Minimum variance importance sampling via Population Monte Carlo

R. Douc*, A. Guillin†, J.-M. Marin‡, and C.P. Robert§

Abstract

Variance reduction has always been a central issue in Monte Carlo experiments. Population Monte Carlo can be used to this effect, in that a mixture of importance functions, called a D-kernel, can be iteratively optimised to achieve the minimum asymptotic variance for a function of interest among all possible mixtures. The implementation of this iterative scheme is illustrated for the computation of the price of a European option in the Cox-Ingersoll-Ross model,

Keywords: Adaptivity, Cox-Ingersoll-Ross model, Euler scheme, importance sampling, mathematical finance, mixtures, population Monte Carlo, variance reduction.

AMS Classification: 60F05, 62L12, 65-04, 65C05, 65C40, 65C60.

1 Introduction

The main bulk of the literature on Monte Carlo methods concentrates on the approximation of integrals

$$\pi(h) = \int_\Omega h(x)\pi(x)\,\mu(dx)\,,$$

where μ is a measure on Ω, π a density and h a π-measurable function on the same set Ω. In particular, the focus of many studies is to reduce the variance of estimators of $\pi(h)$, whether locally, that is, for a given model and a given function, or globally, as in for instance Rao–Blackwellisation, control and antithetic variate or quasi-random

*CMAP, École Polytechnique, Palaiseau douc@cmapx.polytechnique.fr

†CEREMADE, Université Paris Dauphine, and TSI, École Nationale des Télécommunications guillin@ceremade.dauphine.fr

‡CEREMADE, Université Paris Dauphine and Projet SELECT, INRIA FUTURS, Université Paris-Sud marin@ceremade.dauphine.fr

§Corresponding author: CEREMADE, Université Paris Dauphine, and CREST, INSEE, Paris xian@ceremade.dauphine.fr

techniques (see, e.g., Rubinstein, 1981, Robert and Casella, 2004). An illustration of this focus is provided in mathematical Finance by the numerous improvements brought upon the computation of option prices (see, e.g., Lapeyre et al., 1998, Glasserman, 2003, Jackel, 2002).

In the particular case of importance sampling estimators, that is,

$$\hat{\pi}^{IS}_{g,N}(h) = N^{-1} \sum_{i=1}^{N} h(x_i)\pi(x_i)/g(x_i)\,, \qquad x_1, \ldots, x_N \overset{\text{iid}}{\sim} g\,,$$

where g is a distribution dominating π (with density denoted by g), the variance is equal to

$$g\left[(h\pi/g - \pi(h))^2\right]/N\,,$$

if $\pi(h^2\pi/g) < \infty$. A puzzling feature of this class of estimators is the well-known optimality of the importance distribution

$$g^\star(x) = |h(x)|\pi(x) \Big/ \int |h(y)|\pi(y)\,\mu(dy)$$

when aiming at minimising the variance of $\hat{\pi}^{IS}_{g,N}$. This result (Rubinstein, 1981) is paradoxical in that it produces a zero variance estimator when h is either positive or negative (indeed, in both cases, $\hat{\pi}^{IS}_{g^\star,N} = \pi(h)$). The paradox is only superficial, though, in that it points out the fact that, in Monte Carlo settings, there is no ultimate importance function when there is no restriction on the choice of these functions (and when the costs of constructing and simulating these distributions are not taken into account). In particular, $g^\star$ cannot be used in practice because it depends on the integral $\int |h(y)|\pi(y)\,\mu(dy)$. This result is thus rather understood as providing a goal for choosing a importance function g tailored for the approximation of $\pi(h)$.

If the normalizing constants of either the target distribution π or the importance function g are unknown, an alternative to $\hat{\pi}^{IS}_{g,N}$ is the self-normalised importance sampling estimator, that is

$$\hat{\pi}^{SNIS}_{g,N}(h) = \sum_{i=1}^{N} h(x_i)\pi(x_i)/g(x_i) \Big/ \sum_{i=1}^{N} \pi(x_i)/g(x_i)\,. \qquad x_1, \ldots, x_N \overset{\text{iid}}{\sim} g\,,$$

where the sum of the weights normalises the weighted sum. If $g\left((1+h^2)(\pi/g)^2\right) < \infty$, the asymptotic variance of $\hat{\pi}^{SNIS}_{g,N}(h)$ is given by $\pi\left[(h-\pi(h))^2\pi/g\right]$. In this case, $g^\star$ is no longer the best choice: rather,

$$g^\sharp(x) = |h(x) - \pi(h)|\pi(x) \Big/ \int |h(y) - \pi(h)|\pi(y)\,\mu(dy)$$

minimizes (in g) the asymptotic variance of $\hat{\pi}^{SNIS}_{g,N}(h)$. This second optimum is not available either, because it still depends on $\pi(h)$.

The formal aspect of this optimality result may explain why there is little in the literature besides general recommendations that the support of g should be the support of $|h(x)|\pi(x)$ or of $|h(y) - \pi(h)|\pi(y)$, or yet that the tails of g should be at least as thick as those of $|h(x)|\pi(x)$. Note however that a recent reference is the *cross-entropy* method of Rubinstein and Kroese (2004) where the parameter of a family of proposals is optimised, either directly or by an iterative process, to reach minimal variance or maximal entropy against the target $|h(x)|\pi(x)$, the function h being of the specific rare event shape $h(x) = \mathbb{I}(S(x) \leq \gamma)$. The population Monte Carlo methodology studied in this paper encompasses cross-entropy as a special case.

The current paper establishes that the population Monte Carlo (PMC) technique of Cappé et al. (2004) and Douc et al. (2005) can easily be adapted to this purpose and can result in considerable variance reduction. We recall that Cappé et al. (2004) introduced this method, following the denomination of Iba (2000), to advertise the availability of universal adaptive sampling machines that do not encounter the formidable difficulties of designing adaptive MCMC algorithms. Douc et al. (2005) showed in addition that those PMC algorithms can accomodate a progressive adaption to a given target distribution with a diminishing Kullback divergence. We now explain why this is also the case for variance diminution and optimal importance function approximation.

In Section 2, we recall the main features of the PMC algorithm, including the expressions for the asymptotic variances of the unnormalised and self-normalised versions of the PMC estimator. In Section 3, we establish that our updating scheme for the mixture weights in the PMC algorithm does induce a decrease in the asymptotic variance at each step. Section 4 provides an additional improvement through the cumulated estimation of $\pi(h)$. In Section 5, we illustrate the variance reduction for a toy example before launching into the evaluation of a European option price for the Cox-Ingersoll-Ross model.

2 Population Monte Carlo

2.1 Monte Carlo setting

We suppose that the target distribution π is *at least* known up to a normalizing constant, $\pi(x) \propto \tilde{\pi}(x)$ with $\tilde{\pi}$ known. For the importance sampling estimation of $\pi(h)$, the quality of both the unnormalised and the self-normalised approximations to $\pi(h)$ strongly depends on the choice of the proposal distribution g, a choice that is quite delicate for complex distributions like those that occur in high dimensional problems.

We first recall that sampling importance resampling (SIR) (Rubin, 1987, 1988) can be used to reset a given weighted sample from g to a sample from the target distribution π. Once the importance weights are derived, $\omega_i \propto \pi(x_i)/g(x_i)$, a (non-iid) sample from π, $\tilde{x}_1, \ldots, \tilde{x}_M$ can be derived from the instrumental sample $x_1, \ldots, x_N$ by resampling using

the importance weights in $\{x_1, \ldots, x_N\}$, that is,

$$\tilde{x}_i = x_{J_i}\,, \qquad 1 \leq i \leq M\,,$$

where the random variables $J_1, \ldots, J_M$ are distributed as

$$\mathbb{P}\left[J_l = i | x_1, \ldots, x_N\right] = \left(\sum_{j=1}^{N} \frac{\pi(x_j)}{g(x_j)}\right)^{-1} \frac{\pi(x_i)}{g(x_i)}$$

(see, e.g., Robert and Casella, 2004, Section 14.3.5). Multinomial sampling is a possible implementation of the SIR methodology but more efficient alternatives that reduce the variance of the resulting estimators are also available. However, to keep the description of the algorithms as simple as possible, we will use multinomial sampling in the following.

