. According to Janson (2011, Lemma 3), it is sufficient to show that

$$\frac{\sum_n \min(\tau_t^{n, \text{FFBS}}, N)}{N \alpha_N} \xrightarrow{\mathbb{P}} 0$$

for any deterministic sequence α_N such that $\alpha_N / \mathbb{E}[\min(\tau_t^{\infty, \text{FFBS}}, N)] \rightarrow \infty$. By Lemma E.11, we can take expectation with respect to \mathcal{F}_T to derive a sufficient condition, namely

$$\begin{aligned} & \int_{\mathcal{X}_t} \alpha_N^{-1} z^N \left(\frac{r_t^N(x_t)}{M_h} \right) \mathbb{Q}_T^N(dx_t) \xrightarrow{\mathbb{P}} 0 \text{ with } z^N \text{ defined in (65)} \\ \Leftrightarrow & \int_{\mathcal{X}_t} \alpha_N^{-1} z^N \left(\frac{r_t^N(x_t)}{M_h} \right) r_t^N \times G_t \times H_{t:T} d\lambda_t \xrightarrow{\mathbb{P}} 0 \text{ by Proposition 9(b)} \\ \Leftrightarrow & \mathbb{E} \left[\int_{\mathcal{X}_t} \alpha_N^{-1} z^N \left(\frac{r_t^N(x_t)}{M_h} \right) r_t^N \times G_t \times H_{t:T} d\lambda_t \right] \rightarrow 0 \\ \Leftrightarrow & \int_{\mathcal{X}_t} \alpha_N^{-1} z^N \left(\frac{r_t(x_t)}{M_h} \right) r_t \times G_t \times H_{t:T} d\lambda_t \rightarrow 0 \text{ by Lemma E.7} \\ \Leftrightarrow & \frac{\mathbb{E}[\min(\tau_t^{\infty, \text{FFBS}}, N)]}{\alpha_N} \rightarrow 0. \end{aligned}$$

The proof is now complete. \square

Proof of Corollary 2. We have, using the cost-to-go, the z^N functions and the $\tau_t^{\infty, \text{PaRIS}}$ distribution defined respectively in (20), (65) and (59):

$$\begin{aligned} \mathbb{E}[\min(\tau_t^{\infty, \text{FFBS}}, N)] &= \int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{M_h} \right) \mathbb{Q}_T(dx_t) \\ &= [(\mathbb{Q}_{t-1} M_t)(G_t H_{t:T})]^{-1} \int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{M_h} \right) (G_t H_{t:T})(x_t) (\mathbb{Q}_{t-1} M_t)(dx_t) \\ &\leq \|G_t H_{t:T}\|_{\infty} [(\mathbb{Q}_{t-1} M_t)(G_t H_{t:T})]^{-1} \int_{\mathcal{X}_t} z^N \left(\frac{r_t(x_t)}{M_h} \right) (\mathbb{Q}_{t-1} M_t)(dx_t) \\ &= \|G_t H_{t:T}\|_{\infty} [(\mathbb{Q}_{t-1} M_t)(G_t H_{t:T})]^{-1} \mathbb{E}[\min(\tau_t^{\infty, \text{PaRIS}}, N)]. \end{aligned}$$

Proposition 8 then finishes the proof. \square

E.9. Proof of Proposition 4 (MCMC kernel properties).

Proof. Let J_t^n be such that $J_t^n | X_{t-1}^{1:N}, X_t^n, \hat{B}_t^{N, \text{IMH}}(n, \cdot) \sim B_t^{N, \text{IMH}}(n, \cdot)$. Moreover, let K_t^n be such that

$$K_t^n | X_{t-1}^{1:N}, X_t^n, A_t^n, \hat{B}_t^{N, \text{IMH}}(n, \cdot) \sim B_t^{N, \text{IMH}}(n, \cdot).$$

