Biometrika (1996), **83**, 1, *pp.* 81–94
Printed in Great Britain

Rao–Blackwellisation of sampling schemes

BY GEORGE CASELLA
Biometrics Unit, Cornell University, Ithaca, New York 14850, U.S.A.

AND CHRISTIAN P. ROBERT
Département de Mathématique, CNRS-URA 1378, Université de Rouen, BP 118, 76134 Mont Saint-Aignan cedex, France

SUMMARY

This paper proposes a post-simulation improvement for two common Monte Carlo methods, the Accept-Reject and Metropolis algorithms. The improvement is based on a Rao–Blackwellisation method that integrates over the uniform random variables involved in the algorithms, and thus post-processes the standard estimators. We show how the Rao–Blackwellised versions of these algorithms can be implemented and, through examples, illustrate the improvement in variance brought by these new procedures. We also compare the improved version of the Metropolis algorithm with ordinary and Rao–Blackwellised importance sampling procedures for independent and general Metropolis set-ups.

Some key words: Accept-reject; Gibbs sampling; Importance sampling; Metropolis; Monte Carlo algorithm; Polynomial computing time.

1. INTRODUCTION

The Rao–Blackwell Theorem, a well-known result in mathematical statistics, see e.g. Lehmann (1983, p. 50), shows how to improve upon any given estimator under every convex loss function. The improvement is obtained by calculating a conditional expected value, often involving integrating out an ancillary statistic. The appeal of this important theorem has been extended to simulation settings in the case of Markov chain Monte Carlo methods by Gelfand & Smith (1990) and Liu, Wong & Kong (1994, 1995). They worked, for the most part, in the context of Gibbs sampling. The Rao–Blackwell Theorem was used to show that smoothed estimators, using the available conditional distributions, were an improvement over nonsmoothed estimators.

When applied to variance improvement, the core of the Rao–Blackwell Theorem is the well-known relationship between conditional and unconditional variance. Given an estimator $\tau(u, y)$ depending on two variables u and y, its variance satisfies

$$\text{var}(\tau) = \text{var}[E\{\tau(U, Y) \mid Y\}] + E[\text{var}\{\tau(U, Y) \mid Y\}]. \qquad (1\cdot1)$$

Therefore $\tau' = E\{\tau(U, Y) \mid Y\}$, having the same mean as τ, improves upon τ in terms of variance by a factor of $E[\text{var}\{\tau(U, Y) \mid Y\}]$. When τ', the Rao–Blackwellised version of τ, can be computed it thus makes sense from a statistical point of view to use it. There are often many choices for the pair (u, y), but not all are appropriate, and may not be computable. The version we select in this paper is a somewhat natural choice, since many simulation methods rely on the side simulation of uniform $\mathscr{U}[0, 1]$ random variables. These uniform random variables are ancillary to the estimation problem and, when they can be integrated out, the resulting estimator improves on the usual average estimator at no cost in terms of simulation time.

We consider in this paper two general simulation algorithms, the Accept-Reject and the Metropolis algorithms. Estimators that are constructed from these algorithms will typically depend on the ancillary uniform random variables. Our Rao–Blackwellised procedures, analogous to $E\{\tau(U, Y) | Y\}$ explicitly integrate out these random variables, hence use the entire sample of candidate random variables simulated from the algorithm. We thus produce estimators that are independent of the ancillary variables, and are a weighted average of the entire sample. There is some similarity between these procedures and procedures resulting from an Importance Sampling algorithm, although our weight factors are generally more elaborate, but still computable in polynomial time. A full theoretical comparison of the Rao–Blackwell versus importance sampling approach is yet to be undertaken, but we provide some elements of comparison in simulation studies.

The Rao–Blackwellisation methods presented in this paper are essentially nonparametric, in the sense that they depend on neither the form of the density nor the estimated function. In such a nonparametric setting, the Rao–Blackwellised estimator can be perceived as a uniform minimum variance unbiased estimator, being symmetric in the order statistics (Lehmann, 1983, § 2.4). This property should be contrasted with more familiar parametric Rao–Blackwellisation, where an estimator is improved upon by calculating its conditional expected value given a sufficient statistic. However, the only role that sufficiency plays is to ensure that the resulting conditional expectation is, indeed, an estimator, since conditioning on an insufficient statistic may result in a quantity that depends on unknown parameters. Therefore, one can interpret the Rao–Blackwell Theorem as saying that, if the calculation of a conditional expectation can result in an estimator, such an estimator is superior to the original. Parametric Rao–Blackwellisation is often used in Gibbs sampling.

We emphasise that our approach and, in particular, the optimisations involved in the derivation of the improvements, is statistical rather than computational. The overall goal of the statistician is to process samples in an optimal way. We therefore consider our improvements as a post-simulation processing of a generated sample, which is statistically superior, although it may be computationally inferior in taking more computer time. For instance, it may be more efficient in time to generate an Accept-Reject sample of size t^* ($t^* > t$) such that the variance of the estimate based on this augmented sample is smaller than that of the Rao–Blackwellised procedure for sample size t. Although such an approach may be computationally optimal, it will not be statistically optimal.

In § 2 we consider the Accept-Reject algorithm and derive the corresponding weights for the Rao–Blackwellised version of the estimator. An example, illustrating the potential improvement, is also given. Section 3 similarly treats the Metropolis algorithm in the independent case using the same setup as § 2. However, since Accept-Reject leads to the generation of a sample of random size, while Metropolis relies on a fixed sample, comparison of the methods is not straightforward. In § 3 we also provide an example illustrating the magnitude of improvement possible in this case. Section 4 extends the Metropolis Rao–Blackwellisation improvement to the general case, where the simulated random variables are not independent, and § 5 discusses importance sampling based on the sample generated by the general Metropolis algorithm. In the general Metropolis case, we also give a Rao–Blackwellised version of the importance sampling estimator which turns out to have the same order of complexity as the other Rao–Blackwellised procedures. We give an example that shows that the Rao–Blackwellised importance-sampling estimator can dramatically improve upon the Rao–Blackwellised Metropolis estimator. However, the importance sampling approach falls short of providing a true sample from the distribution of interest, contrary to Accept-Reject and Metropolis approaches. Lastly, § 6 contains a discussion and some conclusions.

