The Annals of Statistics
2011, Vol. 39, No. 1, 261–277
DOI: 10.1214/10-AOS838

A VANILLA RAO–BLACKWELLIZATION OF METROPOLIS–HASTINGS ALGORITHMS[1]

BY RANDAL DOUC AND CHRISTIAN P. ROBERT

Telecom SudParis and Université Paris-Dauphine

Casella and Robert [*Biometrika* **83** (1996) 81–94] presented a general Rao–Blackwellization principle for accept-reject and Metropolis–Hastings schemes that leads to significant decreases in the variance of the resulting estimators, but at a high cost in computation and storage. Adopting a completely different perspective, we introduce instead a universal scheme that guarantees variance reductions in all Metropolis–Hastings-based estimators while keeping the computation cost under control. We establish a central limit theorem for the improved estimators and illustrate their performances on toy examples and on a probit model estimation.

1. Introduction. As its accept-reject predecessor, the Metropolis–Hastings simulation algorithm relies in part on the generation of uniform variables to achieve given acceptance probabilities. More precisely, given a target density f with respect to a dominating measure on the space $\mathcal{X}$, if the Metropolis–Hastings proposal is associated with the density $q(x|y)$ (with respect to the same dominating measure), then the acceptance probability of the corresponding Metropolis–Hastings iteration at time t is

$$\alpha\bigl(x^{(t)}, y_t\bigr) = \min\left\{1, \frac{\pi(y_t)}{\pi(x^{(t)})}\frac{q(x^{(t)}|y_t)}{q(y_t|x^{(t)})}\right\},$$

where $y_t \sim q(y_t|x^{(t)})$ is the proposed value for $x^{(t+1)}$. In practice, this means that a uniform $u_t \sim \mathcal{U}(0, 1)$ is first generated and that $x^{(t+1)} = y_t$ if and only if $u_t \le \alpha(x^{(t)}, y_t)$.

Since the uniformity of the u_t's is an extraneous (albeit necessary) noise, in that it does not directly provide information about the target f (but only through its acceptance rate), Casella and Robert (1996) took advantage of this flow of auxiliary variables u_t to reduce the variance of the resulting estimators while preserving their unbiasedness by integrating out the u_t's conditional on all simulated y_t's. This strategy has a nonnegligible cost of $O(N^2)$ for a given sample of size N. While extensions have been proposed in the literature [Casella and Robert (1998),

Received October 2009; revised May 2010.

[1] Supported by the EPSRC and by the Agence Nationale de la Recherche through the 2009–2012 Project BIG MC.

AMS 2000 subject classifications. 62-04, 60F05, 60J22, 60J05, 62B10.

Key words and phrases. Metropolis–Hastings algorithm, Markov chain Monte Carlo (MCMC), probit model, central limit theorem, variance reduction, conditioning.

Perron (1999); see also Delmas and Jourdain (2009) for an analysis of a Rao–Blackwellized version of the estimator when conditioning on the rejected candidates], this solution is therefore not considered in practice, in part due to this very cost. The current paper reproduces the Rao–Blackwellization argument of Casella and Robert (1996) by means of an independent representation that allows the variance to be reduced at a fixed computational cost. Section 2 outlines the Rao–Blackwellization technique and Section 3 validates the resulting variance reduction, including a derivation of the asymptotic variance of the improved estimators, while Section 4 presents some illustrations of the improvement on toy examples.

2. The Rao–Blackwellization solution. When considering the outcome of a Metropolis–Hastings experiment, $(x^{(t)})_t$, and the way it is used in Monte Carlo approximations,

$$\delta = \frac{1}{N}\sum_{t=1}^{N} h(x^{(t)}), \tag{1}$$

alternative representations of this estimator are

$$\delta = \frac{1}{N}\sum_{t=1}^{N}\sum_{j=1}^{t} h(y_j)\mathbb{I}_{x^{(t)}=y_j} \quad \text{and} \quad \delta = \frac{1}{N}\sum_{i=1}^{M} \mathfrak{n}_i h(\mathfrak{z}_i),$$

where the y_j's are the proposed Metropolis–Hastings moves, the $\mathfrak{z}_i$'s are the accepted y_j's, M is the number of accepted y_j's up to time N and $\mathfrak{n}_i$ is the number of times $\mathfrak{z}_i$ appears in the sequence $(x^{(t)})_t$. The first representation is the one used by Casella and Robert (1996), who integrate out the random elements of the outer sum, given the sequence of y_t's. The second representation is also found in Sahu and Zhigljavsky (1998), Gåsemyr (2002), Sahu and Zhigljavsky (2003) and Malefaki and Iliopoulos (2008), and is the basis for our construction.

Let us first recall the basic properties of the pairs $(\mathfrak{z}_i, \mathfrak{n}_i)$, also found in the above references.

LEMMA 1. *The sequence* $(\mathfrak{z}_i, \mathfrak{n}_i)$ *is such that*:

1. $(\mathfrak{z}_i, \mathfrak{n}_i)_i$ *is a Markov chain*;
2. $\mathfrak{z}_{i+1}$ *and* $\mathfrak{n}_i$ *are independent given* $\mathfrak{z}_i$;
3. $\mathfrak{n}_i$ *is distributed as a geometric random variable with probability parameter*

$$p(\mathfrak{z}_i) := \int \alpha(\mathfrak{z}_i, y) q(y|\mathfrak{z}_i)\, dy; \tag{2}$$

4. $(\mathfrak{z}_i)_i$ *is a Markov chain with transition kernel* $\tilde{Q}(\mathfrak{z}, dy) = \tilde{q}(y|\mathfrak{z})\, dy$ *and stationary distribution* $\tilde{\pi}$ *such that*

$$\tilde{q}(\cdot|\mathfrak{z}) \propto \alpha(\mathfrak{z}, \cdot) q(\cdot|\mathfrak{z}) \quad \textit{and} \quad \tilde{\pi}(\cdot) \propto \pi(\cdot) p(\cdot).$$

PROOF. We only prove the last point of the lemma. The transition kernel density $\tilde{q}$ of the Markov chain $(\mathfrak{z}_i)_i$ is obtained by integrating out the geometric waiting time, namely $\tilde{q}(\cdot|\mathfrak{z}_i) = \alpha(\mathfrak{z}_i, \cdot)q(\cdot|\mathfrak{z}_i)/p(\mathfrak{z}_i)$. Thus,

$$\tilde{\pi}(x)\tilde{q}(y|x) = \frac{\pi(x)p(x)}{\int \pi(u)p(u)\,du} \frac{\alpha(x, y)q(y|x)}{p(x)} = \tilde{\pi}(y)\tilde{q}(x|y),$$

where we have used the detailed balance property of the original Metropolis–Hastings algorithm, namely that $\pi(x)q(y|x)\alpha(x, y) = \pi(y)q(x|y)\alpha(y, x)$. This shows that the chain $(\mathfrak{z}_i)_i$ satisfies a detailed balance property with respect to $\tilde{\pi}$, thus that it is $\tilde{\pi}$-reversible, which completes the proof. □

Since the Metropolis–Hastings estimator δ only involves the $\mathfrak{z}_i$'s, that is, the accepted y_t's, an optimal weight for those random variables is the importance weight $1/p(\mathfrak{z}_i)$, leading to the corresponding importance sampling estimator,

$$\delta^* = \frac{1}{N} \sum_{i=1}^{M} \frac{h(\mathfrak{z}_i)}{p(\mathfrak{z}_i)},$$

but this quantity is usually unavailable in closed form and needs to be estimated by an unbiased estimator. The geometric $\mathfrak{n}_i$ is the obvious solution that is used in the original Metropolis–Hastings estimate, but solutions with smaller variance also are available, as shown by the following results.