The *Population Monte Carlo* (PMC) method introduced in Cappé et al. (2004) intrinsically is a form of iterated sampling importance resampling with dynamically adapted importance functions. We refer to Cappé et al. (2004), and to Robert and Casella (2004, Chap. 14) for details on the motivations and foundations of this method, and we simply recall the essential feature of the method: At iteration t of the PMC algorithm, N values are simulated from a proposal distribution and this proposal distribution is based on the $N \times (t-1)$ past realizations, with basically *no constraint* on the form of dependence on the past.

If we define *renormalized* importance weights associated with weights $\omega_{j,t}$ $(1 \leq j \leq N)$ as

$$\overline{\omega}_{i,t} = \omega_{i,t} \Big/ \sum_{j=1}^{N} \omega_{j,t}\,,$$

the generic PMC algorithm reads as follows:

–Generic PMC algorithm–
At time 0,

a) Generate $(x_{i,0})_{1\leq i\leq N} \overset{\text{iid}}{\sim} g_0$ and compute $\omega_{i,0} = \pi(x_{i,0})/g_0(x_{i,0})$;

b) Generate $(J_{i,0})_{1\leq i\leq N} \overset{\text{iid}}{\sim} \mathcal{M}(1, (\overline{\omega}_{i,0})_{1\leq i\leq N})$ and set $\tilde{x}_{i,0} = x_{J_{i,0},0}$ $(1 \leq i \leq N)$.

At time $1 \leq t \leq T$

a) Conditionally on past $x_{i,j}$'s and $\tilde{x}_{i,j}$'s, generate independently $x_{i,t} \sim g_{i,t}$ and compute $\omega_{i,t} = \pi(x_{i,t})/g_{i,t}(x_{i,t})$;

b) Generate $(J_{i,t})_{1\leq i\leq N} \overset{\text{iid}}{\sim} \mathcal{M}(1, (\overline{\omega}_{i,t})_{1\leq i\leq N})$ and set $\tilde{x}_{i,t} = x_{J_{i,t},t}$ $(1 \leq i \leq N)$.

Obviously, the quasi-total freedom in the construction of the above $g_{i,t}$'s has drawbacks, namely that some proposals do not necessarily lead to improvements in terms of variance reduction or target approximation. Therefore, we now restrict the family of proposals from which to select the new $g_{i,t}$'s to mixture of fixed proposals and we establish in the next section that variance improvement does occur within this family. This particular type of algorithm was already shown in Douc et al. (2005) to lead to a reduction in the asymptotic Kullback-Leibler distance between the target and the proposal, for a correct update in the mixture weights.

2.2 D-kernel PMC

We assume from now on that we use in parallel D fixed kernels $Q_d(\cdot,\cdot)$ with densities q_d and that the proposal is a mixture of those kernels

$$g_{i,t}(x) = \sum_{d=1}^{D} \alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x)\,, \qquad \sum_d \alpha_d^{t,N} = 1\,,$$

where the weights $\alpha_d^{t,N} > 0$ can be modified at each iteration. The amount of adaptivity we allow in this version of PMC is thus restricted to a possible modification of the weights $\alpha_d^{t,N}$. The importance weight associated with this mixture proposal is

$$\pi(x_{i,t}) \Big/ \sum_{d=1}^{D} \alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x_{i,t})$$

while simulation from $g_{i,t}$ can be decomposed in the two usual mixture steps: first pick the component d then simulate from the corresponding kernel Q_d:

–Generic D-kernel PMC algorithm–
At time 0, produce the sample $(\tilde{x}_{i,0}, J_{i,0})_{1\le i\le N}$ and set $\alpha_d^{1,N} = 1/D$ for all $1 \le d \le D$.
At time $1 \le t \le T$

a) Conditionally on the $\alpha_d^{t,N}$'s, generate

$$(K_{i,t})_{1\le i\le N} \overset{\text{iid}}{\sim} \mathcal{M}(1, (\alpha_d^{t,N})_{1\le d\le D})$$

b) Conditionally on $(\tilde{x}_{i,t-1}, K_{i,t})_{1\le i\le N}$, generate independently

$$(x_{i,t})_{1\le i\le N} \sim Q_{K_{i,t}}(\tilde{x}_{i,t-1}, \cdot)$$

and set $\omega_{i,t} = \pi(x_{i,t}) \Big/ \sum_{d=1}^{D} \alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x_{i,t})$;

c) Conditionally on $(\tilde{x}_{i,t-1}, K_{i,t}, x_{i,t})_{1\leq i\leq N}$, generate

$$(J_{i,t})_{1\leq i\leq N} \overset{\text{iid}}{\sim} \mathcal{M}(1, (\overline{\omega}_{i,t})_{1\leq i\leq N})$$

and set $(1 \leq i \leq N,\ 1 \leq d \leq D)$

$$\tilde{x}_{i,t} = x_{J_{i,t},t}\,, \qquad \alpha_d^{t+1,N} = \Psi_d\left((\tilde{x}_{i,t-1}, x_{i,t}, K_{i,t})_{1\leq i\leq N}\right)$$

such that $\sum_{d=1}^{D} \alpha_d^{t+1,N} = 1\,.$

In the above algorithm, Ψ_d $(1 \leq d \leq D)$ denotes an update function that depends upon the past iteration. We assume that the individual kernel importance weights are almost surely finite, that is,

$$\forall d \in \{1, \ldots, D\}, \overline{\pi}\left\{q_d(x, x') = 0\right\} = 0, \tag{\textbf{A1}}$$

where $\overline{\pi} = \pi \otimes \pi$. Under **(A1)**, Douc et al. (2005) proved that the updates Ψ_d of the mixture weights given by

$$\alpha_d^{t+1,N} = \sum_{i=1}^{N} \overline{\omega}_{i,t} \mathbb{I}_d(K_{i,t})$$

garantee a systematic decrease of the Kullback-Leibler distance between the target and the D-kernel mixture, a long-term run of the algorithm providing the mixture that is (entropy-) closest to the target. Moreover, Theorem 5.1 of Douc et al. (2005) leads to a LLN (in the number of simulations at a given iteration) for the output of the generic D-kernel PMC algorithm.

Theorem 2.1. *Under **(A1)**, for any function h in L^1_π and for all $t \geq 0$, both the unnormalised and the self-normalized PMC estimators are convergent,*

$$\hat{\pi}_{t,N}^{PMC}(h) = \frac{1}{N}\sum_{i=1}^{N} \omega_{i,t} h(x_{i,t}) \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \pi(h) \quad \textit{and} \quad \hat{\pi}_{t,N}^{SPMC}(h) = \sum_{i=1}^{N} \overline{\omega}_{i,t} h(x_{i,t}) \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \pi(h)\,.$$

As noted earlier, the unnormalised PMC estimator can only be used when π is completely known and even in those instances it is not necessarily improving upon the self-normalized PMC estimator.

A CLT can also be established in this setting, under the additional following integrability condition

$$\overline{\pi}\left\{(1 + h^2(x'))\frac{\pi(x')}{q_{\overline{d}}(x, x')}\right\} < \infty \text{ for a } \overline{d} \in \{1, \ldots, D\}. \tag{\textbf{A2}}$$

Note that this condition must hold only for *one* $1 \leq \overline{d} \leq D$ rather than for *all* d's. Theorem 5.2 of Douc et al. (2005) then provides a CLT for the generic D-kernel PMC algorithm.

