VU Research Portal

A survey of sequential Monte Carlo methods for economics and finance

Creal, D.D.

2009

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Creal, D. D. (2009). *A survey of sequential Monte Carlo methods for economics and finance*. (Research Memorandum; No. 2009-18). Faculteit der Economische Wetenschappen en Bedrijfskunde.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

A survey of sequential Monte Carlo methods for economics and finance

DREW CREAL*

Department of Econometrics, Vrije Universiteit Amsterdam

February 5, 2009

Abstract

This paper serves as an introduction and survey for economists to the field of sequential Monte Carlo methods which are also known as particle filters. Sequential Monte Carlo methods are simulation based algorithms used to compute the high-dimensional and/or complex integrals that arise regularly in applied work. These methods are becoming increasingly popular in economics and finance; from dynamic stochastic general equilibrium models in macro-economics to option pricing. The objective of this paper is to explain the basics of the methodology, provide references to the literature, and cover some of the theoretical results that justify the methods in practice.

Keywords: state space models; sequential Monte Carlo; particle filter; Markov chain Monte Carlo; Kalman filter

JEL Classification: C11; C15; C32.

*Correspondence: Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, email: dcreal@feweb.vu.nl, tel: +31 20 598 6011, fax: +31 20 598 6020.

1 Introduction

Economic theory often prescribes fundamental nonlinear relationships between variables of interest. Nonlinear models for learning and strategic interaction among agents provide the modern foundation for microeconomic models. Building on these microfoundations, macroeconomists formulate their structural models as dynamic stochastic general equilibrium (DSGE) models, which have nonlinear first order conditions. Many important economic time series also exhibit strong patterns of non-Gaussian or time-varying behavior. Regime switching, stochastic volatility, and time-varying parameter models have become increasingly popular over the last decade.

Complex models often lead to integrals that cannot be solved analytically. This has created an increase in the popularity of Bayesian methods that utilize Markov chain Monte Carlo (MCMC) algorithms. Sequential Monte Carlo (SMC) methods are alternative simulation-based algorithms for solving analytically intractable integrals. In these methods, a (partially) continuous probability distribution is approximated by a discrete distribution made of weighted draws termed particles. From one iteration of the algorithm to the next, particles are updated to approximate one distribution after another by changing the particle's location on the support of the distribution and their weights. SMC methods include the particle filter, which generalizes the Kalman filter and hidden Markov model (HMM) filter to nonlinear, non-Gaussian state space models. Particle filters were introduced into the economics literature by Kim et al. (1998) to study the volatility of asset prices. Their popularity has grown in economics since the publication of Fernández-Villaverde and Rubio-Ramírez (2005, 2007), who used them to estimate DSGE models. Particle filters also share a common mathematical structure with genetic algorithms which are popular in economics.

The standard reference for SMC methods is Doucet et al. (2001). A considerable number of advances have taken place since its publication; advances ranging from stimulating new applications, improved algorithms, and new theoretical results. Most of the methodological results have occurred outside economics, where a nice review for engineers is provided by Cappé et al. (2007). This paper provides a guide to the growing literature intended for economists and includes updated references. The presentation given here extends previous reviews by including a discussion of SMC methods applied outside state space models. The methods are also applied to several economic applications. To reach as wide an audience as possible, the survey has been split into two parts. The first half focuses on practical applications of particle filters to general state space models. The second half covers recent developments in the field with more emphasis on Bayesian computation as well as an overview of the theoretical properties of SMC methods.

The theoretical properties of SMC algorithms have been intensely studied since Del Moral

(1996), who provided the first consistency proof for the original particle filter of Gordon et al. (1993). In SMC algorithms, the draws or particles interact and are therefore dependent. Traditional limit theorems for Monte Carlo methods, e.g. Geweke (1989) and Tierney (1994), do not apply. The main theoretical properties that are relevant for applied researchers are reviewed in the paper while additional references are provided for those interested in further study. To make the paper shorter, readers are assumed to be modestly familiar with linear, Gaussian state space models, importance sampling, accept-reject algorithms, and MCMC. Harvey (1989) and Durbin and Koopman (2001) provide introductions to the first while Robert and Casella (2004) and Geweke (2005) are good references for the latter three.

In Section 2, SMC methods are introduced starting with the particle filter and its application to nonlinear, non-Gaussian state space models. This section contains a minimum of technical details and concentrates on best practices that a researcher should consider when implementing them. Some of the theoretical properties of SMC algorithms are reviewed in section 3. In Section 4, more advanced SMC algorithms are discussed which are applicable outside the context of state space models. Both Sections 2 and 4 include economic applications to illustrate the relevance of the methods. The final section concludes.

2 Particle filters for state space and hidden Markov models

State space or hidden Markov models are a convenient means for studying dynamic systems. A state space model consists of two equations: the observation or measurement equation and the transition equation which are respectively given by

$$\mathbf{y}_n = g_n(\mathbf{x}_n, \varepsilon_n), \quad (1)$$

$$\mathbf{x}_n = h_n(\mathbf{x}_{n-1}, \eta_n). \quad (2)$$

The state variables $\mathbf{x}_n \in \mathbf{R}^m$ and observations $\mathbf{y}_n \in \mathbf{R}^d$ may be continuous-valued, discrete-valued, or a combination of the two. The functions h_n and g_n are possibly nonlinear but of known form. Time is denoted by the subscripts n . It is assumed that the distributions of the observations and state variable admit density functions with respect to an appropriate dominating measure. The dominating measure for the state variable will be denoted by $d\mathbf{x}_n$. These densities $p(\mathbf{y}_n | \mathbf{x}_n; \theta)$ and $p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)$, corresponding to (1) and (2) respectively, are traditionally called the observation and transition densities. The latter terminology stems from the fact that $\mathbf{x}_n$ is a Markov process. The densities will typically depend upon a vector of unknown parameters θ that need to be estimated from the observed data $\mathbf{y}_{1:T} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$.

The sequence of state variables $\mathbf{x}_{0:n} = \{\mathbf{x}_0, \dots, \mathbf{x}_n\}$ are generally unobserved and it is the

aim of the researcher to estimate them using the observed data. Uncertainty about the state variable is formulated as a joint conditional probability distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ known as the joint smoothing distribution. Three of its marginal distributions are of interest: the one-step ahead predictive distribution $p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta)$, the filtering distribution $p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta)$, and the smoothing distribution $p(\mathbf{x}_n | \mathbf{y}_{1:T}; \theta)$. Each distribution conditions on a different set of observations. Although θ is unknown, it is traditional in the literature to run filtering and smoothing algorithms assuming a fixed value of θ . Therefore, in the discussion of prediction, filtering, and smoothing algorithms in Sections 2.1-2.5, the value of θ is assumed to be known. The estimation of θ is considered in later sections.

2.1 Prediction and filtering recursions

It is computationally convenient to calculate the one-step ahead predictive and filtering distributions recursively in time. Computing them recursively originates from the fields of signal processing and engineering where these methods are applied in real-time or online. The recursion begins under the assumption that the initial distribution of the state variable $p(\mathbf{x}_0; \theta)$ is known. At a future iteration n , the prediction step projects last period's filtering distribution $p(\mathbf{x}_{n-1} | \mathbf{y}_{1:n-1}; \theta)$ forward using the dynamics of the model (2) and its transition density

$$p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) = \int p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta) p(\mathbf{x}_{n-1} | \mathbf{y}_{1:n-1}; \theta) d\mathbf{x}_{n-1}. \quad (3)$$

This distribution is a one-step ahead forecast of the state variable. With the addition of another observation $\mathbf{y}_n$, the update step computes the filtering distribution by applying Bayes' rule

$$\begin{aligned} p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta) &= \frac{p(\mathbf{y}_n, \mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta)}{p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta)} \\ &= \frac{p(\mathbf{y}_n | \mathbf{x}_n, \mathbf{y}_{1:n-1}; \theta) p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta)}{\int p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) d\mathbf{x}_n} \\ &= \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta)}{\int p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) d\mathbf{x}_n}. \end{aligned} \quad (4)$$

The last step following from the Markovian assumptions of the model. This completes one iteration of the recursion which continues until the end of the dataset.

Difficulty arises in this approach because for most state space models the normalizing constant $p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta) = \int p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) d\mathbf{x}_n$ in (4) cannot be calculated analytically. There are several known cases in which it is possible to solve (4) analytically. The two most useful cases are when (1) and (2) are linear models with Gaussian densities and the recursions can be solved by the Kalman filter, see (Kalman (1960) and Kalman and Bucy (1961)). The other is the HMM filter (Baum and Petrie (1966) and Baum et al. (1970)) when the state

variable $\mathbf{x}_n$ is discrete valued. The latter algorithm being rediscovered and extended to autoregressions by Hamilton (1989) in his influential model for the business cycle. For textbook treatments of these methods see Harvey (1989) and Durbin and Koopman (2001) for the linear, Gaussian state space model while Frühwirth-Schnatter (2006) covers models with discrete state variables. Otherwise, the distributions must be approximated and the particle filter does this by approximating the analytically intractable integrals using Monte Carlo simulation.

2.2 Particle filters

A particle filter recursively approximates at iteration n the entire joint distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ and, as a by-product, it also approximates the marginal distribution $p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta)$. It accomplishes this by simulation and uses an extended version of sequential importance sampling (SIS). Instead of drawing entire sequences $\mathbf{x}_{0:n}$ directly from $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ to form a standard *i.i.d.* Monte Carlo estimate, it employs importance sampling where draws are taken from an importance distribution $q_{0:n}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \psi)$ that approximates the target distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ and whose support includes that of the target. The draws are then reweighted to correct for the fact that they were drawn from the wrong distribution. The parameter vector ψ within the importance distribution denotes a vector of tuning parameters; their choice will be discussed below. The importance weights are defined as the ratio of the target distribution divided by a suitably chosen importance distribution

$$w_n = \frac{p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)}{q_{0:n}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \psi)}. \quad (5)$$

The time subscript on $q_{0:n}(\cdot)$ indicates that the importance distribution or its parameters ψ can potentially be chosen at time n and can change over time.

Computing the entire expression (5) at each time n can be computationally intensive. To avoid evaluating it each period, the importance distribution within a particle filter is factored into two parts

$$q_{0:n}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \psi) \equiv q_n(\mathbf{x}_n | \mathbf{x}_{0:n-1}, \mathbf{y}_{1:n}; \psi) q_{0:n-1}(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \psi). \quad (6)$$

The second distribution $q_{0:n-1}(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \psi)$ is, per particle, a Dirac measure $\delta_{\mathbf{x}_{0:n-1}}$ placing probability one on each path that has already been simulated in the previous iterations up to time $n - 1$. A new set of values $\{\mathbf{x}_n^{(i)}\}_{i=1}^N$ are drawn at time n from the first part of the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{0:n-1}, \mathbf{y}_{1:n}; \psi)$. Consequently, a new sequence of paths is obtained by keeping the trajectories of the old particles fixed and appending the newly simulated values to the end of the old trajectories, $\{\mathbf{x}_{0:n}^{(i)}\}_{i=1}^N = \{\mathbf{x}_{0:n-1}^{(i)}, \mathbf{x}_n^{(i)}\}_{i=1}^N$.

Factoring the joint smoothing distribution in the numerator of (5) as

$$\begin{aligned}
p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta) &= \frac{p(\mathbf{y}_n | \mathbf{x}_{0:n}, \mathbf{y}_{1:n-1}; \theta) p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n-1}; \theta)}{p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta)} \\
&= \frac{p(\mathbf{y}_n | \mathbf{x}_{0:n}, \mathbf{y}_{1:n-1}; \theta) p(\mathbf{x}_n | \mathbf{x}_{0:n-1}, \mathbf{y}_{1:n-1}; \theta) p(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \theta)}{p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta)} \\
&= \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta) p(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \theta)}{p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta)} \tag{7}
\end{aligned}$$

it is possible to compute only one component of the importance weight (5) at each iteration by plugging (6) and (7) into (5) to obtain

$$w_n = \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta) p(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \theta)}{p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta) q_n(\mathbf{x}_n | \mathbf{x}_{0:n-1}, \mathbf{y}_{1:n}; \psi) q_{0:n-1}(\mathbf{x}_{0:n-1} | \mathbf{y}_{1:n-1}; \psi)} \tag{8}$$

$$\propto w_{n-1} \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)}{q_n(\mathbf{x}_n | \mathbf{x}_{0:n-1}, \mathbf{y}_{1:n}; \psi)}. \tag{9}$$

The densities $p(\mathbf{y}_n | \mathbf{x}_n; \theta)$ and $p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)$ are determined by the state space model (1)-(2). The ratio of densities on the right hand side of (9) is referred to as the *incremental importance weight*. The conditioning information in the importance distribution in the denominator of (9) will typically be reduced to $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ for computational convenience. Calculating the incremental weights then does not require the past observations or the entire past trajectories $\{\mathbf{x}_{0:n-2}^{(i)}\}_{i=1}^N$. At the end of each iteration, the algorithm produces N simulated paths and importance weights $\{\mathbf{x}_{0:n}^{(i)}, w_n^{(i)}\}_{i=1}^N$. These provide a discrete distribution that approximates the continuous distribution.

Given the draws $\{\mathbf{x}_{0:n}^{(i)}, w_n^{(i)}\}_{i=1}^N$, it is possible to approximate expectations of a function $f(\mathbf{x}_{0:n})$ of the state variable

$$\mathbb{E}_q[f(\mathbf{x}_{0:n})] = \int f(\mathbf{x}_{0:n}) \frac{p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)}{q_{0:n}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \psi)} q_{0:n}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \psi) d\mathbf{x}_{0:n}. \tag{10}$$

Expectations are taken with respect to the importance distribution instead of the target distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ as in standard importance sampling. The particle filter's estimate of (10) is given by first self-normalizing the importance weights

$$\hat{w}_n = \frac{w_n}{\sum_{i=1}^N w_n^{(i)}}, \tag{11}$$

and taking a weighted average

$$\sum_{i=1}^N f(\mathbf{x}_{0:n}^{(i)}) \hat{w}_n^{(i)} \approx \mathbb{E}_q[f(\mathbf{x}_{0:n})]. \tag{12}$$

The importance weights are self-normalized to increase the stability of the estimator. The empirical distribution function determined by the particles is given by

$$\hat{p}(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta) = \sum_{i=1}^N \hat{w}_n^{(i)} \delta_{\mathbf{x}_{0:n}^{(i)}}(\mathbf{x}_{0:n}^{(i)}) \approx p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta). \tag{13}$$

Figure 1: 30 iterations of a discrete-time log-normal stochastic volatility model. Pictured is the true log-volatility (solid line) and a particle system with $N=12$ particles after: (i) 6 time-steps; (ii) 12 time-steps; (iii) 22 time-steps; and (iv) 30 time-steps.

where $\delta_{\mathbf{x}_{0:n}}$ is the Dirac measure located at $\mathbf{x}_{0:n}$.

To provide some intuition about this distribution, data were simulated from a standard stochastic volatility model

$$\begin{aligned} \mathbf{y}_n &= \exp(\mathbf{x}_n/2)\varepsilon_n, & \varepsilon_n &\sim \mathcal{N}(0, 1), \\ \mathbf{x}_n &= \mu + \phi(\mathbf{x}_{n-1} - \mu) + \sigma_\eta \eta_n, & \eta_n &\sim \mathcal{N}(0, 1). \end{aligned}$$

Figure 1 plots the true value of the state variable $\mathbf{x}_n$ from this model and $N = 12$ particles over the first 30 time periods. The panels show the evolution of a particle system $\left\{ \mathbf{x}_{0:n}^{(i)}, w_n^{(i)} \right\}_{i=1}^{N=12}$ and four of its empirical distributions $\hat{p}(\mathbf{x}_{0:6} | \mathbf{y}_{1:6}; \theta)$, $\hat{p}(\mathbf{x}_{0:12} | \mathbf{y}_{1:12}; \theta)$, $\hat{p}(\mathbf{x}_{0:22} | \mathbf{y}_{1:22}; \theta)$, and $\hat{p}(\mathbf{x}_{0:30} | \mathbf{y}_{1:30}; \theta)$. The graphs indicate how the particle filter approximates a continuous distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ with a discrete distribution. The distributions will be approximated by a much larger number of particles in practice. Another perhaps more accurate perspective is to look at the empirical distribution functions determined by (13). Panel (i) of Figure 2 depicts the particles' approximation of the marginal filtering distribution at time $n = 30$ given by $p(\mathbf{x}_{30} | \mathbf{y}_{1:30}; \theta)$ for $N = 12$ particles. The remaining panels in Figure 2 demonstrate what

Figure 2: Empirical distribution functions created using the particles to approximate the marginal filtering distribution $p(\mathbf{x}_{30}|\mathbf{y}_{1:30};\theta)$ for the stochastic volatility model. (i) $N = 12$ particles; (ii) $N = 100$ particles; (iii) $N = 500$ particles; and (iv) $N = 10000$ particles.

happens as the number of particles increases. In this example where the state variable $\mathbf{x}_n$ is continuous, the particles form a probability mass function that is converging toward a continuous distribution function.

The practical limitation of the method outlined above is that as the number of iterations increases all the probability mass will eventually be allocated to one particle. There will exist one particle whose importance weight equals one while the other particles do not contribute to the estimator (12). This is known as the weight degeneracy problem. In their seminal paper introducing the particle filter, Gordon et al. (1993) added a resampling stage within the SIS algorithm to alleviate this problem. After computing the importance weights, the particles are resampled. This means that particles are replicated in proportion to their normalized importance weight, i.e. draw N random variables with replacement from a multinomial distribution with probabilities $\{\hat{w}^{(i)}\}_{i=1}^N$. Particles with large importance weights are randomly duplicated while particles with small probability are eliminated. Once resampled the particles' weights are set equal to any constant, e.g. $w_n^{(i)} = \frac{1}{N}$ for $i = 1, \dots, N$. This new algorithm called sequential importance sampling with resampling (SISR) extended the sampling importance resampling

Algorithm 1 Sequential Importance Sampling with Resampling (SISR)

At $n = 0$, for $i = 1, \dots, N$

Draw $\mathbf{x}_0^{(i)} \sim q_0(\mathbf{x}_0)$ and set $w_0^{(i)} = \frac{p(\mathbf{x}_0^{(i)})}{q_0(\mathbf{x}_0^{(i)})}$.