Given $X_{t-1}^{1:N}$, X_t^n and A_t^n , the kernel $B_t^{N, \text{IMH}}(n, \cdot)$ applies to A_t^n one or more several MCMC steps keeping invariant $B_t^{N, \text{FFBS}}(n, \cdot)$. Since $A_t^n | X_{t-1}^{1:N}, X_t^n \sim B_t^{N, \text{FFBS}}(n, \cdot)$, it follows that $K_t^n | X_{t-1}^{1:N}, X_t^n \sim B_t^{N, \text{FFBS}}(n, \cdot)$ too. On the other hand, J_t^n and K_t^n share the same distribution given $X_{t-1}^{1:N}$, X_t^n and $\hat{B}_t^{N, \text{IMH}}(n, \cdot)$. Hence they also do, given $X_{t-1}^{1:N}$ and X_t^n only. This implies that $J_t^n | X_{t-1}^{1:N}, X_t^n \sim B_t^{N, \text{FFBS}}(n, \cdot)$, which is the same as $A_t^n | X_{t-1}^{1:N}, X_t^n$. Thus (J_t^n, X_t^n) have the same distribution as (A_t^n, X_t^n) given $X_{t-1}^{1:N}$, as required by Theorem 2.1. The arguments for the kernel $B_t^{N, \text{IMHP}}$ are similar.

To show that a certain kernel B_t^N satisfies (13), we look at two conditionally i.i.d. simulations $J_t^{n,1}$ and $J_t^{n,2}$ from $B_t^N(n, \cdot)$ and lower bound the probability that they are different. For the kernel $B_t^{N, \text{IMH}}$, the variables $J_t^{n,1}$ and $J_t^{n,2}$ both result from one step of MH applied to A_t^n . Let $J_t^{n,1*}$ and $J_t^{n,2*}$ be the corresponding MH proposals. A sufficient condition for $J_t^{n,1} \neq J_t^{n,2}$ is that $J_t^{n,1*} \neq J_t^{n,2*}$ and the two proposals are both accepted. The acceptance rate is at least \bar{M}_ℓ/\bar{M}_h by Assumption 2 and the probability that $J_t^{n,1*} \neq J_t^{n,2*}$ is

$$1 - \sum_{n=1}^N (W_{t-1}^n)^2 \geq 1 - \frac{1}{N} \left(\frac{\bar{G}_h}{\bar{G}_\ell} \right)^2$$

by Assumption 3. Thus (13) is satisfied for $\varepsilon_S = \bar{M}_\ell/2\bar{M}_h$ for N large enough. Similarly, the probability that $J_t^{n,1} \neq J_t^{n,2}$ for the $B_t^{N, \text{IMHP}}$ kernel with $\tilde{N} = 2$ can be lower-bounded via the probability that $\tilde{J}_t^{n,1} \neq \tilde{J}_t^{n,2}$ (where $\tilde{J}_t^{n,1}$ and $\tilde{J}_t^{n,2}$ are defined in (18)). Thus using the same arguments, (13) is satisfied here for $\varepsilon_S = \bar{M}_\ell/4\bar{M}_h$. \square

E.10. Conditional probability of maximal couplings. In general, there exist multiple maximal couplings of two random distributions (i.e. couplings that maximise the probability of equality of the two variables). However, they all satisfy a certain conditional probability property stated in the following lemma. It is closely related to results on the coupling density on the diagonal (see e.g. Wang et al., 2021, Lemma 2 or Douc et al., 2018, Theorem 19.1.6). Its statement, which we were unable to find in the literature in the exact form we need, is obvious in the discrete case but requires lengthier arguments in the continuous one.

Proposition 10. *Let X_1 and X_2 be two random variables with densities f_1 and f_2 with respect to some dominating measure defined on a space \mathcal{X} . Then, the following inequality holds almost surely:*

$$(69) \quad \mathbb{P}(X_2 = X_1 | X_1) \leq 1 \wedge \frac{f_2(X_1)}{f_1(X_1)}.$$

Moreover, the equality occurs almost surely if and only if X_1 and X_2 form a maximal coupling.