Theorem 2.2. *Under* ***(A1)*** *and* ***(A2)****, if for all* $t \geq 1$,

$$\forall 1 \leq d \leq D, \quad \alpha_d^{t,N} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \alpha_d^t > 0\,,$$

then both

$$\sqrt{N}\left(\sum_{i=1}^{N} \overline{\omega}_{i,t} h(x_{i,t}) - \pi(h)\right) \quad \textit{and} \quad \sqrt{N}\left(\frac{1}{N}\sum_{i=1}^{N} \omega_{i,t} h(x_{i,t}) - \pi(h)\right) \tag{1}$$

converge in distribution as n *goes to infinity to normal distributions with variances*

$$\sigma_{1,t}^2 = \overline{\pi}\left((h(x') - \pi(h))^2 \frac{\pi(x')}{\sum_{d=1}^{D} \alpha_d^t q_d(x, x')}\right)$$

and

$$\sigma_{2,t}^2 = \overline{\pi}\left\{\left(\frac{\pi(x')}{\sum_{d=1}^{D} \alpha_d^t q_d(x, x')} h(x') - \pi(h)\right)^2 \frac{\sum_{d=1}^{D} \alpha_d^t q_d(x, x')}{\pi(x')}\right\}.$$

The additional condition in Theorem 2.2 is necessary to ensure a stabilisation of the weights as the number of simulations increases. It is garanteed in cases like those of Douc et al. (2005) updating scheme and we will show below that it also holds for our updating scheme. The quantities $\sigma_{1,t}^2$ and $\sigma_{2,t}^2$ exhibited in this result are thus associated with the limiting set of weights $(\alpha_1^t, \ldots, \alpha_D^t)$, defined on the simplex set of $\mathbb{R}^D$,

$$\mathscr{S}_D = \left\{\alpha = (\alpha_1, \ldots, \alpha_D);\ \forall d \in \{1, \ldots, D\},\ \alpha_d \geq 0 \quad \text{and} \quad \sum_{d=1}^{D} \alpha_d = 1\right\}.$$

We now proceed to exhibit an updating scheme on the weights α_d^t such that the asymptotic variances $\sigma_{1,t}^2$ and $\sigma_{2,t}^2$ are decreasing at each iteration of the D-kernel PMC algorithm.

3 PMC as variance reduction scheme

3.1 Self-normalized PMC estimator

For the estimator $\sum_{i=1}^{N} \overline{\omega}_{i,t} h(x_{i,t})$, we first introduce notations that simplify the study of its asymptotic variance. If ν_h denotes the measure on $\Omega \times \Omega$ defined by

$$\nu_h(dx, dx') = \pi(x') \left(h(x') - \pi(h)\right)^2 \pi(dx)\pi(dx')\,, \tag{2}$$

which naturally appears in $\sigma^2_{1,t}$, we define two functions σ^2_1 and F_1 on $\mathscr{S}_D$ such that

$$\sigma^2_1(\alpha) = \nu_h\left(1\Big/\sum_{d=1}^D \alpha_d q_d(x,x')\right) \quad \text{and} \quad F_1(\alpha) = \left(\nu_h\left(\frac{\alpha_d q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)\Big/\sigma^2_1(\alpha)\right)_{1\le d\le D}$$

Clearly, σ^2_1 is then the asymptotic variance associated with a given set of weights, while F_1 takes its values in $\mathscr{S}_D$ and is thus a transform (or update) of the weights. The central result of this paper is that this particular choice of update induces a reduction of the asymptotic variance at each step of the PMC algorithm:

Proposition 3.1. *Under **(A1)**, for all $\alpha \in \mathscr{S}_D$, we have*

$$\sigma^2_1(F_1(\alpha)) \le \sigma^2_1(\alpha)\,.$$

Proof. We have

$$\begin{aligned}
\sigma^2_1(F_1(\alpha)) &= \nu_h\left(\frac{1}{\sum_{d=1}^D \alpha_d q_d(y,y')\nu_h\left(\frac{q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)}\right)\sigma^2_1(\alpha)\\
&= \nu_h\left(\frac{1}{\sum_{l=1}^D \alpha_l q_l(y,y')}\frac{1}{\sum_{d=1}^D \frac{\alpha_d q_d(y,y')}{\sum_{l=1}^D \alpha_l q_l(y,y')}\nu_h\left(\frac{q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)}\right)\sigma^2_1(\alpha)\\
&\le \nu_h\left(\frac{1}{\sum_{l=1}^D \alpha_l q_l(y,y')}\sum_{d=1}^D \frac{\alpha_d q_d(y,y')}{\sum_{l=1}^D \alpha_l q_l(y,y')}\frac{1}{\nu_h\left(\frac{q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)}\right)\sigma^2_1(\alpha)\\
&= \sum_{d=1}^D \alpha_d \nu_h\left(\frac{q_d(y,y')/(\sum_{l=1}^D \alpha_l q_l(y,y'))^2}{\nu_h\left(q_d(x,x')/(\sum_{l=1}^D \alpha_l q_l(x,x'))^2\right)}\right)\sigma^2_1(\alpha) = \sigma^2_1(\alpha)\,,
\end{aligned}$$

the inequality following from Jensen's inequality. □

We thus take advantage of the diminution of the asymptotic variance to construct a sequence on $\mathscr{S}_D$ such that

$$\begin{cases}\alpha^{1,1} = (1/D,\ldots,1/D)\\ \alpha^{1,t+1} = F_1(\alpha^{1,t}) \qquad \text{for} \quad t \ge 1\,.\end{cases} \tag{3}$$

At each step of the PMC algorithm, the asymptotic variance is therefore decreasing. Since σ^2_1 is a convex function on the connected compact set $\mathscr{S}_D$, it admits a unique minimum. If we denote by

$$\alpha^{1,min} = \arg\min_{\alpha\in\mathscr{S}_D} \sigma^2_1(\alpha)$$

this minimum, we then have the convergence result for this updating mechanism (whose proof is given in Appendix 6.1).

Proposition 3.2. *Under* ***(A1)****,*

$$\lim_{t\to\infty} \alpha^{1,t} = \alpha^{1,min}.$$

Propositions 3.1 and 3.2 together establish the convergence to the minimal variance solution of the ideal algorithm, that is, the one using the update mechanism $\alpha^{1,t+1} = F_1(\alpha^{1,t})$. To complete the validation of a practical algorithm, we now have to replace the ideal updating with a practical updating and to show that the substitution does not jeopardize convergence. In other words, we need to establish the convergence of the mixture weights to the $\alpha^{1,t}$'s and this is sufficient for Theorem 2.2 to apply, i.e., for the asymptotic variance to be a valid assessment of our algorithm.

We thus define, as a substitute to F_1, the following update of the mixture weights

$$\alpha_d^{t+1,N} = \frac{\displaystyle\sum_{i=1}^N \overline{\omega}_{i,t}^2 \left(h(x_{i,t}) - \sum_{j=1}^N \overline{\omega}_{j,t} h(x_{j,t}) \right)^2 \mathbb{I}_d(K_{i,t})}{\displaystyle\sum_{i=1}^N \overline{\omega}_{i,t}^2 \left(h(x_{i,t}) - \sum_{j=1}^N \overline{\omega}_{j,t} h(x_{j,t}) \right)^2}, \tag{4}$$

which also holds when $\overline{\omega}_{i,t}^2$ is replaced with $\omega_{i,t}^2$ in both the numerator and the denominator and is thus independent of the normalising constant to some extent. The convergence of this updating scheme is then ensured by the following result, whose proof is defered to Appendix 6.2

Proposition 3.3. *Under* ***(A1)****, for all* $t \geq 1$ *and* $\forall 1 \leq d \leq D$*,*

$$\alpha_d^{t,N} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \alpha_d^t \tag{5}$$

where the $\alpha_d^{t,N}$*'s are defined by equation* (4) *and the* α_d^t*'s are given in* (3)*.*

Note that in the proof of Proposition 3.2 (see Appendix 6.2), we have shown that

$$\hat{\sigma}_{1,t}^2 = N \sum_{i=1}^N \overline{\omega}_{i,t}^2 \left(h(x_{i,t}) - \sum_{j=1}^N \overline{\omega}_{j,t} h(x_{j,t}) \right)^2 \tag{6}$$

is a consistent estimator of $\sigma_{1,t}^2(\alpha^t)$.