For $n = 1, \dots, T$:

(i) For $i = 1, \dots, N$:

a. Draw: $\mathbf{x}_n^{(i)} \sim q_n(\mathbf{x}_n | \mathbf{x}_{n-1}^{(i)}, \mathbf{y}_n; \psi)$.

b. Compute importance weights: $w_n^{(i)} \propto w_{n-1}^{(i)} \frac{p(\mathbf{y}_n | \mathbf{x}_n^{(i)}; \theta) p(\mathbf{x}_n^{(i)} | \mathbf{x}_{n-1}^{(i)}; \theta)}{q_n(\mathbf{x}_n^{(i)} | \mathbf{x}_{n-1}^{(i)}, \mathbf{y}_n; \psi)}$.

(ii) For $i = 1, \dots, N$:

Normalize the importance weights: $\hat{w}_n^{(i)} = \frac{w_n^{(i)}}{\sum_{j=1}^N w_n^{(j)}}$.

(iii) Compute the filtered estimate: $E[f(\mathbf{x}_n)] \approx \sum_{i=1}^N \hat{w}_n^{(i)} f(\mathbf{x}_n^{(i)})$.

(iv) Resample N particles with probabilities $\{\hat{w}_n^{(i)}\}_{i=1}^N$ and for $i = 1, \dots, N$ set $w_n^{(i)} = \frac{1}{N}$.

(SIR) method of Rubin (1987) to the context of filtering in state space models. The basic SISR particle filter is given as Algorithm 1.

While resampling is a crucial feature to the success of the particle filter, it is important to understand why particles are resampled and what the effects of resampling are. Resampling does not cure the degeneracy problem when it comes to the particle filter's estimate of the entire joint distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$. Repeatedly resampling particles copied from previous generations reduces the number of distinct particles representing the early parts of the joint distribution. The particle filter produces a good approximation of the marginal distribution $p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta)$ and the joint distribution $p(\mathbf{x}_{n-k:n} | \mathbf{y}_{1:n}; \theta)$ when k is small. However, its approximation at time n of the entire joint distribution $p(\mathbf{x}_{0:n} | \mathbf{y}_{1:n}; \theta)$ and the earlier marginal distributions $p(\mathbf{x}_{n-l} | \mathbf{y}_{1:n}; \theta)$ will be poor as n and l increase. Due to this effect, in practice only the most recent generation of particles $\{\mathbf{x}_{n-k:n}^{(i)}\}_{i=1}^N$ are resampled and stored in memory. The purpose of resampling is to prevent future degeneracy by replicating those particles that appear relevant for estimating next period's target distribution. Resampling at time n ensures that next period's marginal distribution $p(\mathbf{x}_{n+1} | \mathbf{y}_{1:n+1}; \theta)$ will be well estimated. Estimates of the state variable (12) should always be calculated before resampling. Resampling introduces additional Monte Carlo variation into the algorithm. In the discussion that follows, we will see that it is better not to resample at every iteration (see Section 2.4).

The tuning parameters ψ of the importance density will depend upon the design of the algorithm. In many particle filters, they are equal to (a subset of) the parameters of the model, i.e. $\psi = \theta_1$ where $\theta = (\theta'_1, \theta'_2)'$. Different particle filtering algorithms are obtained by different choices of the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ and different resampling algorithms, which are both chosen by the user.

The number of particles also does not need to remain constant during the algorithm. Particle size may vary over time either deterministically or at random such that N_n . Alternatively, the number of particles can change within each iteration. For example, one can draw N^+ particles (where $N^+ = \alpha N$ for a positive integer α) from the importance distribution, compute the weights (9) and an estimate of $\mathbb{E}_q[f(\mathbf{x}_{0:n})]$ in (12) using the N^+ particles, and then resample only N out of the N^+ particles. The advantage is that a larger number of particles are used when the estimator is computed.

2.3 Choosing an importance distribution

Designing a good particle filter is analagous to designing a good MCMC algorithm; the appropriate algorithm depends on the problem. In particular, the user should take into account the structure of the state space model when choosing the importance distribution. Selecting a good proposal distribution is important because it results in better balanced importance weights and a more stable estimator. This section covers the major classes of importance distributions. Different distributions will result in different functional forms for the incremental weights in the weight recursion (9). To shorten the survey, detailed derivations of the algorithms are left to the references.

2.3.1 The bootstrap filter

The simplest particle filter uses the transition density as the proposal making it equal to

$$q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi) = p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta). \quad (14)$$

Many authors call this importance distribution the prior kernel or prior distribution given the Bayesian interpretation of the filtering recursions (3) and (4). This was used in the original particle filter of Gordon et al. (1993) called the *bootstrap filter*. The weight recursion (9) simplifies to

$$w_n \propto w_{n-1} p(\mathbf{y}_n | \mathbf{x}_n; \theta). \quad (15)$$

If one resamples each period, then it simplifies further to $w_n \propto p(\mathbf{y}_n | \mathbf{x}_n; \theta)$ because the importance weights from the previous iteration are equal. This particle filter is simple to implement and can perform well for some models but can easily be improved upon. Notice that information in the current observation $\mathbf{y}_n$ is not used in the proposal distribution.

2.3.2 Conditionally optimal importance distribution

The particle filtering literature includes the notion of a *conditionally* optimal importance distribution for any model. The conditionally optimal distribution is defined as the distribution that

minimizes the Monte Carlo variation of the importance weights. The “conditional” portion of this statement emphasizes that the importance distribution is optimal if one only conditions on the current observation $\mathbf{y}_n$ and last period’s particles $\left\{\mathbf{x}_{n-1}^{(i)}\right\}_{i=1}^N$. This idea was introduced by Liu and Chen (1995), although it exists in an earlier literature on SIS algorithms from Zaritskii et al. (1975) and Akashi and Kumamoto (1977). The conditionally optimal importance distribution is given by

$$\begin{aligned} q_n(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \psi) &= p(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \theta), \\ &= \frac{p(\mathbf{y}_n \mid \mathbf{x}_n, \mathbf{x}_{n-1}; \theta) p(\mathbf{x}_n \mid \mathbf{x}_{n-1}; \theta)}{p(\mathbf{y}_n \mid \mathbf{x}_{n-1}; \theta)}, \end{aligned} \quad (16)$$

which implies that the weight update (9) simplifies to

$$w_n \propto w_{n-1} p(\mathbf{y}_n \mid \mathbf{x}_{n-1}; \theta).$$

The incremental weight is interestingly a function of the previous state $\mathbf{x}_{n-1}$ and not the current state $\mathbf{x}_n$. This importance distribution unfortunately requires drawing from $p(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \theta)$ and evaluating $p(\mathbf{y}_n \mid \mathbf{x}_{n-1}; \theta)$ for the importance weights. Using this importance distribution typically will only be possible in special circumstances, e.g. when the measurement equation (1) is linear and its density is Gaussian. However, researchers use this distribution as a benchmark and try to approximate it with sub-optimal choices.

2.3.3 Proposal distributions resulting in i.i.d. samples

When using the conditionally optimal importance distribution (16), the importance weights can be computed before proposals are made as the weights are independent of the time n particles. This suggests that one can calculate the importance weights first, resample last period’s particles, and then extend the paths of only those particles that were resampled. In this case, the order in which resampling and drawing new particles from $q_n(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ are performed is reversed. Resampling the past particles prior to extending their paths can improve the particle filters’ approximation because the resampled particles $\left\{\mathbf{x}_{n-1}^{(i)}\right\}_{i=1}^N$ form a different, improved importance distribution. The importance weights for the newly sampled particles are also equal after sampling $\left\{\mathbf{x}_n^{(i)}\right\}_{i=1}^N$ meaning that the draws produced by the algorithm are *i.i.d.*. This idea originates as a special case of the auxiliary particle filter of Pitt and Shephard (1999, 2001) presented in the next section. It was formalized by Cappé et al. (2005) and has been extended by Johannes et al. (2006). This procedure will only be possible for a limited number of models.

2.3.4 Auxiliary particle filters

A generalization of the *i.i.d.* sampling algorithm just described is the auxiliary particle filter (APF) of Pitt and Shephard (1999, 2001), henceforth (PS). PS introduced a newly defined importance distribution $q_n(\mathbf{x}_n, j \mid \mathbf{y}_{1:n}; \psi)$ where j is an auxiliary variable that indexes the particles in existence from time $n - 1$. This importance distribution will be sampled in two steps. The auxiliary variables $\{j^{(i)}\}_{i=1}^N$ get drawn first (using $\mathbf{y}_n$) in a first-stage resampling step that replicates any particles $\{\mathbf{x}_{n-1}^{(i)}\}_{i=1}^N$ which appear relevant for estimating $\mathbf{x}_n$. Conditional on these indicators $\{j^{(i)}\}_{i=1}^N$, new particles $\{\mathbf{x}_n^{(i)}\}_{i=1}^N$ are simulated from the remainder of the importance distribution using $\{\mathbf{x}_{n-1}^{j^{(i)}}\}_{i=1}^N$. The auxiliary variables are then discarded. The point of introducing the auxiliary variables is to find a way to use the information in the current observation $\mathbf{y}_n$ to find (and replicate) the “good” particles within the existing set $\{\mathbf{x}_{n-1}^{(i)}\}_{i=1}^N$ in order to form a better importance distribution. There are many ways of using this auxiliary variable idea within a particle filter; a basic version is described here.

A simple, two-step procedure for drawing the pair $\{\mathbf{x}_n^{(i)}, j^{(i)}\}_{i=1}^N$ begins by factoring the new importance distribution into two smaller parts

$$q_n(\mathbf{x}_n, j \mid \mathbf{y}_{1:n}; \psi) \equiv q_n(\mathbf{x}_n \mid j, \mathbf{y}_{1:n}; \psi) q_n(j \mid \mathbf{y}_{1:n}; \psi). \quad (17)$$

Writing the importance distribution proportional to the existing information at the beginning of the iteration, we have

$$q_n(\mathbf{x}_n, j \mid \mathbf{y}_{1:n}; \psi) \propto p(\mathbf{x}_n \mid \mathbf{x}_{n-1}; \theta) \tau_n w_{n-1}, \quad (18)$$

where τ_n is a function chosen by the user. If the first term on the right-hand side of (17) is the transition density

$$q_n(\mathbf{x}_n \mid j, \mathbf{y}_{1:n}; \psi) = p(\mathbf{x}_n \mid \mathbf{x}_{n-1}; \theta), \quad (19)$$

then it follows from (18) that the marginal distribution for the indicator must satisfy

$$q_n(j \mid \mathbf{y}_{1:n}; \psi) \propto \tau_n w_{n-1}. \quad (20)$$

These are the (unnormalized) first-stage weights where, because of this expression, PS call τ_n *adjustment multiplier weights*. The auxiliary variables $\{j^{(i)}\}_{i=1}^N$ and the particles they index $\{\mathbf{x}_{n-1}^{(i)}\}_{i=1}^N$ are first drawn with probabilities proportional to (20). PS suggested choosing $\tau_n = p(\mathbf{y}_n \mid \mu_n; \theta)$ where μ_n is the mean or mode of $p(\mathbf{x}_n \mid \mathbf{x}_{n-1}; \theta)$. After sampling the indicators, new particles $\{\mathbf{x}_n^{(i)}\}_{i=1}^N$ are drawn from the transition density using those particles that were resampled. For this choice of τ_n , the recursion for the importance weights in (9), now considered

Algorithm 2 Auxiliary particle filter (APF)

At $n = 0$, for $i = 1, \dots, N$

Draw $\mathbf{x}_0^{(i)} \sim q_0(\mathbf{x}_0)$ and set $w_0^{(i)} = \frac{p(\mathbf{x}_0^{(i)})}{q_0(\mathbf{x}_0^{(i)})}$.

For $n = 1, \dots, T$:

(i) For $i = 1, \dots, N$:

a. Set $\tau_n^{(i)} = p(\mathbf{y}_n | \mu_n^{(i)})$ where $\mu_n^{(i)}$ is equal to the mean/mode of $p(\mathbf{x}_n | \mathbf{x}_{n-1}^{(i)})$.

b. Compute importance weights: $w_{1,n}^{(i)} \propto w_{n-1}^{(i)} \tau_n^{(i)}$.

c. Normalize the importance weights: $\hat{w}_{1,n}^{(i)} = \frac{w_{1,n}^{(i)}}{\sum_{j=1}^N w_{1,n}^{(j)}}$.

(ii) Sample N indicators $\{j^{(i)}\}_{i=1}^N$ with probabilities $\{\hat{w}_{1,n}^{(i)}\}_{i=1}^N$.

(iii) For $i = 1, \dots, N$:

a. Draw: $\mathbf{x}_n^{(i)} \sim p_n(\mathbf{x}_n | \mathbf{x}_{n-1}^{j^{(i)}}; \theta)$.

b. Compute second-stage importance weights: $w_n^{(i)} \propto \frac{p(\mathbf{y}_n | \mathbf{x}_n^{(i)})}{\tau_n^{j^{(i)}}}$.

(iv) Normalize the importance weights: $\hat{w}_n^{(i)} = \frac{w_n^{(i)}}{\sum_{j=1}^N w_n^{(j)}}$.

(v) Compute the filtered estimate: $E[f(\mathbf{x}_n)] \approx \sum_{i=1}^N \hat{w}_n^{(i)} f(\mathbf{x}_n^{(i)})$.

as second-stage weights, reduces to

$$\begin{aligned} w_n &\propto w_{n-1} \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)}{q_n(\mathbf{x}_n, j | \mathbf{y}_{1:n}; \psi)}, \\ &\propto w_{n-1} \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta) p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)}{p(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta) \tau_n w_{n-1}}, \\ &= \frac{p(\mathbf{y}_n | \mathbf{x}_n; \theta)}{p(\mathbf{y}_n | \mu_n; \theta)}. \end{aligned} \tag{21}$$

These weights are used to compute the filtered estimate of the state variable via (12). In their original paper, particles are then resampled a second time using these weights. In their formulation of the algorithm, PS allowed the particle size to vary at each iteration. When resampling in the first stage, they suggest drawing $N^+ = \alpha N$ particles for positive integer α , passing these N^+ particles through the transition density, and then resampling only N particles in the second stage.

Due to the importance of the APF, several recent papers have been written studying its theoretical properties and a number of practical points can be made on how to improve its implementation. In independent work by Johansen and Doucet (2008) and Douc et al. (2007), the authors prove that the second resampling stage in the original algorithm of PS is unnecessary if one keeps the particle size constant at each iteration. Its inclusion increases the asymptotic variance of the corresponding estimator. The APF is widely applicable and easy to implement making it worthwhile to repeat a simple version of it as Algorithm 2.

Secondly, Douc et al. (2007) provide a proof to find the optimal choice of τ_n in (20) that minimizes the variance of the importance weights. The optimal choice will unfortunately be

unavailable for most state space models but their results can still serve as a guideline for its selection. Johansen and Doucet (2008) also demonstrate that the APF can actually degrade the performance of a particle filter even in the case of what PS called perfect adaption. Perfect adaption occurs when τ_n is chosen so that the second-stage importance weights are equal (the special case of *i.i.d.* sampling discussed above). These authors show that the performance of the APF will depend upon the signal to noise ratio in the state space model. If the signal to noise ratio in the state space model is low, first-stage resampling can mislead the cloud of particles away from interesting areas of the support.

2.3.5 Importance distributions built from accept-reject algorithms

Accept-reject algorithms (see Robert and Casella (2004, p. 47)) can also be incorporated within a particle filter. For example, when it is impossible to draw directly from the conditionally optimal importance distribution (16), one can draw N particles from this distribution using an accept-reject algorithm. This idea originates with Hürzeler and Künsch (1998) and also Tanizaki and Mariano (1998). The algorithms have been studied theoretically by Künsch (2005). As an accept-reject algorithm will be run N times per time period, the user needs to find a good proposal distribution within the accept-reject algorithm. If this is not chosen well, a large number of trial simulations may be needed for each particle to be accepted. Finally, it is possible to use the accept-reject algorithm within an APF as in PS.

2.3.6 Local approximations of the conditionally optimal distribution

Doucet et al. (2000) includes an introduction to methods that incorporate the current observation in the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ by approximating the conditionally optimal distribution $p(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \theta)$ through local linearization via Taylor series expansions of the functions within the state space model (1) and (2). After taking the Taylor series expansions, draws are made from a state space model that approximates the true model.

Other closely related proposal distributions exist when either or both of the functions $h_n(\cdot)$ and $g_n(\cdot)$ in (1) and (2) are nonlinear but the disturbances η_n and ε_n are additively Gaussian. These importance distributions, given in van der Merwe et al. (2000) and Guo et al. (2005), make proposals using one-step of the extended or unscented Kalman filter applied to each particle. The extended and unscented Kalman filters are nonlinear filtering algorithms that use analytical approximations; for details, see Anderson and Moore (1979), Julier and Uhlmann (1997), Julier et al. (2000)).

Another means of approximating the optimal distribution $p(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \theta)$ is to choose the parameters ψ of the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ in each time period (and for

each particle) so that its mode matches the mode of the optimal distribution. Finding the mode of the target can be accomplished using Newton-Raphson methods, which will be effective when the target is uni-modal. One can then choose the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ to be a normal or Student t distribution setting the parameters ψ to have this mode and an inflated variance to ensure its support includes the support of the target distribution.

