

On sequential Monte Carlo, partial rejection control and approximate Bayesian computation

G. W. Peters^{*†} Y. Fan^{*} S. A. Sisson^{*‡}

November 11, 2009

Abstract

We present a sequential Monte Carlo sampler variant of the partial rejection control algorithm, and show that this variant can be considered as a sequential Monte Carlo sampler with a modified mutation kernel. We prove that the new sampler can reduce the variance of the incremental importance weights when compared with standard sequential Monte Carlo samplers. We provide a study of theoretical properties of the new algorithm, and make connections with some existing algorithms. Finally, the sampler is adapted for application under the challenging “likelihood free,” approximate Bayesian computation modelling framework, where we demonstrate superior performance over existing likelihood-free samplers.

Keywords: Approximate Bayesian computation; Bayesian computation; Likelihood-free inference; Sequential Monte Carlo samplers; Partial rejection control.

^{*}School of Mathematics and Statistics, University of New South Wales, Sydney, Australia

[†]CSIRO Sydney, Locked Bag 17, North Ryde, NSW, 1670, Australia

[‡]Communicating author: Scott.Sisson@unsw.edu.au

1 Introduction

Sequential Monte Carlo (SMC) methods have emerged out of the fields of engineering, probability and statistics in recent years. Variants of the methods sometimes appear under the names of particle filtering or interacting particle systems (e.g. Arulampalam et al. 2002; Andrieu et al. 2003; Del Moral 2004; Doucet et al. 2001), and their theoretical properties have been extensively studied (Crisan and Doucet 2002; Del Moral 2004; Kunsch 2005).

The standard SMC algorithm involves finding a numerical solution to a set of filtering recursions, such as filtering problems arising from non-linear / non-Gaussian state space models. Under this framework, the SMC algorithm samples from a (often naturally occurring) sequence of distributions π_t , indexed by $t = 1, \dots, T$. Each distribution is defined on the support $E^t = E \times E \times \dots \times E$. Del Moral et al. (2006) (see also Peters 2005) generalize the SMC algorithm to the case where the distributions π_t are all defined on the same support E . This generalization, termed the SMC *sampler*, adapts the SMC algorithm to the more popular setting in which the state space E remains static.

In short, the SMC sampler generates weighted samples (termed *particles*) from a sequence of distributions π_t , for $t = 1, \dots, T$, where π_T may be of particular interest. We refer to π_T as the target distribution. Procedurally, particles obtained from an arbitrary initial distribution π_1 , with a set of corresponding initial weights, are sequentially propagated through each distribution π_t in the sequence via three processes, involving mutation (or move), correction (or importance weighting) and selection (or resampling). The final weighted particles at distribution π_T are considered weighted samples from the target distribution π . The mechanism is similar to sequential importance sampling (resampling) (Liu 2001; Doucet et al. 2001), with

one of the crucial differences being the framework under which the particles are allowed to move, resulting in differences in the calculation of the weights of the particles.

One of the major difficulties with SMC-type algorithms is particle depletion, in which the weights of the majority of the particles gradually decrease to zero, while a few particle weights dominate the population. This severely increases the variability of Monte Carlo estimates of expectations under π . In this article, we develop an algorithm which incorporates the partial rejection control (PRC) strategy of Liu (2001) into the SMC sampler framework. A particular motivation for this stems from the recent developments in “likelihood-free” (or approximate Bayesian) computation (Beaumont et al. 2002; Marjoram et al. 2003; Sisson et al. 2007), where an extremely high proportion of mutated particles are expected to have very small, or exactly zero, posterior weights.

In this article, we develop the SMC samplers PRC algorithm, in which the partial rejection control mechanism is built directly into the mutation kernel of the SMC sampler. In this manner, a particle mutation may be rejected if the resulting importance weight is below a certain threshold. We begin with a brief introduction to the standard sequential Monte Carlo sampler in Section 2, and then present the SMC sampler PRC algorithm. We also discuss implementational issues arising from the inclusion of the PRC stage, including estimation for the resultant kernel normalizing constant. Section 3 provides some theoretical results that justify the addition of PRC in terms of improvements in the variance of the incremental importance weights. We also discuss a central limit theorem and derive a recursive expression for the asymptotic variance of our algorithm. In addition, we make a novel connection between the SMC sampler PRC algorithm and the AliveSMC algorithm from the rare-event literature developed in Le Gland and Oudjane (2004). In Section 4 we adapt the SMC sampler PRC algorithm for application in the likelihood-free modelling framework,

and demonstrate the computational gains achieved over existing likelihood-free algorithms via a simulated example. Finally, we present a stochastic claims reserving analysis using the developed methods in Section 5, and conclude with a discussion.

2 Sequential Monte Carlo and partial rejection

2.1 Sequential Monte Carlo sampler

Del Moral et al. (2006) introduced a modification of the sequential Monte Carlo algorithm, termed the sequential Monte Carlo *sampler*. Consider a sequence of distributions $\pi_t(x), t = 1, \dots, T$, with $x \in E$, where the final distribution π_T is the distribution of interest. By introducing a sequence of backward kernels L_k , a new distribution $\tilde{\pi}_t(x_1, \dots, x_t) = \pi_t(x_t) \prod_{k=1}^{t-1} L_k(x_{k+1}, x_k)$ may be defined for the *path* of a particle $(x_1, \dots, x_t) \in E^t$ through the sequence π_1, \dots, π_t . The only restriction on the backward kernels is that the correct marginal distributions $\int \tilde{\pi}_t(x_1, \dots, x_t) dx_1, \dots, dx_{t-1} = \pi_t(x_t)$ are available.

Within this framework, one may then work with the sequence of distributions, $\tilde{\pi}_t$, under the standard SMC algorithm. In summary, the SMC sampler algorithm involves three stages: *mutation*, whereby the particles are moved from x_{t-1} to x_t via a mutation kernel $M_t(x_{t-1}, x_t)$ as described below (Doucet and Johansen 2009; Del Moral et al. 2006); *correction*, where the particles are reweighted with respect to π_t via the incremental importance weight (1); and *selection*, where according to some measure of particle diversity, commonly the effective sample size (ESS, Kunsch 2005; Kitigawa 1996; Doucet et al. 2001; Liu and Chen 1998), the weighted particles may be resampled in order to reduce the variability of the importance weights.

In more detail, suppose that at time $t - 1$, the distribution $\tilde{\pi}_{t-1}$ can be approximated empirically by $\tilde{\pi}_{t-1}^N$ using N weighted particles. These particles are first prop-

agated to the next distribution $\tilde{\pi}_t$ using a mutation kernel $M_t(x_{t-1}, x_t)$, and then assigned new weights $W_t = W_{t-1} w_t(x_1, \dots, x_t)$, where W_{t-1} is the weight of a particle at time $t - 1$ and w_t is the incremental weight given by

$$w_t(x_1, \dots, x_t) = \frac{\tilde{\pi}_t(x_1, \dots, x_t)}{\tilde{\pi}_{t-1}(x_1, \dots, x_{t-1}) M_t(x_{t-1}, x_t)} = \frac{\pi_t(x_t) L_{t-1}(x_t, x_{t-1})}{\pi_{t-1}(x_{t-1}) M_t(x_{t-1}, x_t)}. \quad (1)$$

The resulting particles are now weighted samples from $\tilde{\pi}_t$. Consequently from (1), under the SMC sampler framework, one may work directly with the marginal distributions $\pi_t(x_t)$ such that $w_t(x_1, \dots, x_t) = w_t(x_{t-1}, x_t)$. While the choice of the backward kernels L_{t-1} is essentially arbitrary, their specification can strongly affect the performance of the algorithm. See Del Moral et al. (2006) for detailed discussion.

2.2 Incorporating partial rejection control

It is well known that the performance of SMC methods are strongly dependent on the mutation kernel (Cornebise et al. 2008). If M_t is poorly chosen, such that it does not place particles in regions of the support of π_t with high density, then many importance sampling weights will be close to zero. This leads to sample degeneracy, as a few well located particles with large weights dominate the particle population, resulting in large variance for estimates made using these samples.

Liu (2001) (see also Liu et al. 1998) introduced the partial rejection control strategy to overcome particle degeneracy in a sequential importance sampling setting. Under this mechanism, when the weight of a particle at distribution π_t falls below a finite threshold, $c_t \geq 0$, the particle is probabilistically discarded. It is replaced with a particle drawn from the previous distribution π_{t-1} which is then mutated to π_t . This new particle's weight is then compared to the threshold, with this process repeating until a particle is accepted. This approach is termed *partial* rejection, as the replacement particle is drawn from π_{t-1} , not π_1 (Liu et al. 1998).

Under the SMC sampler framework we modify this approach and incorporate the partial rejection mechanism directly within the mutation kernel. Hence at time $t - 1$, the particle x_{t-1} is moved via the mutation kernel $M_t(x_{t-1}, x_t)$ and weighted according to (1). This particle is accepted with probability p , determined by the particle's weight and the weight threshold c_t . If rejected, a new particle is obtained via the mutation kernel M_t , until a particle is accepted.