3.2 Unnormalised PMC estimator

The same sequence of results holds for the estimator $\frac{1}{N}\sum_{i=1}^N \omega_{i,t} h(x_{i,t})$. We first define the measure on $\Omega \times \Omega$

$$\rho_h(dx, dx') = \pi(x')\, h(x')^2 \pi(dx) \pi(dx')\,,$$

which relates to $\sigma^2_{2,t}$ since

$$\sigma^2_{2,t} = \overline{\pi}\left(h^2(x')\frac{\pi(x')}{\sum_{d=1}^D \alpha_d^t q_d(x,x')}\right) - 2\pi(h) + \pi(h)^2\,.$$

We also consider two functions σ_2^2 and F_2 on $\mathscr{S}_D$ as

$$\sigma_2^2(\alpha) = \rho_h\left(1\Big/\sum_{d=1}^D \alpha_d q_d(x,x')\right) - 2\pi(h) + \pi(h)^2$$

and

$$F_2(\alpha) = \left\{\rho_h\left(\frac{\alpha_d q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)\Big/\sigma_2^2(\alpha)\right\}_{1\leq d\leq D}\,.$$

Then we can use the same steps as in Section 3.1 and derive convergence from the results there. First, as a corollary to Proposition 3.1, we can derive the decrease in the asymptotic variance for the ideal weights:

Proposition 3.4. *Under **(A1)**, for all $\alpha \in \mathscr{S}_D$, we have*

$$\sigma_2^2(F_2(\alpha)) \leq \sigma_2^2(\alpha)$$

Second, if we set the sequence of ideal weights as

$$\begin{cases}\alpha^{2,1} = (1/D,\ldots,1/D) \\ \alpha^{2,t+1} = F_2(\alpha^{2,t}) \qquad \text{for} \quad t \geq 1\end{cases} \tag{7}$$

then we deduce from Proposition 3.2 that this ideal sequence converges to the optimal set of weights since, as σ_2^2 is convex on $\mathscr{S}_D$, it thus admits a unique minimum

$$\alpha^{2,min} = \arg\min_{\alpha\in\mathscr{S}_D} \sigma_2^2(\alpha)\,.$$

Proposition 3.5. *Under **(A1)**,*

$$\lim_{t\to\infty} \alpha^{2,t} = \alpha^{2,min}.$$

Third, we now exhibit the empirical version of the updating scheme which ensures that the practical version of the algorithm also converges, by virtue of Theorem 2.2 and Propositions 3.4 and 3.5. In the unnormalised case, it is now given by

$$\alpha_d^{t+1,N} = \sum_{i=1}^N \omega_{i,t}^2 h^2(x_{i,t})\mathbb{I}_d(K_{i,t})\Big/\sum_{i=1}^N \omega_{i,t}^2 h^2(x_{i,t})\,. \tag{8}$$

Finally, as a corollary to Proposition 3.3, we then have the overall convergence guarantee:

Proposition 3.6. *Under* ***(A1)****, for all* $t \geq 1$ *and* $\forall 1 \leq d \leq D$,

$$\alpha_d^{t,N} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \alpha_d^t \tag{9}$$

where the $\alpha_d^{t,N}$*'s and the* α_d^t*'s are defined by equations (8) and (7), respectively.*

Note also that

$$\frac{1}{N}\sum_{i=1}^{N}\left(\omega_{i,t}h(x_{i,t}) - N^{-1}\sum_{j=1}^{N}\omega_{j,t}h(x_{j,t})\right)^2$$

is a consistent estimator of $\sigma_{2,t}^2$.

4 A cumulated estimator

While each iteration of the PMC algorithm leads to a decrease in the variance of the PMC estimator compared with the previous iteration estimator, provided updates (4) or (8) are used, ealier iterations can further be taken into account by constructing a weighted cumulated estimator of the PMC estimators over iterations. Since, as noted in Theorem 2.2, each PMC sample is asymptotically independent of the others, the asymptotic variance of the weighted estimator is simply the sum of the cumulated weighted variance and the weights in this cumulated estimator can thus be directly optimised.

Quite naturally, the cumulated self-normalized PMC estimator, $\hat{\pi}_\beta^{CSN}(h)$, of $\pi(h)$ is chosen as

$$\hat{\pi}_\beta^{CSN}(h) = \sum_{t=0}^{T}\beta_t\left(\sum_{i=1}^{N}\overline{\omega}_{i,t}h(x_{i,t})\right), \quad \text{with} \quad \beta \in \mathscr{S}_{T+1}\,.$$

Under assumptions **(A1-2)**, Theorem 2.2, implies that

$$\sqrt{N}\left\{\hat{\pi}_\beta^{CSN}(h) - \pi(h)\right\} \overset{N\to\infty}{\longrightarrow}_{\mathcal{L}} \mathcal{N}\left(0, \sum_{t=0}^{T}\beta_t^2\sigma_{1,t}^2\right).$$

With respect to this asymptotic variance, we can then derive the optimal choice of β, namely $(1 \leq t \leq T)$

$$\beta_t^{min} = \sigma_{1,t}^{-2}/(\sum_{t=0}^{T}\sigma_{1,t}^{-2}) \quad \text{and} \quad \sum_{t=0}^{T}(\beta_t^{min})^2\sigma_{1,t}^2 = \left(\sum_{t=0}^{T}\sigma_{1,t}^{-2}\right)^{-1}$$

Furthermore, this optimal β_t^{min} can be consistently estimated by

$$\hat{\beta}_t^{min} = \frac{\hat{\sigma}_{1,t}^{-2}}{\sum_{t=0}^{T}\hat{\sigma}_{1,t}^{-2}}\,, \tag{10}$$

where $\hat{\sigma}_{1,t}^2$ is defined in (6). Therefore,

Proposition 4.1. *For an arbitrary h in L^1_π, for any T,.*

(i) Under **(A1)**, $\hat{\pi}^{CSN}_{\hat{\beta}^{min}}(h) \xrightarrow{N\to\infty}_{\mathbb{P}} \pi(h)$;

(ii) Under **(A1-2)**, $\sqrt{N}\left\{\hat{\pi}^{CSN}_{\hat{\beta}^{min}}(h) - \pi(h)\right\} \xrightarrow{N\to\infty}_{\mathcal{L}} \mathcal{N}\left\{0, \left(\sum_{t=0}^{T} \sigma_{1,t}^{-2}\right)^{-1}\right\}$.

Note that a corresponding result also holds for the unnormalized PMC estimator.

5 Applications

5.1 A normal toy example

We first consider a toy example where the optimal solution is known: using the $\mathscr{N}(0,1)$ density and $h(x) = x$, the optimal importance distribution which minimises the variance of the unnormalised importance sampling estimator is $g^*(x) \propto |x| \exp -x^2/2$. It actually corresponds to the distribution of the root of an exponential $\mathscr{E}(1/2)$ random variable with random sign, that is, $(-1)^s\sqrt{\mathscr{E}(1/2)}$ where $s \sim \mathscr{B}(0.5)$, a Bernoulli distribution with parameter $1/2$. We then choose g^* as one of $D = 3$ independent kernels, the other kernels being the $\mathscr{N}(0,1)$ and the $\mathscr{C}(0,1)$ (Cauchy) distributions. Note that the fact that the proposals are independent does not modify the validity of the above results. In particular, conditions **(A1-2)** do hold in that case. (The only modification in the algorithm is that the resampling step is no longer necessary.)