2.3.7 Rao-Blackwellization

There are state space models where a subset of the state vector may be integrated out analytically. Analytical integration reduces the Monte Carlo variation of the resulting estimator and is known as Rao-Blackwellization; e.g., see Robert and Casella (2004, p. 130). Separating the state vector into parts $\mathbf{x}_n = (\mathbf{x}'_{1,n}, \mathbf{x}'_{2,n})'$, the model implies that the filtering distribution can be decomposed as $p(\mathbf{x}_{1,n}, \mathbf{x}_{2,n} | \mathbf{y}_{1:n}; \theta) = p(\mathbf{x}_{1,n} | \mathbf{x}_{2,n}, \mathbf{y}_{1:n}; \theta) p(\mathbf{x}_{2,n} | \mathbf{y}_{1:n}; \theta)$. Particles are only simulated randomly for $p(\mathbf{x}_{2,n} | \mathbf{y}_{1:n}; \theta)$ while $p(\mathbf{x}_{1,n} | \mathbf{x}_{2,n}, \mathbf{y}_{1:n}; \theta)$ can be evaluated analytically. One class of state space models amenable to Rao-Blackwellization that is popular in economics is

$$\mathbf{y}_n = Z_n(\mathbf{x}_{2,n}) \mathbf{x}_{1,n} + \varepsilon_n, \quad \varepsilon_n \sim \mathcal{N}(0, H_n(\mathbf{x}_{2,n})), \quad (22)$$

$$\mathbf{x}_{1,n} = T_n(\mathbf{x}_{2,n}) \mathbf{x}_{1,n-1} + \eta_n, \quad \eta_n \sim \mathcal{N}(0, Q_n(\mathbf{x}_{2,n})), \quad (23)$$

$$p_{ij} = p(\mathbf{x}_{2,n} = j | \mathbf{x}_{2,n-1} = i), \quad \mathbf{x}_{2,n} \in \{1, 2, \dots, k\}, \quad (24)$$

which is a linear, Gaussian state space model where the parameters in the state space matrices Z_n, T_n, Q_n, H_n depend upon the value of an additional discrete state variable $\mathbf{x}_{2,n}$. The discrete state variables follow a first-order Markov process as in (24). These models are covered in Kim and Nelson (1999) and Frühwirth-Schnatter (2006).

Rao-Blackwellization may be efficiently employed on both the continuous and discrete state variables. These particle filters are due to Chen and Liu (2000) who named them *mixture Kalman filters*, see also Doucet et al. (2001). Conditional on the discrete state variables $\{\mathbf{x}_{2,n}^{(i)}\}_{i=1}^N$ the resulting system is a linear, Gaussian state space model and $p(\mathbf{x}_{1,n} | \mathbf{x}_{2,n}, \mathbf{y}_{1:n}; \theta)$ can be evaluated by the Kalman filter. de Freitas et al. (2004), Schön et al. (2005), and Bos and Shephard (2006) are additional references which apply some form of this methodology.

Other models that can be Rao-Blackwellized are partially observed Gaussian state space models as in Andrieu and Doucet (2002), which include dynamic probit and Tobit models with unobserved states. State space models where the functions in (1) and (2) are nonlinear but depend on both discrete and continuous-valued states can also be Rao-Blackwellized; e.g., see Andrieu et al. (2003).

2.3.8 Block sampling

When using MCMC, it is well-known that better performing algorithms can be built if one can find a proposal distribution that enables joint sampling of blocks of variables from the target distribution. If the proposal is chosen well, sampling variables in blocks improves the speed by which the Markov chain explores the support of the distribution. Doucet et al. (2006) propose a similar idea for particle filters using some of the simulation methods presented in Section 4. At the beginning of iteration n of a particle filter, the algorithm has already simulated and stored the paths $\left\{ \mathbf{x}_{0:n-1}^{(i)} \right\}_{i=1}^N$. The goal is not only to extend each path at the endpoint but instead returning k time periods into the past (where k is say 5-10) and sample a block $\left\{ \mathbf{x}_{n-k:n}^{(i)} \right\}_{i=1}^N$. Instead of using a proposal distribution $q_n(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$, the proposal distribution is over the path $q_{n-k:n}(\mathbf{x}_{n-k:n} \mid \mathbf{x}_{n-k-1:n-1}, \mathbf{y}_{n-k:n}; \psi)$. In order to implement block sampling, the importance weight recursions (9) need to be rewritten to account for the expanded importance density, see Doucet et al. (2006) for details. This algorithm can potentially improve the particles exploration of the support of the target distribution and simultaneously reduce the degeneracy problem. However, it comes at the expense of additional computing time and places greater demands on the user to design the algorithm.

2.3.9 MCMC and adaptive proposals

It is also possible propose new particle positions each time period by each particle using one iteration of a Metropolis-Hastings or Gibbs sampler. These ideas were proposed by Gilks and Berzuini (2001) under the name of the *resample-move* algorithm, see also Fearnhead (2002). Another recent line of research considers using the past particles $\left\{ \mathbf{x}_{0:n-1}^{(i)}, \hat{w}_n^{(i)} \right\}_{i=1}^N$ to adapt the importance distribution over time. Cornebise et al. (2008) consider selecting the parameters ψ of $q_n(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ each period to minimize an empirical estimate of the Shannon entropy or the coefficient of variation between the empirical distribution of the particles and the target distribution. Their paper contains additional references to work on adaptive methods in particle filters. Using MCMC and adaptive proposals within particle filters are relatively more advanced methods. They are closely related to the algorithms discussed in Section 4 and will be covered in more detail there.

2.4 Resampling and branching algorithms

There are four resampling algorithms that dominate most of the literature: multinomial, residual, systematic, and stratified resampling. All of these algorithms can be performed in $O(N)$

operations.¹ Some of them are also discussed in the literature on genetic algorithms, see Whitley (1994). The main point for applied researchers to note is that some resampling algorithms are preferable because they introduce less Monte Carlo variation into the particle filter's estimator. Douc et al. (2005) compare their efficiency in terms of Monte Carlo variation. They prove that the stratified resampling algorithm of Kitagawa (1996a) and the residual resampling scheme of Liu and Chen (1995) should be preferred to the original multinomial resampling of Gordon et al. (1993). The Monte Carlo variation introduced by these algorithms is strictly smaller.

The residual and stratified resampling algorithms are also unbiased in the sense that the expected number of times a particle $\mathbf{x}_n^{(i)}$ will be resampled is equal to its importance weight. Thus, the algorithms satisfy the condition

$$E \left[\mathbf{x}_n^{(i)} \mid \left\{ \hat{w}_n^{(i)} \right\}_{i=1}^N \right] = N \hat{w}_n^{(i)}.$$

This condition is a maintained assumption in the consistency and asymptotic normality proofs behind most particle filters.

The systematic resampling algorithm of Carpenter et al. (1999) is the easiest to implement. It can also perform well in Monte Carlo studies, see Douc et al. (2005), but does not always dominate multinomial resampling in terms of variance. Other notable resampling algorithms include the optimal resampling algorithm of Fearnhead and Clifford (2003), which should be used for any model whose state variable has a discrete component, e.g. the mixture of linear, Gaussian models (22)-(24). The stopping-time resampling algorithm of Chen et al. (2005) is another recent alternative. The papers by Fearnhead and Clifford (2003) and Chen et al. (2005) illustrate the point that a resampling algorithm can be tailored for specific classes of models or even a specific application.

The original particle filter of Gordon et al. (1993) carries out resampling every time period. To lower the degree of Monte Carlo variation introduced into the estimator, many researchers suggest resampling only after time periods where the importance weights are unstable. There are three commonly used measures of weight instability: the coefficient of variation (CV) of Kong et al. (1994), the effective sample size (ESS) of Liu (1996), and the Shannon entropy (SE) of the weights. During each iteration of the algorithm, one calculates any one of these measures and if it drops above/below a user chosen threshold then resampling is performed. Resampling is therefore performed at random times.

¹Matlab code for each of the resampling algorithms can be found at Nando de Freitas' webpage at <http://www.cs.ubc.ca/~nando/software.html>. R code for some simple particle filters is available at Paul Fearnhead's homepage at <http://www.maths.lans.ac.uk/~fearnhea/PF/>. Simple pseudo-code for the last three resampling algorithms is available at <http://staff.feweb.vu.nl/dcreal/>.

The ESS is given by

$$ESS = \frac{1}{\sum_{i=1}^N \left(\hat{w}_n^{(i)}\right)^2}, \quad (25)$$

and is a number between 1 and N . If the $ESS = N$, the interpretation is that the weights are equally balanced and that all N particles are contributing to the estimator in (12). The threshold for the ESS is typically chosen to be a percentage of the number of particles, say 0.5 to 0.8. If the ESS drops below this level, then one of the resampling algorithms discussed above is applied.

The CV is defined as

$$CV = \left[\frac{1}{N} \sum_{i=1}^N \left(N \hat{w}_n^{(i)} - 1 \right)^2 \right]^{0.5}, \quad (26)$$

and is a number between zero and $\sqrt{N-1}$. If all the weights are equal then $CV = 0$ and if one particle has all the probability mass then $CV = \sqrt{N-1}$.

The SE is

$$SE = - \sum_{i=1}^N \hat{w}_n^{(i)} \log_2 \hat{w}_n^{(i)}, \quad (27)$$

which is minimal at zero when one particle has all the mass. Its largest value is $\log_2 N$ when all the weights are equal. When using the CV and SE criterion to determine when to resample, the threshold will depend upon the model and on the particle size N .

An alternative to resampling algorithms for rejuvenating the particles are “branching” algorithms, which are popular in the theoretical probability literature and are reasonably simple to implement. In most implementations, the number of particles will be random over time N_n and therefore these methods are not as common in applications. For more details on branching algorithms; see, e.g. Crisan et al. (1999) and Del Moral and Miclo (2000).

2.5 Particle smoothing

The marginal smoothing distribution $p(\mathbf{x}_n \mid \mathbf{y}_{1:T}; \theta)$ characterizes the state variable given all the observations in the dataset, where $T \geq n$. Computing the distribution $p(\mathbf{x}_n \mid \mathbf{y}_{1:T}; \theta)$ for all possible n while T is held fixed is the most common form of smoothing in economics. This is known as *fixed-interval smoothing* in the engineering literature, see Anderson and Moore (1979) for further discussion. Fixed-interval smoothing algorithms for state space models are historically based upon one of two frameworks known as *forward-filtering backward-smoothing* or *two-filter formula smoothing*. Both types of algorithms compute the same sequence of marginal distributions $\{p(\mathbf{x}_n \mid \mathbf{y}_{1:T}; \theta)\}_{n=1}^T$ and only differ in how the neighboring states are integrated out of the joint smoothing distribution. Particle smoothing algorithms have been created using both approaches. A good reference for this material is Briers et al. (2004) on which my discussion

is based while Chapter 3 of Cappé et al. (2005) contains a more general, measure-theoretic treatment.

After running a filtering algorithm forward and computing each of the predictive and filtering distributions $\{p(\mathbf{x}_{n+1} | \mathbf{y}_{1:n}; \theta), p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta)\}_{n=1}^T$, the sequence of marginal smoothing distributions can be computed from $n = T - 1, \dots, 1$ using the following backward recursion

$$\begin{aligned}
p(\mathbf{x}_n | \mathbf{y}_{1:T}; \theta) &= \int p(\mathbf{x}_n, \mathbf{x}_{n+1} | \mathbf{y}_{1:T}; \theta) d\mathbf{x}_{n+1}, \\
&= \int p(\mathbf{x}_{n+1} | \mathbf{y}_{1:T}; \theta) p(\mathbf{x}_n | \mathbf{x}_{n+1}, \mathbf{y}_{1:T}; \theta) d\mathbf{x}_{n+1}, \\
&= \int p(\mathbf{x}_{n+1} | \mathbf{y}_{1:T}; \theta) p(\mathbf{x}_n | \mathbf{x}_{n+1}, \mathbf{y}_{1:t}; \theta) d\mathbf{x}_{n+1}, \\
&= p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta) \int \frac{p(\mathbf{x}_{n+1} | \mathbf{y}_{1:T}; \theta) p(\mathbf{x}_{n+1} | \mathbf{x}_n; \theta)}{p(\mathbf{x}_{n+1} | \mathbf{y}_{1:n}; \theta)} d\mathbf{x}_{n+1}. \tag{28}
\end{aligned}$$

The backward recursion is initialized using the last filtering distribution $p(\mathbf{x}_T | \mathbf{y}_{1:T}; \theta)$ and the predictive distribution $p(\mathbf{x}_{T+1} | \mathbf{y}_{1:T}; \theta)$ from the forward filtering recursions. The smoothing algorithms for the linear, Gaussian state space model that are popular in economics, e.g. Harvey (1989), Kim and Nelson (1999), and Durbin and Koopman (2001), are versions of this approach based upon original work by Rauch et al. (1965). Doucet et al. (2000) invented a particle smoother using this framework but it is an $O(N^2T)$ operation making it uncompetitive with MCMC. A second shortcoming is the particles' locations on the support of the distributions are fixed on the forward filtering pass. These particles are then simply reweighted by changing their importance weights on a backwards pass using the information in the future data. Although future observations are available, new particle locations are not simulated on the backwards pass. These one-sided particle locations may not be representative of the marginal smoothing distributions given more data.

Two-filter formula smoothing consists of running two filters that are independent of one another and using their output to construct the marginal smoothing distributions. This method was proposed by Fraser and Potter (1969) for linear, Gaussian models. The first filter calculates the one-step ahead predictive and filtering distributions $\{p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta), p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta)\}_{n=1}^T$ running forward in time and the second filter calculates a series of functions $\{p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta)\}_{n=1}^T$ running backward in time. Together these can compute the marginal smoothing distributions using the forward recursion

$$\begin{aligned}
p(\mathbf{x}_n | \mathbf{y}_{1:T}; \theta) &= p(\mathbf{x}_n | \mathbf{y}_{1:n-1}, \mathbf{y}_{n:T}; \theta), \\
&= \frac{p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) p(\mathbf{y}_{n:T} | \mathbf{y}_{1:n-1}, \mathbf{x}_n; \theta)}{p(\mathbf{y}_{n:T} | \mathbf{y}_{1:n-1}; \theta)}, \\
&\propto p(\mathbf{x}_n | \mathbf{y}_{1:n-1}; \theta) p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta), \\
&\propto p(\mathbf{x}_n | \mathbf{y}_{1:n}; \theta) p(\mathbf{y}_{n+1:T} | \mathbf{x}_n; \theta).
\end{aligned}$$

The set of backward functions $p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta)$ can be computed recursively via

$$p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta) = \int p(\mathbf{y}_{n+1:T} | \mathbf{x}_{n+1}; \theta) p(\mathbf{x}_{n+1} | \mathbf{x}_n; \theta) p(\mathbf{y}_n | \mathbf{x}_n; \theta) d\mathbf{x}_{n+1},$$

which is known as the *backward information filter* and was first proposed by Mayne (1966). Difficulty may sometimes arise with this approach because $p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta)$ is not a probability density. The integral of this function can grow without bound (the integral is infinite). Practical implementations of two-filter formula smoothing are therefore based on normalization of $p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta)$ to ensure that it is a density.

Kitagawa (1996a) proposed the first particle smoother based upon two-filter formula smoothing. However, this algorithm implicitly assumes that $p(\mathbf{y}_{n:T} | \mathbf{x}_n; \theta)$ is integrable. Briers et al. (2004) develop a two-filter formula particle smoother that solves the integrability problem. Their method also simulates fresh particle locations on the backward pass but it remains an $O(N^2T)$ operation. Building on this work, Fearnhead et al. (2008) have recently shown how to apply a two-filter formula particle smoother which is only an $O(NT)$ operation making it competitive with MCMC. This smoother does not solve the problem for all general state space models (1) and (2) but applies to only those models whose state equation is linear and Gaussian. This is typically the case in economics. Details of the implementation of the algorithm are relatively lengthy and therefore I refer to their paper for further discussion.

Godsill et al. (2004) developed a simulation smoothing algorithm for a general nonlinear, non-Gaussian state space model using particle filters that is an $O(NT)$ operation. A simulation smoother is an algorithm that takes random draws of a sequence of state variables $\mathbf{x}_{0:T}$ from the joint smoothing distribution $p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}; \theta)$. Their method can be viewed as an extension of the simulation smoothing algorithms for linear, Gaussian models of Carter and Kohn (1994), Frühwirth-Schnatter (1994), de Jong and Shephard (1995), and Durbin and Koopman (2002). By repeatedly drawing samples from this distribution, smoothed estimates of the state variable can be computed by averaging across the simulations as in standard *i.i.d* Monte Carlo methods. The algorithm is particularly simple and therefore I refer to Godsill et al. (2004) for its implementation.

Particle smoothing can be used to approximate the mean of the marginal or joint smoothing distributions. This is the optimal estimator if the user has a quadratic loss function. Viewing the joint smoothing distribution as a posterior distribution, it is also possible for particle filters to approximate the *maximum a posteriori* (MAP) estimator. This is the sequence $\mathbf{x}_{0:T}$ that maximizes the posterior distribution $p(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}; \theta)$ and is the optimal estimator under a zero-one loss function. Godsill, Doucet, and West (2002) solve this problem by extending the well-known Viterbi (1967) algorithm for discrete-state HMM models to the context of particle filters.

The algorithm is a simple dynamic programming problem and is an $O(N^2T)$ operation.

2.6 Parameter estimation using particle methods

In this subsection, I describe the literature on frequentist estimation of the parameters of general state space models using particle filters. Bayesian estimation of model parameters using SMC is more closely related to the methods presented in Section 4.

