Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

VARIANCE REDUCTION OF MONTE CARLO AND RANDOMIZED QUASI-MONTE CARLO ESTIMATORS FOR STOCHASTIC VOLATILITY MODELS IN FINANCE

Hatem Ben Ameur

École des Hautes Études Commerciales
3000, chemin de la Côte-Ste-Catherine
Montréal, H3T 2A7, CANADA

Pierre L'Ecuyer and Christiane Lemieux

Département d'Informatique et de Recherche Opérationnelle
Université de Montréal, C.P. 6128, Succ. Centre-Ville
Montréal, H3C 3J7, CANADA

ABSTRACT

We illustrate by numerical examples how certain variance reduction methods dramatically improve the efficiency of Monte Carlo simulation for option pricing and other estimation problems in finance, in the context of a geometric Brownian motion model with stochastic volatitity. We consider lookback options and partial hedging strategies, with different models for the volatility process. For variance reduction, we use control variates, antithetic variates, conditional Monte Carlo, and randomized lattice rules coupled with a Brownian bridge technique that reduces the effective dimension of the problem. In some of our examples, the variance is reduced by a factor of more than 100 millions without increasing the work. The examples also illustrate how randomized quasi-Monte Carlo can be effective even if the problems considered involve a large number of dimensions.

1 INTRODUCTION

Monte Carlo (MC) simulation is used on a daily basis by banks and other financial institutions for pricing financial derivatives products. These simulations must provide precise estimates in a very short period of time. Efficiency improvement (e.g., via variance reduction) is therefore quite important in this context. In this paper, we give examples of how efficiency can be improved for pricing options under extended versions of the Black-Scholes model, with stochastic volatility. These models are believed to describe in a more realistic way the behavior of financial markets than the constant volatility model of Black and Scholes (1973).

We price two types of lookback options, for which the payoff depends on the maximal value of the primitive asset over a given time interval. We also estimate the initial gain made by hedging an European option only partially. Analytic formulas are available for these problems when the volatility is constant, but not when the volatility is stochastic. Simulation is an appropriate tool in this case. To reduce the variance, we use control variates, antithetic variates, conditional Monte Carlo, and randomized lattice rules coupled with techniques that reduce the effective dimension of the problem. In our numerical examples, we obtain variance reduction factors of over 2000 for the lookback options pricing, and of up to 100 millions (!) for the partial-hedging example. These variance reductions are obtained with no significant additional work; in fact in some cases (such as for conditional Monte Carlo and lattice rules) the total work (CPU time) is also reduced, so the *efficiency improvement* factors are even larger than the variance reduction factors that we just mentioned.

The use of simulation for pricing financial derivatives is surveyed by Boyle, Broadie, and Glasserman (1997). Among the recent papers where variance reduction methods are studied in that context, we cite Glasserman, Heidelberger, and Shahabuddin (1999), Lemieux and L'Ecuyer (1998) and Willard (1997). For general overviews of variance reduction, see, e.g., Bratley, Fox, and Schrage (1987), Fishman (1996) and L'Ecuyer (1994).

Quasi-Monte Carlo (QMC) methods have also been used successfully to evaluate financial products (e.g., Paskov and Traub 1995; Joy, Boyle, and Tan 1996; Caflish, Morokoff, and Owen 1997; Acworth, Broadie, and Glasserman 1997; Boyle, Broadie, and Glasserman 1997; Willard 1997; Tezuka 1998; Lemieux and L'Ecuyer 1998; Lemieux and L'Ecuyer 1999a). In particular, the empirical results obtained in Paskov and Traub (1995) for problems related to mortgage-backed securities have shown that QMC can work even for large-dimensional problems (up to 360 in their example), in contrast to what was commonly believed previously (Morokoff and Caflisch 1994; Morokoff and Caflish 1995). Their results led to the concept of *effective dimension*, which is roughly the number of "important" variables (or dimensions) of the problem. It is believed that many financial problems have a small associated effective dimension and this somehow explains the success of QMC methods for this type of application.

The QMC methods used in the present paper are *lattice rules* (see, e.g., Sloan and Joe 1994; Lemieux and L'Ecuyer 1999a; Lemieux and L'Ecuyer 1999b). In order to obtain unbiased mean and variance estimators, we randomize the rules via independent random shifts as in Cranley and Patterson (1976), and the method can then be viewed as a variance reduction technique (Lemieux and L'Ecuyer 1999b).