Table 5.1 presents the results of the variance D-kernel PMC scheme with $N = 100,000$ and $T = 20$. At each iteration, the (esatimated) asymptotic variance of the self-normalized PMC estimator decreases and the weights of the mixture proposal correctly concentrate around the correct optimal kernel. The optimal standard deviation in that case is equal to $2/\sqrt{2\pi} = 0.7979$. Figure 1 represents the convergence to this optimal value over 50 iterations.

5.2 The Cox-Ingersoll-Ross model

The Cox-Ingersoll-Ross (CIR) model (Cox et al., 1985) is a diffusion process used to model interest rate changes. The fundamental stochastic differential equation of the CIR model is

$$dr_t = (\eta - kr_t)dt + \sigma\sqrt{r_t}dW_t\,, \tag{11}$$

where $(W_t)_{[0,T]}$ is a Brownian motion under the risk neutral measure $\mathbb{P}$. In the financial application, $(r_t)_{[0,T]}$ represents the short term rate over the measurement period. A quantity of interest is the so-called *European caplet* option, which is an option written on the

t	$\hat{\pi}_{t,N}^{PMC}(x)$	$\alpha_1^{t+1,N}$	$\alpha_2^{t+1,N}$	$\alpha_3^{t+1,N}$	$\sigma_{1,t}$
1	0.0000	0.1000	0.8000	0.1000	0.9524
2	-0.0030	0.1144	0.7116	0.1740	0.9192
3	-0.0017	0.1191	0.6033	0.2776	0.8912
4	-0.0006	0.1189	0.4733	0.4078	0.8608
5	-0.0035	0.1084	0.3545	0.5371	0.8394
6	0.0005	0.0956	0.2546	0.6498	0.8241
7	-0.0007	0.0822	0.1786	0.7392	0.8163
8	-0.0029	0.0696	0.1257	0.8047	0.8107
9	-0.0044	0.0594	0.0887	0.8519	0.8036
10	0.0065	0.0519	0.0622	0.8859	0.8016
11	-0.0030	0.0454	0.0450	0.9096	0.8018
12	-0.0051	0.0405	0.0329	0.9266	0.8000
13	-0.0008	0.0352	0.0236	0.9412	0.7996
14	0.0034	0.0329	0.0175	0.9496	0.7993
15	0.0033	0.0305	0.0136	0.9559	0.7987
16	-0.0025	0.0277	0.0102	0.9621	0.7988
17	0.0007	0.0277	0.0102	0.9621	0.7987
18	0.0026	0.0239	0.0062	0.9699	0.7985
19	0.0018	0.0223	0.0051	0.9726	0.7984
20	-0.0042	0.0204	0.0041	0.9755	0.7984

Table 1: Estimation of $\mathbb{E}[X] = 0$ for a normal variate using the D-kernel PMC algorithm with $D = 3$, $N = 100,000$ and normal, Cauchy and transformed Gamma independent kernels: evolution of the PMC estimates, kernel weights and asymptotic standard deviation estimates over 20 iterations.

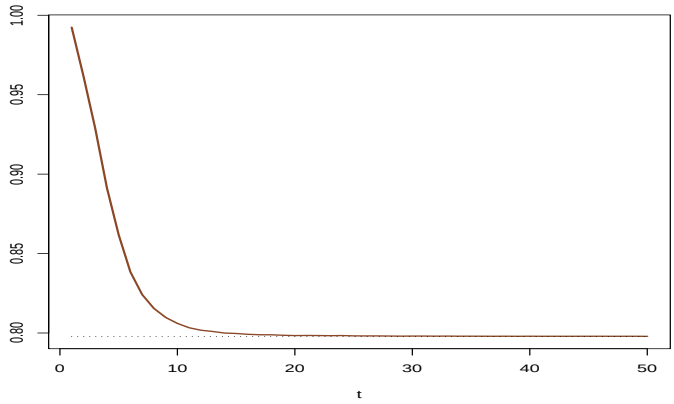

Figure 1: Estimation of $\mathbb{E}[X] = 0$ for a normal variate: decrease of the standard deviation to its optimal value

interest rate with the following payoff function at deadline (or maturity) T:

$$M \max(r_T - K, 0)\,,$$

where K is called the strike rate and M the nominee amount. The actualised price of the caplet at time 0 is therefore given by

$$\mathbb{E}_{\mathbb{P}}\left[\exp\left(-\int_0^T r_t dt\right) M \max(r_T - K, 0)\right]\,. \tag{12}$$

The explicit calculation of (12) is obviously intractable even though the transition density of the diffusion is available (Cox et al., 1985).

The standard approach to processing diffusions is to use a Euler approximation scheme, which consists in discretising the time interval $[0, T]$ into n steps and in studying instead the discrete time process $(r^n_{pT/n})_{0\le p\le n}$ with $r^n_0 = r_0$ and

$$r^n_{(p+1)T/n} = r^n_{pT/n} + (\eta - k r^n_{pT/n})\frac{T}{n} + \sigma\sqrt{r^n_{pT/n}}\left(W_{(p+1)T/n} - W_{pT/n}\right)\,, \tag{13}$$

since the differences $W_{(p+1)T/n} - W_{pT/n}$ are iid $\mathcal{N}(0, T/n)$. The quantity of interest (12) is then approximated by

$$\mathfrak{P} = \mathbb{E}\left[\exp\left\{-(T/n)\left((r^n_0 + r^n_N)/2 + \sum_{i=1}^{n-2} r^n_{(p+1)T/n}\right)\right\} M \max(r^n_T - K, 0)\right]\,, \tag{14}$$

where $r^n_{(p+1)T/n} \sim \mathcal{N}\left(r^n_{pT/n} + (T/n)(\eta - k r^n_{pT/n}), (T/n)\sigma^2 r^n_{pT/n}\right)$. Details on the Euler approximation can be found in Glasserman (2003), Jackel (2002) as well as Talay and Tubaro

(1990) and Bally and Talay (1996a,b). (See also Beskos et al., 2005 for a novel approach on the true distribution of $(r^n_{pT/n})_{0\le p\le n}$ that does not require the Euler scheme.)

Even when using the Euler approximation, the explicit derivation of $\mathfrak{P}$ is impossible and we need to use Monte Carlo simulation to approximate (14). Some importance sampling techniques have already been proposed by Arouna (2003, 2004), Su and Fu (2002) and we now study the behaviour of our D-kernel PMC scheme. While the exact distribution in (13) can be used in the Monte Carlo approximation, it seems rather natural to force the process to end up as much as possible above K for $r_T - K$ to be positive.