2.6.1 Computing the likelihood for a general state space model

The log-likelihood of a time series model is given by the prediction error decomposition

$$\log L(\theta | \mathbf{y}_{1:T}) = \log p(\mathbf{y}_1, \dots, \mathbf{y}_T; \theta) = \sum_{n=1}^T \log p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta).$$

The likelihood of a general state space model is the integral in the denominator of (4), which is typically intractable and can only be approximated. The particle filter's approximation of the likelihood function for a single observation is conveniently the sum of the unnormalized importance weights at that iteration

$$\hat{p}(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta) = \frac{1}{N} \sum_{i=1}^N \frac{p(\mathbf{x}_n^{(i)} | \mathbf{y}_{1:n})}{q_n(\mathbf{x}_n^{(i)} | \mathbf{y}_{1:n})} = \frac{1}{N} \sum_{i=1}^N w_n^{(i)}. \quad (29)$$

This estimator of the likelihood function is unbiased. Taking the log of this approximation and summing over all the observations gives

$$\log L(\theta | \mathbf{y}_{1:T}) \approx \sum_{n=1}^T \log \hat{p}(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta) = \sum_{n=1}^T \log \left[\frac{1}{N} \sum_{i=1}^N w_n^{(i)} \right],$$

which will unfortunately introduce a bias in the estimate due to Jensen's inequality. A first-order Taylor series expansion will provide the following biased-corrected log-likelihood

$$\log L(\theta | \mathbf{y}_{1:T}) \approx \sum_{n=1}^T \log \left[\frac{1}{N} \sum_{i=1}^N w_n^{(i)} \right] + \frac{1}{2} \frac{\hat{\sigma}^2}{N \exp \left(2 \sum_{n=1}^T \log \left[\frac{1}{N} \sum_{i=1}^N w_n^{(i)} \right] \right)}, \quad (30)$$

where $\hat{\sigma}^2 = \frac{1}{N} \sum_{i=1}^N \left(\hat{w}_n^{(i)} \right)^2 - \left(\frac{1}{N} \sum_{i=1}^N \hat{w}_n^{(i)} \right)^2$ is estimated using the particles. These approximations of the likelihood have been used in the work of Kim et al. (1998) and Chib et al. (2002) to compute Bayes factors and likelihood ratio statistics and in the Bayesian analysis of DSGE models; e.g. see An and Schorfheide (2007). Other model diagnostics, e.g. Ljung-Box statistics, can be computed from the output of a particle filter; see Andrieu, Doucet, Singh, and Tadić (2004, p. 429) for details.

2.6.2 Frequentist likelihood-based parameter estimation

Frequentist parameter estimation of nonlinear, non-Gaussian state space models by particle filters remains a current research topic. The two major issues to consider are computing the maximum likelihood (ML) estimator in a computationally efficient way and its statistical properties (i.e. consistency and asymptotic normality) once it is computed. Although the particle filter's approximation of the likelihood function at a point θ is consistent asymptotically in the number of particles, the log-likelihood function is not a continuous function of the parameters. The discontinuity is created from the resampling stage within a particle filter and can cause problems for gradient-based optimizers; e.g., see Hürzeler and Künsch (2001) for a detailed example of the problem.

Pitt (2002) developed a new algorithm called the smooth particle filter to overcome the problem of a non-smooth log-likelihood function. This algorithm replaces the standard resampling algorithm with a new resampling method. It builds a continuous c.d.f. using piecewise linear approximations between particles instead of the discrete c.d.f. used in the standard resampling algorithms. Pitt's algorithm is only viable when the state dimension is equal to one or perhaps two because smoothing the c.d.f. requires the ordering of the state variables during each iteration of the filter. The method becomes an $O(N^2T)$ operation beyond a one-dimensional state vector.

Olsson and Rydén (2008) consider maximization of the log-likelihood and also address the resulting estimator's theoretical properties. They approximate the parameter space using a discrete grid of points and evaluate the log-likelihood function by particle filter at each point. They then prove what conditions are needed on the grid size, the number of particles, and the state space model in order to guarantee consistency and asymptotic normality of the resulting estimator $\hat{\theta}$. This appears to be the first result of this kind.

Otherwise, most work on ML estimation using particle filters has focused on using approaches other than gradient-based optimizers that avoid the discontinuity problem. These methods include stochastic gradient-based methods, recursive maximum likelihood methods (Doucet and Tadić (2003), Poyiadjis et al. (2005a), Poyiadjis et al. (2005b)) and Monte Carlo expectation maximization (MCEM) methods (Cappé et al. (2005), Olsson et al. (2008)). The last paper also analyzes the statistical properties of the estimator. To my knowledge, none of these methods have been applied in the economics literature. Finally, Johansen et al. (2008) use SMC samplers to compute the ML estimator; these are discussed in Section 4.

2.6.3 Alternative methods and online estimation

A number of other proposals have been made for estimating the parameters of general state space models using particle methods. In particular, researchers are interested in estimating the distribution $p(\mathbf{x}_{0:n}, \theta \mid \mathbf{y}_{1:n})$ or $p(\mathbf{x}_n, \theta \mid \mathbf{y}_{1:n})$ online as data arrives. From an economic perspective, this is of practical interest because it provides a Bayesian alternative to the learning algorithms that are popular in macroeconomics; see, e.g. Evans and Honkapohja (2001). Meanwhile, Bayesian statisticians are interested in accounting for parameter uncertainty by approximating the marginal density $p(\mathbf{x}_n \mid \mathbf{y}_{1:n})$ recursively in time instead of the conditional density $p(\mathbf{x}_n \mid \mathbf{y}_{1:n}; \theta)$. Research in this area is still on-going. Some of the earlier methods are reviewed in Doucet et al. (2001). These methods include placing the parameters in the state vector with the variance set to zero, in which case the parameter space is only explored at initialization of the algorithm, see Kitagawa (1996b), and making the parameters dynamic within the state vector by adding artificial noise; e.g., see Liu and West (2001).

Storvik (2002) proposed learning the parameters sequentially in time by storing sufficient statistics related to each of the parameters in θ . One sufficient statistic gets stored for each particle path. After updating the sufficient statistics at each iteration, new parameter values are simulated using the sufficient statistics, which are then resampled along with the current set of parameters. This method is particularly convenient when the measurement and transition densities are in the exponential family because it is easy to summarize the relevant distributions using sufficient statistics. Fearnhead (2002), Polson et al. (2008), and Carvalho et al. (2008) apply more advanced versions of these methods successfully to several applications.

Andrieu et al. (2005) and Künsch (2006) note, however, that the success of these methods will depend upon the mixing properties of the Markov kernels within the algorithm. Past errors produced by the particle filter's approximations need to be forgotten and not accumulated over time. Repeatedly resampling the past particles gives a poor approximation of the joint distribution $p(\mathbf{x}_{0:n} \mid \mathbf{y}_{1:n}; \theta)$ as n increases because few particles represent early parts of the distribution. While estimates of $p(\mathbf{x}_{0:n-k} \mid \mathbf{y}_{1:n}; \theta)$ for large k should contribute information toward estimating the parameters θ , information about θ may not always accumulate if the approximation is poor, see Andrieu et al. (2005) for further discussion.

2.7 Application # 1: forecasting inflation with a time-varying unobserved components model

Forecasting inflation is an important part of monetary policy-making and has a long history in economics. Stock and Watson (2007) recently proposed forecasting inflation π_n using the

following time-varying local level model

$$\pi_n = \mathbf{x}_{1,n} + \varepsilon_n, \quad \varepsilon_n \sim \mathcal{N}(0, \exp(\mathbf{x}_{2,n})), \quad (31)$$

$$\mathbf{x}_{1,n} = \mathbf{x}_{1,n-1} + \eta_n, \quad \eta_n \sim \mathcal{N}(0, \exp(\mathbf{x}_{3,n})), \quad (32)$$

$$\mathbf{x}_{2,n} = \mathbf{x}_{2,n-1} + \omega_{1,n}, \quad \omega_{1,n} \sim \mathcal{N}(0, 0.2), \quad (33)$$

$$\mathbf{x}_{3,n} = \mathbf{x}_{3,n-1} + \omega_{2,n}, \quad \omega_{2,n} \sim \mathcal{N}(0, 0.2), \quad (34)$$

where $\mathbf{x}_{1,n}$ is the unobserved time-varying mean of inflation and $\mathbf{x}_{i,n}$ for $i = 2, 3$ are unobserved log-variances. Stock and Watson (2007) argued that this specification improves forecasting because the model accounts for the structural breaks present in inflation. It can be shown, see e.g. Harvey (1989, p. 68), that the local level model with constant variances is equivalent to an ARIMA(0,1,1) model with additional restrictions on the parameter space. The stochastic variances for the level and irregular components in (31)-(34) imply a time-varying variance and MA parameter in this ARIMA representation. The time-varying MA parameter conveniently summarizes how the model's forecast function changes through time.

Stock and Watson (2007) estimate the state variables of the model by MCMC, while it is (arguably) easier to implement a particle filter when there are no static parameters that need to be estimated. The design of the particle filter should take advantage of the structure of the model, which is linear, Gaussian conditional on the log-variances $\mathbf{x}_{2,n}$ and $\mathbf{x}_{3,n}$. The log-variances can therefore be simulated while the level component $\mathbf{x}_{1,n}$ can be Rao-Blackwellized using the Kalman filter. A good particle filter for this application is an APF version of the mixture Kalman filter; see, Chen and Liu (2000). For illustration purposes, this particle filter is given in detail as Algorithm 3. It includes the Kalman filter within it which is written specifically for the local level model. The notation $\mathbf{x}_{1,n|n-1}$ is the one-step ahead predicted estimate of the level component while $\mathbf{x}_{1,n|n}$ denotes the filtered estimate. The filtering algorithm was implemented with $N = 10,000$ particles and systematic resampling. Smoothed estimates of each of the components were computed by taking 1000 draws using the simulation smoothing algorithm of Godsill et al. (2004) discussed in Section 2.5.

In this APF algorithm, the function τ_n in (18) is the normal density with mean and variance equal to the prediction errors v_n and prediction error variances F_n produced by the Kalman filter. Forecasts are computed using the importance weights from the previous iteration, which are equal. The data are quarterly U.S. CPI inflation from Q1:1959-Q7:2008 constructed from the “real-time” price indices available from the U.S. Federal Reserve Bank of Philadelphia.

The one-step ahead forecast of inflation and the filtered and smoothed estimates of the volatilities are pictured in Figure 3. These estimates largely confirm the results of Stock and

Algorithm 3 Rao-Blackwellized APF for the Stock Watson (2007) model

At $n = 0$, for $i = 1, \dots, N$

Draw $\mathbf{x}_{1,0}^{(i)}, P_{1,0|0}^{(i)} \sim p(\mathbf{x}_{1,0}), \mathbf{x}_{2,0}^{(i)} \sim p(\mathbf{x}_{2,0}), \mathbf{x}_{3,0}^{(i)} \sim p(\mathbf{x}_{3,0})$ and set $w_0^{(i)} = \frac{1}{N}$.

For $n = 1, \dots, T$:

(i) For $i = 1, \dots, N$:

a. Kalman prediction step: $\mathbf{x}_{1,n|n-1}^{(i)} = \mathbf{x}_{1,n-1|n-1}^{(i)}$,
 $P_{1,n|n-1}^{(i)} = P_{1,n-1|n-1}^{(i)} + \exp(\mathbf{x}_{3,n-1}^{(i)})$.

(ii) Compute the forecast of inflation: $E[f(\mathbf{x}_{1,n|n-1})] \approx \sum_{i=1}^N w_{n-1}^{(i)} f(\mathbf{x}_{1,n|n-1}^{(i)})$.

(iii) For $i = 1, \dots, N$:

a. Compute importance weights: $v_n^{(i)} = \mathbf{y}_n - \mathbf{x}_{1,n|n-1}^{(i)}$,
 $F_n^{(i)} = P_{1,n|n-1}^{(i)} + \exp(\mathbf{x}_{2,n-1}^{(i)})$,
 $w_n^{(i)} \propto \mathcal{N}(v_n^{(i)}, F_n^{(i)})$.

b. Normalize the importance weights: $\hat{w}_n^{(i)} = \frac{w_n^{(i)}}{\sum_{j=1}^N w_n^{(j)}}$.

(iv) Resample N particles $\left\{ \mathbf{x}_{1,n-1|n-1}^{(i)}, P_{1,n-1|n-1}^{(i)}, \mathbf{x}_{2,n-1}^{(i)}, \mathbf{x}_{3,n-1}^{(i)} \right\}_{i=1}^N$ with probabilities

$\left\{ \hat{w}_n^{(i)} \right\}_{i=1}^N$ and for $i = 1, \dots, N$ set $w_n^{(i)} = \frac{1}{N}$.

(v) For $i = 1, \dots, N$:

a. Draw $\mathbf{x}_{2,n}^{(i)} \sim \mathcal{N}(\mathbf{x}_{2,n-1}^{(i)}, 0.2)$ and $\mathbf{x}_{3,n}^{(i)} \sim \mathcal{N}(\mathbf{x}_{3,n-1}^{(i)}, 0.2)$.

b. Kalman filter: $\mathbf{x}_{1,n|n-1}^{(i)} = \mathbf{x}_{1,n-1|n-1}^{(i)}$,
 $P_{1,n|n-1}^{(i)} = P_{1,n-1|n-1}^{(i)} + \exp(\mathbf{x}_{3,n}^{(i)})$,
 $v_n^{(i)} = \mathbf{y}_n - \mathbf{x}_{1,n|n-1}^{(i)}$,
 $F_n^{(i)} = P_{1,n|n-1}^{(i)} + \exp(\mathbf{x}_{2,n}^{(i)})$,
 $\mathbf{x}_{1,n|n}^{(i)} = \mathbf{x}_{1,n|n-1}^{(i)} + P_{1,n|n-1}^{(i)} F_n^{-1,(i)} v_n^{(i)}$,
 $P_{1,n|n}^{(i)} = P_{1,n|n-1}^{(i)} - P_{1,n|n-1}^{(i)} F_n^{-1,(i)} P_{1,n|n-1}^{(i)}$,

(vi) Compute the filtered estimate: $E[f(\mathbf{x}_{i,n})] \approx \sum_{i=1}^N \hat{w}_n^{(i)} f(\mathbf{x}_{i,n}^{(i)})$ for $i = 1$ to 3.

Watson (2007). The volatility of the level or permanent component $\exp(\mathbf{x}_{3,n}/2)$ increased during the period of high-inflation in the 1970's, while the volatility of the irregular component $\exp(\mathbf{x}_{2,n}/2)$ was relatively more stable. Filtered and smoothed estimates of the implied MA(1) parameter are shown in panel (ii) and they indicate that it also increased during this period. The forecastability of inflation appears to have changed over time as argued by Stock and Watson (2007). This data set includes five additional years of inflation beyond that analyzed by these authors. The volatility of inflation has recently increased beginning in the middle of 2007. It appears to be concentrated in the irregular or transitory component.

Figure 3: *Estimates from the time-varying local level model applied to quarterly U.S. inflation Q1:1959-Q3:2008: (i) inflation and its one-step ahead forecast; (ii) filtered and smoothed estimates of the implied MA(1) parameter; (iii) filtered and smoothed estimates of the irregular volatility $\exp(\mathbf{x}_{2,n}/2)$; (iv) filtered smoothed estimates of the state volatility $\exp(\mathbf{x}_{3,n}/2)$. NBER recession dates are indicated by the vertical bars.*

3 Theoretical properties

Early reviews of the theoretical properties of particle filters can be found in Chapters 2-3 of Doucet et al. (2001) and Crisan and Doucet (2002) while Chapter 9 of Cappé et al. (2005) includes a nice introduction to consistency and asymptotic normality for several particle filtering algorithms. More recent papers on consistency and asymptotic normality are Douc et al. (2007) and Douc and Moulines (2008). The goal of this section is to discuss some of the main results at an intuitive level.

3.1 Consistency and asymptotic normality

At each iteration, a particle filter produces samples $\{\mathbf{x}_{0:n}^{(i)}, w_n^{(i)}\}_{i=1}^N$ that can be used to approximate the expectation of a function f with respect to the joint importance distribution $q_{0:n}(\mathbf{x}_{0:n}|y_{1:n}; \psi)$ given by

$$\mathbb{E}_q[f(\mathbf{x}_{0:n})] = \int f(\mathbf{x}_{0:n}) \frac{p(\mathbf{x}_{0:n}|y_{1:n}; \theta)}{q_{0:n}(\mathbf{x}_{0:n}|y_{1:n}; \psi)} q_n(\mathbf{x}_{0:n}|y_{1:n}; \psi) d\mathbf{x}_{0:n}. \quad (35)$$

The particle filter's estimator is given again as

$$\mathbb{E}_q[f(\mathbf{x}_{0:n})] \approx \sum_{i=1}^N f(\mathbf{x}_{0:n}^{(i)}) \hat{w}_n^{(i)}, \quad (36)$$

where $\hat{w}_n$ are the normalized importance weights. The exact conditions for consistency and asymptotic normality of the estimator (36) depend upon the particle filter one implements. Proofs in the literature vary accordingly with different types of regularity conditions favored by different authors. Assuming that the integral of interest (35) is finite, the regularity conditions ensure that the particle filter's estimator remains finite and that a suitable law of large numbers (LLN) can be applied to show that the sum (36) converges as N grows large. Under additional regularity conditions, a central limit theorem (CLT) holds where additional conditions are needed to ensure the asymptotic variance remains finite. The latter guarantees that as the number of particles increases the estimator grows more accurate and the errors produced by the approximation become smaller.

The use of standard importance sampling algorithms requires some simple technical conditions on both the functions $f(\mathbf{x}_{0:n})$ within the integral and on the importance weights; see, e.g. Geweke (2005, p. 114). These are that the importance weights remain bounded so that the estimator remains well-behaved. In addition, the function $f(\mathbf{x}_{0:n})$ within the integral (35) must have finite variance $\mathbb{V}_p[f(x)] < \infty$ when evaluated under the target distribution. There are similar conditions for particle filters limiting both the set of functions $f(\mathbf{x}_{0:n})$ that are valid and conditions to ensure that the importance weights are not too variable. In a particle filter, the importance weights are determined recursively through the weight recursion (9). The variability of the importance weights w_n depends on the Monte Carlo variation introduced at the current iteration as well as any variability that is carried over from previous periods. This is due to the fact that particles simulated at previous iterations form part of the future joint importance distribution through the Dirac measure on past paths, see (6).