For the sequence of distributions π_t , $t = 1, \dots, T$, the mutation and backward kernels M_t and L_{t-1} , a sequence of weight thresholds c_t , and PRC normalizing constants $r(c_t, x_{t-1})$ (defined below), the SMC sampler PRC algorithm is given by:

SMC sampler PRC algorithm

Initialization: Set $t = 1$.

For $i = 1, \dots, N$, sample $x_1^{(i)} \sim \pi_1(x)$, and set weights $W_1(x_1^{(i)}) = \frac{1}{N}$.

Resample: Normalize the weights $\sum_i W_t(x_t^{(i)}) = 1$. If $[\sum_i W_t(x_t^{(i)})^2]^{-1} < H$ resample N particles with respect to $W_t(x_t^{(i)})$ and set $W_t(x_t^{(i)}) = \frac{1}{N}$, $i = 1, \dots, N$.

Mutation and correction: Set $t = t + 1$ and $i = 1$:

(a) Sample $x_t^{(i)} \sim M_t(x_{t-1}^{(i)}, x_t)$ and set weight for $x_t^{(i)}$ to

$$W_t(x_t^{(i)}) = W_{t-1}(x_{t-1}^{(i)}) \frac{\pi_t(x_t^{(i)}) L_{t-1}(x_t^{(i)}, x_{t-1}^{(i)})}{\pi_{t-1}(x_{t-1}^{(i)}) M_t(x_{t-1}^{(i)}, x_t^{(i)})}.$$

(b) With probability $1 - p^{(i)} = 1 - \min\{1, W_t(x_t^{(i)})/c_t\}$, reject $x_t^{(i)}$ and go to (a).

(c) Otherwise, accept $x_t^{(i)}$ and set $W_t(x_t^{(i)}) = W_t(x_t^{(i)}) r(c_t, x_{t-1}^{(i)})/p^{(i)}$.

(d) Increment $i = i + 1$. If $i \leq N$, go to (a).

(e) If $t < T$ go to Resample.

The above algorithm without the mutation and correction steps (b) and (c) is equivalent to the standard SMC sampler algorithm (Del Moral et al. 2006). In the resample stage, the degeneracy of the particle approximation is quantified through the

usual estimate of the effective sample size, $1 \leq [\sum_i W_t(x_t^{(i)})^2]^{-1} \leq N$ (Liu and Chen 1998).

We discuss the choice of the thresholds, c_t , in later Sections.

The addition of a rejection step at each time t effectively modifies the mutation kernel M_t . We denote the resulting kernel by M_t^* , where

$$M_t^*(x_{t-1}, x_t) = r(c_t, x_{t-1})^{-1} \min \left[\left\{ 1, W_{t-1}(x_{t-1}) \frac{w_t(x_{t-1}, x_t)}{c_t} \right\} M_t(x_{t-1}, x_t) \right]. \quad (2)$$

The quantity $r(c_t, x_{t-1})$ denotes the normalizing constant for particle x_{t-1} , given by

$$r(c_t, x_{t-1}) = \int \min \left\{ 1, W_{t-1}(x_{t-1}) \frac{w_t(x_{t-1}, x_t)}{c_t} \right\} M_t(x_{t-1}, x_t) dx_t. \quad (3)$$

Note that $0 < r(c_t, x_{t-1}) \leq 1$ if (w.l.o.g.) the mutation kernel M_t is normalized, so that $\int M_t(x_{t-1}, x_t) dx_t = 1$, and if the PRC threshold $0 \leq c_t < \infty$ is finite. Thus the SMC sampler PRC algorithm can be considered as an SMC sampler algorithm with the mutation kernel $M_t^*(x_{t-1}, x_t)$, and the correction weight

$$W_t(x_t) = W_{t-1}(x_{t-1}) \frac{\pi_t(x_t) L_{t-1}(x_t, x_{t-1})}{\pi_{t-1}(x_{t-1}) M_t^*(x_{t-1}, x_t)}. \quad (4)$$

2.3 Estimation of the normalizing constant

As the normalizing constant $r(c_t, x_{t-1})$ in the weight calculation (4) in general depends on x_{t-1} , it must be evaluated. Where no analytic solution can be found, approximating (3) may be achieved by, for example, quadrature methods if the sample space E is relatively low dimensional or Monte Carlo methods if E is high dimensional. For example, for $j = 1, \dots, m$ independent samples $x_t^{*(j)}$ sampled from $M_t(x_{t-1}, x_t)$

$$\hat{r}(c_t, x_{t-1}) \approx \frac{1}{m} \sum_{j=1}^m \min \left\{ 1, W_{t-1}(x_{t-1}) \frac{w(x_{t-1}, x_t^{*(j)})}{c_t} \right\}.$$

An alternative, computationally more efficient approach is to select kernels M_t and L_{t-1} such that $r(c_t, x_{t-1}) = r(c_t)$ will be constant for all particles x_{t-1} . In this case,

the value of $r(c_t)$ may be absorbed into the proportionality constant of the weights, and safely ignored. Equation (3) suggests that this can be achieved if $M_t(x_{t-1}, x_t)$, $W_{t-1}(x_{t-1})$ and $w(x_{t-1}, x_t)$ are independent of x_{t-1} .

Specifying mutation kernels M_t such that $M_t(x_{t-1}, x_t) = M_t(x_t)$ amounts to choosing a *global* kernel which is the same for all particles x_{t-1} . This is common in practice (e.g. West 1993). The particle dependent weight $W_{t-1}(x_{t-1})$ can be set to $1/N$ for all particles following a resampling (or preselection) step. Finally, consider for a moment the backward kernel of the form

$$L_{t-1}^{opt}(x_t, x_{t-1}) = \frac{\pi_{t-1}(x_{t-1})M_t(x_{t-1}, x_t)}{\int \pi_{t-1}(x_{t-1})M_t(x_{t-1}, x_t)dx_{t-1}}. \quad (5)$$

This backward kernel is an approximation of the optimal backward kernel, in the sense of the choice of L_{t-1} that minimizes the variance of the importance sampling weights (Del Moral et al. 2006). Under the backward kernel (5), the incremental weight can be approximated by

$$\begin{aligned} w_t(x_{t-1}, x_t) &= \pi_t(x_t) / \int \pi_{t-1}(x_{t-1})M_t(x_{t-1}, x_t)dx_{t-1} \\ &\approx \pi_t(x_t) / \sum_{i=1}^N W_{t-1}(x_{t-1}^{(i)})M_t(x_{t-1}^{(i)}, x_t). \end{aligned}$$

Under a global mutation kernel $M_t(x_t)$, and following a resampling step, then the incremental weight under this backward kernel reduces to $w_t(x_{t-1}, x_t) = \pi_t(x_t)/M_t(x_t)$, which is independent of x_{t-1} . Thus, the weight calculation in (4) becomes

$$\begin{aligned} W_t(x_t) &\propto \pi_t(x_t) / \left[\min \left\{ 1, \frac{w(x_{t-1}, x_t)}{Nc_t} \right\} M_t(x_t) \right] \\ &= \begin{cases} \pi_t(x_t)/M_t(x_t) & \text{if } \min \left\{ 1, \frac{w(x_{t-1}, x_t)}{Nc_t} \right\} = 1 \\ Nc_t & \text{otherwise.} \end{cases} \end{aligned}$$

Note that under this setting, the SMC sampler PRC algorithm can be considered as a sequence of importance sampling strategies with partial rejection control.

3 SMC Sampler PRC algorithm analysis

In this section we study theoretical properties of the SMC sampler PRC algorithm. We firstly bound the variance of the importance weights, and then present a central limit theorem for the sampler with a recursive expression for the asymptotic variance. Finally, via a connection with an existing algorithm, we establish a condition for which the number of rejection steps under the PRC mechanism is almost surely finite.

3.1 Variance of the incremental weights

We begin this section by establishing a bound on the variance of the importance weights of the SMC sampler PRC algorithm.

Theorem 3.1 *Let $W_t(x_t)$ denote the importance sampling weight at time t from a standard SMC sampler with mutation kernel M_t , and let $W_t^*(x_t)$ denote the equivalent weight following a partial rejection control step under the SMC sampler PRC algorithm, with resulting mutation kernel M_t^* . Then*

$$\mathbb{V}ar_{M_t^*}[W_t^*(x_t)] \leq \mathbb{V}ar_{M_t}[W_t(x_t)].$$

Proof: See Appendix A.1.

Hence, applying partial rejection control within the SMC sampler framework will not worsen, and may improve the variance of the importance weights, by reducing the χ^2 -distance between the sampling and target distributions at each stage, t .

In the case where $\min \left\{ 1, \frac{W_t(x_t)}{c_t} \right\} = 1$ for all x_t , which is achieved when $c_t \leq \inf_{x_t} \{W_t(x_t)\}$, then from (3) we have $r(c_t, x_{t-1}) = 1$ for all x_{t-1} . From (2), this results in $M_t^*(x_{t-1}, x_t) = M_t(x_{t-1}, x_t)$ and hence $\mathbb{V}ar_{M_t^*}[W_t^*(x_t)] = \mathbb{V}ar_{M_t}[W_t(x_t)]$. That is, the SMC sampler PRC algorithm reduces to the standard SMC sampler when $c_t \leq \inf_{x_t} \{W_t(x_t)\}$, and in this case, the variance of the importance weights is maximised. When $W_t(x_t) \in [0, \infty)$ this is realized for $c_t = 0$ where we define $0/0 := 1$.