Our alternatives to the Euler distribution in (13) are based on the introduction of a location drift on the Brownian motion $(W_t)_t$. More precisely, for $\theta \in \mathbb{R}$, we define the family of all equivalent probability measures $Q(\theta)$ with respect to $\mathbb{P}$ that follow from introducing a drift of θ in (W_t). By Girsanov's theorem, we know that, under the measure $Q(\theta)$,

$$dr_t = (\eta - kr_t + \theta\sigma\sqrt{r_t})dt + \sigma\sqrt{r_t}d\tilde{W}_t\,,$$

where $\left(\tilde{W}_t\right)_{[0,T]}$ is a Brownian motion and the change of measure process is given by

$$\frac{d\mathbb{P}}{dQ(\theta)} = \exp\left(-\theta W_T - 0.5\theta^2 T\right)\,. \tag{15}$$

Reverting to the Euler approximation, we can then define (and simulate) the associated process

$$r^n_{(p+1)T/n} = r^n_{pT/n} + (\eta - kr^n_{pT/n} + \theta\sigma\sqrt{r^n_{pT/n}})\frac{T}{n} + \sigma\sqrt{r^n_{pT/n}}\epsilon_p\,, \tag{16}$$

where the $\epsilon'_p s$ are iid $\mathcal{N}(0,1)$, and compute the importance sampling weight associated with the simulation. (In the discrete case, the change of measure is the same as (15).) Obviously, the importance weights are based on the comparison between (13) and (16) and do not take into account the Euler approximation. Note that the idea of a location drift is already present in the literature (see, e.g., Arouna, 2003, 2004, Su and Fu, 2002), with Su and Fu (2002) deriving optimal choices of θ towards variance minimising. As in the toy example, we are thus using independent proposals and thus do not require the resampling step in the algorithm.

The choice of θ being open, a D-kernel scheme can be used to select efficient values of θ towards the approximation of (14). Let us stress once more that the underlying idea is to force r^n_T to be larger than K in order to decrease the variance of the Monte Carlo estimator and so positive values of θ are called for. Note that the case $\theta = 0$ corresponds to a crude Monte Carlo approximation. Figure 2 compares the range of the proposed processes $r^n_{(p+1)T/n}$ with the range of $1,000$ resampled processes using the importance weights. While the range of the proposed values is obviously larger, the decrease in the range due to resampling is quite limited, which shows a good fit between the target distribution and the optimised mixture.

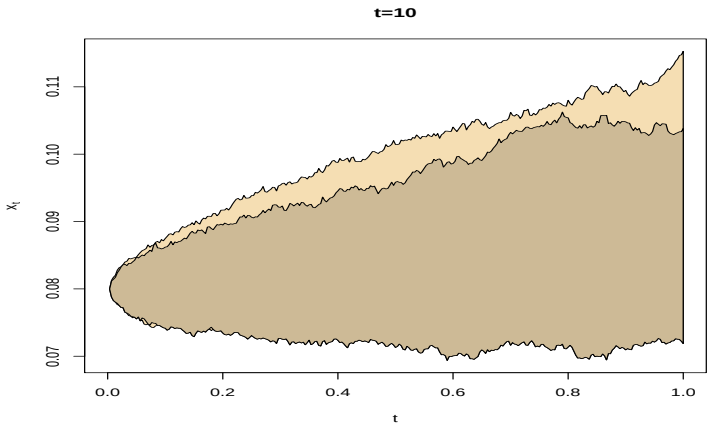

Figure 2: Cox-Ingersoll-Ross European option: range of the simulated process $r^n_{(p+1)T/n}$ *(lighter background)* compared with the resampled process $\tilde{r}^n_{(p+1)T/n}$ *(darker background)* for $K = 0.09$, using a D-kernel PMC algorithm with $\theta = 0$, $\theta = 1$ and $\theta = 2$, $T = 10$ iterations and $N = 10,000$ simulations.

We thus ran a simulation experiment with the following parameters: $\eta = 0.016$, $k = 0.2$, $\sigma = 0.02$, $M = 1000$, $r0 = 0.08$, $T = 1$ and $n = 299$, For simplicity's sake, we only took three values of θ, $\theta = 0, 1, 2$, the value $\theta = 0$ acting as a stabilising factor in the importance weight (since it ensures that assumptions **(A1-2)** hold). Obviously, a finer grid of values of θ with D larger than 3 could also have been chosen as well. Tables 2, 3 and 4 present the results of this experiment over $N = 100,000$ simulations and $T = 10$ iterations of the weight update for $K = 0.7$, $K = 0.8$ and $K = 0.9$. The larger the bound K, the larger the weight on the larger value of θ. In the three cases, the decrease in variance from the equaly weighted D-kernel proposal is quite appreciable. (A phenomenon that is quite common is the quick decrease of the variance in the very first iterations, followed by a much slower decline.)

The gain compared with the naive Monte Carlo approximation to $\mathfrak{P}$ is quite important: for $K = 0.7$, $K = 0.8$ and $K = 0.9$, the variances are 21.59, 7.914 and 0.1937, respectively. Note that Su and Fu (2002) derived optimal values for θ in exactly the same setting, obtaining $\theta = 0.487$, $\theta = 1.077$ and $\theta = 1.234$ in the three cases, respectively. An interesting remark is that, while $\theta = 0.487$ does lead to a smaller variance when $K = 0.7$, 2.318, compared with 6.88, the second case leads to almost the same variance, 1.126 versus 0.9939, when $K = 0.8$ since the optimal value is $\theta = 1.077$, and, surprisingly, the case $K = 0.9$ produces a much smaller variance, 0.0037 versus 0.0112, the reason being that Su and Fu (2002) then ran a fixed number of iterations of their optimisation algorithm, rather than to wait for the minimum: $\theta = 2$ then produces a much smaller variance than the proposed value $\theta = 1.234$.

t	$\widehat{\mathfrak{P}}_{t,N}^{PMC}$	$\alpha_1^{t,N}$	$\alpha_2^{t,N}$	$\alpha_3^{t,N}$	$\sigma_{1,t}^2$
1	9.2635	0.3333	0.3333	0.3334	27.0664
2	9.2344	0.4748	0.3703	0.1549	13.4474
3	9.2785	0.5393	0.3771	0.0836	9.7458
4	9.2495	0.5672	0.3835	0.0493	8.5258
5	9.2444	0.5764	0.3924	0.0312	7.8595
6	9.2400	0.5780	0.4014	0.0206	7.5471
7	9.2621	0.5765	0.4098	0.0137	7.2214
8	9.2435	0.5727	0.4183	0.0090	7.1354
9	9.2553	0.5682	0.4260	0.0058	7.0289
10	9.2602	0.5645	0.4320	0.0035	6.8854

Table 2: Cox-Ingersoll-Ross European option: approximation of the price $\mathfrak{P}$ for $K = 0.07$ using a D-kernel PMC algorithm with $\theta = 0$, $\theta = 1$ and $\theta = 2$

t	$\widehat{\mathfrak{P}}_{t,N}^{PMC}$	$\alpha_1^{t,N}$	$\alpha_2^{t,N}$	$\alpha_3^{t,N}$	$\sigma_{1,t}^2$
1	1.8784	0.3333	0.3333	0.3334	2.1781
2	1.8791	0.2458	0.4187	0.3355	1.9287
3	1.8793	0.1797	0.5078	0.3125	1.7329
4	1.8848	0.1279	0.5924	0.2797	1.5670
5	1.8877	0.0878	0.6704	0.2418	1.4374
6	1.8881	0.0589	0.7340	0.2071	1.3303
7	1.8892	0.0359	0.7873	0.1768	1.2530
8	1.8853	0.0229	0.8275	0.1496	1.2010
9	1.8860	0.0137	0.8613	0.1250	1.1593
10	1.8879	0.0079	0.8883	0.1038	1.1262

Table 3: Same table as Table 2 for $K = 0.08$

t	$\widehat{\mathfrak{P}}_{t,N}^{PMC}$	$\alpha_1^{t,N}$	$\alpha_2^{t,N}$	$\alpha_3^{t,N}$	$\sigma_{1,t}^2$
1	0.0555	0.3333	0.3333	0.3334	0.0114
2	0.0559	0.0333	0.2474	0.7193	0.0053
3	0.0554	0.0026	0.1108	0.8866	0.0043
4	0.0558	0.0001	0.0443	0.9556	0.0039
5	0.0557	0.0000	0.0164	0.9836	0.0038
6	0.0559	0.0000	0.0059	0.9941	0.0038
7	0.0559	0.0000	0.0028	0.9972	0.0038
8	0.0555	0.0000	0.0010	0.9990	0.0038
9	0.0558	0.0000	0.0003	0.9997	0.0038
10	0.0556	0.0000	0.0002	0.9998	0.0037

Table 4: Same table as Table 2 for $K = 0.09$

Acknowledgments

This work was partially supported by an ACI *"Nouvelles Interfaces des Mathématiques"* grant from the Ministère de la Recherche.