LEMMA 2. *If $(y_j)_j$ is an i.i.d. sequence with distribution $q(y|\mathfrak{z}_i)$, then the quantity*

$$\hat{\xi}_i = 1 + \sum_{j=1}^{\infty} \prod_{\ell \le j} \{1 - \alpha(\mathfrak{z}_i, y_\ell)\}$$

is an unbiased estimator of $1/p(\mathfrak{z}_i)$, the variance of which, conditional on $\mathfrak{z}_i$, is lower than the conditional variance of $\mathfrak{n}_i$, $\{1 - p(\mathfrak{z}_i)\}/p^2(\mathfrak{z}_i)$.

PROOF. Since $\mathfrak{n}_i$ can be written as

$$\mathfrak{n}_i = 1 + \sum_{j=1}^{\infty} \prod_{\ell \le j} \mathbb{I}\{u_\ell \ge \alpha(\mathfrak{z}_i, y_\ell)\},$$

where the u_j's are i.i.d. $\mathcal{U}(0, 1)$, given that the sum actually stops with the first pair (u_j, y_j) such that $u_j \le \alpha(\mathfrak{z}_i, y_j)$, a Rao–Blackwellized version of $\mathfrak{n}_i$ consists in its expectation conditional on the sequence $(y_j)_j$:

$$\begin{aligned}
\hat{\xi}_i &= 1 + \sum_{j=1}^{\infty} \mathbb{E}\Big[\prod_{\ell \le j} \mathbb{I}\{u_\ell \ge \alpha(\mathfrak{z}_i, y_\ell)\}\Big|(y_t)_{t\ge 1}\Big] \\
&= 1 + \sum_{j=1}^{\infty} \prod_{\ell \le j} \mathbb{P}\big(u_\ell \ge \alpha(\mathfrak{z}_i, y_\ell)|(y_t)_{t\ge 1}\big) \\
&= 1 + \sum_{j=1}^{\infty} \prod_{\ell \le j} \{1 - \alpha(\mathfrak{z}_i, y_\ell)\}.
\end{aligned}$$

Therefore, since $\hat{\xi}_i$ is a conditional expectation of $\mathfrak{n}_i$, its variance is necessarily smaller. □

We note that this unbiased estimate of $1/p(\mathfrak{z}_i)$ can be related to the Bernoulli factory approach of Latuszynski et al. (2010), in that we are only using Bernoulli events in this derivation.

Given that $\alpha(\mathfrak{z}_i, y_j)$ involves a ratio of probability densities, $\alpha(\mathfrak{z}_i, y_j)$ takes the value 1 with positive probability and the sum $\hat{\xi}_i$ is therefore almost surely finite. This may, however, require far too many iterations to be realistically computed or it may involve too much variability in the number of iterations thus required. An intermediate estimator with a fixed computational cost is fortunately available.

PROPOSITION 1. *If $(y_j)_j$ is an i.i.d. sequence with distribution $q(y|\mathfrak{z}_i)$ and $(u_j)_j$ is an i.i.d. uniform sequence, for any $k \geq 0$, the quantity*

$$\hat{\xi}_i^k = 1 + \sum_{j=1}^{\infty} \prod_{1 \leq \ell \leq k \wedge j} \{1 - \alpha(\mathfrak{z}_i, y_j)\} \prod_{k+1 \leq \ell \leq j} \mathbb{I}\{u_\ell \geq \alpha(\mathfrak{z}_i, y_\ell)\} \tag{3}$$

is an unbiased estimator of $1/p(\mathfrak{z}_i)$ with an almost sure finite number of terms. Moreover, for $k \geq 1$,

$$\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i] = \frac{1-p(\mathfrak{z}_i)}{p^2(\mathfrak{z}_i)} - \frac{1-(1-2p(\mathfrak{z}_i)+r(\mathfrak{z}_i))^k}{2p(\mathfrak{z}_i)-r(\mathfrak{z}_i)}\left(\frac{2-p(\mathfrak{z}_i)}{p^2(\mathfrak{z}_i)}\right)(p(\mathfrak{z}_i)-r(\mathfrak{z}_i)),$$

where p is defined in (2) and $r(\mathfrak{z}_i) := \int \alpha^2(\mathfrak{z}_i, y)q(y|\mathfrak{z}_i)\,dy$. Therefore, we have

$$\mathbb{V}[\hat{\xi}_i|\mathfrak{z}_i] \leq \mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i] \leq \mathbb{V}[\hat{\xi}_i^0|\mathfrak{z}_i] = \mathbb{V}[\mathfrak{n}_i|\mathfrak{z}_i].$$

The truncation at the kth proposal thus allows for a calibration of the computational effort since $\hat{\xi}_i^k$ costs on average k additional simulations of y_j and computations of $\alpha(\mathfrak{z}_i, y_j)$ to compute $\hat{\xi}_i^k$, when compared with the regular Metropolis–Hastings weight $\mathfrak{n}_i$.

PROOF OF PROPOSITION 1. Define $y = (y_j)_{j\geq 1}$ and $u_{k:\infty} = (u_\ell)_{\ell \geq k}$. Note that $\hat{\xi}_i^0 = \mathfrak{n}_i$ and therefore the conditional variance of $\hat{\xi}_i^0$ is the variance of a geometric variable. Now, obviously, $\hat{\xi}_i^{k+1} = \mathbb{E}[\hat{\xi}_i^k|\mathfrak{z}_i, y, u_{k+2:\infty}]$; thus, we have

$$\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i] = \mathbb{V}[\hat{\xi}_i^{k+1}|\mathfrak{z}_i] + \mathbb{E}[\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i, y, u_{k+2:\infty}]|\mathfrak{z}_i].$$

To get a closed-form expression for the second term on the right-hand side, we first introduce a geometric random variable T_k defined by

$$T_k = 1 + \sum_{j=1}^{\infty} \prod_{\ell \leq j} \mathbb{I}\{u_{k+\ell} \geq \alpha(\mathfrak{z}_i, y_{k+\ell})\}.$$

Then, by straightforward algebra, $\hat{\xi}_i^k$ may be rewritten as

$$\hat{\xi}_i^k = C + \left(\prod_{\ell=1}^{k}\{1-\alpha(\mathfrak{z}_i, y_j)\}\right) T_{k+2}\mathbb{I}\{u_{k+1} > \alpha(\mathfrak{z}_i, y_{k+1})\},$$

where C does not depend on $u_1, \dots, u_{k+1}$. Thus,

$$\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i, y, u_{k+2:\infty}] = \left(\prod_{\ell=1}^{k}\{1-\alpha(\mathfrak{z}_i, y_j)\}^2\right) T_{k+2}^2 \alpha\{\mathfrak{z}_i, y_{k+1})\bigl(1-\alpha(\mathfrak{z}_i, y_{k+1})\bigr).$$