Proof. Let h be any non-negative test function from \mathcal{X} to \mathbb{R} . Putting

$$\begin{aligned} A_1 &:= \{x \in \mathcal{X} \mid f_1(x) \geq f_2(x)\} \\ A_2 &:= \{x \in \mathcal{X} \mid f_2(x) \geq f_1(x)\}, \end{aligned}$$

we have

$$\begin{aligned} \mathbb{E}[\mathbb{P}(X_2 = X_1 | X_1)h(X_1)] &= \mathbb{E}[\mathbb{1}_{X_2=X_1}h(X_1)] \\ &= \mathbb{E}[\mathbb{1}_{X_2=X_1} \mathbb{1}_{X_1 \in A_1}h(X_1)] + \mathbb{E}[\mathbb{1}_{X_2=X_1} \mathbb{1}_{X_1 \in A_2}h(X_1)] \\ &= \mathbb{E}[\mathbb{1}_{X_2=X_1} \mathbb{1}_{X_2 \in A_1}h(X_2)] + \mathbb{E}[\mathbb{1}_{X_2=X_1} \mathbb{1}_{X_1 \in A_2}h(X_1)] \\ &\leq \mathbb{E}[\mathbb{1}_{X_2 \in A_1}h(X_2)] + \mathbb{E}[\mathbb{1}_{X_1 \in A_2}h(X_1)] \\ &= \int h(x)f_2 \wedge f_1(x)dx = \mathbb{E} \left[\left(1 \wedge \frac{f_2(X_1)}{f_1(X_1)} \right) h(X_1) \right]. \end{aligned}$$

The inequality (69) is now proved almost-surely. As a result, almost-sure equality occurs if and only if the expectation of the two sides of (69) are equal, which means, via Lemma A.1, that the two variables are maximally coupled. \square

The following lemma establishes the symmetry of Assumption 8. Again, its statement is obvious in the discrete case, though some work is needed to rigorously justify the continuous one.

Lemma E.12. *Let X_1 and X_2 be two random variables of densities f_1 and f_2 w.r.t. some dominating measure defined on some space \mathcal{X} . Suppose that almost-surely*

$$\mathbb{P}(X_2 = X_1 | X_1) \geq \varepsilon \left(1 \wedge \frac{f_2(X_1)}{f_1(X_1)} \right)$$

for some $\varepsilon > 0$. Then almost-surely,

$$\mathbb{P}(X_1 = X_2 | X_2) \geq \varepsilon \left(1 \wedge \frac{f_1(X_2)}{f_2(X_2)} \right).$$

Proof. We introduce a non-negative test function $h_2 : \mathcal{X} \rightarrow \mathbb{R}$ and write

$$\begin{aligned} \mathbb{E}[\mathbb{P}(X_1 = X_2 | X_2)h(X_2)] &= \mathbb{E}[\mathbb{1}_{X_1=X_2}h(X_2)] = \mathbb{E}[\mathbb{1}_{X_2=X_1}h(X_1)] \\ &= \mathbb{E}[\mathbb{P}(X_2 = X_1 | X_1)h(X_1)] \\ &\geq \mathbb{E}\left[\varepsilon \left(1 \wedge \frac{f_2(X_1)}{f_1(X_1)}\right) h(X_1)\right] \\ &= \int \varepsilon f_1 \wedge f_2(x)h(x)dx \\ &= \mathbb{E}\left[\varepsilon \left(1 \wedge \frac{f_1(X_2)}{f_1(X_2)}\right) h(X_2)\right] \end{aligned}$$

which implies the desired result. \square

E.11. Proof of Theorem 4.1 (intractable kernel properties).

Proof. Let J_t^n be a random variable such that

$$J_t^n | X_{t-1}^{1:N}, X_t^n, \hat{B}_t^{N,\text{ITR}}(n, \cdot) \sim B_t^{N,\text{ITR}}(n, \cdot).$$

By construction of Algorithm 7, given $X_{t-1}^{1:N}$, the couple (J_t^n, X_t^n) has the same distribution as $(A_t^{n,L}, X_t^{n,L})$. Thus, $B_t^{N,\text{ITR}}$ satisfies the hypotheses of Theorem 2.1. To verify (13), we define the variables $J_t^{n,1:2}$ accordingly and write:

$$\begin{aligned} &\mathbb{P}\left(J_t^{n,1} \neq J_t^{n,2} \mid X_t^n = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &= \frac{1}{2} \mathbb{P}\left(X_t^{n,1} = X_t^{n,2}, A_t^{n,1} \neq A_t^{n,2} \mid X_t^n = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &= \mathbb{P}\left(X_t^{n,1} = X_t^{n,2}, A_t^{n,1} \neq A_t^{n,2}, L = 1 \mid X_t^n = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &= \frac{1}{2} \mathbb{P}\left(X_t^{n,1} = X_t^{n,2}, A_t^{n,1} \neq A_t^{n,2} \mid L = 1, X_t^n = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \text{ (by symmetry)} \\ &= \frac{1}{2} \mathbb{P}\left(X_t^{n,1} = X_t^{n,2}, A_t^{n,1} \neq A_t^{n,2} \mid X_t^{n,1} = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &= \frac{1}{2} \sum_{a_t^{n,1} \neq a_t^{n,2}} \mathbb{P}\left(X_t^{n,2} = x_t \mid A_t^{n,1} = a_t^{n,1}, A_t^{n,2} = a_t^{n,2}, X_t^{n,1} = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \times \\ &\quad \times \mathbb{P}\left(A_t^{n,1} = a_t^{n,1}, A_t^{n,2} = a_t^{n,2} \mid X_t^{n,1} = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &\geq \frac{1}{2} \varepsilon_D \frac{\bar{M}_\ell}{M_h} \sum_{a_t^{n,1} \neq a_t^{n,2}} \mathbb{P}\left(A_t^{n,1} = a_t^{n,1}, A_t^{n,2} = a_t^{n,2} \mid X_t^{n,1} = x_t, X_{t-1}^{1:N} = x_{t-1}^{1:N}\right) \\ &\quad \text{(by Assumptions 8 and 2)} \\ &\geq \frac{1}{2} \varepsilon_D \left(\frac{\bar{M}_\ell}{M_h}\right)^2 \varepsilon_A \text{ by Lemma E.13.} \end{aligned}$$

The proof is complete. \square

Lemma E.13. *We use the notations of Algorithm 7. Under Assumptions 2 and 7, we have*

$$\mathbb{P}\left(A_t^{n,1} \neq A_t^{n,2} \mid X_t^{n,1}, X_{t-1}^{1:N}\right) \geq \frac{\bar{M}_\ell}{M_h} \varepsilon_A.$$

Proof. We write (and define new notations along the way):

$$\begin{aligned} \pi(a_t^{n,1}, a_t^{n,2}) &:= p(a_t^{n,1}, a_t^{n,2} | X_t^{n,1}, X_{t-1}^{1:N}) \\ &\propto p(a_t^{n,1}, a_t^{n,2} | X_{t-1}^{1:N}) m_t(X_{t-1}^{a_t^{n,1}}, X_t^{n,1}) \\ &=: p(a_t^{n,1}, a_t^{n,2} | X_{t-1}^{1:N}) \phi(a_t^{n,1}) \\ &=: \pi_0(a_t^{n,1}, a_t^{n,2}) \phi(a_t^{n,1}). \end{aligned}$$

Thus

$$\begin{aligned} \mathbb{P}\left(A_t^{n,1} \neq A_t^{n,2} \mid X_t^{n,1}, X_{t-1}^{1:N}\right) &= \int \mathbb{1}\{a_t^{n,1} \neq a_t^{n,2}\} \pi(a_t^{n,1}, a_t^{n,2}) \\ &= \frac{\int \mathbb{1}\{a_t^{n,1} \neq a_t^{n,2}\} \pi_0(a_t^{n,1}, a_t^{n,2}) \phi(a_t^{n,1})}{\int \pi_0(a_t^{n,1}, a_t^{n,2}) \phi(a_t^{n,1})} \\ &\geq \int \mathbb{1}\{a_t^{n,1} \neq a_t^{n,2}\} \pi_0(a_t^{n,1}, a_t^{n,2}) \frac{\bar{M}_\ell}{\bar{M}_h} \end{aligned}$$

by the boundedness of the function ϕ between \bar{M}_ℓ and \bar{M}_h . From this, we get the desired result by virtue of Assumption 7. \square

E.12. Validity of Algorithm 14 (modified Lindvall-Rogers coupler). Recall that generating a random variable is equivalent to uniformly simulating under the graph of its density (see e.g. [Robert and Casella, 2004](#), The Fundamental Theorem of Simulation, chapter 2.3.1). Algorithm 14's correctness is thus a direct corollary of the following intuitive lemma.