It is not possible to cover all the results in the literature and the different types of regularity conditions. Instead, the discussion here is limited to Theorem 10 from Douc and Moulines (2008), which covers both the SISR and APF algorithms with multinomial and residual resampling which can be performed at random times via the coefficient of variation (CV). The analysis by these authors can also be applied to some of the algorithms in Section 4. First, some of the simpler assumptions used in the proofs are described and then the main results are given in an intuitive manner.

The initial iteration of a particle filter is a standard importance sampling iteration. Therefore, the standard importance sampling assumptions apply to the first iteration $n = 0$ and these

are given by

- (i) $\mathbb{E}_p[f(\mathbf{x}_0)]$ exists;
- (ii) $\mathbb{V}_p[f(\mathbf{x}_0)]$ exists;
- (iii) The support of the initial importance distribution $q_0(\mathbf{x}_0)$ includes the target $p(\mathbf{x}_0)$;
- (iv) The initial importance weights w_0 are bounded.

Let X denote the state space of the Markov chain $\mathbf{x}_n$. Denote by C_n, A_n , and W_n , three proper sets of functions². C_n is set the set of functions f for which the particle filter will be consistent A_n is the set for which it is asymptotically normal. Additional conditions are:

- (v) For all iterations $n > 0$, the support of the incremental importance distribution $q_n(\mathbf{x}_n|\mathbf{x}_{n-1}\mathbf{y}_{1:n};\psi)$ includes the support of the target $p(\mathbf{y}_n|\mathbf{x}_n;\theta)p(\mathbf{x}_n|\mathbf{x}_{n-1};\theta)$. This follows from the weight recursion (9);
- (vi) For all iterations $n > 0$, the incremental importance weights are bounded;
- (vii) The initial sets of functions C_0, A_0 , and W_0 are proper sets;
- (viii) For all iterations $n > 0$, $\int_X p(\mathbf{y}_n|\mathbf{x}_n;\theta)p(\mathbf{x}_n|\mathbf{x}_{n-1};\theta)d\mathbf{x}_n > 0$ for all $\mathbf{x}_n \in X$.

Before proving consistency and asymptotic normality for a particle filter, Douc and Moulines (2008) (see also Cappé et al. (2005)) prove preliminary theorems showing that if one starts with a sample $\left\{\mathbf{x}_{0:n-1}^{(i)}, w_{n-1}^{(i)}\right\}_{i=1}^N$ that produces a consistent and asymptotically normal estimator for a function f , then one iteration of the sampling and resampling operations produce a new sample $\left\{\mathbf{x}_{0:n}^{(i)}, w_n^{(i)}\right\}_{i=1}^N$ whose estimator is also consistent and asymptotically normal for the function f . These preliminary theorems also govern consistency and asymptotic normality for the SIR algorithm of Rubin (1987). Therefore, an additional assumption is required

- (ix) Estimates produced by the initial particles $\left\{\mathbf{x}_0^{(i)}, w_0^{(i)}\right\}_{i=1}^N$ are consistent and asymptotically normal for a function $f \in C_0$ and the target $p(\mathbf{x}_0)$.

This assumption makes clear that the first importance sampling iteration needs to produce an estimator that is consistent and asymptotically normal.

Given assumptions (i)-(ix), then by induction Theorem 10 of Douc and Moulines (2008) states that the estimator computed from the samples produced by the particle filter is consistent at

²Douc and Moulines (2008) define a *proper set* of functions C as: (i) for any functions f and g in C and real numbers α and β , $\alpha f + \beta g \in C$; (ii) if $g \in C$ and f is measurable with $|f| \leq |g|$, then $|f| \in C$; (iii) the constant function belongs to C .

iteration n for any function $f \in C_n$ meaning that as $N \rightarrow \infty$

$$\begin{aligned} \sum_{i=1}^N \hat{w}_n^{(i)} f(\mathbf{x}_{0:n}) &\xrightarrow{p} \mathbb{E}_q[f(\mathbf{x}_{0:n})], \\ \max_{1 \leq i \leq N} w_n^{(i)} &\xrightarrow{p} 0. \end{aligned} \quad (37)$$

They are asymptotically normal for any function $f \in A_n$ and $\gamma \in W_n$ meaning that as $N \rightarrow \infty$

$$\begin{aligned} \sqrt{N} \left[\sum_{i=1}^N \hat{w}_n^{(i)} f(\mathbf{x}_n) - \mathbb{E}_q[f(\mathbf{x}_n)] \right] &\xrightarrow{d} \mathcal{N}(0, \sigma_n^2), \\ N \sum_{i=1}^N \left(\hat{w}_n^{(i)} \right)^2 f(\mathbf{x}_{0:n}) &\xrightarrow{p} \gamma, \\ \sqrt{N} \max_{1 \leq i \leq N} w_n^{(i)} &\xrightarrow{p} 0. \end{aligned} \quad (38)$$

Due to the recursive nature of the algorithm, both the asymptotic variance σ_n^2 in (38) and the sets of functions for which these results will hold are determined recursively.³ The class of functions are restricted to ensure the particle filters estimator is bounded.⁴ Integrals approximated at iteration n are functions of integrals approximated at previous iterations. Loosely speaking, the regularity conditions on the set of functions and on the model ensure that integrals from previous iterations remain well-defined in order to make integrals at the current iteration well-defined. Chopin (2004), Künsch (2005), Cappé et al. (2005), Douc, Moulines, and Olsson (2007), and Cornebise et al. (2008) provide further discussion on these conditions. Chapter 11 of Del Moral (2004) also provides consistency and asymptotic normality results for the particle filters' estimator of the likelihood $p(\mathbf{y}_n | \mathbf{y}_{1:n-1}; \theta)$ of the state space model.

Under the simplifying assumptions that particles are resampled each period using multinomial resampling, Johansen and Doucet (2008) show that it is possible to write the asymptotic variance expression explicitly for the SISR and APF algorithms. Their expression for the SISR

³The function γ is a term within the asymptotic variance σ_n^2 . Exact expressions for the asymptotic variance recursion can be found in Douc and Moulines (2008).

⁴The set C_n is the space L^1 of integrable functions with respect to the joint filtering distribution at time n . The sets A_n and W_n are defined recursively. W_n is the space L^1 of integrable functions with respect to the joint filtering distribution at time n . The set also includes those functions $w_{n-1}^2 |f|$ which are integrable with respect to the time $n-1$ importance distribution. A_n is the space L^2 of square integrable functions with respect to the joint filtering distribution at time n . The set also includes those functions $f \in A_{n-1}$ that were integrable with respect to the joint target distribution at time $n-1$ and those functions $w_{n-1}^2 f^2 \in W_{n-1}$ integrable with respect to the time $n-1$ importance distribution.

algorithm is given by

$$\begin{aligned}
\sigma_n^2 = & \int \frac{p(\mathbf{x}_{0:n}|\mathbf{y}_{1:n})^2}{p(\mathbf{x}_{0:n-1}|\mathbf{y}_{1:n-1})q_n(\mathbf{x}_n|\mathbf{x}_{n-1},\mathbf{y}_n)} (f(\mathbf{x}_{0:n}) - \mathbb{E}_q[f(\mathbf{x}_{0:n})])^2 d\mathbf{x}_{0:n} \\
& + \sum_{k=1}^{n-1} \int \frac{p(\mathbf{x}_{0:k}|\mathbf{y}_{1:n})^2}{p(\mathbf{x}_{0:k-1}|\mathbf{y}_{1:k-1})q_k(\mathbf{x}_k|\mathbf{x}_{k-1},\mathbf{y}_{1:n})} \\
& \left(\int f(\mathbf{x}_{0:n})p(\mathbf{x}_{k+1:n}|\mathbf{y}_{k+1:n},\mathbf{x}_k)d\mathbf{x}_{k+1:n} - \mathbb{E}_q[f(\mathbf{x}_{0:n})] \right)^2 d\mathbf{x}_{0:k} \\
& + \int \frac{p(\mathbf{x}_0)^2}{q_0(\mathbf{x}_0)} \left(\int f(\mathbf{x}_{0:n})p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n},\mathbf{x}_0)d\mathbf{x}_{1:n} - \mathbb{E}_q[f(\mathbf{x}_{0:n})] \right)^2 d\mathbf{x}_0. \quad (39)
\end{aligned}$$

where the parameter vectors θ and ψ have been omitted from the notation. The first integral on the right hand side is the asymptotic variance of a standard importance sampling estimator for the joint distribution $p(\mathbf{x}_{0:n}|\mathbf{y}_{1:n})$ using as importance distribution $p(\mathbf{x}_{0:n-1}|\mathbf{y}_{1:n-1})q_n(\mathbf{x}_n|\mathbf{x}_{n-1},\mathbf{y}_n)$. The importance distribution is equivalent to (6) because in theory after the resampling step particles that form the importance distribution $q_{0:n-1}(\mathbf{x}_{0:n-1}|\mathbf{y}_{1:n-1})$ have distribution roughly equal to $p(\mathbf{x}_{0:n-1}|\mathbf{y}_{1:n-1})$. However, this condition only holds approximately in practice. Particles that form the importance distribution $q_{0:n-1}(\mathbf{x}_{0:n-1}|\mathbf{y}_{1:n-1})$ were originally simulated in previous periods using less information. For example in a standard particle filter, particle locations used to approximate the marginal $p(\mathbf{x}_{n-k}|\mathbf{y}_{1:n-1})$ were simulated k -periods ago and have had their locations fixed since that period. The remaining terms in the asymptotic variance express the fact that these earlier quantities are approximated and not computed exactly. Expressions for the asymptotic variance become more complicated than (39) when using different resampling schemes and when resampling at random times via measures like the ESS.

Another important theoretical result for particle filters is that the asymptotic variance in the CLT can be proven to remain bounded over time. Bounds on the asymptotic variance have been obtained by several authors. These results generally require additional assumptions on the ergodic properties of the transition equation within the state space model. Künsch (2005), Chapter 9 of Cappé et al. (2005), and Douc et al. (2007) discuss these results for different types of particle filtering algorithms. Results that bound the asymptotic variance for the particle filters' approximation of the marginal $p(\mathbf{x}_n|\mathbf{y}_{1:n})$ are the most important. This is because they contrast sharply with other methods that try to approximate the filtering recursions in Section 2.1, i.e. regular importance sampling or deterministic methods such as the extended or unscented Kalman filters. As noted by Künsch (2001), the approximation error for other algorithms will generally accumulate asymptotically and the algorithms' estimates may diverge from the true value as more observations are included.

3.2 Some additional references and comments

Del Moral (2004) includes more advanced coverage of particle systems including probabilistic properties other than consistency and asymptotic normality. More recent work addressing theoretical aspects of SMC since the publication of this book are Douc et al. (2007), Douc and Moulines (2008), Bercu et al. (2008b) and Bercu et al. (2008a) which contain additional references.

4 Recent Developments in Sequential Monte Carlo

This section covers two more recent developments that extend SMC outside the context of traditional particle filtering. In the first extension, researchers working in Monte Carlo methods recognized that particle filters could be used to simulate from sequences of distributions other than the filtering distributions defined by a state space model. These methods are particularly applicable to Bayesian inference problems because they provide an alternative to MCMC for simulating from complex distributions. They can also be applied to models for cross-sectional data. Sections 4.1-4.3 review this research in detail. The second extension of standard particle filtering uses the particle filter to approximate the likelihood function within a standard Metropolis-Hastings algorithm. This type of algorithm is currently being used in the macroeconomics literature on Bayesian estimation of DSGE models. Section 4.5 covers this material.

4.1 SMC samplers and Population Monte Carlo

Recognizing that the particles form a collection of interacting Markov chains on a sequence of general state spaces is the key to building other types of SMC algorithms. Leading references in this field include Gilks and Berzuini (2001), Chopin (2002), Liang (2002), Cappé et al. (2004), and Carvalho et al. (2008). Del Moral, Doucet, and Jasra (2006b, 2006) and Jasra et al. (2008) built a framework titled *SMC samplers* that encompasses a number of the algorithms in the literature. A special case of an SMC sampler that is simpler to implement and conceptually easier to understand are the Population Monte Carlo (PMC) algorithms developed in a series of papers by Cappé et al. (2004), Celeux et al. (2006), Douc et al. (2007a), Douc et al. (2007b).

Research in this area of Monte Carlo methods is on-going. There are several key themes in this research: (i) an emphasis on building adaptive Monte Carlo algorithms that learn from their previous draws; (ii) understanding the practical circumstances where allowing the Markov chains to interact is beneficial relative to MCMC; (iii) developing the necessary limit theory to justify the methods in practice.

To connect a particle filter with an SMC sampler, it is necessary to formalize some ideas

from section 2. At time n of a particle filter, a particle (the state variable) $\mathbf{x}_n$ takes values in a measurable space $(X_n, \mathcal{X}_n)$ where at each iteration the state space X_n is simply $\mathbf{R}^m$ and the σ -algebra $\mathcal{X}_n$ is the corresponding Borel σ -algebra $\mathcal{B}(\mathbf{R}^m)$. The sequence of “target” densities are the marginal filtering densities $\{p(\mathbf{x}_n|\mathbf{y}_{1:n})\}_{n=1}^T$ defined by the state space model. This implies that the sequence of joint smoothing densities $\{p(\mathbf{x}_{0:n}|\mathbf{y}_{1:n})\}_{n=1}^T$ are defined on a sequence of product spaces $X_{0:n} = X_0 \times X_1 \times \cdots \times X_n = \mathbf{R}^{\otimes(m+1)}$. The dimensionality of the larger state spaces $X_{0:n}$ is clearly increasing from one period to the next as new observations are added. The particle filter approximates the joint densities, although the marginal densities are actually of primary interest.

An SMC sampler uses the same idea but the sequence of “low-dimensional” measurable spaces $(X_n, \mathcal{X}_n)$ does not have to be the same at each iteration as in a standard particle filter nor are they determined by the model. This sequence is instead chosen by the user when they design an algorithm for a specific application. This added flexibility requires some additional, more generic notation not needed for standard particle filters. In particular, it is helpful to index both each target density and the particle by its iteration n . The iteration number n in the sequence is a counter that may or may not represent calendar time. The random variable or particle $\mathbf{x}_n$ is no longer restricted to denote a state variable in a state space model as in Section 2. It is simply a quantity of interest with its interpretation depending upon the application. For example, it can represent a parameter vector, a sequence of latent variables, or a combination of the two. Let $\{p_n(\mathbf{x}_n)\}_{n=1}^J$ be a sequence of probability distributions defined on a sequence of measurable spaces $\{(X_n, \mathcal{X}_n)\}_{n=1}^J$. The number of observations in the researcher’s sample (not necessarily a time series) is denoted by T while J is the number of distributions in the sequence. Each density in the sequence is defined as

$$p_n(\mathbf{x}_n) = \frac{\gamma_n(\mathbf{x}_n)}{Z_n}, \quad (40)$$

where $\gamma_n(\mathbf{x}_n)$ is the unnormalized density which can be calculated for any realization of $\mathbf{x}_n$. The normalizing constant Z_n in the denominator of (40) typically includes integrals that cannot be solved analytically.

An SMC sampler begins by drawing N particles $\{\mathbf{x}_1^{(i)}\}_{i=1}^N$ from an initial importance density $q_1(\mathbf{x})$ and reweighting the particles using standard importance weights. Importance weights at the first iteration are

$$w_1 = \frac{\gamma_1(\mathbf{x}_1)}{q_1(\mathbf{x}_1)}, \quad (41)$$

which can be computed explicitly because the user knows the initial importance density $q_1(\mathbf{x}_1)$. Beginning at the second iteration and continuing forward, each particle is sampled from a forward

nonhomogenous Markov transition kernel $\mathbf{x}_n^{(i)} \sim K_n(\mathbf{x}_{n-1}^{(i)}, \cdot)$. This Markov kernel is simply a generalization of the Markovian importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$ within a standard particle filter from Section 2. The marginal distribution of the particles after drawing from the transition kernel K_n is

$$q_n(\mathbf{x}_n) = \int q_{n-1}(d\mathbf{x}_{n-1}) K_n(\mathbf{x}_{n-1}, \mathbf{x}_n). \quad (42)$$

The importance weights at the n -th iteration are the ratio of the target density to the importance density and are given by

$$w_n = \frac{\gamma_n(\mathbf{x}_n)}{q_n(\mathbf{x}_n)}. \quad (43)$$

Unfortunately, the integral in (42) cannot usually be solved analytically for an arbitrary choice of the transition kernel K_n . This makes it impossible to directly calculate the importance weights.

Del Moral et al. (2006b) solve the problem of having to evaluate the unknown importance density $q_n(\mathbf{x}_n)$ to compute importance weights beyond the first iteration by introducing new artificial target densities $p_{1:n}(\mathbf{x}_{1:n})$. The sequence of artificial targets $\{p_{1:n}(\mathbf{x}_{1:n})\}_{n=1}^J$ are defined on the product spaces $X_{1:n} = X_1 \times X_2 \times \cdots \times X_n$ along with their respective product σ -algebra.⁵ The artificial joint densities in an SMC sampler are not of interest in themselves but their introduction allows the importance weights to be computed. An artificial target must be defined up to a normalizing constant

$$p_{1:n}(\mathbf{x}_{1:n}) = \frac{\gamma_{1:n}(\mathbf{x}_{1:n})}{Z_n}, \quad (44)$$

where the new target is intentionally designed to admit $p_n(\mathbf{x}_n)$ as a marginal density. The expanded target is similar to the earlier presentation of the particle filter which operated on the joint smoothing distributions to approximate the marginal filtering distributions. By sampling in a larger space, estimates of the marginal using the particles' locations and importance weights can be computed as a by-product.