Taking the expectation of the above expression, we obtain

$$\mathbb{E}(\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i, y, u_{k+2:\infty}]) = \bigl(1-2p(\mathfrak{z}_i)+r(\mathfrak{z}_i)\bigr)^k \left(\frac{2-p(\mathfrak{z}_i)}{p^2(\mathfrak{z}_i)}\right)\bigl(p(\mathfrak{z}_i)-r(\mathfrak{z}_i)\bigr),$$

which completes the proof. □

3. Convergence properties. Using those Rao–Blackwellized versions of δ brings about an asymptotic improvement for the estimation of $\mathbb{E}_\pi[h(X)]$, as shown by the following result which, for any $M > 0$, compares the estimators ($k \ge 0$)

$$\delta_M^k = \frac{\sum_{i=1}^M \hat{\xi}_i^k h(\mathfrak{z}_i)}{\sum_{i=1}^M \hat{\xi}_i^k}.$$

For any positive function φ, we denote by $\mathcal{C}_\varphi = \{h; |h/\varphi|_\infty < \infty\}$ the set of functions bounded by φ up to a constant and we assume that the reference importance sampling estimator is sufficiently well behaved, in that there exist positive functions $\varphi \ge 1$ and ψ such that

$$\forall h \in \mathcal{C}_\varphi \qquad \frac{\sum_{i=1}^M h(\mathfrak{z}_i)/p(\mathfrak{z}_i)}{\sum_{i=1}^M 1/p(\mathfrak{z}_i)} \xrightarrow{\mathbb{P}} \pi(h), \tag{4}$$

$$\forall h \in \mathcal{C}_\psi \qquad \sqrt{M}\left(\frac{\sum_{i=1}^M h(\mathfrak{z}_i)/p(\mathfrak{z}_i)}{\sum_{i=1}^M 1/p(\mathfrak{z}_i)} - \pi(h)\right) \xrightarrow{\mathcal{L}} \mathcal{N}(0, \Gamma(h)). \tag{5}$$

THEOREM 1. *Under the assumption that $\pi(p) > 0$, the following convergence properties hold*:

(i) *if h is in $\mathcal{C}_\varphi$, then*

$$\delta_M^k \xrightarrow[M\to\infty]{\mathbb{P}} \pi(h);$$

(ii) *if, in addition, $h^2/p \in \mathcal{C}_\varphi$ and $h \in \mathcal{C}_\psi$, then*

$$\sqrt{M}\bigl(\delta_M^k - \pi(h)\bigr) \xrightarrow[M\to\infty]{\mathcal{L}} \mathcal{N}\bigl(0, V_k[h-\pi(h)]\bigr), \tag{6}$$

where $V_k(h) := \pi(p)\int \pi(d\mathfrak{z})\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}]h^2(\mathfrak{z})p(\mathfrak{z}) + \Gamma(h)$.

PROOF. We will prove that for all $g \in \mathcal{C}_\varphi$,

$$M^{-1} \sum_{i=1}^{M} \hat{\xi}_i^k g(\mathfrak{z}_i) \xrightarrow{\mathbb{P}} \pi(g)/\pi(p). \tag{7}$$

Then, (i) directly follows from (7) applied to both $g = h$ and $g = 1$. Now, denote by $\mathcal{F}_i$ the σ-field $\mathcal{F}_i := \sigma(\mathfrak{z}_1, \dots, \mathfrak{z}_{i+1}, \hat{\xi}_1^k, \dots, \hat{\xi}_i^k)$. Since $\mathbb{E}[\hat{\xi}_i^k g(\mathfrak{z}_i)|\mathcal{F}_{i-1}] = g(\mathfrak{z}_i)/p(\mathfrak{z}_i)$, we have

$$M^{-1} \sum_{i=1}^{M} \hat{\xi}_i^k g(\mathfrak{z}_i) = \left(\sum_{i=1}^{M} U_{M,i} - \mathbb{E}[U_{M,i}|\mathcal{F}_{i-1}] \right) + M^{-1} \sum_{i=1}^{M} g(\mathfrak{z}_i)/p(\mathfrak{z}_i)$$

with $U_{M,i} := M^{-1} \hat{\xi}_i^k g(\mathfrak{z}_i)$. First, consider the second term on the right-hand side. Since $\varphi \geq 1$, the function p is in $\mathcal{C}_\varphi$; equation (4) then implies that $M/\{\sum_{i=1}^{M} 1/p(\mathfrak{z}_i)\} \xrightarrow{\mathbb{P}} \pi(p) > 0$ and therefore that

$$\forall g \in \mathcal{C}_\varphi \qquad M^{-1} \sum_{i=1}^{M} g(\mathfrak{z}_i)/p(\mathfrak{z}_i) \xrightarrow{\mathbb{P}} \pi(g)/\pi(p). \tag{8}$$

It remains to check that $\sum_{i=1}^{M} U_{M,i} - \mathbb{E}[U_{M,i}|\mathcal{F}_{i-1}] \xrightarrow{\mathbb{P}} 0$. We use asymptotic results for conditional triangular arrays of random variables given in Douc and Moulines (2008), Theorem 11. Obviously, since $|g| \in \mathcal{C}_\varphi$, we have

$$\sum_{i=1}^{M} \mathbb{E}[|U_{M,i}||\mathcal{F}_{i-1}] = M^{-1} \sum_{i=1}^{M} |g(\mathfrak{z}_i)|/p(\mathfrak{z}_i) \xrightarrow{\mathbb{P}} \pi(|g|)/\pi(p)$$

and we only need to show that $\sum_{i=1}^{M} \mathbb{E}[|U_{M,i}|\mathbb{I}\{|U_{M,i}| > \varepsilon\}|\mathcal{F}_{i-1}] \xrightarrow{\mathbb{P}} 0$. Let $C > 0$ and note that $\{|U_{M,i}| > \varepsilon\} \subset \{|g(\mathfrak{z}_i)| > (\varepsilon M)/C\} \cup \{\hat{\xi}_i^k > C\}$. Again, using $\mathbb{E}[\hat{\xi}_i^k g(\mathfrak{z}_i)|\mathcal{F}_{i-1}] = g(\mathfrak{z}_i)/p(\mathfrak{z}_i)$, we have

$$\begin{aligned} &\sum_{i=1}^{M} \mathbb{E}[|U_{M,i}|\mathbb{I}\{|U_{M,i}| > \varepsilon\}|\mathcal{F}_{i-1}] \\ &\quad \leq \frac{1}{M} \sum_{i=1}^{M} \frac{|g(\mathfrak{z}_i)|\mathbb{I}\{|g(\mathfrak{z}_i)| > (\varepsilon M)/C\}}{p(\mathfrak{z}_i)} + \frac{1}{M} \sum_{i=1}^{M} \frac{F_C(\mathfrak{z}_i)}{p(\mathfrak{z}_i)} \end{aligned} \tag{9}$$

with $F_C(\mathfrak{z}_i) := |g(\mathfrak{z}_i)|\mathbb{E}[\hat{\xi}_i^k \mathbb{I}\{\xi_i^k > C\}|\mathfrak{z}_i]p(\mathfrak{z}_i)$. Since $F_C \leq |g|$, we have $F_C \in \mathcal{C}_\varphi$. Then, again using (8), we have

$$\frac{1}{M} \sum_{i=1}^{M} \frac{|g(\mathfrak{z}_i)|\mathbb{I}\{|g(\mathfrak{z}_i)| > (\varepsilon M)/C\}}{p(\mathfrak{z}_i)} \xrightarrow{\mathbb{P}} 0,$$

$$\frac{1}{M} \sum_{i=1}^{M} \frac{F_C(\mathfrak{z}_i)}{p(\mathfrak{z}_i)} \xrightarrow{\mathbb{P}} \pi(F_C)/\pi(p),$$

which can be arbitrarily small when taking C sufficiently large. Indeed, using Lebesgue's theorem in the definition of F_C, for any fixed $\mathfrak{z}$, $\lim_{C\to\infty} F_C(\mathfrak{z}) = 0$ and then, again using Lebesgue's theorem, $\lim_{C\to\infty} \pi(F_C) = 0$. Finally, (7) is proved. The proof of (i) follows.