Lemma E.14. *Let S_A and S_B be two subsets of \mathbb{R}^d with finite Lebesgue measures. Let A and B be two not necessarily independent random variables distributed according to $\text{Uniform}(S_A)$ and $\text{Uniform}(S_B)$ respectively. Denote by S_0 the intersection of S_A and S_B ; and by C a certain $\text{Uniform}(S_A)$ -distributed random variable that is independent from (A, B) . Define A^* and B^* as*

$$A^* = \begin{cases} C & \text{if } (A, C) \in S_0 \times S_0 \\ A & \text{otherwise} \end{cases}$$

and

$$B^* = \begin{cases} C & \text{if } (B, C) \in S_0 \times S_0 \\ B & \text{otherwise} \end{cases}$$

Then $A^* \sim \text{Uniform}(S_A)$ and $B^* \sim \text{Uniform}(S_B)$.

Proof. Given $(A, C) \in S_0 \times S_0$, the two variables A and C have the same distribution (which is $\text{Uniform}(S_0)$). Thus, the definition of A^* implies that A and A^* have the same (unconditional) distribution. The same argument applies to B and B^* notwithstanding the asymmetry in the definition of C . \square

E.13. Hoeffding inequalities. This section proves a Hoeffding inequality for ratios, which helps us to bound (28). It is essentially a reformulation of [Douc et al. \(2011, Lemma 4\)](#) in a slightly more general manner.

Definition 2. A real-valued random variable X is called (C, S) -sub-Gaussian if

$$\mathbb{P}\left(\frac{|X|}{S} > t\right) \leq 2Ce^{-t^2/2}, \forall t \geq 0.$$

This definition is close to other sub-Gaussian definitions in the literature, see e.g. [Vershynin \(2018, Chapter 2.5\)](#). It basically means that the tails of X decreases at least as fast as the tails of the $\mathcal{N}(0, S^2)$ distribution, which is itself $(1, S)$ -sub-Gaussian. The following result is classic.

Theorem E.15 (Hoeffding's inequality). *Let X_1, \dots, X_N be N i.i.d. random variables with mean μ and almost surely contained between a and b . Then $N^{1/2}(\sum X_i/N - \mu)$ is $(1, (b-a)/2)$ -sub-Gaussian.*

The following lemma is elementary from Definition 2. The proof is omitted.

Lemma E.16. *Let X and Y be two (not necessarily independent) random variables. If X is (C_1, S_1) -sub-Gaussian and Y is (C_2, S_2) -sub-Gaussian, then $X + Y$ is $(C_1 + C_2, S_1 + S_2)$ -sub-Gaussian.*

We are ready to state the main result of this section.

Proposition 11 (Hoeffding's inequality for ratios). *Let a_N, b_N, a^*, b^* be random variables such that $\sqrt{N}(a_N - a^*)$ is (C_a, S_a) -sub-Gaussian and $\sqrt{N}(b_N - b^*)$ is (C_b, S_b) -sub-Gaussian. Then $\sqrt{N}(a_N/b_N - a^*/b^*)$ is sub-Gaussian with parameters (C^*, S^*) where*

$$\begin{cases} C^* &= C_a + C_b \\ S^* &= \left\| \frac{1}{b^*} \right\|_\infty (S_a + S_b \left\| \frac{a_N}{b_N} \right\|_\infty). \end{cases}$$

The terms with inf-norm can be infinite if the corresponding random variables are unbounded.

Proof. We have

$$\begin{aligned} \left| \sqrt{N} \left(\frac{a_N}{b_N} - \frac{a^*}{b^*} \right) \right| &\leq \left| \sqrt{N} \left(\frac{a_N}{b_N} - \frac{a_N}{b^*} \right) \right| + \left| \sqrt{N} \left(\frac{a_N}{b^*} - \frac{a^*}{b^*} \right) \right| \\ &= \left| \frac{a_N}{b_N} \right| \left| \frac{1}{b^*} \right| \left| \sqrt{N}(b_N - b^*) \right| + \left| \frac{1}{b^*} \right| \left| \sqrt{N}(a_N - a^*) \right| \end{aligned}$$

by which the proposition follows from Lemma [E.16](#). □

CREST-ENSAE, INSTITUT POLYTECHNIQUE DE PARIS

Email address: `nicolas.chopin@ensae.fr`