3.2 A central limit theorem

Central Limit Theorems (CLTs) for SMC and particle filtering algorithms have been derived in various literatures (Del Moral 2004; Kunsch 2005; Chopin 2004; Del Moral et al. 2006; Johansen and Doucet 2008). They are based on the observation that an SMC algorithm introduces local errors (fluctuations) as a result of the approximations introduced by sampling numerically from the transitions. Hence, at each stage t , one can decompose the error between the target distribution π_t and the N -particle approximation π_t^N . This turns out to be a sum of the local sampling fluctuations at each discrete time in the past, propagated forwards in time to t .

In the setting of the SMC sampler algorithm, the existence of a CLT is established by Del Moral (2004). Explicitly, under the assumption of multinomial resampling at each stage of the algorithm, and the integrability conditions given in Chopin (2004) [Theorem 1] and Del Moral (2004) [Section 9.4, pp.300-306], then for a suitable continuous and bounded test function $\varphi \in C_b(E)$ we have

$$N^{1/2} \left(\mathbb{E}_{\pi_t^N}(\varphi) - \mathbb{E}_{\pi_t}(\varphi) \right) \rightarrow \mathcal{N}(0, V_{SMC,t}(\varphi)) \quad (6)$$

as $N \rightarrow \infty$, for each $t = 1, \dots, T$. Del Moral et al. (2006) obtain a recursive expression for the asymptotic variance $V_{SMC,t}(\varphi)$ as an explicit function of the backward kernels L_{t-1} and the sequence of distributions on path space, $\tilde{\pi}_t$.

Following Del Moral et al. (2006), we obtain an analogous result for the SMC sampler PRC algorithm. Under the same assumptions as the above, we have the CLT (6) with asymptotic variance given by

$$\begin{aligned} V_{SMC-PRC,t}(\varphi) &= \int I_1 \frac{\tilde{\pi}_t^2(x_1)}{\pi_1(x_1)} \left(\int \varphi(x_1) \tilde{\pi}_1(x_t|x_1) dx_t - \mathbb{E}_{\pi_t}(\varphi) \right)^2 dx_1 \\ &+ \sum_{k=2}^{t-1} \int I_k \frac{(\tilde{\pi}_t(x_k) L_{k-1}(x_k, x_{k-1}))^2}{\pi_{k-1}(x_{k-1}) M_k(x_{k-1}, x_k)} \left(\int \varphi(x_t) \tilde{\pi}_t(x_t|x_k) dx_t - \mathbb{E}_{\pi_t}(\varphi) \right)^2 dx_{k-1} dx_k \\ &+ \int I_t \frac{(\pi_t(x_t) L_{t-1}(x_t, x_{t-1}))^2}{\pi_{t-1}(x_{t-1}) M_t(x_t, x_{t-1})} (\varphi(x_t) - \mathbb{E}_{\pi_t}(\varphi))^2 dx_{t-1} dx_t \end{aligned} \quad (7)$$

with $I_k = \left[r(c_k, x_{k-1})^{-1} \min \left\{ 1, \frac{1}{N} \frac{\pi_k(x_k) L_{k-1}(x_k, x_{k-1})}{\pi_{k-1}(x_{k-1}) M_k(x_{k-1}, x_k) c_k} \right\} \right]^{-1}$ and $I_1 = 1$.

The contribution of the PRC stage to the asymptotic variance is encapsulated in the I_k terms. Under the standard SMC sampler algorithm we have $c_k \leq \inf_{x_t} \{W_t(x_t)\}$ so that $I_k = 1$ for all $k = 1, \dots, t$. In this setting, (7) reduces to the asymptotic variance expression obtained by Del Moral et al. (2006).

3.3 Connections to an existing SMC algorithm

In rare event applications there is a high probability of generating particles with exactly zero weights. The AliveSMC algorithm (Le Gland and Oudjane 2004; Le Gland and Oudjane 2005) was developed to ensure that a particle population of a desired size persists at each iteration of a standard SMC algorithm (see Del Moral et al. 2001; Johansen et al. 2006 for related methods). In this setting, the number of particles at each time t is considered as a random variable N_{c_t} . That is, N_{c_t} is the number of particles required to generate exactly N non-zero weighted particles. In this Section we reinterpret the SMC sampler PRC algorithm in terms of the AliveSMC algorithm. As a consequence, in Section 3.4 we are able to establish a condition under which the PRC resampling stage will require a finite number of rejection attempts.

In Le Gland and Oudjane (2005), at iteration t , a fitness function is applied to select particles satisfying a desired criteria. Those particles not satisfying the criteria receive a zero weight. In an SMC sampler PRC setting, the fitness function can be interpreted as selecting those particles with a weight that is immediately accepted under the PRC acceptance probability. As such, we may rewrite the SMC sampler PRC algorithm with a modified mutation and correction step:

Reinterpreted SMC sampler PRC/AliveSMC algorithm

Mutation and correction: Set $t = t + 1$ and $j = 1$. For $i = 1, \dots, N_{c_t}$:

(a) Sample $x_t^{(i)} \sim M_t(x_{t-1}^{(j)}, x_t)$ and calculate $W = W_{t-1}(x_{t-1}^{(j)}) \frac{\pi_t(x_t^{(i)})L_{t-1}(x_t^{(i)}, x_{t-1}^{(j)})}{\pi_{t-1}(x_{t-1}^{(j)})M_t(x_{t-1}^{(j)}, x_t^{(i)})}$.

(b) Set weight for $x_t^{(i)}$ as

$$W_t(x_t^{(i)}) \propto \begin{cases} W_{t-1}(x_{t-1}^{(j)})/p^{(i)} & \text{with probability } p^{(i)} = \min\{1, W/c_t\} \\ 0 & \text{otherwise.} \end{cases}$$

(c) If $W_t(x_t^{(i)}) \neq 0$, increment $j = j + 1$.

If $t < T$ go to Resample.

Note that $j = 1, \dots, N$ indexes the particles $x_{t-1}^{(j)}$ at time $t - 1$ such that particle mutations from $x_{t-1}^{(j)}$ generate a non-zero weight exactly once. Also, under the fitness function, particle $x_t^{(i)}$ has a probability $1 - p^{(i)} = 1 - \min\{1, W/c_t\}$ of being exactly zero. Given that it is possible to express the SMC sampler PRC algorithm within the AliveSMC framework, we may adapt the results of Le Gland and Oudjane (2004) and Le Gland and Oudjane (2005), to obtain a condition under which the SMC sampler PRC algorithm is guaranteed to require a finite number of attempts, $N_{c_t} < \infty$, to obtain exactly N non-zero weighted particles.

3.4 Analysis of the number of rejection attempts

Following Le Gland and Oudjane (2004) and Le Gland and Oudjane (2005), we define the random variable N_{c_t} as

$$N_{c_t} \triangleq \inf \left\{ N^* \geq 1 : \sum_{i=1}^{N^*} W_t^{(i)}(x_t) \geq N \sup_{x_t \in E} W_t(x_t) \right\}$$

Le Gland and Oudjane (2005) proved for the AliveSMC algorithm that the random number of particles N_{c_t} is almost surely finite with $N_{c_t} \geq N$, under the condition that $\langle \pi_{t-1} M_t, W_t \rangle = \frac{\int W_t(x_t) \int \pi_{t-1}(x_{t-1}) M_t(x_{t-1}, x_t) dx_{t-1} dx_t}{\int \pi_{t-1}(x_{t-1}) M_t(x_{t-1}, x_t) dx_{t-1}} > 0$. A sufficient condition for

this to hold is $\mathbb{E}_{M_t} [W_t(x_t) \mid x_{t-1} = x] > 0$, for all $x \in E$. Thus for the SMC sampler PRC algorithm, we have the following theorem:

Theorem 3.2 *Under the SMC sampler PRC algorithm, the number of rejection attempts at each stage of the algorithm, $N_{c_t} \geq N$, is almost surely finite if $c_t < \infty$.*

Proof: See Appendix A.2.

Corollary 5.1 *The following convergence in probability holds, with a rate of $1/\sqrt{N}$:*

$$\frac{N_{c_t}}{N} \rightarrow \frac{\sup_{x_t \in E} W_t(x_t)}{\langle \pi_{t-1} M_t, W_t \rangle} < \infty.$$

See Le Gland and Oudjane (2005) for further details. Hence, the SMC sampler PRC algorithm possesses an almost surely finite number of rejection attempts if the PRC threshold c_t is finite, with the above rate of convergence.