We now consider (ii). Without loss of generality, we assume that $\pi(h) = 0$. Write

$$\sqrt{M}\delta_M^k = \frac{M^{-1/2}\sum_{i=1}^M \hat{\xi}_i^k h(\mathfrak{z}_i)}{M^{-1}\sum_{i=1}^M \hat{\xi}_i^k}.$$

By (7), the denominator of the right-hand side converges in probability to $1/\pi(p)$. Thus, by Slutsky's lemma, we only need to prove a central limit theorem for the numerator of the right-hand side. Define $U_{M,i} := M^{-1/2}\hat{\xi}_i^k h(\mathfrak{z}_i)$ and write

$$M^{-1/2}\sum_{i=1}^M \hat{\xi}_i^k h(\mathfrak{z}_i) = \left(\sum_{i=1}^M U_{M,i} - \mathbb{E}[U_{M,i}|\mathcal{F}_{i-1}]\right) + M^{-1/2}\sum_{i=1}^M h(\mathfrak{z}_i)/p(\mathfrak{z}_i).$$

Since $h \in \mathcal{C}_\psi$ and $M^{-1}\sum_{i=1}^M 1/p(\mathfrak{z}_i) \xrightarrow{\mathbb{P}} 1/\pi(p)$, the second term, thanks again to Slutsky's lemma and equation (5), converges in distribution to $\mathcal{N}(0, \Gamma(h)/\pi^2(p))$. Now, consider the first term on the right-hand side. We will once again use asymptotic results on triangular arrays of random variables [as in Douc and Moulines (2008), Theorem 13]. We have

$$\begin{aligned}
&\sum_{i=1}^M \mathbb{E}[U_{M,i}^2|\mathcal{F}_{i-1}] - (\mathbb{E}[U_{M,i}|\mathcal{F}_{i-1}])^2 \\
&\qquad = M^{-1}\sum_{i=1}^M (h^2(\mathfrak{z}_i)\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i]p(\mathfrak{z}_i))/p(\mathfrak{z}_i) \\
&\qquad \xrightarrow{\mathbb{P}} \pi[\mathbb{V}[\hat{\xi}_i^k|\cdot]h^2(\cdot)p(\cdot)]/\pi(p),
\end{aligned}$$

by (8) applied to the nonnegative function $\mathfrak{z}_i \mapsto h^2(\mathfrak{z}_i)\mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i]p(\mathfrak{z}_i)$ which is in $\mathcal{C}_\varphi$ since it is bounded from above by $h^2/p \in \mathcal{C}_\varphi$. It remains to show that, for any $\varepsilon > 0$,

$$\sum_{i=1}^M \mathbb{E}\big[|U_{M,i}|^2\mathbb{I}_{|U_{M,i}|>\varepsilon}|\mathcal{F}_{i-1}\big] \xrightarrow{\mathbb{P}} 0. \tag{10}$$

Following the same lines as in the proof of (i), note that for any $C > 0$, we have $\{|U_{M,i}| > \varepsilon\} \subset \{|h(\mathfrak{z}_i)| > (\varepsilon\sqrt{M})/C\} \cup \{\hat{\xi}_i^k > C\}$. Using the fact that

$$\mathbb{E}[(\hat{\xi}_i^k)^2|\mathcal{F}_{i-1}] = \mathbb{V}[\hat{\xi}_i^k|\mathfrak{z}_i] + (\mathbb{E}[\hat{\xi}_i^k|\mathfrak{z}_i])^2 \le 2/p^2(\mathfrak{z}_i),$$

we have

$$\sum_{i=1}^{M} \mathbb{E}[|U_{M,i}|\mathbb{I}\{|U_{M,i}| > \varepsilon\}|\mathcal{F}_{i-1}]$$
$$\leq \frac{2}{M}\sum_{i=1}^{M} \frac{h^2(\mathfrak{z}_i)\mathbb{I}\{|h(\mathfrak{z}_i)| > (\varepsilon\sqrt{M})/C\}}{p^2(\mathfrak{z}_i)} + \frac{1}{M}\sum_{i=1}^{M} \frac{F_C(\mathfrak{z}_i)}{p(\mathfrak{z}_i)}$$

with $F_C(\mathfrak{z}_i) := h^2(\mathfrak{z}_i)\mathbb{E}[(\hat{\xi}_i^k)^2\mathbb{I}\{\xi_i^k > C\}|\mathfrak{z}_i]p(\mathfrak{z}_i)$. Since $F_C \leq (2h^2)/p$ and $h^2/p \in \mathcal{C}_\varphi$, we have $F_C \in \mathcal{C}_\varphi$. Then, again using (8),

$$\frac{1}{M}\sum_{i=1}^{M} \frac{(h^2(\mathfrak{z}_i)/p(\mathfrak{z}_i))\mathbb{I}\{|h(\mathfrak{z}_i)| > (\varepsilon\sqrt{M})/C\}}{p(\mathfrak{z}_i)} \xrightarrow{\mathbb{P}} 0,$$
$$\frac{1}{M}\sum_{i=1}^{M} \frac{F_C(\mathfrak{z}_i)}{p(\mathfrak{z}_i)} \xrightarrow{\mathbb{P}} \pi(F_C)/\pi(p),$$

which can be made arbitrarily small by taking C sufficiently large. Indeed, as in the proof of (i), one can use Lebesgue's theorem in the definition of F_C so that for any fixed $\mathfrak{z}$, $\lim_{C\to\infty} F_C(\mathfrak{z}) = 0$. Then, again using Lebesgue's theorem, $\lim_{C\to\infty} \pi(F_C) = 0$. Finally, (10) is proved. The proof of (ii) follows. □

The main consequence of this central limit theorem is thus that, asymptotically, the correlation between the ξ_i's vanishes, hence that the variance ordering on the ξ_i's extends to the same ordering on the δ_M's.

It remains to link the central limit theorem of the usual Markov chain Monte Carlo (MCMC) estimator (1) with the central limit theorem expressed in (6), with $k = 0$ associated with the accepted values. We will need some additional assumptions, starting with a maximal inequality for the Markov chain $(\mathfrak{z}_i)_i$: there exists a measurable function ζ such that for any starting point x,

$$\forall h \in \mathcal{C}_\zeta \qquad \mathbb{P}_x\left(\left|\sup_{0\leq i\leq N}\sum_{j=0}^{i}[h(\mathfrak{z}_j) - \tilde{\pi}(h)]\right| > \varepsilon\right) \leq \frac{NC_h(x)}{\varepsilon^2}, \tag{11}$$

where $\mathbb{P}_x$ is the probability measure induced by the Markov chain $(\mathfrak{z}_i)_{i\geq 0}$ starting from $\mathfrak{z}_0 = x$.