6 Proofs

6.1 Proof of Proposition 3.2

The proof of this result follows the same lines as the proof of Proposition 4.3 in Douc et al. (2005). The only condition to check is the equivalent of Proposition 4.2 in Douc et al. (2005). For every $\alpha \in \mathscr{S}_D, \alpha \neq \alpha^{1,min}$, we now show that there exists a neighborhood V_α of α such that if $\alpha^{t_0} \in V_\alpha$ then $(\alpha_t)_{t \geq t_0}$ leaves V_α within a finite time. Then, by continuity of $\sigma_1^2(\alpha)$, there exists $\epsilon > 0$ such that

$$
\begin{aligned}
-\epsilon\sigma_1^2(\alpha) \geq \sigma_1^2(\alpha^{1,min}) - \sigma_1^2(\alpha) &= \nu_h\left(\frac{1}{\sum_{d=1}^D \alpha_d^{1,min} q_d(x,x')} - \frac{1}{\sum_{d=1}^D \alpha_d q_d(x,x')}\right) \\
&\geq \nu_h\left(\frac{\sum_{d=1}^D (\alpha_d - \alpha_d^{1,min}) q_d(x,x')}{(\sum_{d=1}^D \alpha_d q_d(x,x'))^2}\right) = \sum_{d=1}^D \alpha_d^{1,min}\left[\sigma_1^2(\alpha) - \nu_h\left(\frac{q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right)\right]
\end{aligned}
$$

Thus, there exists $1 \leq d \leq D$ such that

$$
\nu_h\left(\frac{q_d(x,x')}{(\sum_{l=1}^D \alpha_l q_l(x,x'))^2}\right) \geq (1+\epsilon)\sigma_1^2(\alpha)
$$

which implies that $[F_1(\alpha)]_d \geq (1+\epsilon)\alpha_d$. Since $0 \leq \alpha_d \leq 1$, it follows that if $\alpha^{t_0} \in V_\alpha$ then $(\alpha_t)_{t \geq t_0} = (F_1^{t-t_0}(\alpha_{t_0}))_{t \geq t_0}$ will leave V_α within a finite time. The proof is completed.

6.2 Proof of Proposition 3.3

The case $t = 1$ is obvious. Now, assume (5) holds for some $t \geq 1$. Recall that ν_h is defined in (2). We now prove that the following convergence results

$$\frac{1}{N}\sum_{i=1}^{N}\omega_{i,t}^2\left(h(x_{i,t}) - \sum_{j=1}^{N}\overline{\omega}_{j,t}h(x_{j,t})\right)^2 \mathbb{I}_d(K_{i,t}) \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \nu_h\left(\frac{\alpha_d^t q_d(x,x')}{(\sum_{l=1}^{D}\alpha_l^t q_l(x,x'))^2}\right), \tag{17}$$

$$\frac{1}{N}\sum_{i=1}^{N}\omega_{i,t}^2\left(h(x_{i,t}) - \sum_{j=1}^{N}\overline{\omega}_{j,t}h(x_{j,t})\right)^2 \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \sigma_1^2(\alpha^t). \tag{18}$$

Only the first convergence needs be considered since the latter can be easily deduced from the former by summing over d. To prove Eq. (17), we will show that

$$\frac{1}{N}\sum_{i=1}^{N}\omega_{i,t}^2 H(x_{i,t})\mathbb{I}_d(K_{i,t}) \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \nu_H\left(\frac{\alpha_d^t q_d(x,x')}{(\sum_{l=1}^{D}\alpha_l^t q_l(x,x'))^2}\right)$$

for any function H satisfying $\pi(H) = 0$ and

$$\overline{\pi}\left\{(1+H(x'))\pi(x')/q_{\overline{d}}(x,x')\right\} < \infty \quad \text{for some} \quad \overline{d} \in \{1,\ldots,D\}. \tag{19}$$

We apply Theorem A.1 of Douc et al. (2005) with

$$\mathcal{G}_N = \sigma\left((\tilde{x}_{i,t-1})_{1\leq i\leq N}, (\alpha_d^{t,N})_{1\leq d\leq D}\right) \quad \text{and} \quad U_{N,i} = N^{-1}\omega_{i,t}^2 H(x_{i,t})\mathbb{I}_d(K_{i,t}).$$

Conditionally on $\mathcal{G}_N$, $(K_{i,t}, x_{i,t})_{1\leq i\leq N}$ are independent and for all (d, A) in $\{1,\ldots,D\}\times\mathcal{A}$,

$$\mathbb{P}\left(K_{i,t} = d, x_{i,t} \in A \middle| \mathcal{G}_N\right) = \alpha_d^{t,N} Q_d(\tilde{x}_{i,t-1}, A)$$

To apply Theorem A.1 of Douc et al. (2005), we just need to check condition (iii). We have

$$\begin{aligned}
&\mathbb{E}\left(\sum_{i=1}^{N}\frac{\omega_{i,t}^2 H(x_{i,t})\mathbb{I}_d(K_{i,t})}{N}\mathbb{I}_{\left\{\omega_{i,t}^2 H(x_{i,t})\mathbb{I}_d(K_{i,t})>C\right\}}\middle|\mathcal{G}_N\right)\\
&\quad\leq \sum_{j=1}^{D}\frac{1}{N}\sum_{i=1}^{N}\int\pi(dx)\frac{\pi(x)H(x)\alpha_d^{t,N}q_d(\tilde{x}_{i,t-1},x)}{(\sum_{l=1}^{D}\alpha_l^{t,N}q_l(\tilde{x}_{i,t-1},x))^2}\mathbb{I}_{\left\{\frac{\pi(x)^2H(x)}{D^{-2}q_j^2(\tilde{x}_{i,t-1},x)}>C\right\}}\\
&\quad\leq \sum_{j=1}^{D}\frac{1}{N}\sum_{i=1}^{N}\int\pi(dx)\frac{\pi(x)H(x)}{\alpha_d^{t,N}q_{\overline{d}}(\tilde{x}_{i,t-1},x)}\mathbb{I}_{\left\{\frac{\pi(x)^2H(x)}{D^{-2}q_j^2(\tilde{x}_{i,t-1},x)}>C\right\}}\\
&\quad\overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \sum_{j=1}^{D}\int\pi(dx')\pi(dx)\frac{\pi(x)H(x)}{\alpha_{\overline{d}}^t q_{\overline{d}}(x',x)}\mathbb{I}_{\left\{\frac{\pi(x)^2H(x)}{D^{-2}q_j^2(x',x)}>C\right\}}
\end{aligned}$$

by the LLN stated in Theorem 2.1 and since the induction assumption implies that $\alpha_d^{t,N} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \alpha_d^t$ which is positive by the updating formula of α_d^t. The rhs converges to 0 as C gets to infinity using (19) and $\overline{\pi}\{H(x') = \infty \text{ or } q_j(x,x') = 0\} = 1$. Thus, Theorem A.1 of Douc et al. (2005) applies and

$$\frac{1}{N}\sum_{i=1}^{N} \omega_{i,t}^2 H(x_{i,t}) \mathbb{I}_d(K_{i,t}) - \mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N} \omega_{i,t}^2 H(x_{i,t}) \mathbb{I}_d(K_{i,t}) \middle| \mathcal{G}_N \right) \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} 0.$$