2. The Accept-Reject algorithm

The Accept-Reject algorithm is based on the following lemma.

LEMMA 2·1. *If f and g are two densities, and there exists $M < \infty$ such that $f(x) \leqslant Mg(x)$ for every x, the random variable X provided by the algorithm:*

1. *simulate $Y \sim g(y)$;*
2. simulate $U \sim \mathscr{U}[0, 1]$ *and take $X = Y$ if $U \leqslant f(Y)/Mg(Y)$, otherwise, repeat step* 1;

is distributed according to f.

This algorithm is widely used for simulation, often with some refinements (Devroye, 1985, § 2.3; Gilks & Wild, 1992) to increase the probability of acceptance at each step. Nonetheless, this method leads to the rejection of a part of the sample simulated from g, that is, although we simulate the values $Y_1, \ldots, Y_n$, the Y_i's for which $U_i > f(Y_i)/Mg(Y_i)$ are eliminated. We now propose an improvement upon the original Accept-Reject procedure which makes use of every simulated value.

We consider the distributions f and g to be given; the choice of g is not considered in this paper. A sequence $Y_1, Y_2, \ldots$ of independent random variables is generated from g along with a corresponding sequence $U_1, U_2, \ldots$ of uniform random variables. Given a function h, the Accept-Reject estimator of $E^f\{h(X)\}$, based upon a sample $X_1, \ldots, X_t$ generated according to Lemma 2·1, is given by

$$\hat{\tau}_1 = \frac{1}{t} \sum_{i=1}^{t} h(X_i). \tag{2·1}$$

For a fixed sample size t, that is for a fixed number of accepted random variables, the number of generated Y_i's is a random integer N satisfying

$$\sum_{i=1}^{N} I(U_i \leqslant w_i) = t, \quad \sum_{i=1}^{N-1} I(U_i \leqslant w_i) = t - 1,$$

where we define $w_i = f(Y_i)/Mg(Y_i)$. Since $\hat{\tau}_1$ can be written as

$$\hat{\tau}_1 = \frac{1}{t} \sum_{i=1}^{N} I(U_i \leqslant w_i) h(Y_i),$$

the conditional expectation

$$\hat{\tau}_2 = \frac{1}{t} E\left\{ \sum_{i=1}^{N} I(U_i \leqslant w_i) h(Y_i) \,\middle|\, N, Y_1, \ldots, Y_N \right\} \tag{2·2}$$

improves upon (2·1).

For completeness, we first present some of the necessary distributions for deriving a manageable formula for the estimator $\hat{\tau}_2$. The joint distribution of

$(N, Y_1, \ldots, Y_N, U_1, \ldots, U_N)$ is given by

$$\mathrm{pr}(N=n, Y_1 \leqslant y_1, \ldots, Y_n \leqslant y_n, U_1 \leqslant u_1, \ldots, U_n \leqslant u_n)$$

$$= \int_{-\infty}^{y_n} g(t_n)(u_n \wedge w_n)\, dt_n \int_{-\infty}^{y_1} \cdots \int_{-\infty}^{y_{n-1}} g(t_1) \ldots g(t_{n-1})$$

$$\times \sum_{(i_1,\ldots,i_{t-1})} \prod_{j=1}^{t-1} (w_{i_j} \wedge u_{i_j}) \prod_{j=t}^{n-1} (u_{i_j} - w_{i_j})^+ \, dt_1 \ldots dt_{n-1}.$$

The last sum is over all subsets of $\{1, \ldots, n-1\}$ of size $t-1$. The conditional density of the U's is given by

$$f(u_1, \ldots, u_n \mid N=n, Y_1, \ldots, Y_n) = \left\{ \sum_{(i_1,\ldots,i_{t-1})} \prod_{j=1}^{t-1} w_{i_j} \prod_{j=t}^{n-1} (1-w_{i_j}) \right\}^{-1}$$

$$\times \left\{ \sum_{(i_1,\ldots,i_{t-1})} \prod_{j=1}^{t-1} I(u_{i_j} \leqslant w_{i_j}) \prod_{j=t}^{n-1} I(u_{i_j} > w_{i_j}) \right\}$$

$$\times \frac{I(u_n \leqslant w_n)}{w_n}.$$

Using this distribution we can calculate, conditional on $(N, Y_1, \ldots, Y_N)$, the probability ρ_i of the event $\{U_i < w_i\}$ and thus derive the weights of $h(Y_i)$ in the estimator $\hat{\tau}_2$. The calculations involve averaging over permutations of the realised sample and yield, for $i < n$,

$$\rho_i = w_i \sum_{(i_1,\ldots,i_{t-2})} \prod_{j=1}^{t-2} w_{i_j} \prod_{j=t-1}^{n-2} (1-w_{i_j}) \Bigg/ \sum_{(i_1,\ldots,i_{t-1})} \prod_{j=1}^{t-1} w_{i_j} \prod_{j=t}^{n-1} (1-w_{i_j}), \qquad (2{\cdot}3)$$

while $\rho_n = 1$. The numerator sum is over all subsets of $\{1, \ldots, i-1, i+1, \ldots, n-1\}$ of size $t-2$, and the denominator sum is over all subsets of size $t-1$. From a statistical point of view, $(N, Y_1, \ldots, Y_N)$ is sufficient since the conditional distribution of the U_i's does not depend on the density function f.

The computation of $\hat{\tau}_2$ now follows quickly from equations (2·2) and (2·3), with the resulting estimator being an average over all the possible permutations of the realised sample, the permutations being weighted by their probabilities. The Rao–Blackwellised estimator is then a function only of $(N, Y_{(1)}, \ldots, Y_{(N-1)}, Y_N)$.