Moreover, we assume that there exists a measurable function $\phi \geq 1$ such that for any starting point x,

$$\forall h \in \mathcal{C}_\phi \qquad \tilde{Q}^n(x,h) \xrightarrow{\mathbb{P}} \tilde{\pi}(h) = \pi(ph)/\pi(p), \tag{12}$$

where $\tilde{Q}$ is the transition kernel of $(\mathfrak{z}_i)_i$ expressed in Lemma 1.

THEOREM 2. *In addition to the assumptions of Theorem 1, assume that h is a measurable function such that $h/p \in \mathcal{C}_\zeta$ and $\{C_{h/p}, h^2/p^2\} \subset \mathcal{C}_\phi$. Assume, moreover, that*

$$\sqrt{M}(\delta_M^0 - \pi(h)) \xrightarrow{\mathcal{L}} \mathcal{N}(0, V_0[h - \pi(h)]).$$

Then, for any starting point x,

$$\sqrt{M_N}\left(\frac{\sum_{t=1}^N h(x^{(t)})}{N} - \pi(h)\right) \underset{N\to\infty}{\xrightarrow{\mathcal{L}}} \mathcal{N}(0, V_0[h - \pi(h)]),$$

where M_N is defined by

$$\sum_{i=1}^{M_N} \hat{\xi}_i^0 \le N < \sum_{i=1}^{M_N+1} \hat{\xi}_i^0. \tag{13}$$

PROOF. Without loss of generality, we assume that $\pi(h) = 0$. In this proof, we will denote by $\mathbb{P}_x$ (resp., $\mathbb{E}_x$) the probability (resp., expectation) associated with the Markov chain $(x^{(t)})_{t\ge 0}$ starting from a fixed point x. Using (7) with $g = 1$, one may divide (13) by M_N and let N go to infinity. This yields that $M_N/N \xrightarrow{\mathbb{P}} \pi(p) > 0$. Then, by Slutsky's lemma, Theorem 2 will be proven if we are able to show that

$$\sqrt{N}\left(\frac{\sum_{t=1}^N h(x^{(t)})}{N} - \pi(h)\right) \underset{N\to\infty}{\xrightarrow{\mathcal{L}}} \mathcal{N}(0, V_0[h - \pi(h)]/\pi(p)).$$

To that end, consider the decomposition

$$N^{-1/2} \sum_{t=1}^N h(x^{(t)}) := \Delta_{N,1} + \Delta_{N,2} + \Delta_{N,3},$$

where $M_N^\star := \lfloor N\pi(p) \rfloor$,

$$\Delta_{N,1} := N^{-1/2}\left(N - \sum_{i=1}^{M_N} \hat{\xi}_i^0\right) h(\mathfrak{z}_{M_N+1}),$$

$$\Delta_{N,2} := N^{-1/2}\left(\sum_{i=1}^{M_N} \hat{\xi}_i^0 h(\mathfrak{z}_i) - \sum_{i=1}^{M_N^\star} \hat{\xi}_i^0 h(\mathfrak{z}_i)\right),$$

$$\Delta_{N,3} := N^{-1/2} \sum_{i=1}^{M_N^\star} \hat{\xi}_i^0 h(\mathfrak{z}_i).$$

Using the fact that $0 \le N - \sum_{i=1}^{M_N} \hat{\xi}_i^0 \le \hat{\xi}_{M_N+1}^0$ and Markov's inequality, we have

$$\mathbb{P}_x(|\Delta_{N,1}| > \varepsilon) \le \frac{\mathbb{E}_x(\hat{\xi}_{M_N+1}^0 |h(\mathfrak{z}_{M_N+1})|)}{\varepsilon\sqrt{N}} = \frac{\tilde{Q}^{M_N+1}(x, |h|/p)}{\varepsilon\sqrt{N}},$$

which converges in probability to 0 using the facts that $|h|/p \le h^2/p^2 + 1$ and $\{h^2/p^2, 1\} \subset \mathcal{C}_\phi$. Thus, $\Delta_{N,1} \xrightarrow{\mathbb{P}} 0$. We now consider $\Delta_{N,2}$. Note that

$$\mathbb{P}_x(|\Delta_{N,2}| > \varepsilon) \le \mathbb{P}_x\big(|A_N| > \varepsilon\sqrt{N}/2\big) + \mathbb{P}_x\big(|B_N| > \varepsilon\sqrt{N}/2\big) \tag{14}$$

with

$$A_N = \sum_{i=M_N \wedge M_N^\star}^{M_N \vee M_N^\star} h(\mathfrak{z}_i)/p(\mathfrak{z}_i) \quad \text{and} \quad B_N = \sum_{i=M_N \wedge M_N^\star}^{M_N \vee M_N^\star} (\hat{\xi}_j^0 - 1/p(\mathfrak{z}_i))h(\mathfrak{z}_i).$$

Now, pick an arbitrary $\alpha \in (0, 1)$ and set $\underline{M}_N := M_N^\star(1-\alpha)$ and $\overline{M}_N := M_N^\star(1+\alpha)$. Since $M_N/N \xrightarrow{\mathbb{P}} \pi(p)$ for all $\eta > 0$, there exists N_0 such that for all $N \ge N_0$, $\mathbb{P}_x(\underline{M}_N \le M_N \le \overline{M}_N) \ge 1 - \eta$. Then, obviously for $N \ge N_0$, the first term on the right-hand side of (14) is bounded by

$$\begin{aligned} &\mathbb{P}_x\big(|A_N| > \varepsilon\sqrt{N}/2\big) \\ &\quad \le \eta + \mathbb{P}_x\left(\sup_{M_N^\star \le i \le \overline{M}_N} \left|\sum_{j=M_N^\star}^{i} h(\mathfrak{z}_j)/p(\mathfrak{z}_j)\right| > \varepsilon\sqrt{N}/2\right) \\ &\qquad + \mathbb{P}_x\left(\sup_{\underline{M}_N^\star \le i \le M_N^\star} \left|\sum_{j=i}^{M_N^\star} h(\mathfrak{z}_j)/p(\mathfrak{z}_j)\right| > \varepsilon\sqrt{N}/2\right). \end{aligned} \tag{15}$$