4 Approximate Bayesian computation

With the aim of posterior simulation from $\pi(x|\mathcal{D}) \propto \pi(\mathcal{D}|x)\pi(x)$ for parameters x and observed data \mathcal{D} , “likelihood-free,” approximate Bayesian computation (ABC) methods are often utilised when the likelihood function, $\pi(\mathcal{D}|x)$, is computationally intractable or when its evaluation is computationally prohibitive. ABC methods can be based on rejection sampling (Tavaré et al. 1997; Beaumont et al. 2002), Markov chain Monte Carlo (Marjoram et al. 2003; Bortot et al. 2007; Ratmann et al. 2009) and SMC-type samplers (Sisson et al. 2007; Toni et al. 2009; Beaumont et al. 2009; Del Moral et al. 2008). While currently among the most efficient ABC methods, the underlying practical issue with SMC-type algorithms is in avoiding sample degeneracy through extreme numbers of particles with low or exactly zero weights. In this Section, we will demonstrate that the SMC sampler PRC algorithm applied within the ABC

framework can achieve significant performance gains and greater modelling flexibility over existing SMC-type ABC methods.

The underlying approach of ABC methods is to augment the (intractable) posterior to $\pi(x, \mathcal{D}'|\mathcal{D}) \propto \pi(\mathcal{D}|\mathcal{D}', x)\pi(\mathcal{D}'|x)\pi(x)$, where the auxiliary parameter is an artificial data set distributed according to the model $\mathcal{D}' \sim \pi(\cdot|x)$. An approximation of the target posterior $\pi(x|\mathcal{D})$ is then given by

$$\pi_{ABC}(x|\mathcal{D}) \propto \int \pi(\mathcal{D}|\mathcal{D}', x)\pi(\mathcal{D}'|x)\pi(x)d\mathcal{D}'. \quad (8)$$

The weighting distribution $\pi(\mathcal{D}|\mathcal{D}', x)$ takes high density in regions where the datasets \mathcal{D} and \mathcal{D}' are similar, and low density otherwise. Comparison of the datasets is usually achieved through low-dimensional summary statistics $T(\cdot)$, so that, for example

$$\pi(\mathcal{D}|\mathcal{D}', x) \propto \begin{cases} 1 & \text{if } \rho(T(\mathcal{D}), T(\mathcal{D}')) \leq \epsilon \\ 0 & \text{else,} \end{cases} \quad (9)$$

for some small tolerance value $\epsilon > 0$ and distance measure ρ . If $T(\cdot)$ are sufficient statistics, and $\epsilon \rightarrow 0$ so that $\pi(\mathcal{D}|\mathcal{D}', x)$ reduces to a point mass at $T(\mathcal{D}) = T(\mathcal{D}')$ then $\pi_{ABC}(x|\mathcal{D}) = \pi(x|\mathcal{D})$ is recovered exactly, otherwise the ABC approximation to $\pi(x|\mathcal{D})$ is of the form (8), with greater accuracy for smaller ϵ . The computational overhead of all ABC samplers increases as ϵ decreases, producing a trade off between computation and accuracy. ABC methods either sample from the joint density $\pi(x, \mathcal{D}'|\mathcal{D})$ by arranging to cancel out the intractable likelihood in a weight or acceptance probability, or sample from $\pi_{ABC}(x|\mathcal{D})$ directly via Monte Carlo integration

$$\pi_{ABC}(x|\mathcal{D}) \approx \frac{\pi(x)}{S} \sum_{s=1}^S \pi(\mathcal{D}|\mathcal{D}'_s, x), \quad (10)$$

where $\mathcal{D}'_1, \dots, \mathcal{D}'_S \sim \pi(\mathcal{D}'|x)$ are draws from the likelihood given x . Almost all current ABC methods have the weighting density (9) written directly into the algorithm.

We apply the SMC sampler PRC algorithm in the ABC framework as follows: The target $\pi_t(x_t) = \pi_{ABC,t}(x_t|\mathcal{D})$ is given by (8), with the weighting function $\pi_t(\mathcal{D}|\mathcal{D}', x)$

parameterized by a different scaling parameter ϵ_t for each t , where $\infty = \epsilon_1 \geq \dots \geq \epsilon_T$, produces increasing accuracy at each step, t . The ϵ_t sequence and its length, T , may be determined *a priori* or dynamically. Evaluation of $\pi_t(x_t)$ is defined by (10) through S Monte Carlo draws. Given the high computational overheads of ABC methods, we avoid evaluating the PRC normalizing constant (as demonstrated in Section 2.3), through a global mutation kernel $M_t(x_t)$, the backward kernel L_{t-1}^{opt} (c.f. 5) and enforced resampling.

SMC sampler PRC-ABC algorithm

Initialization: Set $t = 1$.

For $i = 1, \dots, N$, sample $x_1^{(i)} \sim \mu(x)$, and set weights $W_t(x_1^{(i)}) = \pi_{ABC,1}(x_1^{(i)}|\mathcal{D})/\mu(x_1^{(i)})$.

Resample: Resample N particles with respect to $W_t(x_t^{(i)})$ and set $W_t(x_t^{(i)}) = \frac{1}{N}$,

$i = 1, \dots, N$.

Mutation and correction: Set $t = t + 1$ and $i = 1$:

(a) Sample $x_t^{(i)} \sim M_t(x_t)$ and set weight for $x_t^{(i)}$ to

$$W_t(x_t^{(i)}) = \pi_{ABC,t}(x_t^{(i)}|\mathcal{D})/M_t(x_t^{(i)}).$$

(b) With probability $1 - p^{(i)} = 1 - \min\{1, W_t(x_t^{(i)})/c_t\}$, reject $x_t^{(i)}$ and go to (a).

(c) Otherwise, accept $x_t^{(i)}$ and set $W_t(x_t^{(i)}) = W_t(x_t^{(i)})/p^{(i)}$.

(d) Increment $i = i + 1$. If $i \leq N$, go to (a).

(e) If $t < T$ then go to Resample.

The density $\mu(x)$ is an initial sampling distribution, from which direct sampling is available. As with the tolerance ϵ_t , the PRC thresholds c_t may also be determined dynamically (see below for an illustration). Note that as the resampled particles in the Resample step play no subsequent part in the sampler, in practice this step can be omitted. The path of each particle $(x_1^{(i)}, \dots, x_T^{(i)}) \in E^T$ can be recon-

structured post-simulation, if required, by resampling the recorded marginal populations $(W_t(x_t^{(i)}), x_t^{(i)})$.

The above algorithm has a number of benefits over existing SMC-type ABC samplers (Sisson et al. 2007; Toni et al. 2009; Beaumont et al. 2009; Del Moral et al. 2008). Firstly, the weighting density $\pi(\mathcal{D}|\mathcal{D}', x)$ can take any form – we suggest any smoothing kernel, following Blum (2009). Existing samplers in the literature are restricted to the uniform function (9). Secondly, there is complete control over the PRC threshold, c_t , unlike Sisson et al. (2007) who impose a specific value. Thirdly, in estimating $\pi_{ABC}(x_t|\mathcal{D})$, as long as the L_{t-1}^{opt} backward kernel is used, any number $S \geq 1$ of Monte Carlo draws can be used in (10). Existing samplers only use $S = 1$, and so there is less control over the variability of the weights. Finally, providing that the computation required to estimate the PRC normalizing constants, $r(c_t, x_{t-1})$, is acceptable, a form of the SMC sampler PRC-ABC sampler may be constructed which uses arbitrary mutation and backward kernels, allowing the user to select the most appropriate tools for a given problem.

4.1 Simulation study

We now demonstrate the superior performance of the SMC sampler PRC-ABC algorithm through a controlled study. Specifically, we specify the true posterior $\pi(x|\mathcal{D})$ as $N(0, 1)$ by defining the likelihood and prior as $\mathcal{D} \sim N(x, 1)$ and $\pi(x) \propto 1$, with a single observed datum, $\mathcal{D} = 0$. For this model, a sufficient statistic is $T(\mathcal{D}) = \mathcal{D}$. From (8), for the uniform weighting density (9) with $\rho(a, b) = |a - b|$ and $\epsilon = \epsilon_u$, or for $\pi(\mathcal{D}' | \mathcal{D}, x) = N(\mathcal{D}, \epsilon_g^2)$, then $\pi_{ABC}(x|\mathcal{D})$ may be obtained in closed form as

$$\pi_{ABC}(x|\mathcal{D}) \propto \frac{\Phi(\epsilon_u - x) - \Phi(-\epsilon_u - x)}{2\epsilon_u} \quad \text{or} \quad \pi_{ABC}(x|\mathcal{D}) = N(0, 1 + \epsilon_g^2)$$

respectively, where $\Phi(\cdot)$ denotes the standard Gaussian CDF. In both cases $\pi_{ABC}(x|\mathcal{D}) \rightarrow N(0, 1)$ as $\epsilon \rightarrow 0$. In order to directly compare the two approximate posteriors we impose equal variances on the two weighting functions, so that $\epsilon_g = \sqrt{3}\epsilon_u$.