PROPOSITION 2·1. *For $N=n$, the Rao–Blackwellised version of* (2·1) *is given by*

$$\hat{\tau}_2 = \frac{1}{t} \sum_{i=1}^{n} \rho_i h(Y_i),$$

where ρ_i is provided by equation (2·3).

Note that the computation of the ρ_i's may appear formidable but these weights are easily derived from a recurrence relation which is of order n^2. For example, if we define

$$S_k(m) = \sum_{(i_1,\ldots,i_k)} \prod_{j=1}^{k} w_{i_j} \prod_{j=k+1}^{m} (1-w_{i_j}) \quad (k \leqslant m < n),$$

with $\{i_1, \ldots, i_m\} = \{1, \ldots, m\}$, $S_k(m) = 0$ for $k > m$ and $S_k^i(i) = S_k(i-1)$, we can recursively calculate

$$S_k(m) = w_m S_{k-1}(m-1) + (1 - w_m) S_k(m-1),$$

$$S_k^i(m) = w_m S_{k-1}^i(m-1) + (1 - w_m) S_k^i(m-1) \quad (m > i)$$

and note that weight ρ_i of (2·3) is given by

$$\rho_i = w_i S_{t-2}^i(n-1)/S_{t-1}(n-1) \quad (i < n).$$

The estimator $\hat{\tau}_2$ may be regarded as an importance sampling estimator which uses weights based on every variable generated in the process. Compared with the classical importance sampling procedure,

$$\hat{\tau}_3 = \frac{1}{N} \sum_{i=1}^{N} \frac{f(Y_i)}{g(Y_i)} h(Y_i), \tag{2·4}$$

$\hat{\tau}_2$ is more involved. However, given that the sample $Y_1, \ldots, Y_N$ is produced via Accept-Reject and therefore that N is random, having a negative binomial distribution with parameters t and $1/M$, $\hat{\tau}_3$ lacks its original motivation of unbiasedness. An unbiased version of $\hat{\tau}_3$ would have modified weights which are of the same degree of complexity of the ρ_i's.

Example 2·1. The target distribution is a gamma distribution $\mathcal{G}a(\alpha, \beta)$ with $\alpha > 1$. We set $\beta = 2\alpha$ so that the mean of the distribution is $\frac{1}{2}$. The candidate distribution we select is the gamma $\mathcal{G}a(a, b)$ distribution with $a = [\alpha]$ and $b = \beta a/\alpha$. We require $a < \alpha$ in order for M in Lemma 2·1 to be finite. The choice $b = 2a$ improves the fit between the two distributions since both means match. We consider two cases which reflect different acceptance rates for the Accept-Reject algorithm. In Case 1 we set $\alpha = 2{\cdot}434$, $a = 2$ and $1/M = 0{\cdot}9$, and, in Case 2, $\alpha = 20{\cdot}62$, $a = 2$ and $1/M = 0{\cdot}3$.

For each case we estimate the mean, chosen to be $\frac{1}{2}$, and a tail probability, chosen to be 5%, using both the simple Accept-Reject algorithm and its Rao–Blackwellised version. The averages of both the Accept-Reject estimators and their Rao–Blackwellised counterparts are presented in Tables 1 and 2. We also include mean squared error estimates for the Accept-Reject estimator and the improvement brought by Rao–Blackwellising. This improvement is measured by the percentage decrease in mean squared error. From both tables, it can be seen that the Rao–Blackwellisation provides a substantial decrease in mean squared error, reaching 60% in the case where the acceptance rate of the algorithm is 0·3. The improvement is better at the lower Accept-Reject acceptance rate partially because the Rao–Blackwellised sample is about three times bigger, with approximately two-thirds of the sample being discarded by the Accept-Reject algorithm. Another interesting observation is that the percent improvement in mean squared error remains constant as the Accept-Reject sample size increases, implying that the variance of the original Accept-Reject estimator does not approach the variance of the Rao–Blackwellised estimator even as the sample size increases.

3. The Metropolis algorithm in the independent case

Similarly to the Accept-Reject algorithm, the independent Metropolis algorithm constructs a sample $Z_1, \ldots, Z_n$ with distribution f from a sample $Y_1, \ldots, Y_n$, generated according to a distribution g, by discarding some of the Y_i's. Now, however, the Z_i are not necessarily independent, and the ratio f/g need not be bounded although Mengersen & Tweedie (1996) show that boundedness is desirable for good convergence properties. As proposed by Metropolis et al. (1953) and Hastings (1970), the Metropolis algorithm

Table 1. *Estimation of a gamma mean, chosen to be* $\frac{1}{2}$, *using the Accept-Reject algorithm, with acceptance rates* 0·9 *and* 0·3, *based on* 7500 *simulations*

	Acceptance rate 0·9				Acceptance rate 0·3			
AR sample size	AR estimate $\hat{\tau}_1$	RB estimate $\hat{\tau}_2$	AR MSE	Decrease in MSE (%)	AR estimate $\hat{\tau}_1$	RB estimate $\hat{\tau}_2$	AR MSE	Decrease in MSE (%)
10	0·5002	0·5007	0·0100	15·83	0·5005	0·5004	0·0012	53·28
25	0·5001	0·4999	0·0041	19·66	0·4997	0·5000	0·0005	59·28
50	0·4996	0·4997	0·0020	20·70	0·4998	0·5001	0·0002	59·28
100	0·4996	0·4997	0·0010	22·73	0·4995	0·5001	0·0001	60·46

AR, Accept-Reject; RB, Rao–Blackwellised; MSE, mean squared error.