Using (11), the second term of the right-hand side is bounded by

$$4\overline{M}_N - M_N^\star \mathbb{E}_x[C_{h/p}(\mathfrak{z}_{M_N^\star})]/\varepsilon^2 N,$$

which converges to $4\alpha\pi(p)\tilde{\pi}(C_{h/p})/\varepsilon^2$ as N goes to infinity, using the fact that $C_{h/p} \in \mathcal{C}_\phi$. The resulting bound can thus be arbitrarily small as α goes to 0. Similarly, one can bound the third term on the right-hand side of (15) and let N go to infinity. Again letting α go to 0, we obtain that $A_N/\sqrt{N} \xrightarrow{\mathbb{P}} 0$. Similarly, the second term of the right-hand side of (14) is bounded by

$$\begin{aligned} &\mathbb{P}_x\big(|B_N| > \varepsilon\sqrt{N}/2\big) \\ &\quad \le \eta + \mathbb{P}_x\left(\sup_{M_N^\star \le i \le \overline{M}_N} \left|\sum_{j=M_N^\star}^{i} \left(\hat{\xi}_j^0 - \frac{1}{p(\mathfrak{z}_j)}\right)h(\mathfrak{z}_j)\right| > \varepsilon\sqrt{N}/2\right) \\ &\qquad + \mathbb{P}_x\left(\sup_{\underline{M}_N^\star \le i \le M_N^\star} \left|\sum_{j=i}^{M_N^\star} \left(\hat{\xi}_j^0 - \frac{1}{p(\mathfrak{z}_j)}\right)h(\mathfrak{z}_j)\right| > \varepsilon\sqrt{N}/2\right). \end{aligned} \tag{16}$$

We write $R_N = \sum_{\ell=1}^N (\hat{\xi}_\ell^0 - \frac{1}{p(\mathfrak{z}_\ell)})h(\mathfrak{z}_\ell)$. Clearly, (R_N) is a $\mathcal{F}$-martingale where $\mathcal{F} = (\mathcal{F}_i)_{i\ge 1}$ and $\mathcal{F}_i$ is the σ-field $\mathcal{F}_i := \sigma(\mathfrak{z}_1, \dots, \mathfrak{z}_{i+1}, \hat{\xi}_1^0, \dots, \hat{\xi}_i^0)$. Then, by Kolmogorov's inequality, one can bound the second term of (16) in the following way:

$$\begin{aligned}
&\mathbb{P}_x\Big(\sup_{M_N^\star \le i \le \overline{M}_N} |R_i - R_{M_N}| > \varepsilon\sqrt{N}/2\Big) \\
&\quad \le 4\frac{\mathbb{E}_x[(R_{M_N^\star} - R_{M_N})^2]}{\varepsilon^2 N} = \frac{4}{\varepsilon^2 N}\mathbb{E}_x\Bigg[\sum_{i=M_N^\star}^{\overline{M}_N} \frac{1-p(\mathfrak{z}_i)}{p^2(\mathfrak{z}_i)} h^2(\mathfrak{z}_i)\Bigg] \\
&\quad = \frac{4(\overline{M}_N - M_N^\star + 1)}{\varepsilon^2 N} \frac{\sum_{i=M_N^\star}^{\overline{M}_N} \tilde{Q}^i(x, (1-p)/p^2h^2)}{\overline{M}_N - M_N^\star + 1} \\
&\quad \xrightarrow{\mathbb{P}} \frac{4\alpha\pi((1-p)/ph^2)}{\varepsilon^2},
\end{aligned}$$

which can be arbitrarily small as α goes to 0. Similarly, one can bound the third term of (16) and let N go to infinity. Finally, letting α go to 0, we obtain that $B_N/\sqrt{N} \xrightarrow{\mathbb{P}} 0$. Thus, $\Delta_{N,2} \xrightarrow{\mathbb{P}} 0$. Finally, by Slutsky's lemma,

$$\Delta_{N,3} := (N/M_N^\star)^{-1/2} \frac{\sum_{i=1}^{M_N^\star} \hat{\xi}_i^0 h(\mathfrak{z}_i)}{\sqrt{M_N^\star}} \xrightarrow{\mathcal{L}} \mathcal{N}\big(0, V_0[h - \pi(h)]/\pi(p)\big).$$

The proof is thus complete. □

Note that the above analysis also provides us with a universal control variate for Metropolis–Hastings algorithms. Indeed, while Lemma 2 shows that

$$\hat{\xi}_i = 1 + \sum_{j=1}^{\infty} \prod_{\ell \le j} \{1 - \alpha(\mathfrak{z}_i, y_\ell)\}$$

is an unbiased estimator of $1/p(\mathfrak{z}_i)$, a simple independent estimator of $p(\mathfrak{z}_i)$ is provided by $\alpha(\mathfrak{z}_i, y_0)$ when y_0 is an independent draw from $q(Y|\mathfrak{z}_i)$. While the variation in this estimate may result in a negligible improvement in the control variate estimation, it is nonetheless available for free in all settings and should thus be exploited.

4. Illustrations. We first consider a series of toy examples to assess the possible gains brought about by the essentially free Rao–Blackwellization. Our initial example is a random walk Metropolis–Hastings algorithm with target the $\mathcal{N}(0, 1)$ distribution and with proposal $q(y|x) = \varphi(x - y; \tau)$, a normal random walk with scale τ. The acceptance probability is then the ratio of the targets, and Figure 1 illustrates the gain provided by the Rao–Blackwellization scheme by repeating the simulation 250 times and by representing the 90% range as well as the whole range of both estimators. The gain provided by the Rao–Blackwellization is not

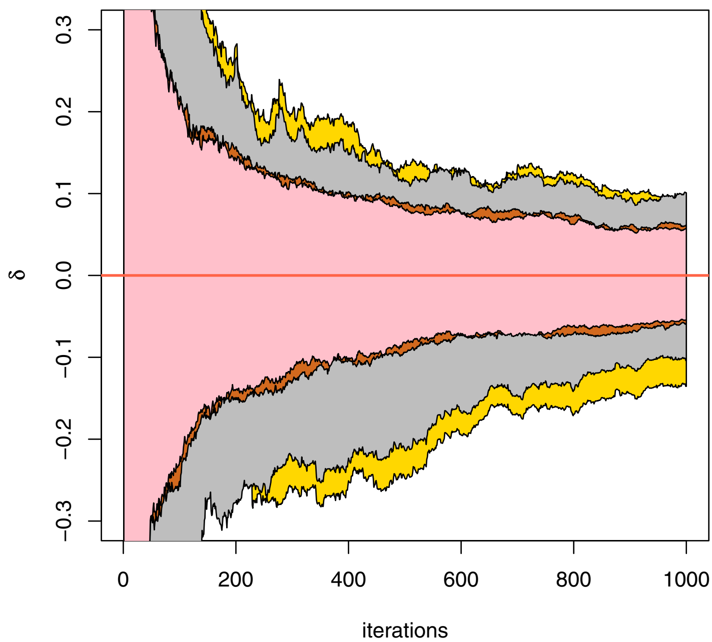

FIG. 1. *Overlay of the variations of* 250 *i.i.d. realizations of the estimates* δ (gold) *and* δ^∞ (grey) *of* $\mathbb{E}[X] = 0$ *for* 1000 *iterations*, *along with the* 90% *interquantile range for the estimates* δ (brown) *and* δ^∞ (pink), *in the setting of a random walk Gaussian proposal with scale* $\tau = 10$.

huge with respect to the overlap of both estimates, but one must consider that the variability of the estimator δ is due to two sources of randomness, one due to the $\mathfrak{n}_i$'s and the other due to the $\mathfrak{z}_i$'s. In addition, the gain forecasted by the above developments is in terms of variance, not of tails, and this gain is illustrated in Table 1. In this table, we provide the ratio of the empirical variances of the terms $\mathfrak{n}_i h(\mathfrak{z}_i)$ and $\hat{\xi}_i h(\mathfrak{z}_i)$ for several functions h. The minimal gains when $\tau = 0.1$ are explained by the fact that the acceptance probability is almost 1 with such a small