We adopt the following sampler specifications: A particle population of size $N = 1000$ was drawn from the initial sampling distribution $\mu(x) \sim U(-5, 5)$, and the sequence of distributions, π_1, \dots, π_{10} , is defined by $\{\epsilon_t\} = \{\infty, 10, 5, 2, 1, 0.5, 0.2, 0.1, 0.05, 0.05\}$, on the ϵ_g scale. The mutation kernel $M_t(x_t) = \sum_{j=1}^N W_{t-1}^{(j)} \psi(x_t|x_{t-1}^{(j)}, \tau^2)$ is taken as a Normal kernel density estimate of $\pi_{t-1}(x_{t-1})$, with $\tau^2 = 1$ and where $\psi(x|\mu, \sigma^2)$ denotes the PDF of a $N(\mu, \sigma^2)$ distribution evaluated at x . We initially use $S = 1$ Monte Carlo draws to approximate $\pi_{ABC}(x|\mathcal{D})$ (c.f. 10).

Figure 1 examines the effect of PRC on the effective sample size (ESS), the variance of the importance weights and the mean number of rejections per particle. The PRC threshold was determined dynamically at each iteration as $c_t = Q(W_t^+(x_t), q)$, the q -th quantile of the *non-zero* weights at time t (obtained by mutating all x_{t-1} particles under M_t before implementing the PRC stage), for $q = 0, 0.5, 0.75, 0.85, 0.9, 0.95, 0.99, 0.995, 0.999$. Results are shown using the Gaussian (left plots) and uniform (right plots) weighting density, based on 250 sampler replications. Note that the PRC threshold with $q = 0$ approximately corresponds to a standard SMC sampler (“No-PRC”) only for the Gaussian weighting function, as the uniform weighting function permits exactly zero importance weights. Setting $q = 0$ for the uniform weighting density corresponds to existing SMC-type ABC samplers (Sisson et al. 2007; Toni et al. 2009; Beaumont et al. 2009).

For both weighting densities, the effective sample size increases as c_t increases, and the variance of the importance weights decreases. Naturally, the higher the PRC threshold, the more rejections occur, quantifying the extra computation required for the gains in sampler performance. However, there is a notable difference in the

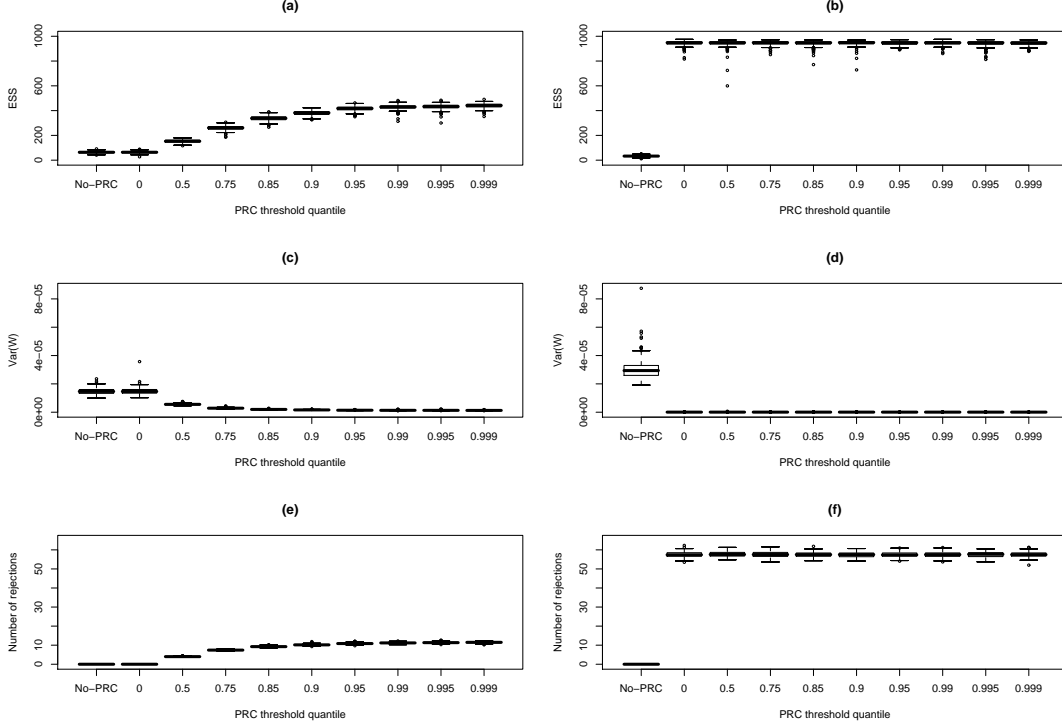


Figure 1: Effective sample size (a,b), variance of normalized importance weights (c,d) and mean number of rejections per particle (e,f) as functions of PRC threshold c_t . PRC threshold is defined dynamically as a quantile of the non-zero weights at time t (x-axis). Left plots (a,c,e) and right plots (b,d,f) are obtained under the Gaussian and uniform weighting densities $\pi(\mathcal{D}|\mathcal{D}', x)$ respectively. Boxplots are based on 250 sampler replications.

transition from poor (no PRC) to improved (under PRC) performance between the two different weighting densities. This occurs as the uniform weighting density only permits 0/1 weights, compared to the smoother scale under the Gaussian. As a result, the uniform weighting density (which is the only choice under existing ABC samplers) has a fixed, albeit strong, performance gain over not implementing PRC, but at a very high computational cost (Figure 1,f). Comparison with panels (a,c,e) suggests that considerable computational gains can be achieved with alternative weighting functions, without sacrificing sampler performance. This is easily permitted under

the SMC sampler PRC framework.

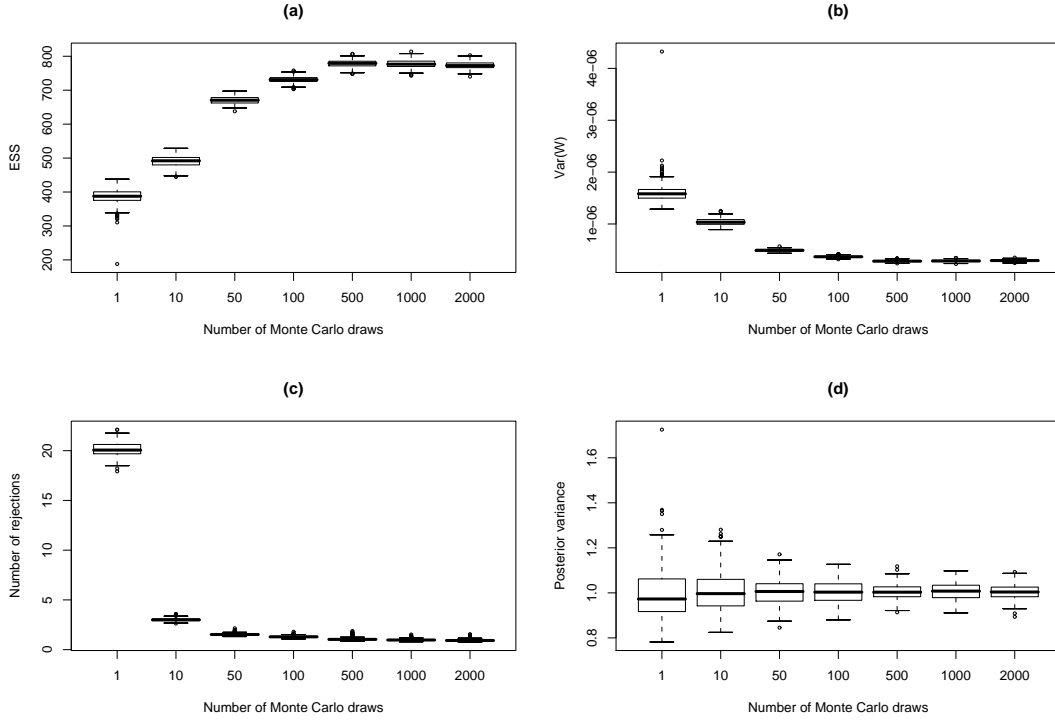


Figure 2: The effect of the number of Monte Carlo draws, S , on sampler performance. Panels show (a) effective sample size (ESS), (b) variance of the importance weights, (c) the mean number of rejection attempts, and (d) estimates of the posterior variance (true value ≈ 1), as a function of the number of Monte Carlo draws in the estimation of $\pi_{ABC}(x|\mathcal{D})$.

When using the L_{t-1}^{opt} backward kernel (5), any number $S \geq 1$ of Monte Carlo draws may be used to approximate $\pi_{ABC}(x|\mathcal{D})$ via (10). While $S = 1$ is near universal under existing ABC algorithms, one would expect to realize less variable importance weights for $S > 1$. Figure (2) illustrates the effect of increasing S , under the Gaussian density function, based on the PRC threshold $c_t = Q(W_t^+(x_t), 0.95)$. An increase in the effective sample size (panel a) is reflected by the reduction in the variability of the importance weights (panel b), as is the variability in the estimates of the posterior variance (panel d). This in turn results in lower numbers of rejections at the PRC

stage (panel c). Of course, these performance gains are again balanced by the strong increases in computation required for $S > 1$. It would appear that unless the data-generation procedure $\mathcal{D}'_s \sim \pi(\mathcal{D}|x)$ is computationally inexpensive, $1 \leq S \leq 10$ would seem to be the most useful choice in practice. Regardless, the greatest gains in sampler performance under the SMC sampler PRC algorithm are achieved for $S = 1$.