Del Moral et al. (2006b) provide a framework for choosing both the artificial target densities $p_{1:n}(\mathbf{x}_{1:n})$ as well as the forward Markov kernels. As in Jarzynski (1997) and Neal (2001), they suggest defining the artificial targets as a sequence of artificial backward Markov kernels $L_n(\mathbf{x}_{n+1}, \mathbf{x}_n)$ which can be written as

$$\gamma_{1:n}(\mathbf{x}_{1:n}) = \gamma_n(\mathbf{x}_n) \prod_{k=1}^{n-1} L_k(\mathbf{x}_{k+1}, \mathbf{x}_k). \quad (45)$$

Given particles $\{w_{n-1}^{(i)}, \mathbf{x}_{1:n-1}^{(i)}\}_{i=1}^N$ that approximate the artificial target $\gamma_{1:n-1}(\mathbf{x}_{1:n-1})$, the next artificial target $\gamma_{1:n}(\mathbf{x}_{1:n})$ can be approximated by sampling from the forward Markov kernel.

⁵In a standard particle filter from Section 2, the joint smoothing densities are analagous to the artificial joint densities described here.

The (unweighted) particles' joint distribution after n transitions is

$$q_{1:n}(\mathbf{x}_{1:n}) = q_1(\mathbf{x}_1) \prod_{j=2}^n K_j(\mathbf{x}_{j-1}, \mathbf{x}_j). \quad (46)$$

Reweightings the particles using the importance weights changes their distribution from $q_{1:n}(\mathbf{x}_{1:n})$ to $p_{1:n}(\mathbf{x}_{1:n})$.

The unnormalized importance weights w_n for the joint distribution are defined as the ratio of the (unnormalized) joint target density (45) to the joint importance density (46) and are given by

$$w_n = \frac{\gamma_{1:n}(\mathbf{x}_{1:n})}{q_{1:n}(\mathbf{x}_{1:n})}. \quad (47)$$

These can be written recursively such that at each iteration one only calculates the incremental importance weights $\tilde{w}_n$ given by

$$w_n = w_{n-1} \tilde{w}_n.$$

where

$$\tilde{w}_n = \frac{\gamma_n(\mathbf{x}_n) L_{n-1}(\mathbf{x}_n, \mathbf{x}_{n-1})}{\gamma_{n-1}(\mathbf{x}_{n-1}) K_n(\mathbf{x}_{n-1}, \mathbf{x}_n)}. \quad (48)$$

Notice the similarities between this recursion and (9). These unnormalized weights then lead to normalized importance weights

$$\hat{w}_n = \frac{w_n}{\sum_{i=1}^N w_n^{(i)}}. \quad (49)$$

Once the normalized importance weights are calculated, estimates of a marginal target distribution can be calculated as

$$\hat{p}_n(\mathbf{x}_n) = \sum_{i=1}^N \hat{w}_n^{(i)} \delta_{\mathbf{x}_n^{(i)}}(\mathbf{x}_n^{(i)}) \approx p_n(\mathbf{x}_n).$$

In addition, estimates of the ratio of normalizing constants can be computed as

$$\frac{\hat{Z}_n}{Z_{n-1}} = \sum_{i=1}^N w_{n-1}^{(i)} \tilde{w}_n^{(i)} \quad (50)$$

If the user chooses an initial distribution where the normalizing constant Z_1 can be calculated, then they obtain an estimate of the normalizing constant for any distribution in the sequence including the final iteration $\hat{Z}_J$. For example, this could be the marginal likelihood in a Bayesian context or the likelihood of a general state space model.

Like the standard particle filter described previously, it is usually not optimal to resample the particles at each iteration of an SMC sampler. Instead, particles should only be resampled

Algorithm 4 Sequential Monte Carlo sampler

At $n = 1$, for $i = 1, \dots, N$

Draw $\mathbf{x}_1^{(i)} \sim q_1(\mathbf{x}_1)$ and set $w_1^{(i)} = \frac{p(\mathbf{x}_1^{(i)})}{q_1(\mathbf{x}_1^{(i)})}$.

For $n = 2, \dots, J$:

(i) For $i = 1, \dots, N$:

a. Draw: $\mathbf{x}_n^{(i)} \sim K_n(\mathbf{x}_{n-1}^{(i)}, \cdot)$.

b. Compute importance weights: $w_n^{(i)} \propto w_{n-1}^{(i)} \tilde{w}_n^{(i)}$.

(ii) For $i = 1, \dots, N$:

Normalize the importance weights: $\hat{w}_n^{(i)} = \frac{w_n^{(i)}}{\sum_{j=1}^N w_n^{(j)}}$.

(iii) Calculate the effective sample size (ESS).

(iv) If $\text{ESS} < \text{threshold}$, resample N particles with probabilities $\left\{ \hat{w}_n^{(i)} \right\}_{i=1}^N$ and

for $i = 1, \dots, N$ set $w_n^{(i)} = \frac{1}{N}$.

when the variance of the importance weights grows and becomes unstable. This can be measured by any of the criterion described in Section 2.4. A standard SMC sampler is given by Algorithm 4.

Although an SMC sampler is simply a particle filter in a more general context, it requires more input and experience from the user. In a standard particle filter, the sequence of target densities (and implicitly the backwards kernels) are already defined for the user by their state space model. This leaves only the choice of the forward Markov kernel (i.e. the importance distribution $q_n(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{y}_n; \psi)$) which is relatively easy to select. Conversely in an SMC sampler, the user will have to define the sequence of target densities and choose the forward and backward Markov kernels. Different choices for the forward and backward Markov kernels also determine how challenging it is to compute the incremental weight (48) in an SMC sampler. Del Moral et al. (2006b) provide suggestions to users for choosing each of these quantities in practice. Although these authors consider many options, the easiest algorithms to implement for practitioners with experience using MCMC will be to choose $K_n(\mathbf{x}_{n-1}, \mathbf{x}_n)$ to be a Gibbs or Metropolis-Hastings kernel. Del Moral et al. (2006a, 2006b) give the equations to calculate the incremental weights (48) when using these kernels, see their paper for details. Many of these only involve evaluating the unnormalized target density (40) as in a standard MCMC algorithm.

4.2 PMC algorithms

A special case of the SMC sampling framework that may be easier to implement in practice are the PMC algorithms developed by Cappé et al. (2004), Douc et al. (2007a), and Douc et al. (2007b), see also Cappé et al. (2008). In these algorithms, the sequence of target densities are the same at each iteration, $p_n(\mathbf{x}_n) = p(\mathbf{x}) \forall n$. The researcher does not need to formally

consider the sequence of measurable spaces $\{(X_n, \mathcal{X}_n)\}_{n=1}^J$ nor the backward Markov kernels because these are implicitly determined by the target density.

The algorithm begins by drawing N particles from an initial importance density $\mathbf{x}_1 \sim q_1(\mathbf{x}_1)$ and computing importance weights

$$w_1 = \frac{p_1(\mathbf{x}_1)}{q_1(\mathbf{x}_1)}.$$

At each future iteration, particles are moved according to a forward Markov transition kernel. These authors place more structure on the forward Markov kernels $K_n(\mathbf{x}_{n-1}, \mathbf{x}_n)$ than in the more generic SMC sampler. They suggest choosing the forward Markov kernel as a mixture of D Markov kernels $q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_{n,j})$ for $j = 1, \dots, D$ with mixture weights $\alpha_{n,j}$ where $\sum_{j=1}^D \alpha_{n,j} = 1$. For example, the D different kernels in the mixture might be a collection of $D = 4$ multivariate Student t distributions with different covariance matrices and degrees of freedom for the parameters $\psi_{n,j}$. Formally, the importance density of a PMC algorithm at iteration n is defined as

$$K_n(\mathbf{x}_{n-1}, \mathbf{x}_n) = \sum_{j=1}^D \alpha_{n,j} q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_{n,j}). \quad (51)$$

At iteration n , each particle is moved with one of the D kernels $q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_{n,j})$ in the mixture, which is selected at random by drawing N indicators $\{d_n^{(i)}\}_{i=1}^N$ from a multinomial distribution using the mixture weights as probabilities. Note that the particles from the previous iteration will usually serve as the location parameter for $q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_{n,j})$. The importance weights at iteration n are then defined as the target distribution $p(\mathbf{x}_n)$ divided by the relevant component of the mixture importance density (51) that is

$$w_n^{(i)} \propto \frac{p(\mathbf{x}_n^{(i)})}{\alpha_{n,d_n^{(i)}}^{(i)} q_{n,d_n^{(i)}}(\mathbf{x}_{n-1}^{(i)}, \mathbf{x}_n^{(i)}; \psi_{d_n^{(i)}}^{(i)})}. \quad (52)$$

where the mixture indicator $d_n^{(i)}$ is used to compute the weight. Particles are then resampled using these normalized importance weights to form a new mixture importance density for the next iteration of the algorithm. A Rao-Blackwellized-version of the PMC algorithm is given as Algorithm 5. Notice that the PMC algorithm is a special case of an SMC sampler for a specific forward/backward Markov kernel.

The reason iterations are introduced into the importance sampling algorithm is to adapt or “tune” the importance distribution over time. The resampling mechanism in the algorithm is only one form of adaptation. There are other ways that one could consider adapting the importance density. The original algorithm in Cappé et al. (2004) and also Robert and Casella (2004, p. 562) was intended to adapt the mixture weights $\{\alpha_{n,j}\}_{j=1}^D$ over iterations while the

Algorithm 5 Rao-Blackwellized Population Monte Carlo sampler with adaption

At $n = 1$, for $i = 1, \dots, N$

- (i) Draw $\mathbf{x}_1^{(i)} \sim q_1(\mathbf{x}_1)$, set $w_1^{(i)} = \frac{p(\mathbf{x}_1^{(i)})}{q_1(\mathbf{x}_1^{(i)})}$, and
- (ii) For $j = 1, \dots, D$, set $\alpha_{2,j} = 1/D$.

For $n = 2, \dots, J$:

- (i) For $i = 1, \dots, N$:
 - a. Draw an indicator $d_n^{(i)}$ from $\mathcal{M}(D, \{\alpha_{n,j}\}_{j=1}^D)$.
 - b. Conditional on the indicator, draw $\mathbf{x}_n^{(i)} \sim q_{n,d_n^{(i)}}(\mathbf{x}_{n-1}^{(i)}, \mathbf{x}_n; \psi_{n,d_n^{(i)}})$.
 - c. Compute importance weights: $w_n^{(i)} \propto \frac{p(\mathbf{x}_n^{(i)})}{\sum_{j=1}^D \alpha_{n,j} q_{n,j}(\mathbf{x}_{n-1}^{(i)}, \mathbf{x}_n^{(i)}; \psi_{n,j})}$.
 - (ii) For $i = 1, \dots, N$:

Normalize the importance weights: $\hat{w}_n^{(i)} = \frac{w_n^{(i)}}{\sum_{j=1}^N w_n^{(j)}}$.
 - (iii) Using the particles, adapt the mixture weights $\{\alpha_{n+1,j}\}_{j=1}^D$ and/or parameters $\{\psi_{n+1,j}\}_{j=1}^D$.
 - (iv) Resample N particles with probabilities $\{\hat{w}_n^{(i)}\}_{i=1}^N$.
-

distributions in the mixture $q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_{n,j})$ and their parameters $\psi_{n,j}$ remained the same. The mixture weights are updated using the indicators $\{d_n^{(i)}\}_{i=1}^N$ and the normalized importance weights to compute $\{\alpha_{n+1,j}\}_{j=1}^D$. As the algorithm progresses, more particles are hopefully simulated from the component in the mixture that is most like the target density. Using the expected Kullback-Leibler (K-L) distance (where the expectation is taken with respect to the target) between the target density $p(\mathbf{x}_n)$ and the importance density (51) as a criterion, Douc et al. (2007a) proved that as $N \rightarrow \infty$ the mixture weights will only converge to minimize the expected K-L distance if the importance weights are Rao-Blackwellized. In other words, in an algorithm where the mixture probabilities $\{\alpha_{n,j}\}_{j=1}^D$ are adapted based upon the past simulations, improvement over a standard IS algorithm only occurs if one computes the Rao-Blackwellized importance weights

$$w_n^{(i)} \propto \frac{p(\mathbf{x}_n^{(i)})}{\sum_{j=1}^D \alpha_{n,j} q_{n,j}(\mathbf{x}_{n-1}^{(i)}, \mathbf{x}_n^{(i)}; \psi_j)}, \quad (53)$$

instead of (52). Their theoretical result is for a PMC algorithm where the kernels $q_{n,j}(\mathbf{x}_{n-1}, \mathbf{x}_n; \psi_j)$ and the parameters $\psi_{n,j}$ for $j = 1, \dots, D$ remain the same over iterations. In a related paper, Douc et al. (2007b) considered minimizing the asymptotic variance of the IS algorithm for a specific function of interest instead of the expected K-L distance.

Cappé et al. (2008) consider improving the algorithm of Douc et al. (2007a) by adapting not only the mixture probabilities $\{\alpha_{n,j}\}_{j=1}^D$ but also the parameters $\{\psi_{n,j}\}_{j=1}^D$ of the distributions within the mixture. They suggest using the expected K-L distance as a criterion for updating the parameters at each iteration. They provide closed-form formulas to update the parameters

$\{\psi_{n,j}\}_{j=1}^D$ when the importance densities within the mixture are either Normal or Student t distributions. The formulas are similar to those one would use in an Expectation-Maximization (EM) algorithm for these distributions. Although similar in structure to an SMC algorithm, the adaptive-IS algorithm described in this last paper is not strictly an SMC algorithm because it does not use Markov kernels within the mixture distribution. Instead of using the past particle's as the location parameter within each of the distributions in the mixture, these authors update the mean of the Normal/Student t distribution over iterations.

4.3 Application # 2: Monetary DSGE Model

In this section, the SMC and PMC algorithms are applied to a simple DSGE model and are compared to MCMC and importance sampling on simulated data. The goal is to improve upon standard importance sampling algorithms using some of the adaptive methods discussed above. The model is a monetary DSGE model from An and Schorfheide (2007). The main equations of the model are an Euler equation describing the evolution of output y_n , a Phillips curve for inflation π_n , and a Taylor Rule for the short-term interest rate R_n along with specifications for the exogenous variables driving the economy. These variables are a shock z_n to the growth rate of technology, a shock to preferences g_n , and a serially uncorrelated monetary policy shock $\epsilon_{R,n}$, while c_n denotes consumption. This forms a linear rational expectations model for the state vector $[\hat{y}_n, \hat{c}_n, \hat{p}_n, \hat{R}_n, \epsilon_{R,n}, \epsilon_{g,n}, \epsilon_{z,n}]'$. Hats over the variables denote percentage deviations from their steady state values. The log-linearized model is given by

$$\begin{aligned}
\hat{y}_n &= E_n[\hat{y}_{n+1}] + \hat{g}_n - E[\hat{g}_{n+1}] - \frac{1}{\tau}(\hat{R}_n - E_n[\pi_{n+1}] - E[\hat{z}_{n+1}]), \\
\hat{\pi}_n &= \beta E_n[\hat{\pi}_{n+1}] + \kappa(\hat{y}_n - \hat{g}_n), \\
\hat{R}_n &= \rho_R \hat{R}_{n-1} + (1 - \rho_R)\psi_1 \hat{\pi}_n + (1 - \rho_R)\psi_2(\hat{y}_n - \hat{g}_n) + \epsilon_{R,n}, \\
\hat{c}_n &= \hat{y}_n - \hat{g}_n, \\
\hat{g}_n &= \rho_g \hat{g}_{n-1} + \epsilon_{g,n}, & \epsilon_{g,n} &\sim \mathcal{N}(0, \sigma_g^2), \\
\hat{z}_n &= \rho_z \hat{z}_{n-1} + \epsilon_{z,n}, & \epsilon_{z,n} &\sim \mathcal{N}(0, \sigma_z^2), \\
\epsilon_{R,n} &= \epsilon_{R,n}, & \epsilon_{R,n} &\sim \mathcal{N}(0, \sigma_R^2).
\end{aligned}$$

where $\kappa = \frac{\tau(1-\nu)}{\nu\pi^2\phi}$. More details on the interpretation of the structural parameters of the model can be found in Section 2 of An and Schorfheide (2007).

These authors compare first and second-order solution methods of the above system. A first-order solution is implemented here, which results in a linear transition equation for the state

vector. The measurement equation for the model is

$$\begin{aligned} ygr_n &= \gamma^{(Q)} + 100(\hat{y}_n - \hat{y}_{n-1} + \hat{z}_n), \\ infl_n &= \pi^{(A)} + 400\hat{\pi}^n, \\ int_n &= \pi^{(A)} + r^{(A)} + 4\gamma^{(Q)} + 400\hat{R}_n. \end{aligned}$$

The likelihood for the model can be computed by the Kalman filter. Only determinate solutions of the model are considered. The same priors for the parameters are used as in Table 2 on page 129 of An and Schorfheide (2007), which for convenience are given in Table 2 here. In this example, 80 observations were simulated from the model, where the parameters of the DGP are also given in Table 2.⁶

	DGP	Dist.	Prior		DGP	Dist.	Prior
τ	2.00	gamma	2.00 (0.50)	$r^{(A)}$	0.40	gamma	0.50 (0.50)
κ	0.15	gamma	0.20 (0.10)	$\pi^{(A)}$	4.00	gamma	7.00 (2.00)
ψ_1	1.50	gamma	1.50 (0.25)	$\gamma^{(Q)}$	0.50	normal	0.40 (0.20)
ψ_2	1.00	gamma	0.50 (0.25)	$100\sigma_R$	0.20	inv. gamma	0.11 (0.08)
ρ_R	0.60	beta	0.50 (0.20)	$100\sigma_g$	0.80	inv. gamma	0.66 (0.47)
ρ_g	0.95	beta	0.80 (0.10)	$100\sigma_z$	0.45	inv. gamma	0.17 (0.12)
ρ_z	0.65	beta	0.66 (0.15)				

Table 1: *Parameter settings used in the DGP as well as prior means and standard deviations.*

⁶The MCMC algorithm used as a comparison was a standard random-walk Metropolis-Hastings (M-H) algorithm as described in An and Schorfheide (2007). The posterior mode and the negative inverse Hessian at the mode were first found by optimization. The latter was used as the covariance matrix to generate proposals for the random-walk with normally distributed increments in the M-H algorithm. The covariance matrix was scaled to accept roughly 40% of the draws. The method of Chib and Jeliazkov (2001) was used to compute the marginal likelihood. An importance sampling algorithm was also implemented using as importance distribution a Student t distribution with 3 degrees of freedom, mean equal to the posterior mode, and covariance matrix equal to the negative inverse Hessian at the mode. The log marginal likelihood can be computed from the importance sampling weights; see, e.g. Geweke(2005, p. 256). The number of draws for each of the algorithms was selected so that they were run for roughly the same computational time. After a burn-in of 10,000 iterations, the MCMC algorithm was run for 500,000 draws and retained every 100 draws due to their correlation. The importance sampling algorithm was run for 500,000 draws.