The paper is organized as follows. In Section 2, we describe the financial models and their assumptions. In the third section, we recall briefly what are lattice rules and how they can be used for simulation problems. In Section 4, we review the basic idea of the *Brownian bridge* method, which reduces the effective dimension of problems involving the simulation of Brownian motion paths. Finally, in Section 5 we introduce the two examples, lookback options and partial hedging, and we present numerical results giving the estimated variance reduction factors obtained with the variance reduction techniques.

2 FINANCE MODELS

2.1 Dynamic of the Primitive Asset

We assume a frictionless and arbitrage-free market with two traded assets in which trading takes place continuously over the period $[0, T]$. The first asset is a *bank account* $B(\cdot)$, with initial value 1, and growing with a constant interest rate r, i.e., $B(t) = e^{rt}$ for $0 \le t \le T$. The discount factor process is defined as $\gamma(t) = 1/B(t)$ for $0 \le t \le T$. We also assume the existence of a *risk-neutral* probability measure Q (as usual) for which the second asset $S(\cdot)$ (called the *primitive asset*) verifies the stochastic differential equation (SDE):

$$\begin{aligned} dS(t) &= rS(t)dt + \sigma(t)S(t)\cdot \\ &\qquad \left[\sqrt{1-\rho^2}dW_1(t) + \rho dW_2(t)\right], \end{aligned} \tag{1}$$

for $S(0) > 0$ and $0 \le t \le T$, where $(W_1(t), W_2(t))$ for $0 \le t \le T$ is a two-dimensional Brownian motion, $\rho \in [0, 1]$ is a constant, and $\sigma(\cdot)$ is the *volatility process*, to be discussed in a moment. All stochastic processes are assumed to verify enough assumptions to ensure the existence of a unique solution to (1) and to be adapted to the (Q-augmented) natural filtration of $(W_1(t), W_2(t))$ for $0 \le t \le T$, denoted by $\{\mathcal{F}_t, 0 \le t \le T\}$. The standard Black-Scholes (BS) model is obtained when $\sigma(t) = \sigma$ (a constant) and $\rho = 0$.

Formally, an *option* is an $\mathcal{F}_T$-measurable random variable $Y \ge 0$ with $E^Q(Y) < \infty$, where E^Q denotes the expectation under the probability measure Q. It is a privilege whose holder can exercise at the time horizon T when $Y > 0$ (i.e., when the contract expires in the money). The rational price of this product is the expected value of its (future) discounted payoff under Q:

$$v = E^Q\left[\gamma(T)Y\right]. \tag{2}$$

This expectation is denoted by $E\left[\gamma(T)Y\right]$ in the following. Details about arbitrage-free markets and risk-neutral evaluation of options can be found in Duffie (1996).

2.2 Dynamics for the Volatility

We consider three different models for the volatility process $\sigma(\cdot)$. For further details on stochastic volatility models in general, we refer the reader to Detemple and Osakwe (1997).

The first model is the *Geometric Brownian Motion Process* (GBMP), for which $\sigma(\cdot)$ follows the SDE

$$d\sigma(t) = \alpha\sigma(t)dt + \theta\sigma(t)dW_2(t), \qquad 0 \le t \le T,$$

where the *appreciation rate* α and the *volatility of the volatility* θ are constants. The ρ defined in (1) is the correlation between the innovations of the volatility and the price of the primitive asset. For a GBMP, $\sigma(T)/\sigma(0)$ is lognormal with parameters $(\alpha - \theta^2/2)T$ and $\theta\sqrt{T}$.

The second is the *Mean Reverting Proportional Process* (MRPP), for which

$$d\sigma(t) = \kappa\left[\bar{\sigma} - \sigma(t)\right]dt + \theta\sigma(t)dW_2(t), \qquad 0 \le t \le T,$$

where κ is the *reverting rate*, $\bar{\sigma}$ is the *long-term volatility*, and θ is the *volatility of the volatility*. To the best of our knowledge, the distribution of $\sigma(T)$ is unknown when the volatility process is a MRPP.