5 A stochastic claims reserving analysis

We present an analysis of an important and popular class of statistical models in actuarial science using stochastic claims reserving. We consider a time series formulation of the distribution-free chain ladder model (Mack 1993; Gisler and Wüthrich 2008; Peters et al. 2008). For a claim on an insurance company for an accident in year i , $C_{i,j}$ denotes the cumulative claim in subsequent years $j \geq i$. Cumulative claims can refer to payments, claims incurred and other expenses. At time I , we have observations $\mathcal{D}_I = \{C_{i,j}; i + j \leq I\}$, and for reserving against future claims we wish to predict $\mathcal{D}_I^c = \{C_{i,j}; i + j > I, i \leq I\}$. One such dataset is illustrated in Table 1.

Under a time series formulation, cumulative claims $C_{i,j}$ in different accident years i are independent and satisfy, for $j = 0, \dots, I - 1$,

$$C_{i,j+1} = f_j C_{i,j} + \sigma_j \sqrt{C_{i,j}} \varepsilon_{i,j+1}, \quad (11)$$

where $\mathbf{f} = (f_0, \dots, f_{I-1})$ and $\boldsymbol{\sigma} = (\sigma_0, \dots, \sigma_{I-1})$ are respectively the chain ladder factors and standard deviations, and the residuals $\varepsilon_{i,j}$ are i.i.d. with mean 0 and variance 1. The model is constrained such that $P(C_{i,j} > 0 | \{C_{k,0}\}_{k=1}^j, \mathbf{f}, \boldsymbol{\sigma}) = 1$ for all i, j (see Peters et al. 2008). If distributional assumptions are made on the residuals $\varepsilon_{i,j}$ (e.g. Yao 2008), the posterior distribution can be made computationally tractable. However, a primary intention of this model is to work with distribution-free assumptions on the residuals, and therefore on the cumulative claims. Within this distribution-free

context one wishes to quantify popular risk metrics such as value-at-risk and expected-shortfall to be calculated for the predicted claims distribution, both of which are highly relevant to regulatory reporting. Alternative approaches, based on credibility results, can relax such distributional assumptions, but can only provide statements on posterior first and second moments in limiting cases (Gisler and Wüthrich 2008).

Previously, actuaries have proceeded by predicting claims via a deterministic model known as the classical chain ladder algorithm. This approach predicts unobserved future cumulative claims by the recursion $\hat{C}_{i,I-i} = C_{i,I-i}$, and for $j > I - i$

$$\hat{C}_{i,j} = \hat{C}_{i,j-1} \hat{f}_{j-1}^{(CL)} \quad \text{where} \quad \hat{f}_{j-1}^{(CL)} = \frac{\sum_{i=0}^{I-j} C_{i,j}}{\sum_{i=0}^{I-j} C_{i,j-1}}, \quad (12)$$

and where, in the time series formulation, the variances are estimated by

$$\hat{\sigma}_j^{2(CL)} = \frac{1}{I-j-1} \sum_{i=0}^{I-j-1} C_{i,j} \left(\frac{C_{i,j+1}}{C_{i,j}} - \hat{f}_j^{(CL)} \right)^2.$$

See Mack (1993) for an estimator of $\hat{\sigma}_{I-1}^{2(CL)}$. As this algorithm is deterministic there is strong interest in stochastic chain ladder models, which naturally allow the quantification of uncertainty, such as the mean square error of prediction. In the claims reserving setting the most popular stochastic models are those with estimators which recover the classical chain ladder estimators. We consider one such Bayesian stochastic model which has the property that as the diffusivity of the priors $\pi(\mathbf{f}, \boldsymbol{\sigma})$ tends to infinity $\hat{\mathbf{f}}^{(MMSE)} \rightarrow \hat{\mathbf{f}}^{(CL)}$ where $MMSE$ denotes the posterior mean (Gisler and Wüthrich 2008).

Hence by (12) the posterior mean $E[C_{i,J}|\mathcal{D}_I] = \hat{C}_{i,J}$ recovers the classical estimators, thereby justifying the classical model. We sample from the intractable posterior $\pi_{ABC}(\mathbf{f}, \boldsymbol{\sigma}|\mathcal{D}_I)$ using the SMC sampler PRC-ABC algorithm.

5.1 Analysis and results

This model is interesting as the intractability of the likelihood directly impacts the ability to generate synthetic data sets, \mathcal{D}'_I , from the model. That is, if the distri-

butional form of the residuals were known, data-generation from the model would be trivial. To retain a distribution-free setting we alternatively utilise a conditional bootstrap approach (Peters et al. 2008). Conditional upon proposed parameters \mathbf{f} and $\boldsymbol{\sigma}$, the residuals $\tilde{\epsilon}_{i,j}|\mathbf{f},\boldsymbol{\sigma}$ are iteratively obtained by inversion of (11). Then, by independently drawing resampled residuals from the empirical conditional residual distribution, a bootstrap sample of the cumulative claims \mathcal{D}'_I is then available through recursion on (11).

In analyzing the real claims reserving data in Table 1 we specify independent priors $f_j \sim \text{Gamma}(\alpha_j, \beta_j)$ with mean $\alpha_j\beta_j = \hat{f}_j^{(CL)}$ and $\sigma_j \sim \text{IGamma}(a_j, b_j)$ with mean $b_j/(a_j - 1) = \hat{\sigma}_j^{(CL)}$ for $j = 0, \dots, I - 1$, each with large variance. For summary statistics we adopt $T(\mathcal{D}') = (\mathcal{D}'_I, \mu'(\tilde{\epsilon}), s'(\tilde{\epsilon}))$ where $\mu'(\tilde{\epsilon})$ and $s'(\tilde{\epsilon})$ denote the sample mean and standard deviation of the conditionally resampled residuals $\tilde{\epsilon}_{i,j}|\mathbf{f},\boldsymbol{\sigma}$. The observed summary statistics are given by $T(\mathcal{D}) = (\mathcal{D}_I, 0, 1)$ following the zero mean and unit variance assumptions on the true residuals.

We implement the SMC sampler PRC-ABC algorithm with uniform weighting density (9) and $\rho(T(\mathcal{D}_I), T(\mathcal{D}'_I)) = [(T(\mathcal{D}_I) - T(\mathcal{D}'_I))^\top \Sigma^{-1} (T(\mathcal{D}_I) - T(\mathcal{D}'_I))]^{1/2}$ defined as Mahalanobis distance, where the covariance Σ is estimated following Peters et al. (2008). We use $N = 5000$ particles, PRC threshold $c_t = Q(W_t^+(x_t), 0)$ and a deterministic distribution schedule $\{\epsilon_t\} = \{\infty, 10, \dots, 0.00001\}$ with $T = 22$. The mutation kernel $M_t(x_t) = \sum_{i=1}^N W_{t-1}^{(i)}(x_{t-1}^{(i)}) \text{Gamma}(a(x_{t-1}^{(i)}), b(x_{t-1}^{(i)}))$ is a mixture of gamma densities, with mean $a(x_{t-1}^{(i)})b(x_{t-1}^{(i)}) = x_{t-1}^{(i)}$ and large variance.

Table 2 presents a comparison of the parameter estimates $\hat{\mathbf{f}}$ and $\hat{\boldsymbol{\sigma}}$, and predicted cumulative claims, $\hat{C}_{i,j}$ under classical and Bayesian models. Given the uninformative priors, the posterior mean estimates and resulting predicted claims agree well with those obtained under the classical model. This provides some validation for the deterministic classical model estimates under the Bayesian stochastic interpretation.

Perhaps more usefully for inference, the full posterior $\pi_{ABC}(\mathbf{f}, \boldsymbol{\sigma} | \mathcal{D}_I)$ is available. Figure 3 illustrates how the estimated marginal densities of the first chain ladder factor, $\pi_{ABC,t}(f_0 | \mathcal{D}_I)$, and the associated standard deviation, $\pi_{ABC,t}(\sigma_0 | \mathcal{D}_I)$, evolve as ϵ_t decreases. The precision of the densities clearly improves, as decreasing ϵ_t imposes stricter restrictions on the permissible deviations of the ABC approximate posterior $\pi_{ABC}(\mathbf{f}, \boldsymbol{\sigma} | \mathcal{D})$ from the target posterior $\pi(\mathbf{f}, \boldsymbol{\sigma} | \mathcal{D})$. A full predictive analysis may now follow, including upper and lower credible bounds on predicted future claims.

6 Discussion

When used in challenging settings, sequential Monte Carlo samplers often suffer from severe particle degeneracy. In this article we have provided a practical approach to tackling this problem, by incorporating the partial rejection control mechanism of Liu (2001) directly into the mutation kernel of the SMC sampler. The resulting sampler will not worsen, and can improve the variance of the importance weights, sometimes substantially so. By establishing clear relationships with existing samplers (Del Moral et al. 2006; Le Gland and Oudjane 2004), many theoretical properties may be extended to the SMC sampler PRC algorithm, including a central limit theorem, and a proof of an almost sure finite number of PRC rejection attempts.