4.3.1 SMC sampler for a DSGE model

The SMC sampler used all of the observed data $\mathbf{y}_{1:T}$ at each iteration with the densities in the sequence defined up to a normalizing constant as

$$p_n(\mathbf{x}_n) \propto [p(\mathbf{y}_{1:T}|\theta_n) p(\theta_n)]^{\zeta_n} [\mu(\theta_n)]^{1-\zeta_n}. \quad (54)$$

In this example, a particle at iteration n is a parameter vector in existence at that iteration, i.e. $\mathbf{x}_n = \theta_n$.⁷ Indexing the parameter vector θ_n by n in this notation does not mean that θ is time-varying. The densities in the sequence (54) differed only based on a simulated tempering sequence $0 = \zeta_1 < \dots < \zeta_J = 1$ known as a cooling schedule; see, e.g. Liu (2001, p. 210). Simulated tempering raises a function to a power less than one. At iteration $n = 1$, $\zeta_1 = 0$ and the target density is $\mu(\theta_1)$. As ζ_n gradually gets larger, the densities within the sequence get closer to the (unnormalized) posterior density and will be equal to the posterior density at the last iteration when $\zeta_J = 1$. In this example, the initial distribution was the same as the IS algorithm. Therefore, particles were initialized by drawing from a Student t distribution with 3 degrees of freedom, mean equal to the posterior mode, and covariance matrix equal to the negative inverse Hessian at the mode.

The forward Markov kernel $K_n(\mathbf{x}_{n-1}, \mathbf{x}_n)$ was a random-walk M-H kernel where the random walk has a normal distribution. Del Moral et al. (2006b) show that for an M-H transition kernel, the backwards Markov kernel should be selected as

$$L_n(\mathbf{x}_n, \mathbf{x}_{n-1}) = \frac{p_n(\mathbf{x}_{n-1}) K_n(\mathbf{x}_{n-1}, \mathbf{x}_n)}{p_n(\mathbf{x}_n)}. \quad (55)$$

Plugging this expression into (48), the incremental weights are

$$\tilde{w}_n = \frac{\gamma_n(\mathbf{x}_{n-1})}{\gamma_{n-1}(\mathbf{x}_{n-1})}, \quad (56)$$

which in this application with target density (54) reduces to

$$\tilde{w}_n = \frac{[p(\mathbf{y}_{1:T}|\theta_{n-1}) p(\theta_{n-1})]^{\zeta_n} [\mu(\theta_{n-1})]^{1-\zeta_n}}{[p(\mathbf{y}_{1:T}|\theta_{n-1}) p(\theta_{n-1})]^{\zeta_{n-1}} [\mu(\theta_{n-1})]^{1-\zeta_{n-1}}}. \quad (57)$$

Without the tempering parameters, this expression is similar to evaluating the acceptance probability in an independence M-H algorithm with $\mu(\theta_1)$ as the proposal distribution. However, the incremental weights are independent of the time n particles. This means that (57) can be computed, and the old particles resampled before drawing new particles. Finally, contributions to the log-marginal likelihood can be calculated using (50).

⁷An alternative sequence of densities that could be considered are $p_n(\mathbf{x}_n) \propto (\mathbf{y}_{1:n}|\theta_n) p(\theta_n)$ with $\mathbf{x}_n = \theta_n$. Once again, θ is not time-varying. By adding an additional observation at each iteration, one is performing sequential Bayesian estimation.

For the sequence of cooling parameters $\{\zeta_j\}_{j=1}^J$, a linear schedule was chosen with differentials $\zeta_n - \zeta_{n-1} = 1/(J-1)$ and $\zeta_1 = 0$, $\zeta_J = 1$. Many applications of simulated tempering have the tempering parameters changing slower at the beginning and then gradually increasing at a faster rate. The complexity of the tempering schedule and most importantly the number of densities J will depend on how “close” the initial density $p_1 = \mu(\theta_1)$ is from the final target density $p_J \propto p(\mathbf{y}_{1:T}|\theta_J)p(\theta_J)$. The number of densities may need to be quite large if the user cannot find a reasonable initial approximation. In this example, the algorithm was run with $J = 50$ densities, $N = 10,000$ particles, and systematic resampling.

4.3.2 PMC sampler for a DSGE model

The target density for the Rao-Blackwellized PMC sampler is the same at each iteration

$$p_n(\mathbf{x}_n) \propto p(\mathbf{y}_{1:T}|\theta_n)p(\theta_n). \quad (58)$$

For the forward Markov kernel, I follow Cappé et al. (2008) and select a mixture of $D = 4$ Student t distributions. The mean $\mu_{n,j}$, and covariance matrices $\Sigma_{n,j}$ within each component $q_{n,j}(\mathbf{x}_n; \psi_{n,j}) = \mathcal{T}(\mathbf{x}_n; \nu_j, \mu_{n,j}, \Sigma_{n,j})$ of the mixture are updated at each iteration. The four degrees of freedom parameters were kept fixed over iterations at $\nu_j = (3, 7, 15, 50)'$, respectively. For convenience, the equations for updating the parameters and probabilities from Cappé et al. (2008) are repeated here. They are

$$\alpha_{n+1,j} = \sum_{i=1}^N \hat{w}_n^{(i)} \rho_j^{(i)}, \quad (59)$$

$$\mu_{n+1,j} = \frac{\sum_{i=1}^N \hat{w}_n^{(i)} \rho_j^{(i)} \kappa_j^{(i)} \mathbf{x}_n^{(i)}}{\sum_{i=1}^N \hat{w}_n^{(i)} \rho_j^{(i)} \kappa_j^{(i)}}, \quad (60)$$

$$\Sigma_{n+1,j} = \frac{\sum_{i=1}^N \hat{w}_n^{(i)} \rho_j^{(i)} \kappa_j^{(i)} (\mathbf{x}_n^{(i)} - \mu_{n,j})(\mathbf{x}_n^{(i)} - \mu_{n,j})'}{\sum_{i=1}^N \hat{w}_n^{(i)} \rho_j^{(i)}}, \quad (61)$$

where

$$\rho_j^{(i)} = \frac{\alpha_{n,j} \mathcal{T}(\mathbf{x}_n^{(i)}; \nu_j, \mu_{n,j}, \Sigma_{n,j})}{\sum_{k=1}^D \alpha_{n,k} \mathcal{T}(\mathbf{x}_n^{(i)}; \nu_k, \mu_{n,k}, \Sigma_{n,k})}, \quad (62)$$

$$\kappa_j^{(i)} = \frac{\nu_j + \dim(\mathbf{x}_n)}{\nu_j + (\mathbf{x}_n^{(i)} - \mu_{n,j})' \Sigma_{n,j}^{-1} (\mathbf{x}_n^{(i)} - \mu_{n,j})}. \quad (63)$$

In these expressions, $\hat{w}_n^{(i)}$ are the normalized importance weights. The algorithm was initialized by drawing particles from the same Student t importance density as in the importance sampling algorithm. The PMC algorithm is not intended to be run for many iterations because the parameters (59)-(61) will eventually converge and stop adapting. The algorithm was run for $J = 10$ iterations and $N = 50,000$ particles.

4.3.3 Results

Results from each of the four algorithms are presented in Table 2. Point estimates from each of the methods are reasonably close, including their estimates of the log-marginal likelihood. Like the burn-in phase of an MCMC algorithm, the SMC and PMC algorithms require taking an initial number of draws before storing any draws for an estimate. This means that the total number of draws used to construct the estimates for each algorithm are not the same. The tradeoff between a standard importance sampling algorithm and these more elaborate algorithms comes in a potential increase in the stability of the estimator at the expense of more time spent coding. Stability of the estimator can be measured by the number of effective draws. Table 2 contains an estimate of the effective sample size (ESS) from each method, see equation (25). As both the PMC and SMC algorithms are initialized with the same importance density as the standard importance sampling algorithm, it is interesting to see a substantial increase in the stability of the estimators as measured by the ESS. Draws from the SMC sampler are close to being equally weighted.

The relative numerical efficiency (RNE) of the MCMC, IS, and PMC algorithms can be computed by standard methods, e.g. Geweke (2005, p. 114, 149). These are reported in Table 2. The RNE from the PMC algorithm offers a substantial improvement over the standard IS algorithm. The reason statistics like the ESS and coefficient of variation (CV) are used in the particle filtering literature instead of the RNE is because expressions for the asymptotic variance in the CLT are complicated. They are also dependent on the algorithm that is implemented (see Section 3). This makes it challenging to compute the RNE in practice. Particles were only resampled twice over the 50 iterations of the SMC sampler, at iterations $n = 20$ and $n = 38$ respectively. The correlations between the particles were reasonably low. Therefore for illustrative purposes, RNE estimates for the SMC sampler assuming the sample is uncorrelated are reported although this is not formally correct. Particles within a standard particle filter applied to a state space model will typically resample much more often. It is not recommended that users regularly report this statistic. However, given that it only resampled twice with the last resample occurring at iteration $n = 38$, the SMC sampler's performance is similar to or slightly better than the MCMC algorithm. This is not surprising given that both algorithms use a similar random walk Metropolis Markov kernel within them.

4.4 Some additional references and comments

Gilks and Berzuini (2001) and Chopin (2002) both considered sequential Bayesian estimation under the SMC approach. Chopin (2002) introduced the concept for static parameter estimation;

	MCMC	RNE	IS	RNE	SMC	RNE	PMC	RNE
τ	2.092 (0.460)	0.839	2.082 (0.498)	0.010	2.106 (0.487)	0.864	2.105 (0.489)	0.233
κ	0.170 (0.041)	0.592	0.170 (0.041)	0.003	0.173 (0.045)	0.672	0.172 (0.044)	0.230
ψ_1	1.520 (0.242)	0.852	1.540 (0.248)	0.010	1.526 (0.249)	0.902	1.527 (0.252)	0.319
ψ_2	0.475 (0.197)	0.717	0.495 (0.236)	0.006	0.511 (0.233)	0.771	0.478 (0.206)	0.245
ρ_R	0.437 (0.079)	0.777	0.441 (0.081)	0.005	0.429 (0.085)	0.792	0.436 (0.082)	0.255
ρ_g	0.830 (0.070)	0.724	0.835 (0.069)	0.005	0.837 (0.069)	0.878	0.831 (0.069)	0.264
ρ_z	0.516 (0.048)	0.716	0.519 (0.049)	0.003	0.520 (0.053)	0.747	0.516 (0.049)	0.176
$r^{(A)}$	0.578 (0.328)	0.947	0.568 (0.307)	0.004	0.584 (0.335)	0.882	0.578 (0.332)	0.216
$\pi^{(A)}$	3.994 (0.033)	0.761	3.994 (0.031)	0.003	3.994 (0.034)	0.795	3.995 (0.034)	0.145
$\gamma^{(Q)}$	0.414 (0.094)	0.860	0.415 (0.092)	0.003	0.413 (0.097)	0.851	0.413 (0.095)	0.189
$100\sigma_R$	0.212 (0.021)	0.854	0.214 (0.022)	0.007	0.214 (0.022)	0.858	0.212 (0.022)	0.261
$100\sigma_g$	0.792 (0.062)	0.756	0.799 (0.067)	0.008	0.800 (0.066)	0.865	0.795 (0.066)	0.325
$100\sigma_z$	0.558 (0.078)	0.699	0.554 (0.089)	0.002	0.554 (0.086)	0.766	0.558 (0.083)	0.146
log-marg like	-216.38		-217.15		-217.17		-217.11	
ESS	-		2173.1		8625.6		14915.2	
time (sec)	498.4		494.4		497.5		540.1	

Table 2: *Posterior estimates from the monetary DSGE model in An Schorfheide (2007) based on 80 observations simulated from the model. Point estimates of the marginals are the mean with standard deviations given in parenthesis beneath them. ESS denotes the effective sample size. The RNE is reported in the column to the right of the estimator.*

his applications also included cross-sectional data. Chopin and Pelgrin (2004) and Chopin (2007) estimate discrete-state HMM models with the unique ability to estimate the number of states in the HMM as the data-set gets processed. Johansen et al. (2006) consider applications to rare event simulation. Carvalho et al. (2008) focus on learning the parameters sequentially in time. Jasra et al. (2008) use adaptive SMC samplers to estimate Lévy-driven SV models. Del Moral et al. (2008) design new types of adaptive schemes that determine the tuning parameters of the algorithm internally.

The theoretical analysis of adaptive SMC algorithms is a current area of research. Del Moral et al. (2006b) provide a LLN and a CLT for their SMC sampler under some simplifying assumptions. Papers from the theoretical probability literature studying these methods under the name of self-interacting Metropolis-Hastings algorithms or alternatively non-linear MCMC are Del Moral and Doucet (2004) and Chapter 5 of Del Moral (2004). More recent work includes Bercu et al. (2008b) and the papers cited therein.

4.5 Particle Filters within Metropolis-Hastings algorithms

The discussion in Sections 4.1-4.4 assumed that SMC algorithms are used as an alternative to MCMC. Another possibility is to use an SMC algorithm as a proposal distribution within a standard MCMC algorithm. This computational method has been used within the macroeconomics literature for the Bayesian analysis of second-order approximations to DSGE models; see, e.g. Fernández-Villaverde and Rubio-Ramírez (2007) and An and Schorfheide (2007). In these papers, a particle filter is used to approximate the likelihood function, see (29), of the DSGE model, which is a nonlinear state space model. The log-likelihood approximation is then used within a standard random-walk Metropolis algorithm.

Recently, Andrieu et al. (2007) have given a formal proof for the convergence of the algorithm. These authors prove that as long as the estimate of the likelihood function is unbiased then the estimation error produced by the approximation does not change the equilibrium distribution of the Markov chain being simulated. These authors label their algorithms Particle Markov chain Monte Carlo (PMCMC). In addition to providing convergence results for the random-walk Metropolis algorithm currently used in the DSGE literature, they also establish the results for a particle-filter based Gibbs sampler and an independent Metropolis-Hastings algorithm. They note that the Particle Gibbs sampler is not a standard Gibbs sampler. Additional care needs to be used when implementing an MCMC algorithm that uses a particle filter within it and has steps other than random-walk Metropolis. Flury and Shephard (2008) apply the methodology to several simple economic models to demonstrate its applicability.

5 Summary

This paper surveyed SMC methods that are applicable for economics and finance. The methods were applied to simple economic examples to illustrate their relevance on practical economic problems. From either a frequentist or Bayesian perspective, particle filters enable researchers to perform prediction and filtering in nonlinear, non-Gaussian state space models easily. Particle filters and other SMC methods may play a larger role in risk management, option pricing, and high-frequency financial econometrics. Following recent trends in macroeconomics, particle filters are appearing more frequently to estimate structural models. In a frequentist setting, econometricians can use the particle filter in testing situations (i.e. to compute likelihood-ratio statistics or Ljung-Box statistics). Maximum likelihood estimation of nonlinear, non-Gaussian state space models using particle filters still remains an open research area. No single method has demonstrated an overwhelming computational or theoretical advantage for a reasonably large class of models. Work remains to be done on the statistical properties of the estimators as well.

SMC methods are likely to have a continued impact on Bayesian inference. SMC opens many new research avenues for estimating challenging models. These include trans-dimensional models, models that result in multimodal posteriors, and models with potentially a large number of parameters. The emphasis in this literature is currently on developing adaptive Monte Carlo algorithms that learn from previously simulated samples. Understanding how the algorithms should be built in practice to make adaption work and its comparison with MCMC is part of this research. The limit theorems needed to justify their use is another. A second theme is the introduction of particle filters within MCMC algorithms.

Acknowledgements

I would like to thank Charles Bos, Siem Jan Koopman, Michael Massmann, Herman van Dijk, Eric Zivot and participants at the Emerging Methods in Bayesian Econometrics Workshop at Erasmus Universiteit Rotterdam for helpful comments. I would also like to acknowledge financial support from the Grover and Creta Ensley Fellowship, which funded part of this research while I was a graduate student at the University of Washington. All the computations reported in this paper were carried out using the OxMetrix 5.0 programming environment of Doornik (2007). I used the SsfPack functions of Koopman et al. (1999) and solved the DSGE model using the LiRE package of Mavroeidis and Zwols (2007). Ox code used in the examples is available upon request from the author.