TABLE 1
Ratios of the empirical variances of the components of the estimators δ^∞ *and* δ *of* $\mathbb{E}[h(X)]$ *for* 100 *MCMC iterations over* 10^3 *replications*, *in the setting of a random walk Gaussian proposal with scale* τ, *when started with a normal simulation*

$h(x)$	x	x^2	$\mathbb{I}_{X>0}$	$p(x)$
$\tau = 0.1$	0.971	0.953	0.957	0.207
$\tau = 2$	0.965	0.942	0.875	0.861
$\tau = 5$	0.913	0.982	0.785	0.826
$\tau = 7$	0.899	0.982	0.768	0.820

TABLE 2
Evaluations of the additional computing effort due to the use of the Rao–Blackwell correction: median and mean numbers of additional iterations, 80% and 90% quantiles for the additional iterations, and ratio of the average R computing times obtained over 10^5 simulations in the same setting as Table 1

	Median	**Mean**	$q_{0.8}$	$q_{0.9}$	**Time**
$\tau = 0.1$	1.0	6.49	5.0	11	2.33
$\tau = 2$	0.0	7.06	4.3	11	6.5
$\tau = 5$	0.0	9.02	4.6	13	8.4
$\tau = 7$	0.0	9.47	4.8	13	3.5

scale, while the higher rejection rate of 82% when $\tau = 7$ leads to more improvement in the variances because of a higher variability in the original $\mathfrak{n}_i$'s. Note that the last column of Table 1 estimates $\mathbb{E}[p(x)]$ via an additional draw from $q(Y|\mathfrak{z}_i)$, as pointed out at the end of the previous section. Table 2 gives an evaluation of the additional time required by the Rao–Blackwellization, even though this should not be overinterpreted. As shown by both the difference between the median and the mean additional times and the variability of the increase in the R computing time, despite the use of 10^5 replications, the occurrence of a few very lengthy runs accounts for the apparently much higher computing times. Note that this difficulty with very long runs can be completely bypassed when using a truncated version δ^k instead of the unconstrained version δ^∞.

Our second example is an independent Metropolis–Hastings algorithm with target the $\mathcal{N}(0, 1)$ distribution and with proposal a Cauchy $\mathcal{C}(0, 0.25)$ distribution. The outcome is quite similar, but producing a slightly superior improvement, as shown in Figure 2. Table 3 also indicates much more clearly that the gains in variance can be substantial. Once again, Table 4 shows that the computing time may vary quite widely due to a few outlying instances of late acceptance.

Our third example is an independent Metropolis–Hastings algorithm with target the $\mathcal{E}\text{xp}(\lambda)$ distribution and with proposal the $\mathcal{E}\text{xp}(\mu)$ distribution. In this case, the probability functions $p(x)$ in (2) and $r(x)$ in Proposition 1 can be derived in closed form as

$$p(x) = 1 - \frac{\lambda - \mu}{\lambda} e^{-\mu x} \quad \text{and} \quad r(x) = 1 - \frac{2(\lambda - \mu)}{2\lambda - \mu} e^{-\mu x}.$$

This special case means that we can compare the variability of the original Metropolis–Hastings estimator with its Rao–Blackwellized version δ_M^∞, but also with the optimal importance sampling version shown in (4). As illustrated by Table 5, the gain brought about by the Rao–Blackwellization is significant, even when compared with the reduction in variance of the optimal importance sampling version. Obviously, the most extreme case of $\mu = 0.1$ shows that the ideal

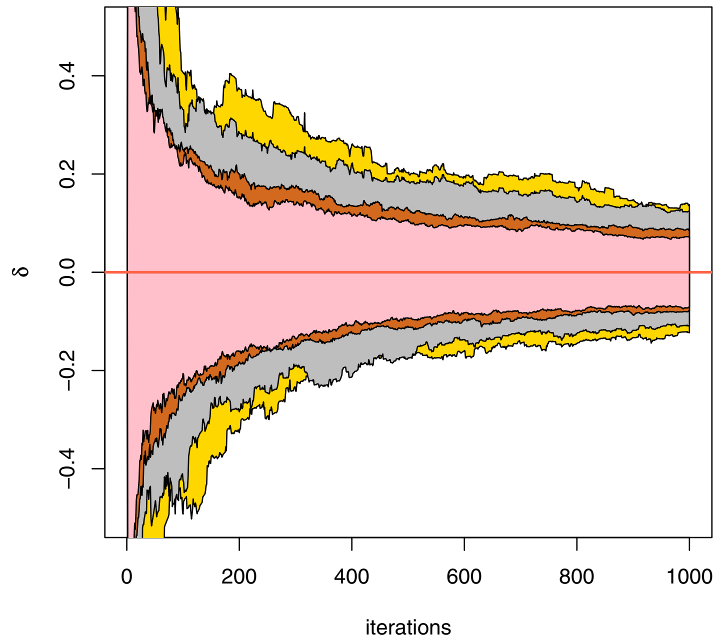

FIG. 2. *Overlay of the variations of* 250 *i.i.d. realizations of the estimates* δ (gold) *and* δ^∞ (grey) *of* $\mathbb{E}[X]=0$ *for* 1000 *iterations*, *along with the* 90% *interquantile range for the estimates* δ (brown) *and* δ^∞ (pink), *in the setting of an independent Cauchy proposal with scale* 0.25.

importance sampling estimator (4) could bring considerable improvement, were it available.

Our fourth and final toy example is a geometric $\mathcal{G}eo(\beta)$ target associated with a one-step random walk proposal:

$$\pi(x)=\beta(1-\beta)^x \quad \text{and} \quad 2q(y|x)=\begin{cases}\mathbb{I}_{|x-y|=1}, & \text{if } x>0,\\ \mathbb{I}_{|y|\le 1}, & \text{if } x=0.\end{cases}$$

TABLE 3

Ratios of the empirical variances of the components of the estimators δ^∞ *and* δ *of* $\mathbb{E}[h(X)]$ *for* 100 *MCMC iterations over* 10^3 *replications, in the setting of an independent Cauchy proposal with scale* τ *started with a normal simulation*

$h(x)$	x	x^2	$\mathbb{I}_{X>0}$	$p(x)$
$\tau=0.25$	0.677	0.630	0.663	0.599
$\tau=0.5$	0.790	0.773	0.716	0.603
$\tau=1$	0.937	0.945	0.889	0.835
$\tau=2$	0.781	0.771	0.694	0.591

TABLE 4
Evaluations of the additional computing effort due to the use of the Rao–Blackwell correction: median and mean numbers of additional iterations, 80% and 90% quantiles for the additional iterations, and ratio of the average R computing times obtained over 10^5 simulations in the same setting as Table 3

	Median	**Mean**	$q_{0.8}$	$q_{0.9}$	**Time**
$\tau = 0.25$	0.0	8.85	4.9	13	4.2
$\tau = 0.50$	0.0	6.76	4	11	2.25
$\tau = 1.0$	0.25	6.15	4	10	2.5
$\tau = 2.0$	0.20	5.90	3.5	8.5	4.5

For this problem,

$$p(x) = 1 - \beta/2 \quad \text{and} \quad r(x) = 1 - \beta + \beta^2/2.$$

We can therefore compute the gain in variance

$$\frac{p(x) - r(x)}{2p(x) - r(x)} \frac{2 - p(x)}{p^2(x)} = 2\frac{\beta(1-\beta)(2+\beta)}{(2-\beta^2)(2-\beta)^2},$$

which is optimal for $\beta = 0.174$, leading to a gain of 0.578, while the relative gain in variance is

$$\frac{p(x) - r(x)}{2p(x) - r(x)} \frac{2 - p(x)}{1 - p(x)} = \frac{(1-\beta)(2+\beta)}{(2-\beta^2)},$$

which is decreasing in β.