Table 2. *Estimation of a gamma tail probability, chosen to be* 0·05, *using the Accept-Reject algorithm, with acceptance rates* 0·9 *and* 0·3, *based on* 7500 *simulations*

	Acceptance rate 0·9				Acceptance rate 0·3			
AR sample size	AR estimate $\hat{\tau}_1$	RB estimate $\hat{\tau}_2$	AR MSE	Decrease in MSE (%)	AR estimate $\hat{\tau}_1$	RB estimate $\hat{\tau}_2$	AR MSE	Decrease in MSE (%)
10	0·0512	0·0503	0·0049	20·45	0·0495	0·0505	0·0048	63·02
25	0·0508	0·0501	0·0019	22·37	0·0499	0·0506	0·0019	69·80
50	0·0506	0·0507	0·0009	21·06	0·0491	0·0498	0·0009	72·17
100	0·0504	0·0503	0·0005	21·14	0·0487	0·0498	0·0005	73·77

AR, Accept-Reject; RB, Rao–Blackwellised; MSE, mean squared error.

starts with a random variable Z_0 generated from f, and at iteration $n+1$, generates $Z_{n+1} \mid Z_n$ as

$$Z_{n+1} = \begin{cases} Z_n & \text{with probability } 1-\rho_{n+1}, \\ Y_{n+1} \sim g(y) & \text{with probability } \rho_{n+1}, \end{cases} \tag{3·1}$$

where

$$\rho_{n+1} = \frac{f(Y_{n+1})/g(Y_{n+1})}{f(Z_n)/g(Z_n)} \wedge 1.$$

The assumption that Z_0 is generated from f is not very restrictive since the Metropolis algorithm converges, in the ergodic sense, to the distribution f (Tierney, 1994). Therefore, after a 'burn-in' period, the current simulation from the Metropolis algorithm can be considered to be approximately generated from the true distribution f.

A major difference between the Metropolis algorithm and the Accept-Reject algorithm of § 2 is that the sample size n is now fixed. To estimate a quantity $E^f\{h(Z)\}$, the usual estimator, justified by the Ergodic Theorem, is

$$\hat{\tau}_4 = \frac{1}{n+1} \sum_{i=0}^{n} h(Z_i),$$

which only involves the Y_i's accepted by the Metropolis algorithm. Using the full sample of Y_i's, the estimator $\hat{\tau}_4$ can be written in the form

$$\hat{\tau}_4 = \frac{1}{n+1}\left[h(Z_0) + \sum_{i=1}^{n} \{I(Z_i = Y_i)h(Y_i) + I(Z_i = Z_{i-1})h(Z_{i-1})\}\right]$$

$$= \frac{1}{n+1}\sum_{i=0}^{n} h(Y_i) \sum_{j=i}^{n} I(Z_j = Y_i). \quad (3{\cdot}2)$$

With the convention $Y_0 = Z_0$, this form incorporates all the Y_i's. Since the value of $\hat{\tau}_4$ is determined by the ancillary uniform variables linked with (3·1), we can apply a Rao–Blackwellisation to integrate over the U_i's. If we define

$$w_i = f(Y_i)/g(Y_i), \quad \rho_{ij} = (w_j/w_i) \wedge 1, \quad \delta_i = \mathrm{pr}(Z_i = Y_i \mid Y_1, \ldots, Y_i),$$

$$\xi_{ii} = 1, \quad \xi_{ij} = \prod_{t=i+1}^{j} (1 - \rho_{it}) \quad (i < j)$$

we get the following improvement upon $\hat{\tau}_4$.

PROPOSITION 3·1. *The Rao–Blackwellised version of* $\hat{\tau}_4$ *is*

$$\hat{\tau}_5 = \frac{1}{n+1}\sum_{i=0}^{n} \varphi_i h(Y_i),$$

where φ_i *is the expected number of times* Y_i *occurs in the sample, given by*

$$\varphi_i = \delta_i \sum_{j=i}^{n} \xi_{ij},$$

with

$$\delta_i = \sum_{j=0}^{i-1} \delta_j \xi_{j(i-1)} \rho_{ji} \quad (i > 0).$$

Proof. The probability that Y_i occurs at least once in the chain is

$$\delta_i = \sum_{j=0}^{i-1} \mathrm{pr}(Z_i = Y_i \mid Z_{i-1} = Y_j)\, \mathrm{pr}(Z_{i-1} = Y_j)$$

$$= \sum_{j=0}^{i-1} \rho_{ji}\, \mathrm{pr}(Z_{i-1} = Y_j) = \sum_{j=0}^{i-1} \rho_{ji}\delta_j\xi_{j(i-1)},$$

since the probability $\mathrm{pr}(Z_i = Y_j)$ $(i > j)$ is given by

$$\mathrm{pr}(Z_i = Y_j) = \mathrm{pr}(Z_i = Y_j \mid Z_{i-1} = Y_j) \ldots \mathrm{pr}(Z_j = Y_j)$$

$$= (1 - \rho_{ji}) \ldots \delta_j.$$

Moreover, once $Y_i = y_i$ is accepted, y_i remains in the sequence as the value of $z_i, \ldots, z_t$ until a new y_{t+1} is accepted. Therefore, using (3·2), the expected number of times y_j occurs in the sample $z_0, \ldots, z_n$ is indeed

$$\delta_j(1 + \xi_{j(j+1)} + \xi_{j(j+2)} + \ldots + \xi_{jn}). \qquad \square$$

Despite its intricate form, the improved estimator $\hat{\tau}_5$ only requires the computation of $n(n-1)/2$ ξ_{ij}'s and n probabilities δ_i. An open question about $\hat{\tau}_5$ is the improvement, not only over $\hat{\tau}_4$, but also over the importance sampling estimator (2·4).

Example 3·1. The target distribution is a Student's t distribution with 3 degrees of freedom, from which we estimate the mean and a 5% tail probability. Estimation is based on an independent Metropolis algorithm, with candidate distribution a Cauchy distribution, which is both easy to simulate from and results in a finite supremum of the ratio $f(y)/g(y)$.