To complete the proof, it remains to show that

$$\mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N} \omega_{i,t}^2 H(x_{i,t}) \mathbb{I}_d(K_{i,t}) \middle| \mathcal{G}_N \right)$$
$$= \frac{1}{N}\sum_{i=1}^{N} \int \pi(dx) \frac{\pi(x) H(x) \alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x)}{(\sum_{l=1}^{D} \alpha_l^{t,N} q_l(\tilde{x}_{i,t-1}, x))^2} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \nu_H\left(\frac{\alpha_d^t q_d(x,x')}{(\sum_{l=1}^{D} \alpha_l^t q_l(x,x'))^2} \right) \quad (20)$$

Using again the LLN stated in Theorem 2.1,

$$\frac{1}{N}\sum_{i=1}^{N} \int \pi(dx) \frac{\pi(x) H(x) \alpha_d^{t} q_d(\tilde{x}_{i,t-1}, x)}{(\sum_{l=1}^{D} \alpha_l^{t} q_l(\tilde{x}_{i,t-1}, x))^2} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \nu_H\left(\frac{\alpha_d^t q_d(x,x')}{(\sum_{l=1}^{D} \alpha_l^t q_l(x,x'))^2} \right) \quad (21)$$

Thus, to prove (20), we use (21) and check that the difference between both terms converges to 0 in probability. To see this, first note that for all $t \geq 1$, for all d in $\{1,\ldots,D\}$, $\alpha_d^t > 0$ and thus, by the induction assumption, for all d in $\{1,\ldots,D\}$, $\frac{\alpha_d^{t,N} - \alpha_d^t}{\alpha_d^t} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} 0$. It has been shown in Douc et al. (2005) that

$$\left| \frac{\alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x)}{\sum_{l=1}^{D} \alpha_l^{t,N} q_l(\tilde{x}_{i,t-1}, x)} - \frac{\alpha_d^{t} q_d(\tilde{x}_{i,t-1}, x)}{\sum_{l=1}^{D} \alpha_l^{t} q_l(\tilde{x}_{i,t-1}, x)} \right| \leq 2 \sup_{l\in\{1,\ldots,D\}} \left| \frac{\alpha_l^{t,N} - \alpha_l^t}{\alpha_l^t} \right|.$$

Combining with $\left|\frac{A}{B^2} - \frac{C}{D^2}\right| \leq \left|\frac{A}{B}\right| \left|\frac{D-B}{BD}\right| + \left|\frac{1}{D}\right| \left|\frac{A}{B} - \frac{C}{D}\right|$ yields by straightforward algebra,

$$\left| \frac{\alpha_d^{t,N} q_d(\tilde{x}_{i,t-1}, x)}{(\sum_{l=1}^{D} \alpha_l^{t,N} q_l(\tilde{x}_{i,t-1}, x))^2} - \frac{\alpha_d^{t} q_d(\tilde{x}_{i,t-1}, x)}{(\sum_{l=1}^{D} \alpha_l^{t} q_l(\tilde{x}_{i,t-1}, x))^2} \right|$$
$$\leq \frac{1}{\sum_{l=1}^{D} \alpha_l^{t,N} q_l(\tilde{x}_{i,t-1}, x)} \left(\sup_{l\in\{1,\ldots,D\}} \left| \frac{\alpha_l^{t,N} - \alpha_l^t}{\alpha_l^t} \right| \right)$$
$$+ \frac{1}{\sum_{l=1}^{D} \alpha_l^{t} q_l(\tilde{x}_{i,t-1}, x)} \left(2 \sup_{l\in\{1,\ldots,D\}} \left| \frac{\alpha_d^{t,N} - \alpha_d^t}{\alpha_d^t} \right| \right)$$
$$\leq \left(\frac{1}{\alpha_{\overline{d}}^{t,N}} + \frac{2}{\alpha_{\overline{d}}^{t}} \right) \frac{1}{q_{\overline{d}}(\tilde{x}_{i,t-1}, x)} \left(\sup_{l\in\{1,\ldots,D\}} \left| \frac{\alpha_l^{t,N} - \alpha_l^t}{\alpha_l^t} \right| \right).$$

The proof follows from $\frac{\alpha_d^{t,N}-\alpha_d^t}{\alpha_d^t} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} 0$ and

$$\frac{1}{N}\sum_{i=1}^{N}\int \pi(dx)\frac{\pi(x)H(x)}{q_{\overline{d}}(\tilde{x}_{i,t-1},x))} \overset{N\to\infty}{\longrightarrow}_{\mathbb{P}} \overline{\pi}\left\{\frac{H(x')\pi(x')}{q_{\overline{d}}(x,x')}\right\}.$$

References

Arouna, B. (2003). Robbins-Monro algorithms and variance reduction in finance. *The Journal of Computational Finance*, 7(2):1245–1255.

Arouna, B. (2004). Adaptative Monte Carlo Method, A variance Reduction Technique. *Monte Carlo Methods and Appl.*, 10(1):1–24.

Bally, V. and Talay, D. (1996a). The law of the Euler scheme for stochastic differential equations (i): convergence rate of the distribution function. *Probability Theory and Related Fields*, 104:43–60.

Bally, V. and Talay, D. (1996b). The law of the Euler scheme for stochastic differential equations (ii): convergence rate of the density. *Probability Theory and Related Fields*, 104:98–128.

Beskos, A., Papaspiliopoulos, O., Roberts, G., and Fearnhead, P. (2005). Exact and efficient likelihood based estimation for discretely observed diffusion processes. Technical report, Department of Statistics, University of Lancaster.

Cappé, O., Guillin, A., Marin, J.-M., and Robert, C. (2004). Population Monte Carlo. *J. Comput. Graph. Statist.*, 13(4):907–929.

Cox, J., Ingersoll, J., and Ross, A. (1985). A Theory of The Term Structure of Interest Rate. *Econometrica*, 53(2):385–408.

Douc, R., Guillin, A., Marin, J.-M., and Robert, C. (2005). Convergence of adaptive sampling schemes. Technical Report 2005-6, University Paris Dauphine.

Glasserman, P. (2003). *Monte Carlo methods in financial engineering*. Springer-Verlag.

Iba, Y. (2000). Population-based Monte Carlo algorithms. *Trans. Japanese Soc. Artificial Intell.*, 16(2):279–286.

Jackel, P. (2002). *Monte Carlo methods in finance*. John Wiley and Sons.

Lapeyre, B., Pardoux, E., and Sentis, R. (1998). *Méthodes de Monte Carlo pour les équations de transport et de diffusion*. Mathématiques et Applications, Vol. 29. Springer Verlag.

Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer-Verlag, New York, 2 edition.

Rubin, D. (1987). A noniterative sampling importance resampling alternative to the data augmentation algorithm for creating a few imputations when fractions of missing information are modest: the SIR algorithm. (In the discussion of Tanner and Wong paper). *J. American Statist. Assoc.*, 82:543–546.

Rubin, D. (1988). Using the SIR algorithm to simulate posterior distributions. In Bernardo, J., Degroot, M., Lindley, D., and Smith, A., editors, *Bayesian Statistics 3: Proceedings of the Third Valencia International Meeting, June 1-5, 1987*. Clarendon Press.

Rubinstein, R. (1981). *Simulation and the Monte Carlo Method*. J. Wiley, New York.

Rubinstein, R. and Kroese, D. (2004). *The Cross-Entropy Method*. Springer-Verlag, New York.

Su, Y. and Fu, M. (2002). Optimal importance sampling in securities pricing. *The Journal of Computational Finance*, 5(4):27–50.

Talay, D. and Tubaro, L. (1990). Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic Analysis and Applications*, 8(4):94–120.