References

- Akashi, H. and H. Kumamoto (1977). Random sampling approach to state estimation in switching environment. *Automatica* 13, 429–434.
- An, S. and F. Schorfheide (2007). Bayesian analysis of DSGE models. *Econometric Reviews* 26(2-4), 113–172.
- Anderson, B. D. O. and J. B. Moore (1979). *Optimal Filtering*. New Jersey: Prentice-Hall.
- Andrieu, C., M. Davy, and A. Doucet (2003). Efficient particle filtering for jump Markov systems. Application to time-varying autoregressions. *IEEE Transactions on Signal Processing* 51(7), 1762–1770.
- Andrieu, C. and A. Doucet (2002). Particle filtering for partially observed Gaussian state space models. *Journal of the Royal Statistical Society, Series B* 64(4), 827–836.
- Andrieu, C., A. Doucet, and R. Holenstein (2007). Particle Markov chain Monte Carlo. Working paper, Department of Mathematics, University of Bristol.
- Andrieu, C., A. Doucet, S. S. Singh, and V. B. Tadić (2004). Particle methods for change detection, system identification, and control. *Proceedings of the IEEE* 92(3), 423–438.
- Andrieu, C., A. Doucet, and V. B. Tadić (2005). On-line simulation-based algorithms for parameter estimation in general state-space models. *Proc. Conf. Dec. Control*. New York: Institute of Electrical and Electronics Engineers.
- Baum, L. E. and T. P. Petrie (1966). Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics* 37, 1554–1563.
- Baum, L. E., T. P. Petrie, G. Soules, and N. Weiss (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics* 41, 164–171.
- Bercu, B., P. Del Moral, and A. Doucet (2008a). Fluctuations of interacting Markov chain Monte Carlo. Working paper, Centre de recherche, INRIA, Bordeaux, France.
- Bercu, B., P. Del Moral, and A. Doucet (2008b). A functional central limit theorem for a class of interacting Markov chain Monte Carlo methods. Working paper, Centre de recherche, INRIA, Bordeaux, France.
- Bos, C. S. and N. Shephard (2006). Inference for adaptive time series models: stochastic volatility and conditionally Gaussian state space form. *Econometric Reviews* 25(2-3), 219–244.

- Briers, M., A. Doucet, and S. Maskell (2004). Smoothing algorithms for state space models. Technical Report, University of Cambridge, Department of Engineering.
- Cappé, O., S. Godsill, and E. Moulines (2007). An overview of existing methods and recent advances in sequential Monte Carlo. *IEEE Proceedings in Signal Processing* 95(5), 899–924.
- Cappé, O., A. Guillin, J.-M. Marin, and C. P. Robert (2004). Population Monte Carlo. *Journal of Computational and Graphical Statistics* 13(4), 907–929.
- Cappé, O., A. Guillin, J.-M. Marin, and C. P. Robert (2008). Adaptive importance sampling in general mixture classes. *Statistics and Computing* (forthcoming).
- Cappé, O., E. Moulines, and T. Rydén (2005). *Inference in Hidden Markov Models*. New York: Springer Press.
- Carpenter, J., P. Clifford, and P. Fearnhead (1999). An improved particle filter for non-linear problems. *IEE Proceedings. Part F: Radar and Sonar Navigation* 146(1), 2–7.
- Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika* 81(3), 541–553.
- Carvalho, C. M., M. Johannes, H. F. Lopes, and N. G. Polson (2008). Particle learning and smoothing. Working Paper, Booth School of Business, University of Chicago.
- Celeux, G., J.-M. Marin, and C. P. Robert (2006). Iterated importance sampling in missing data problems. *Computational Statistics and Data Analysis* 50(12), 3386–3404.
- Chen, R. and J. S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B* 62(3), 493–508.
- Chen, Y., J. Xie, and J. S. Liu (2005). Stopping-time resampling for sequential Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* 67(2), 199–217.
- Chib, S. and I. Jeliazkov (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association* 96(453), 270–281.
- Chib, S., F. Nardari, and N. Shephard (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics* 108(2), 281–316.
- Chopin, N. (2002). A sequential particle filter method for static models. *Biometrika* 89(3), 539–551.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo and its application to Bayesian inference. *The Annals of Statistics* 32(6), 2385–2411.

- Chopin, N. (2007). Inference and model choice for sequentially ordered hidden Markov models. *Journal of the Royal Statistical Society, Series B* 69(2), 269–284.
- Chopin, N. and F. Pelgrin (2004). Bayesian inference and state number determination for hidden Markov models: an application to the information content of the yield curve about inflation. *Journal of Econometrics* 123(2), 327–344.
- Cornebise, J., E. Moulines, and J. Olsson (2008). Adaptive methods for sequential importance sampling with application to state space models. Unpublished manuscript, Ecole Nationale Supérieure des Télécommunications, Paris, France.
- Crisan, D., P. Del Moral, and T. Lyons (1999). Discrete filtering using branching and interacting particle systems. *Markov Processes and Related Fields* 5(3), 293–318.
- Crisan, D. and A. Doucet (2002). A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing* 50(3), 736–746.
- de Freitas, N., R. Dearden, F. Hutter, R. Morales-Menendez, J. Mutch, and D. Poole (2004). Diagnosis by a waiter and a Mars explorer. *Proceedings of the IEEE* 92(3), 1–11.
- de Jong, P. and N. Shephard (1995). The simulation smoother for time series models. *Biometrika* 82(2), 339–350.
- Del Moral, P. (1996). Nonlinear filtering: interacting particle solution. *Markov Processes and Related Fields* 2, 555–579.
- Del Moral, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. New York, NY: Springer.
- Del Moral, P. and A. Doucet (2004). On a class of genealogical and interacting Metropolis models. *Séminaire de Probabilités XXXVII, Ed Azéma, J., Emery, M., Ledoux, M., and Yor, M. 1832*, 415–446. Lecture Notes in Mathematics, Berlin, Springer.
- Del Moral, P., A. Doucet, and A. Jasra (2006a). Sequential Monte Carlo for Bayesian computation. *Bayesian Statistics 8, Oxford University Press, Oxford, UK*.
- Del Moral, P., A. Doucet, and A. Jasra (2006b). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society, Series B* 68(3), 1–26.
- Del Moral, P., A. Doucet, and A. Jasra (2008). An adaptive sequential Monte Carlo method for approximate Bayesian computation. Working paper, Department of Statistics, University of British Columbia.
- Del Moral, P. and L. Miclo (2000). Branching and interacting particle systems approximations of Feynman-Kac formulae with applications to non-linear filtering. *Séminaire de*

- Probabilités XXXVII 1729*, 1–145. Lecture Notes in Mathematics, Berlin, Springer.
- Doornik, J. A. (2007). *Ox: An Object-Oriented Matrix Programming Language*. London: Timberlake Consultants Press.
- Douc, R., O. Cappé, and E. Moulines (2005, September). Comparison of resampling schemes for particle filtering. *4th International Symposium on Image and Signal Processing and Analysis (ISPA)*. Zagreb, Croatia.
- Douc, R., A. Guillin, J.-M. Marin, and C. P. Robert (2007a). Convergence of adaptive sampling schemes. *The Annals of Statistics* 35(1), 1–35.
- Douc, R., A. Guillin, J.-M. Marin, and C. P. Robert (2007b). Minimum variance importance sampling via population Monte Carlo. *ESAIM: Probability and Statistics* 11, 420–448.
- Douc, R. and E. Moulines (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *The Annals of Statistics* (forthcoming).
- Douc, R., E. Moulines, and J. Olsson (2007). On the auxiliary particle filter. Unpublished manuscript, Ecole Nationale Supérieure des Télécommunications, Paris, France.
- Doucet, A., M. Briers, and S. Sénécal (2006). Efficient block sampling strategies for sequential Monte Carlo. *Journal of Computational and Graphical Statistics* 15(3), 693–711.
- Doucet, A., N. de Freitas, and N. Gordon (2001). *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag.
- Doucet, A., S. J. Godsill, and C. Andrieu (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing* 10(3), 197–208.
- Doucet, A., N. J. Gordon, and V. Krishnamurthy (2001). Particle filters for state estimation of jump Markov linear systems. *IEEE Transactions on Signal Processing* 49, 613–624.
- Doucet, A. and V. Tadić (2003). Parameter estimation in general state space models using particle methods. *Annals of the Institute of Statistical Mathematics* 55(2), 409–422.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford, UK: Oxford University Press.
- Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89(3), 603–616.
- Evans, G. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton, New Jersey: Princeton University Press.
- Fearnhead, P. (2002). MCMC, sufficient statistics, and particle filters. *Journal of Computational and Graphical Statistics* 11(4), 848–862.

- Fearnhead, P. and P. Clifford (2003). On-line inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society, Series B* 65(4), 887–899.
- Fearnhead, P., D. Wyncoll, and J. Tawn (2008). A sequential smoothing algorithm with linear computational cost. Working Paper. Department of Statistics, University of Lancaster.
- Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2005). Estimating dynamic equilibrium economies: linear versus nonlinear likelihood. *Journal of Applied Econometrics* 20, 891–910.
- Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007). Estimating macroeconomic models: a likelihood approach. *The Review of Economic Studies* 74(4), 1059–1087.
- Flury, T. and N. Shephard (2008). Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models. Unpublished Paper, Department of Economics, University of Oxford.
- Fraser, D. and J. Potter (1969). The optimum linear smoother as a combination of two optimum linear filters. *IEEE Transactions on Automatic Control* 4, 387–390.
- Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis* 15, 183–202.
- Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. New York, NY: Springer Press.
- Geweke, J. F. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* 57(6), 1317–1139.
- Geweke, J. F. (2005). *Contemporary Bayesian Econometrics and Statistics*. New Jersey: Wiley Press.
- Gilks, W. and C. Berzuini (2001). Following a moving target - Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society, Series B* 63(1), 126–146.
- Godsill, S. J., A. Doucet, and M. West (2002). Maximum a posteriori sequence estimation using Monte Carlo particle filters. *Annals of the Institute of Statistical Mathematics* 53(1), 82–96.
- Godsill, S. J., A. Doucet, and M. West (2004). Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association* 99(465), 156–168.
- Gordon, N. J., D. J. Salmond, and A. F. M. Smith (1993). A novel approach to nonlinear and non-Gaussian Bayesian state estimation. *IEE Proceedings. Part F: Radar and Sonar*

- Navigation* 140(2), 107–113.
- Guo, D., X. Wang, and R. Chen (2005). New sequential Monte Carlo methods for nonlinear dynamic systems. *Statistics and Computing* 15(2), 135–147.
- Hamilton, J. D. (1989). A new approach to the econometric analysis of nonstationary time series and the business cycle. *Econometrica* 57(2), 357–384.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge, UK: Cambridge University Press.
- Hürzeler, M. and H. R. Künsch (1998). Monte Carlo approximations for general state space models. *Journal of Computational and Graphical Statistics* 7(2), 175–193.
- Hürzeler, M. and H. R. Künsch (2001). Approximating and maximising the likelihood for a general state space model. A. Doucet, N. de Freitas, and N. Gordon, Eds. *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer.
- Jarzynski, C. (1997). Nonequilibrium equality for free energy differences. *Physical Review Letters* 78, 2690–2693.
- Jasra, A., A. Doucet, D. Stephens, and C. Holmes (2008). Interacting sequential Monte Carlo samplers for trans-dimensional simulation. *Computational Statistics and Data Analysis* 52(4), 1765–1791.
- Jasra, A., D. A. Stephens, A. Doucet, and T. Tsagaris (2008). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. Unpublished manuscript, Institute for Statistical Mathematics, Tokyo, Japan.
- Johannes, M., N. G. Polson, and J. R. Stroud (2006). Exact particle filtering and parameter learning. Working paper, Graduate School of Business, Columbia University.
- Johansen, A. M., P. Del Moral, and A. Doucet (2006). Sequential Monte Carlo samplers for rare events. *Proc. 6th International Workshop on Rare Event Simulation*.
- Johansen, A. M. and A. Doucet (2008). A note on auxiliary particle filters. *Statistics and Probability Letters*, (forthcoming).
- Johansen, A. M., A. Doucet, and M. Davy (2008). Particle methods for maximum likelihood parameter estimation in latent variable models. *Statistics and Computing* 18, 47–57.
- Julier, S. J. and J. K. Uhlmann (1997). A new method for nonlinear transformation of means and covariances in filters and estimators. *The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Multi Sensor Fusion, Tracking and Resource Management II*. SPIE.

- Julier, S. J., J. K. Uhlmann, and H. F. Durrant-White (2000). A new method for nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control* 45(3), 477–482.
- Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, Transactions of the ASME* 82(1), 35–46.
- Kalman, R. and R. Bucy (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering, Transactions of the ASME* 83(3), 95–108.
- Kim, C.-J. and C. R. Nelson (1999). *State Space Models with Regime Switching Classical and Gibbs Sampling Approaches with Applications*. Cambridge, MA: MIT Press.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies* 65(3), 361–393.
- Kitagawa, G. (1996a). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics* 5(1), 1–25.
- Kitagawa, G. (1996b). A self-organizing state-space model. *Journal of the American Statistical Association* 93(443), 1203–1215.
- Kong, A., J. S. Liu, and W. Wong (1994). Sequential imputation and Bayesian missing data problems. *Journal of the American Statistical Association* 89, 278–288.
- Koopman, S. J., N. Shephard, and J. A. Doornik (1999). Statistical algorithms for models in state space using SsfPack 2.2. *The Econometrics Journal* 2, 107–166.
- Künsch, H. R. (2001). State space and hidden Markov models. in O. E. Barndorff-Nielsen, D. R. Cox, and C. Klüppelburg, Eds. *Complex Stochastic Systems*, 109–173. Chapman Hall CRC, London.
- Künsch, H. R. (2005). Recursive Monte Carlo filters: algorithms and theoretical analysis. *The Annals of Statistics* 33(5), 1983–2021.
- Künsch, H. R. (2006). Filtering, smoothing and parameter estimation for general state space models. Lecture notes, Troisième Cycle Romand, Course at EPFL, <http://stat.ethz.ch/kuensch>.
- Liang, F. (2002). Dynamically weighted importance sampling in Monte Carlo computation. *Journal of the American Statistical Association* 93, 807–821.
- Liu, J. and M. West (2001). Combined parameter and state estimation in simulation-based filtering. A. Doucet, N. de Freitas, and N. Gordon, Eds. *Sequential Monte Carlo Methods in Practice*, Springer, New York, NY..

- Liu, J. S. (1996). Metropolized independent sampling with comparison to rejection sampling and importance sampling. *Statistics and Computing* 6, 113–119.
- Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. New York: Springer Press.
- Liu, J. S. and R. Chen (1995). Blind deconvolutions via sequential imputations. *Journal of the American Statistical Association* 90(430), 567–576.
- Mavroeidis, S. and Y. Zwols (2007). LiRE: An Ox package for solving linear rational expectations models. Unpublished Manuscript, Department of Economics, Brown University.
- Mayne, D. Q. (1966). A solution of the smoothing problem for linear dynamic systems. *Automatica* 4, 73–92.
- Neal, R. (2001). Annealed importance sampling. *Statistics and Computing* 11(2), 125–139.
- Olsson, J., O. Cappé, R. Douc, and E. Moulines (2008). Sequential Monte Carlo smoothing with application to parameter estimation in non-linear state space models. *Bernoulli* 14(1), 155–179.
- Olsson, J. and T. Rydén (2008). Asymptotic properties of particle filter-based maximum likelihood estimators for state space models. *Stochastic Processes and their Applications* 118, 649–680.
- Pitt, M. K. (2002). Smooth particle filters for likelihood evaluation and maximisation. Unpublished manuscript, Department of Economics, Warwick University.
- Pitt, M. K. and N. Shephard (1999). Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association* 94(446), 590–599.
- Pitt, M. K. and N. Shephard (2001). Auxiliary variable based particle filters. A. Doucet, N. de Freitas, and N. Gordon, Eds. *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer.
- Polson, N. G., J. R. Stroud, and P. Müller (2008). Practical filtering with sequential parameter learning. *Journal of the Royal Statistical Society, Series B* 70(2), 413–428.
- Poyiadjis, G., A. Doucet, and S. S. Singh (2005a). Maximum likelihood parameter estimation in general state space models using particle methods. Unpublished manuscript, Department of Engineering, Cambridge University.
- Poyiadjis, G., A. Doucet, and S. S. Singh (2005b, March). Particle methods for optimal filter derivative: application to parameter estimation. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 925–928.

- Rauch, H., F. Tung, and C. Striebel (1965). Maximum likelihood estimation of linear dynamic systems. *AIAA Journal* 3, 1445–1450.
- Robert, C. P. and G. Casella (2004). *Monte Carlo Statistical Methods* (Second ed.). New York: Springer Press.
- Rubin, D. B. (1987). A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when the fraction of missing information is modest: the SIR algorithm (discussion of Tanner and Wong). *Journal of the American Statistical Association* 82, 543–546.
- Schön, T., F. Gustafsson, and P.-J. Nordlund (2005). Marginalized particle filters for mixed linear/nonlinear state-space models. *IEEE Transactions on Signal Processing* 53(7), 2279–2289.
- Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit, and Banking* 39(1), 3–34.
- Storvik, G. (2002). Particle filters in state space models with the presence of unknown static parameters. *IEEE Transactions on Signal Processing* 50, 281–289.
- Tanizaki, H. and R. S. Mariano (1998). Nonlinear and non-Gaussian state-space modeling with Monte-Carlo simulations. *Journal of Econometrics* 83(1-2), 263–290.
- Tierney, L. (1994). Markov chains for exploring posterior distributions (with discussion). *The Annals of Statistics* 22(4), 1701–1762.
- van der Merwe, R., A. Doucet, and N. de Freitas (2000). The unscented particle filter. T. K. Leen, T. G. Dietterich, and V. Tresp, Eds. *Advances in Neural Information Processing Systems* 13. MIT Press.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimal decoding algorithm. *IEEE Transactions in Information Theory* 13, 260–269.
- Whitley, D. (1994). A genetic algorithm tutorial. *Statistics and Computing* 4, 65–85.
- Zaritskii, V., V. Svetnik, and L. Shimelevich (1975). Monte Carlo techniques in problems of optimal data processing. *Automation and Remote Control* 12, 95–103.