We compare the usual Metropolis estimate $\hat{\tau}_4$ with its Rao–Blackwellised improvement $\hat{\tau}_5$ as well as an importance sampling estimate. We do not use the importance sampling estimator $\hat{\tau}_3$, but the often-used alternative

$$\hat{\tau}_6 = \sum_{i=1}^{n} \frac{f(Y_i)}{g(Y_i)} h(Y_i) \bigg/ \sum_{i=1}^{n} \frac{f(Y_i)}{g(Y_i)},$$

which, in our simulations, uniformly improved upon $\hat{\tau}_3$. Note that $\hat{\tau}_6$ does not depend on ancillary random variables, so Rao–Blackwellisation will not improve it.

The results, presented in Table 3, are similar to those of Example 2·1. For the sample sizes examined, the Rao–Blackwellised estimator yields a 40–50% decrease in mean squared error over the ordinary Metropolis mean. What is most surprising is that the importance sampling estimator yields an improvement that is comparable to the Rao–Blackwellised Metropolis estimator. This, perhaps, indicates that, as an estimation technique, the Metropolis mean may not be very desirable. Of course, the Metropolis algorithm has other uses, such as providing a sample from the target distribution. In our simulations the unbiased importance sampling estimator (2·4) did not provide an improvement comparable to its Rao–Blackwellised Metropolis counterpart. Thus, if it is desired to retain unbiasedness, the Rao–Blackwellised estimator is the choice.

Table 3. *Estimation of the mean* (0) *and tail probability* (0·05) *of a Student's t distribution with* 3 *degrees of freedom based on an independent Metropolis sample using a Cauchy distribution,* 7500 *simulations*

(*a*) *Mean* (0)

Sample size	Metrop. estimate $\hat{\tau}_4$	RB Metrop. $\hat{\tau}_5$	Import. sampl. $\hat{\tau}_6$	MSE Metrop.	MSE decrease $\hat{\tau}_5$ over $\hat{\tau}_4$ (%)	MSE decrease $\hat{\tau}_6$ over $\hat{\tau}_4$ (%)
10	−0·0035	−0·0058	−0·0064	0·3622	50·11	47·80
25	−0·0037	−0·0015	−0·0023	0·1468	49·39	51·05
50	−0·0020	0·0012	0·0004	0·0724	48·27	51·59
100	−0·0027	−0·0007	−0·0009	0·0361	46·68	51·19

(*b*) *Tail probability* (0·05)

Sample size	Metrop. estimate $\hat{\tau}_4$	RB Metrop. $\hat{\tau}_5$	Import. sampl. $\hat{\tau}_6$	MSE Metrop.	MSE decrease $\hat{\tau}_5$ over $\hat{\tau}_4$ (%)	MSE decrease $\hat{\tau}_6$ over $\hat{\tau}_4$ (%)
10	0·0490	0·0487	0·0505	0·0056	42·20	38·91
25	0·0486	0·0490	0·0498	0·0024	44·75	45·90
50	0·0488	0·0491	0·0497	0·0012	45·44	48·68
100	0·0494	0·0496	0·0498	0·0006	44·57	49·16

RB, Rao–Blackwellised; MSE, mean squared error.

4. The Metropolis algorithm in the general case

We now consider the general Metropolis algorithm, where Y_i is generated according to a conditional distribution $g(y|Z_{i-1})$. The transition from Z_n to Z_{n+1} is given by

$$Z_{n+1} = \begin{cases} Z_n & \text{with probability } 1-\rho_{n+1}, \\ Y_{n+1} \sim g(y_{n+1}|Z_n) & \text{with probability } \rho_{n+1}, \end{cases}$$

where

$$\rho_{n+1} = \frac{f(Y_{n+1})/g(Y_{n+1}|Z_n)}{f(Z_n)/g(Z_n|Y_{n+1})} \wedge 1.$$

When the support of $g(.|z)$ contains the support of f, convergence is guaranteed (Geyer, 1992).

The Y_i's do not form a Markov chain, the distribution of Y_i depending on $Y_1, \ldots, Y_{i-1}$. This more complex structure does not prevent us from representing the Metropolis algorithm as the construction of a sample $Z_0, Z_1, \ldots, Z_n$ distributed according to f if $Z_0 \sim f$. The sample is derived from the generation of two samples $Y_0 = Z_0$, $Y_1, \ldots, Y_n$ and $U_1, \ldots, U_n$, the second sample being ancillary for the estimation of $E^f\{h(Z)\}$.

The joint distribution of the two samples $(Y_0, \ldots, Y_n)$ and $(U_1, \ldots, U_n)$ is rather involved. The dependence between the Y_i's creates a dependence between the U_i's, conditionally on the Y_i's, and this complicates the derivation of the weights $\text{pr}(Z_i = Y_j)$. For example, if $n=4$, a typical weight is

$$\text{pr}(Z_3 = Y_2 | Y_0, \ldots, Y_4) \propto \{\rho_{01} g(Y_2|Y_1)\rho_{12} g(Y_3|Y_1) + (1-\rho_{01}) g(Y_2|Y_0)\rho_{02} g(Y_3|Y_1)\}(1-\rho_{23}) g(Y_4|Y_2),$$

where

$$\rho_{ij} = \text{pr}(Z_j = Y_j | Z_{j-1} = Y_i) = \frac{f(Y_j)/g(Y_j|Y_i)}{f(Y_i)/g(Y_i|Y_j)} \wedge 1.$$