There is much opportunity for the specification of the sequence of PRC thresholds to be further automated, if desired. For example, by dynamically determining $c_t > c_{t-1}$ if the effective sample size of the particle population at time $t - 1$ falls too low, and conversely allowing $c_t < c_{t-1}$ if the level of PRC resampling is too high, in order to reduce computational overheads.

As the SMC sampler PRC algorithm allows practical inference in challenging situations in which particle weights are highly variable, we anticipate that a primary

application of the sampler will be within the rapidly developing “likelihood-free” approximate Bayesian computation framework. The presented sampler is more flexible and efficient than existing SMC-type ABC samplers, allowing a previously unavailable degree of control over the computation utilised for a given analysis. Perhaps more importantly, the extra flexibility achieved by allowing arbitrary weighting densities (unlike existing ABC samplers) enables the analysis of improved models within the ABC framework, in line with recent non-parametric interpretations (Blum 2009).

7 Acknowledgements

YF and SAS are supported by the Australian Research Council through the Discovery Project scheme (DP0664970, DP0877432). GWP is supported by an APA scholarship and by the School of Mathematics and Statistics, UNSW and CSIRO CMIS. The authors thank P. Shevchenko , A. Johansen and M. Briers for thoughtful discussions.

References

- Andrieu, C., N. D. Freitas, A. Doucet, and M. Jordan (2003). An introduction to MCMC for machine learning. *Machine Learning* 50, 5–43.
- Arulampalam, S., S. Maskell, N. Gorodon, and T. Clapp (2002). A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* 50(2), 174 – 188.
- Beaumont, M. A., J.-M. Cornuet, J.-M. Marin, and C. P. Robert (2009). Adaptivity for ABC algorithms: the ABC-PMC scheme. *Biometrika*, in press.
- Beaumont, M. A., W. Zhang, and D. J. Balding (2002). Approximate Bayesian computation in population genetics. *Genetics* 162, 2025 – 2035.

- Blum, M. G. B. (2009). Approximate Bayesian computation: a non-parametric perspective. Technical report, Université Joseph Fourier, Grenoble, France.
- Bortot, P., S. G. Coles, and S. A. Sisson (2007). Inference for stereological extremes. *Journal of the American Statistical Association* 102, 84–92.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Stat.* 32, 2385 – 2411.
- Cornebise, J., E. Moulines, and J. Olsson (2008). Adaptive methods for sequential importance sampling with application to state space models. In *Proc. 16th European Sig. Proc. Conference, Lausanne*.
- Crisan, D. and A. Doucet (2002). A survey of convergence results on particle filtering for practitioners. *IEEE Trans. Signal Processing* 50(3), 736–746.
- Del Moral, P. (2004). *Feynman-Kac formulae: Genealogical and interacting particle systems with applications*. Springer, New York.
- Del Moral, P., A. Doucet, and A. Jasra (2006). Sequential Monte Carlo samplers. *J. R. Statist. Soc. B* 68, 411 – 436.
- Del Moral, P., A. Doucet, and A. Jasra (2008). Adaptive sequential Monte Carlo samplers. Technical report, University of Bordeaux.
- Del Moral, P., J. Jacod, and P. Protter (2001). The Monte Carlo method for filtering with discrete-time observations. *Prob. Theor. Related Fields* 120(3), 346–368.
- Doucet, A., N. de Freitas, and N. Gordon (Eds.) (2001). *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York.
- Doucet, A. and A. M. Johansen (2009). A tutorial on particle filtering and smoothing: Fiteen years later. In *Oxford Handbook of Nonlinear Filtering*, D. Crisan and B. Rozovsky (eds.). Oxford University Press. To appear.

- Gisler, A. and M. V. Wüthrich (2008). Credibility for the chain ladder reserving method. *Astin Bulletin* 38, 565–600.
- Johansen, A. and A. Doucet (2008). A note on auxiliary particle filters. *Statistics and Probability Letters* 78(12), 1498–1504.
- Johansen, A. M., P. Del Moral, and A. Doucet (2006). Sequential Monte Carlo samplers for rare events. In *Proc. 6th Int. Workshop on Rare Event Simulation*.
- Kitigawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian, non-linear state space models. *Journal of Computational and Graphical Statistics* 5, 1–25.
- Kunsch, H. R. (2005). Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann. Statist.* 33, 1983 – 2021.
- Le Gland, F. and N. Oudjane (2004). Stability and uniform approximation of non-linear filters using the Hilbert metric, and applications to particle filters. *The Annals of Applied Probability* 14(1), 144–187.
- Le Gland, F. and N. Oudjane (2005). A sequential particle algorithm that keeps the particle system alive. In *Proc. 13th European Sig. Proc. Conference, Antalya*.
- Liu, J. and R. Chen (1998). Sequential Monte Carlo for dynamic systems. *Journal of the American Statistical Association* 93, 1032 – 1044.
- Liu, J., R. Chen, and W. Wong (1998). Rejection control and sequential importance sampling. *Journal of the American Statistical Association* 93, 1022–1031.
- Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag.
- Mack, T. (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. *Astin Bulletin* 23, 213–225.
- Marjoram, P., J. Molitor, V. Plagnol, and S. Tavaré (2003). Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* 100, 15324 – 15328.

- Peters, G. W. (2005). Topics in Sequential Monte Carlo samplers. Master’s thesis, University of Cambridge.
- Peters, G. W., M. V. Wüthrich, and P. V. Shevchenko (2008). Chain ladder method: Bayesian bootstrap versus classical bootstrap. Technical report, UNSW.
- Ratmann, O., C. Andrieu, T. Hinkley, C. Wiuf, and S. Richardson (2009). Model criticism based on likelihood-free inference, with an application to protein network evolution. *Proc. Natl. Acad. Sci. USA* *106*, 10576–10581.
- Sisson, S. A., Y. Fan, and M. M. Tanaka (2007). Sequential Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci.* *104*, 1760–1765. Errata (2009).
- Tavaré, S., D. J. Balding, R. C. Griffiths, and P. Donnelly (1997). Inferring coalescence times from DNA sequence data. *Genetics* *145*, 505–518.
- Toni, T., D. Welch, N. Strelkowa, A. Ipsen, and M. P. H. Stumpf (2009). Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. R. Soc. Interface* *6*, 187–202.
- West, M. (1993). Approximating posterior distributions by mixtures. *J. R. Statist. Soc. B* *55*, 409–422.
- Wüthrich, M. V. and M. Merz (2008). *Stochastic claims reserving methods in insurance*. Wiley.
- Yao, J. (2008). Bayesian approach for prediction error in chain ladder claims reserving. 38th International ASTIN Colloquium, July.

Appendix

A.1: Proof of Theorem 3.1

The proof follows the arguments presented by Liu et al. (1998). In particular we study the SMC sampler PRC algorithm in terms of χ^2 distance between the sampling distribution and the target distribution at stage t . Let

$$W_t(x_t) \propto W_{t-1}(x_{t-1}) \frac{\pi_t(x_t) L_{t-1}(x_t, x_{t-1})}{\pi_{t-1}(x_{t-1}) M_t(x_{t-1}, x_t)} \quad \text{and} \quad W_t^*(x_t) \propto W_{t-1}(x_{t-1}) \frac{\pi_t(x_t) L_{t-1}(x_t, x_{t-1})}{\pi_{t-1}(x_{t-1}) M_t^*(x_{t-1}, x_t)}.$$

Recall that the normalising constant for the mutation kernel M_t^* at time t is

$$r(c_t, x_{t-1}) = \int \min \left\{ 1, \frac{W_t(x_t)}{c_t} \right\} M_t(x_{t-1}, x_t) dx_t = \frac{1}{c_t} \mathbb{E}_{M_t} [\min \{c_t, W_t(x_t)\}].$$

The variance of the importance weight at time t from a standard SMC sampler, with respect to M_t , is given by

$$\text{Var}_{M_t}[W_t(x_t)] = \int [W_t(x_t)]^2 M_t(x_{t-1}, x_t) dx_t - \mu^2,$$

and similarly, the variance of the equivalent importance weight at time t following a PRC step under the SMC sampler PRC algorithm, with respect to M_t^* , is given by

$$\text{Var}_{M_t^*}[W_t^*(x_t)] = \int [W_t^*(x_t)]^2 M_t^*(x_{t-1}, x_t) dx_t - \mu^2,$$

where

$$\mu = \mathbb{E}_{M_t^*}[W_t^*(x_t)] = \mathbb{E}_{M_t}[W_t(x_t)].$$

We also have that

$$\begin{aligned}
\int [W_t^*(x_t)]^2 M_t^*(x_{t-1}, x_t) dx_t &= \int \left[W_{t-1}(x_{t-1}) \frac{\pi(x_t) L_{t-1}(x_t, x_{t-1})}{\pi_{t-1}(x_{t-1}) M_t^*(x_{t-1}, x_t)} \right]^2 M_t^*(x_{t-1}, x_t) dx_t \\
&= r(c_t, x_{t-1}) \int \frac{W_{t-1}^2(x_{t-1})}{\min\{1, \frac{W_t(x_t)}{c_t}\}} \frac{\pi_t^2(x_t) L_{t-1}^2(x_{t-1})}{\pi_{t-1}^2(x_{t-1}) M_t^2(x_{t-1}, x_t)} M_t(x_{t-1}, x_t) dx_t \\
&= r(c_t, x_{t-1}) \int \max\{W_t^2(x_t), c_t W_t(x_t)\} M_t(x_{t-1}, x_t) dx_t \\
&= r(c_t, x_{t-1}) \mathbb{E}_{M_t} [\max\{W_t(x_t), c_t\} W_t(x_t)] \\
&= \frac{1}{c_t} \mathbb{E}_{M_t} [\min\{c_t, W_t(x_t)\}] \mathbb{E}_{M_t} [\max\{W_t(x_t), c_t\} W_t(x_t)] \\
&\leq \frac{1}{c_t} \mathbb{E}_{M_t} [\min\{c_t, W_t(x_t)\} \max\{W_t(x_t), c_t\} W_t(x_t)] \\
&= \frac{1}{c_t} \mathbb{E}_{M_t} [c_t W_t^2(x_t)] = \mathbb{E}_{M_t} [W_t^2(x_t)].
\end{aligned}$$