TABLE 5
Ratios of the empirical variances of the components of the estimators δ and δ^∞ of $\mathbb{E}[h(X)]$ for 100 MCMC iterations over 10^3 replications, in the setting of an independent exponential proposal with scale μ started with an exponential $\mathcal{E}\text{xp}(1)$ simulation from the target distribution; the second row is the optimal gain obtained by using $1/p(\mathfrak{z}_i)$ as importance weight, that is, the importance sampling estimator (4)

$h(x)$	x	x^2	$\mathbb{I}_{X>1}$	$p(x)$
$\mu = 0.9$	0.933	0.953	0.939	0.238
	0.787	0.774	0.859	0.106
$\mu = 0.5$	0.722	0.807	0.759	0.591
	0.291	0.394	0.418	0.285
$\mu = 0.3$	0.671	0.738	0.705	0.657
	0.131	0.175	0.263	0.295
$\mu = 0.1$	0.641	0.700	0.676	0.703
	0.0561	0.0837	0.159	0.289

We now apply the Rao–Blackwellization to a probit modeling of the Pima Indian diabetes study [Venables and Ripley (2002)]. The data set we consider covers a population of 332 women who were at least 21 years old, of Pima Indian heritage and living near Phoenix, Arizona. These women were tested for diabetes according to World Health Organization (WHO) criteria. The data were collected by the US National Institute of Diabetes and Digestive and Kidney Diseases, and is available with the basic R package. The goal is to explain the diabetes variable in terms of the body mass index. We use a standard representation of the diabetes binary variables y_i as indicators $y_i = \mathbb{I}_{z_i>0}$ of latent variables z_i, $z_i|\beta \sim \mathcal{N}(\mathbf{x}_i^{\mathrm{T}}\beta, 1)$, associated with a standard regression model, that is, where the $\mathbf{x}_i$'s are p-dimensional covariates and β is the vector of regression coefficients. Given β, the y_i's are independent Bernoulli random variables with $\mathbb{P}(y_i = 1|\beta) = \Phi(\mathbf{x}_i^{\mathrm{T}}\beta)$, where Φ is the standard normal cumulative distribution function. The choice of a prior distribution for the probit parameter β is open to debate [Marin and Robert (2007)], but, for the purposes of illustration, we opt for a flat prior. The Metropolis–Hastings algorithm associated with the posterior is a simple two-dimensional random walk proposal with a single scale τ, due to the normalization of the body mass index. Simulations based on different scales τ show significant improvements in the variance of the terms of δ and δ_∞ by a factor of 2. If we consider, in addition, the possible improvement brought about by the control variate indicated at the end of the previous section, the regression coefficient can be obtained by a simple regression of $\hat{\xi}_i h(\mathfrak{z}_i)$ over $\hat{\xi}_i \alpha(\mathfrak{z}_i, y_0)$ and Table 6 shows that this additional step brings about a significant improvement over the Rao–Blackwellized version.

TABLE 6

Ratios of the empirical variances of the components of the estimators δ and δ^∞ of $\mathbb{E}[h(\beta)]$ for 10^4 MCMC iterations, in the setting of a random walk proposal with scale τ started from the MLE estimate of β applied to the Pima Indian diabetes study; the second row for each value of τ is the additional improvement in the empirical variances resulting from using the control variate

$h(\boldsymbol{\beta})$	$\boldsymbol{\beta_1}$	$\boldsymbol{\beta_2}$	$\mathbb{I}_{\boldsymbol{\beta_2>0.5}}$
$\tau = 0.01$	0.523	0.516	0.944
	0.999	0.999	0.996
$\tau = 0.05$	0.481	0.518	0.877
	0.864	0.888	0.929
$\tau = 0.1$	0.550	0.555	0.896
	0.749	0.748	0.765
$\tau = 0.2$	0.562	0.568	0.845
	0.532	0.527	0.620
$\tau = 0.5$	0.556	0.565	0.778
	0.412	0.433	0.479

Acknowledgments. Both authors are grateful to Elke Thonnes and Gareth Roberts for organizing the MCMC workshop in Warwick, March 2009, that led to the completion of this work. Suggestions from the editorial team led to an improved presentation, for which the authors are most grateful.

REFERENCES

CASELLA, G. and ROBERT, C. (1996). Rao-Blackwellisation of sampling schemes. *Biometrika* **83** 81–94. MR1399157

CASELLA, G. and ROBERT, C. (1998). Post-processing accept-reject samples: Recycling and rescaling. *J. Comput. Graph. Statist.* **7** 139–157. MR1649370

DELMAS, J. F. and JOURDAIN, B. (2009). Does waste recycling really improves the multi-proposal Metropolis Hastings algorithm? An analysis based on control variates. *J. Appl. Probab.* **46** 938–959. MR2582699

DOUC, R. and MOULINES, E. (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *Ann. Statist.* **36** 2344–2376. MR2458190

GÅSEMYR, J. (2002). Markov chain Monte Carlo algorithms with independent proposal distribution and their relation to importance sampling and rejection sampling. Technical Report 2, Dept. Statistics, Univ. Oslo.

LATUSZYNSKI, K., KOSMIDIS, I., PAPASPILIOPOULOS, O. and ROBERTS, G. (2010). Simulating event of unknown probabilities via reverse time martingales. *Random Structures Algorithms*. To appear.

MALEFAKI, S. and ILIOPOULOS, G. (2008). On convergence of importance sampling and other properly weighted samples to the target distribution. *J. Statist. Plann. Inference* **138** 1210–1225. MR2381076

MARIN, J.-M. and ROBERT, C. (2007). *Bayesian Core*. Springer, New York. MR2289769

PERRON, F. (1999). Beyond accept–reject sampling. *Biometrika* **86** 803–813. MR1741978

SAHU, S. and ZHIGLJAVSKY, A. (1998). Adaptation for self regenerative MCMC. Technical report, Univ. Wales, Cardiff.

SAHU, S. and ZHIGLJAVSKY, A. (2003). Self regenerative Markov chain Monte Carlo with adaptation. *Bernoulli* **9** 395–422. MR1997490

VENABLES, W. and RIPLEY, B. (2002). *Modern Applied Statistics with S-PLUS*, 4th ed. Springer, New York. MR1337030

TELECOM SUDPARIS
9 RUE CHARLES FOURIER
91011 EVRY CEDEX
FRANCE
E-MAIL: randal.douc@it-sudparis.eu

CEREMADE
UNIVERSITÉ PARIS-DAUPHINE
75775 PARIS CEDEX 16
CREST, INSEE
92245 MALAKOFF CEDEX
FRANCE
E-MAIL: xian@ceremade.dauphine.fr