However, we are still able to derive a Rao–Blackwellised version of the Metropolis estimator,

$$\hat{\tau}_7 = \frac{1}{n+1} \sum_{i=0}^{n} h(Z_i),$$

as shown by the following result. We first define the quantities

$$\rho_{ij}^* = \rho_{ij} g(Y_{j+1}|Y_j), \quad \rho_{*_{ij}} = (1-\rho_{ij}) g(Y_{j+1}|Y_i) \quad (i<j<n),$$

$$\xi_{jj} = 1, \quad \xi_{jt} = \prod_{l=j+1}^{t} \rho_{*_{jl}} \quad (j<t<n),$$

$$\delta_0 = 1, \quad \delta_j = \sum_{t=0}^{j-1} \delta_t \xi_{t(j-1)} \rho_{tj}^*, \quad \delta_n = \sum_{t=0}^{n-1} \delta_t \xi_{t(n-1)} \rho_{tn} \quad (j<n),$$

$$\omega_n^j = 1, \quad \omega_i^j = \rho_{ji}^* \omega_{i+1}^i + \rho_{*_{ji}} \omega_{i+1}^j \quad (0 \leqslant j < i < n).$$

Proposition 4·1. *The Rao–Blackwellised version of the general Metropolis algorithm estimator is*

$$\hat{\tau}_8 = \frac{\sum_{i=0}^{n} \varphi_i h(Y_i)}{\sum_{i=0}^{n-1} \delta_i \xi_{i(n-1)}},$$

where $\varphi_n = \delta_n$ *and, for* $i < n$,

$$\varphi_i = \delta_i \left\{ \sum_{j=i}^{n-1} \xi_{ij}\omega_{j+1}^i + \xi_{i(n-1)}(1-\rho_{in}) \right\}.$$

Proof. As in the independent case, $\hat{\tau}_7$ can be written

$$\frac{1}{n+1} \sum_{i=0}^{n} h(Y_i) \sum_{j=i}^{n} I(Z_j = Y_i).$$

The conditional expectation of the above indicator variables is then

$$\begin{aligned} \mathrm{pr}(Z_j = Y_i \mid Y_0, Y_1, \ldots, Y_n) &= E\{I(Z_i = Y_i) I(U_{i+1} > \rho_{i(i+1)}) \ldots I(U_j > \rho_{ij}) \mid Y_0, Y_1, \ldots, Y_n\} \\ &= E[\{I(Z_{i-1} = Y_{i-1}) I(U_i \leqslant \rho_{(i-1)i}) + I(Z_{i-1} = Y_{i-2}) I(U_i \leqslant \rho_{(i-2)i}) + \ldots\} \\ &\quad \ldots I(U_{i+1} > \rho_{i(i+1)}) \ldots I(U_j > \rho_{ij}) \mid Y_0, Y_1, \ldots, Y_n]. \end{aligned} \tag{4·1}$$

Therefore, conditionally on $Y_0, Y_1, \ldots, Y_n$, the event $\{Z_j = Y_i\}$ consists of the set of all the possible sequences of $(U_1, \ldots, U_i)$ leading to the acceptance of Y_i, of the sequences $(U_{i+1}, \ldots, U_j)$ corresponding to the rejection of $Y_{i+1}, \ldots, Y_j$ and of the sequences $(U_{j+1}, \ldots, U_n)$ constrained by $Z_j = Y_i$. That is,

$$\{Z_j = Y_i\} = \bigcup_{k=0}^{i-1} [B_k^{i-1}(U_1, \ldots, U_{i-1}) \cap \{U_i < \rho_{ki}, U_{i+1} > \rho_{i(i+1)}, \ldots, U_j > \rho_{ij}\}],$$

with, for $0 \leqslant k \leqslant t$,

$$B_k^t(U_1, \ldots, U_t) = \bigcup_{m=0}^{k-1} [B_m^{k-1}(U_1, \ldots, U_{k-1}) \cap \{U_k < \rho_{mk}, U_{k+1} > \rho_{k(k+1)}, \ldots, U_t > \rho_{kt}\}],$$

$$B_0^1 = \{U_1 > \rho_{01}\}, \quad B_1^1 = \{U_1 < \rho_{01}\}.$$

Taking the expectation in (4·1) with respect to the U_i's leads to a quantity proportional to

$$\delta_i \prod_{t=i+1}^{j} \rho_{*_{it}} \omega_{j+1}^i$$

since the probability that $U_i < \rho_{ji}$ is proportional to ρ_{ji}^*, while the probability that $U_i > \rho_{ji}$ is proportional to $\rho_{*_{ji}}$. The weighting factor in $\hat{\tau}_8$ is derived from the following expression:

$$1 = \sum_{i=0}^{n} \mathrm{pr}(Z_n = Y_i) \propto \left\{ \sum_{i=0}^{n-1} \delta_i \xi_{i(n-1)}(1-\rho_{in}) + \delta_n \right\} = \sum_{i=0}^{n-1} \delta_i \xi_{i(n-1)}. \qquad \square$$

Thus, despite the correlation between the Y_i's, the conditional expectation of the Metropolis procedure approximately has the same formal structure as in the independent case, and requires an amount of computation of the order of n^2.

5. Importance sampling estimators for general Metropolis samples

This section considers Rao–Blackwellisation of an importance sampling estimator based on the Metropolis algorithm. While it is possible to implement an importance sampling formula in this set-up, as the true marginal distribution of the Y_i's can be derived explicitly, its expression is involved. A simpler approach is to consider that the Y_i's $(i > 1)$ are actually generated from the conditional distributions $g(y | Z_{i-1})$ and then use the weights

$$\omega_i = f(Y_i)/g(Y_i | Z_{i-1}).$$

Although formally correct, this solution seems unsatisfactory since the resulting estimator

$$\hat{\tau}_9 = \frac{1}{n+1} \sum_{j=0}^{n} \omega_j h(Y_j)$$

still depends on the ancillary uniform random variables through the Z_i's.