The above inequality holds since the random variables $\min\{c_t, W_t(x_t)\}$ and $\max\{W_t(x_t), c_t\} W_t(x_t)$ are positively correlated (see Liu 2001), and so

$$\begin{aligned}
&\mathbb{E}_{M_t} [\min\{c_t, W_t(x_t)\} \max\{W_t(x_t), c_t\} W_t(x_t)] \\
&\quad - \mathbb{E}_{M_t} [\min\{c_t, W_t(x_t)\}] \mathbb{E}_{M_t} [\max\{W_t(x_t), c_t\} W_t(x_t)] \geq 0.
\end{aligned}$$

Hence

$$\mathbb{V}ar_{M_t^*}[W_t^*(x_t)] \leq \mathbb{E}_{M_t} [W_t^2(x_t)] - \mu^2 = \mathbb{V}ar_{M_t}[W_t(x_t)].$$

□

A.2: Proof of Theorem 3.2

To satisfy the condition $\mathbb{E}_{M_t} [W_t(x_t) \mid x_{t-1} = x] > 0, \forall x \in E$, for the SMC sampler PRC algorithm, we require

$$\int_E W_{t-1}(x) w_t(x, x_t) r(c_t, x) \left[\min \left\{ 1, \frac{W_{t-1}(x) w_t(x, x_t)}{c_t} \right\} \right]^{-1} M_t(x, x_t) dx_t > 0. \quad (13)$$

For a particle $x_{t-1} = x$, the proposed state x_t can take values in a support which can be split into two regions, $A(x)$ and $A^c(x)$, such that $A(x) \cup A^c(x) = E$. These respectively correspond to when $\min \left\{ 1, \frac{W_{t-1}(x) w_t(x, x_t)}{c_t} \right\} = 1$ (i.e. particle acceptance probability under PRC is 1) and when $\min \left\{ 1, \frac{W_{t-1}(x) w_t(x, x_t)}{c_t} \right\} = W_{t-1}(x) w_t(x, x_t) / c_t$

(i.e. the particle may be rejected under PRC). Note that in the extreme cases, $c_t \leq W_{t-1}(x)w_t(x, x_t) \Rightarrow A^c(x) = \emptyset$ reduces to the SMC sampler algorithm, and $c_t > \sup_{x_t} W_{t-1}(x)w_t(x, x_t) \Rightarrow A(x) = \emptyset$. More generally, (13) may be expanded as

$$\frac{r(c_t, x)W_{t-1}(x)}{\pi_{t-1}(x)} \int_{A(x)} \pi_t(x_t) L_{t-1}(x_t, x) dx_t + c_t r(c_t, x) \int_{A^c(x)} M_t(x, x_t) dx_t > 0$$

which is always greater than zero for finite $c_t < \infty$ as $0 < r(c_t, x) \leq 1$.

□

Accident Year, i	Development Year, j									
	0	1	2	3	4	5	6	7	8	9
0	594.6975	372.1236	89.5717	20.7760	20.6704	6.2124	6.5813	1.4850	1.1130	1.5813
1	634.6756	324.6406	72.3222	15.1797	6.7824	3.6603	5.2752	1.1186	1.1646	
2	626.9090	297.6223	84.7053	26.2768	15.2703	6.5444	5.3545	0.8924		
3	586.3015	268.3224	72.2532	19.0653	13.2976	8.8340	4.3329			
4	577.8885	274.5229	65.3894	27.3395	23.0288	10.5224				
5	618.4793	282.8338	57.2765	24.4899	10.4957					
6	560.0184	289.3207	56.3114	22.5517						
7	528.8066	244.0103	52.8043							
8	529.0793	235.7936								
9	567.5568									

to be predicted $Y_{i,j}$

Table 1: A claims development triangle. Upper triangle denotes observed annual claims $Y_{i,j}$ from which $C_{i,j} = \sum_{k=0}^j Y_{i,k} \in \mathcal{D}_I$ may be obtained; lower triangle denotes annual $Y_{i,j}$ and cumulative claims $C_{i,j} \in \mathcal{D}_I^c$ to be predicted. Data are real insurance figures in units of \$10,000 (c.f. Wüthrich and Merz, 2008). The triangle assumes that the number of accident years is equal to the number of observed development periods.


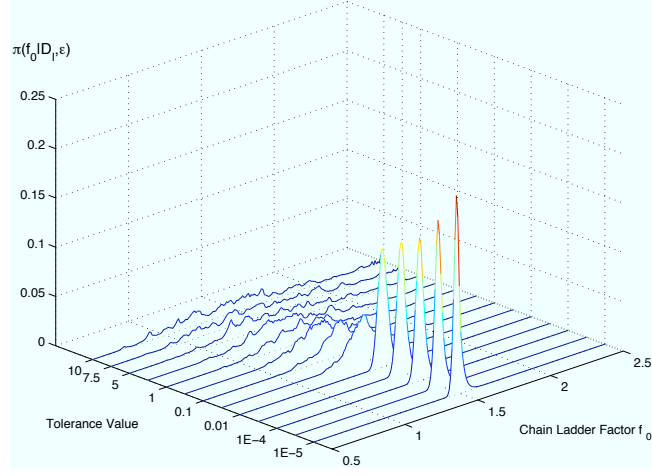


Figure 3: Evolution of the marginal posterior density estimates of the chain ladder factor $\pi_t(f_0|\mathcal{D}_I)$ (left) and the associated standard deviation $\pi_t(\sigma_0|\mathcal{D}_I)$ (right) as a function of ϵ_t .

Parameters	Year	0	1	2	3	4	5	6	7	8	9	$\widehat{C}_{i,I} - C_{i,I-i}$
$(\mathbf{f}^{(CL)})$	0											0
$(\mathbf{f}^{(MMSE)})$												0
$(\mathbf{f}^{(CL)})$	1											15, 126
$(\mathbf{f}^{(MMSE)})$												15, 972
$(\mathbf{f}^{(CL)})$	2											26, 257
$(\mathbf{f}^{(MMSE)})$												25, 540
$(\mathbf{f}^{(CL)})$	3											34, 538
$(\mathbf{f}^{(MMSE)})$												36, 023
$(\mathbf{f}^{(CL)})$	4											85, 302
$(\mathbf{f}^{(MMSE)})$												90, 283
$(\mathbf{f}^{(CL)})$	5											156, 494
$(\mathbf{f}^{(MMSE)})$												175, 886
$(\mathbf{f}^{(CL)})$	6											286, 121
$(\mathbf{f}^{(MMSE)})$												306, 953
$(\mathbf{f}^{(CL)})$	7											449, 167
$(\mathbf{f}^{(MMSE)})$												471, 761
$(\mathbf{f}^{(CL)})$	8											1, 043, 242
$(\mathbf{f}^{(MMSE)})$												1, 050, 760
$(\mathbf{f}^{(CL)})$	9											3, 950, 814
$(\mathbf{f}^{(MMSE)})$												3, 966, 655
$\widehat{f}_j^{(CL)}$		1.4925	1.0778	1.0229	1.0148	1.0070	1.0051	1.0011	1.0010	1.0014		6, 047, 061
$\sigma_j^{(CL)}$		135.253	33.803	15.760	19.847	9.336	2.001	0.823	0.219	0.059		
$\widehat{f}_j^{(MMSE)}$		1.4937	1.0759	1.0233	1.0151	1.0084	1.0055	1.0013	1.0009	1.0015		6, 139, 834
$\sigma_j^{(MMSE)}$		132.917	34.566	14.742	21.972	8.547	2.736	0.789	0.159	0.061		

Table 2: Predicted parameter estimates, $\widehat{\mathbf{f}}, \widehat{\boldsymbol{\sigma}}$, cumulative chain ladder claims, $\widehat{C}_{i,j}$, and estimated chain ladder reserves under the classical (CL) and Bayesian ($MMSE$) models.