The developments of the previous section can be exploited to build an improved version of $\hat{\tau}_9$, by integrating out the U_i's in $\hat{\tau}_9$. In fact, from the proof of Proposition 4·1,

$$E\{g(Y_i | Z_{i-1})^{-1} | Y_0, Y_1, \ldots, Y_n\} = \sum_{j=0}^{i-1} E\{g(Y_i | Y_j)^{-1} I(Z_{i-1} = Y_j) | Y_0, Y_1, \ldots, Y_n\},$$

$$E\{I(Z_{i-1} = Y_j) | Y_0, Y_1, \ldots, Y_n\} \propto \delta_j \xi_{j(i-1)} \omega_i^j \quad (0 < i \leqslant n, 0 \leqslant j < n).$$

Therefore, the Rao–Blackwellised version of $\hat{\tau}_9$ is

$$\hat{\tau}_{10} = \frac{1}{n+1} \left\{ h(Z_0) + \frac{f(Y_1)}{g(Y_1 | Y_0)} h(Y_1) + \frac{\sum_{i=2}^{n} \sum_{j=0}^{i-1} f(Y_i) \delta_j \xi_{j(i-2)} (1 - \rho_{j(i-1)}) \omega_i^j h(Y_i)}{\sum_{i=0}^{n-1} \delta_i \xi_{i(n-1)}} \right\}.$$

Note the strong similarity between $\hat{\tau}_8$ and $\hat{\tau}_{10}$. Both estimators take advantage of all the random variables which have been simulated, although in slightly different ways. We illustrate the performances of these procedures with an example.

Example 5·1. The target distribution is again a Student's t distribution with 3 degrees of freedom for which we wish to estimate the mean and the 5% tail probability. The estimation of these quantities is based on a Metropolis algorithm, where the candidate distribution is a Cauchy distribution centered at the previous random variable Z_{n-1}, with scale parameter σ^2. This is an inefficient sampling scheme, but the results are interesting.

Because of the somewhat involved form of the estimators, an extended simulation study was done. We performed 50 000 simulations, using both $\sigma = 0{\cdot}4$, giving an average acceptance rate of 0·327, and $\sigma = 3$, giving 0·749. We compare the usual Metropolis estimate $\hat{\tau}_7$ and its Rao–Blackwellised improvement with the Rao–Blackwellised importance sampling estimate $\hat{\tau}_{10}$. As expected, the performance of $\hat{\tau}_9$ is not as good as that of $\hat{\tau}_{10}$. The results presented in Tables 4 and 5 are rather surprising. For the high acceptance rate of 0·749, the mean squared error improvement upon $\hat{\tau}_7$ by Rao–Blackwellisation is only 0·4% for the mean estimation and 7% for the tail probability estimation. This improvement becomes more substantial for the lower acceptance rate, and the relative decrease in mean squared error reaches 25% in the best case. The importance sampling estimate is an improvement on the corresponding Metropolis and Rao–Blackwellised Metropolis estimates, the accompanying decrease in mean squared error improving as the sample size increases, to about 95% for $n = 100$. Therefore, in this case, importance sampling appears to be a significant improvement upon its Metropolis counterpart.

6. Conclusion

We have seen that the outputs of simulation schemes such as Accept-Reject and Metropolis algorithms can be improved by use of the entire set of simulated random

Table 4. *Estimation of the mean* (0) *of a Student's t distribution with* 3 *degrees of freedom based on a dependent Metropolis sample using a Cauchy distribution,* 50 000 *simulations*

(*a*) *Acceptance rate* 0·327

Sample size	Metrop. estimate $\hat{\tau}_7$	RB Metrop. $\hat{\tau}_8$	RB import. sampl. $\hat{\tau}_{10}$	MSE Metrop.	MSE decrease $\hat{\tau}_8$ over $\hat{\tau}_7$ (%)	MSE decrease $\hat{\tau}_{10}$ over $\hat{\tau}_7$ (%)
10	−0·0002	−0·0016	−0·0028	1·517	10·71	87·17
25	−0·0012	−0·0004	−0·0008	0·9841	8·78	92·02
50	0·0032	0·0021	−0·0004	0·6252	7·68	93·63
100	0·0012	0·0004	0·0005	0·3002	7·89	93·40

(*b*) *Acceptance rate* 0·749

Sample size	Metrop. estimate $\hat{\tau}_7$	RB Metrop. $\hat{\tau}_8$	RB import. sampl. $\hat{\tau}_{10}$	MSE Metrop.	MSE decrease $\hat{\tau}_8$ over $\hat{\tau}_7$ (%)	MSE decrease $\hat{\tau}_{10}$ over $\hat{\tau}_7$ (%)
10	0·0014	0·0001	0·0005	2·2849	0·1751	77·93
25	0·0003	0·0030	0·0008	1·7698	0·1526	85·98
50	0·0017	0·0017	0·0018	1·3066	0·1071	90·31
100	0·0014	0·0014	0·0003	0·8681	0·0691	92·85

RB, Rao–Blackwellised; MSE, mean squared error.

Table 5. *Estimation of a tail probability, chosen to be* 0·05, *of a Student's t distribution with* 3 *degrees of freedom based on a dependent Metropolis sample using a Cauchy distribution,* 50 000 *simulations*

(*a*) *Acceptance rate* 0·327

Sample size	Metrop. estimate $\hat{\tau}_7$	RB Metrop. $\hat{\tau}_8$	RB import. sampl. $\hat{\tau}_{10}$	MSE Metrop.	MSE decrease $\hat{\tau}_8$ over $\hat{\tau}_7$ (%)	MSE decrease $\hat{\tau}_{10}$ over $\hat{\tau}_7$ (%)
10	0·0500	0·0499	0·0499	0·0203	23·64	88·66
25	0·0504	0·0503	0·0501	0·0111	25·22	92·70
50	0·0505	0·0504	0·0500	0·0062	25·80	93·87
100	0·0501	0·0502	0·0500	0·0032	25·00	94·37

(*b*) *Acceptance rate* 0·749

Sample size	Metrop. estimate $\hat{\tau}_7$	RB Metrop. $\hat{\tau}_8$	RB import. sampl. $\hat{\tau}_{10}$	MSE Metrop.	MSE decrease $\hat{\tau}_8$ over $\hat{\tau}_7$ (%)	MSE decrease $\hat{\tau}_{10}$ over $\hat{\tau}_7$ (%)
10	0·0501	0·0503	0·0502	0·0301	0·9967	75·08
25	0·0505	0·0505	0·0500	0·0212	0·9433	84·43
50	0·0503	0·0503	0·0501	0·0140	0·7142	87·86
100	0·0500	0·0501	0·0498	0·0084	1·190	90·24

RB, Rao–Blackwellised; MSE, mean squared error.

variables, by application of a conditioning argument used in the Rao–Blackwell Theorem. This improvement relies on the recycling of the 'wasted' simulated random variables and hence represents a statistically better management of resources. Although the computational implementation may seem involved, the Rao–Blackwellised versions can be easily programmed via recursion relations with computing times that are quadratic in the sample size.

The fact that the Rao–Blackwell post-processing increases computation time should not be of concern. If we are interested in providing the best statistical solution, then post-processing to remove the uniform random variables is essential. No matter how large a sample is simulated, and no matter how good the simple estimators are, dependence on the uniform random variables is a statistical deficiency. Thus, comparisons of computation time are somewhat irrelevant. Moreover, such comparisons will reflect not only the increase in computation time due to the post-processing, but also the increase in computation time due to our programming skills, or lack thereof. However, we looked at a few timing runs and found that the post-processing will increase computation time from an almost negligible 1·15 times, in the Metropolis case, to as much as a 7 times increase in the Accept-Reject case. These runs were on a 486DX2 running at 66 megahertz, using the Gauss_{TM} programming language (Aptec Systems, 1992).

Our comparisons show that Rao–Blackwellisation is a viable method that may yield substantial improvement in mean squared error. This improvement is not associated with the property of unbiasedness, as Rao–Blackwellisation leaves the expectation untouched, so a biased estimator will remain biased. As pointed out by a referee, in practice, experimenters will often make many runs, using stopping criteria that may impart bias into all of the estimators. Such practice does not invalidate the post-processing improvement.

More important is the statistical comparison of Rao–Blackwellised importance sampling and of Rao–Blackwellised Metropolis, $\hat{\tau}_8$ and $\hat{\tau}_{10}$. Our single experiment shows an advantage for the importance sampling estimate. These results cast some doubt on the value of the Metropolis estimator, since it can be so dramatically improved upon. Further work is necessary to assess these improvements theoretically, but it seems that the use of importance sampling in practical Markov chain Monte Carlo environments should bring an even greater improvement than in the above simulations. This is because the chain does not usually start from the stationary distribution, and importance sampling automatically corrects for the simulation from an incorrect distribution. Since using importance sampling in a Metropolis environment does not require additional calculations other than those of the weight, we advise the use of this estimate, either as the only estimate of the quantity of interest, or at least as a control estimate to help to assess whether the Metropolis estimate has actually reached stationarity.

The 'corrected' importance sampling estimator $\hat{\tau}_6$ is the preferred choice of an importance sampling estimator. This is because without the correction achieved by dividing through by the sum of the weights the estimator will certainly be dominated by either Metropolis or Accept-Reject for functions h that are nearly constant. If h is exactly constant, the Metropolis and Accept-Reject estimators, and their Rao–Blackwellised versions, will have zero variance. This is not true for the importance sampling estimator $\hat{\tau}_3$. Thus, the requirement of unbiasedness in importance sampling estimation needs relaxing.

There exist set-ups where importance sampling estimates cannot be applied in practice, because the ratio $f(y)/g(y \mid z)$ has infinite variance under g, although the point estimates still converge. In these situations we are simulating a distribution f from a distribution g with lighter tails, as occurs in Gibbs sampling. While this does not formally prevent the corresponding Metropolis algorithm from converging, Mengersen & Tweedie (1996) have shown that convergence to the stationary distribution cannot be geometric in such cases. Hence, such schemes should only be used when better candidate densities are unavailable.

Acknowledgement

We are grateful to the Editor, Associate Editor and a referee for their careful reading and detailed comments which resulted in an improved version of the paper. This research was supported by the National Science Foundation and Unité de Recherche associée du Centre National de la Recherche Scientifique (CNRS). The work of C. P. Robert was done at Centre de Recherche en Economie et Statistique Théorique (CREST), Institut National de la Statistique et des Etudes (Insee), Paris.

References

Aptech Systems (1992). *Gauss System, Version 3·0.* Maple Valley, Washington: Aptech Systems Inc.

Devroye, L. (1985). *Non-uniform Random Variate Generation.* New York: Springer-Verlag.

Gelfand, A. & Smith, A. F. M. (1990). Sampling based approaches to calculating marginal densities. *J. Am. Statist. Assoc.* **85**, 398–409.

Geyer, C. J. (1992). Practical Monte Carlo Markov chain (with Discussion). *Statist. Sci.* **7**, 473–511.

Gilks, W. & Wild, P. (1992). Adaptive rejection sampling for Gibbs sampling. *Appl. Statist.* **41**, 337–48.

Hastings, W. K. (1970). Monte-Carlo sampling methods using Markov chains and their application. *Biometrika* **57**, 97–109.

Lehmann, E. L. (1983). *Theory of Point Estimation.* Pacific Grove: Wadsworth.

Liu, J. S., Wong, W. H. & Kong, A. (1994). Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes. *Biometrika* **81**, 27–40.

Liu, J. S., Wong, W. H. & Kong, A. (1995). Covariance structure and convergence rate of the Gibbs sampler with various scans. *J. R. Statist. Soc.* B **57**, 157–69.

Mengersen, K. L. & Tweedie, R. L. (1996). Exact convergence rates for the Hastings and Metropolis sampling algorithms. *Ann. Statist.* **24.** To appear.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. (1953). Equations of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–92.

Tierney, L. (1994). Markov chains for exploring posterior distributions (with Discussion). *Ann. Statist.* **22**, 1701–62.

[*Received August* 1994. *Revised July* 1995]