The third model we use is the *Square-Root Mean Reverting Process* (SRMRP) (Cox, Ingersoll, and Ross 1985) which evolves according to the SDE

$$d\sigma(t) = \kappa\left[\bar{\sigma} - \sigma(t)\right]dt + \theta\sqrt{\sigma(t)}dW_2(t), \qquad 0 \le t \le T$$

for $2\kappa\bar{\sigma} \ge \theta^2$. For a SRMRP, $\sigma(T)$ follows a non-central chi-square distribution with $4\kappa\bar{\sigma}/\theta^2$ degrees of freedom and parameter of non-centrality $4\kappa\sigma(0)e^{-\kappa T}/\left[\theta^2\left(1-e^{-\kappa T}\right)\right]$. The unconditional mean and variance of $\sigma(T)$ are

$$E\left[\sigma(T)\right] = \sigma(0)e^{-\kappa T} + \bar{\sigma}\left(1 - e^{-\kappa T}\right),$$

and

$$V\left[\sigma(T)\right] = \sigma(0)\frac{\theta^2}{\kappa}\left(e^{-\kappa T} - e^{-2\kappa T}\right) + \bar{\sigma}\frac{\theta^2}{2\kappa}\left(1 - e^{-\kappa T}\right)^2.$$

3 LATTICE RULES IN SIMULATION

Stochastic simulations normally draw their randomness from a so-called random number generator (RNG), which provides sequences of real numbers in the interval $[0, 1)$. These numbers are viewed as realizations of i.i.d. $U(0, 1)$ random variables, and transformed as needed to generate variates from other distributions and, ultimately, to compute the desired estimator. Thus, a simulation that requires s such uniforms can be viewed as computing a function f defined over the s-dimensional unit hypercube $[0, 1)^s$, at some point $\mathbf{x}$ determined by the RNG, to estimate the quantity

$$\mu = \int_{[0,1)^s} f(\mathbf{x}) d\mathbf{x}. \tag{3}$$

This can be repeated n times (by doing n independent simulation runs) and the estimator of μ becomes

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^{n} f(\mathbf{x}_i),$$

where $\mathbf{x}_i$ is the vector of s uniforms used for run i. This is the MC method. Variance reduction techniques essentially amount to changing the function f, replacing it by a "smoother" one.

The idea of QMC is to use in (3) a point set $P_n = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ that is more evenly distributed over $[0, 1)^s$ than typical random points. *Lattice rules* do this by taking P_n as the intersection of an integration lattice with the hypercube $[0, 1)^s$ (Sloan and Joe 1994; Lemieux and L'Ecuyer 1999b). A simple and convenient special case of this is to take P_n as the set of all successive overlapping s-tuples obtained from a given initial seed, for a multiplicative linear congruential generator (LCG) with modulus n and (maximal) period length $n - 1$, and then add the zero vector. With this approach, P_n is deterministic, and so are the estimator $\hat{\mu}_n$ and the integration error $|\hat{\mu}_n - \mu|$, which is usually hard to estimate.

A practical way of assessing the error is to randomize the lattice rule as suggested by Cranley and Patterson (1976): Generate a random vector $\mathbf{U}$ uniformly over $[0, 1)^s$, and shift P_n by adding $\mathbf{U}$ to each point, modulo 1. That is, each $\mathbf{x}_i$ is replaced by $\tilde{\mathbf{x}}_i = (\mathbf{x}_i + \mathbf{U}) \bmod 1$. The estimator $\tilde{\mu}_n$ thus obtained by replacing $\mathbf{x}_i$ by $\tilde{\mathbf{x}}_i$ in (3) is unbiased for μ, and its variance turns out to be equal to the sum of squares of the Fourier coefficients of f over the dual of the integration lattice (Lemieux and L'Ecuyer 1999b). By repeating this procedure m times, with m independent random shifts, one obtains m i.i.d. copies of $\tilde{\mu}_n$, whose sample variance provides an unbiased estimator of the variance of $\tilde{\mu}_n$. In our numerical experiments in Section 5, we compare this sample variance with the sample variance of the MC estimator with sample size mn (i.e., the same total number of simulation runs).

The variance expression, given as a sum over the dual lattice, motivates selection criteria for lattice rules (Hickernell et al. 1999; Lemieux and L'Ecuyer 1999a; Lemieux and L'Ecuyer 1999b). For our numerical experiments reported in Section 5, we use lattice rules that correspond to LCGs with prime modulus n and multiplier a. The multipliers a have been selected via the general figure of merit $M_{t,d}$ defined by Lemieux and L'Ecuyer (1999b). We use $(t, d) = (32, 3)$ in Section 5.1, and $(t, d) = (8, 3)$ in Section 5.2. This general figure of merit selects the lattice rule in terms of the quality of the projections of P_n over low-dimensional subspaces of $[0, 1)^t$, and over the subspaces determined by successive dimensions. A different class of lattice rules, defined over polynomial spaces, is discussed by L'Ecuyer and Lemieux (1999).

4 BROWNIAN BRIDGE TECHNIQUE

As we said earlier, QMC typically works better when the effective dimension of the problem is small, e.g., when the first few components of the vectors $\mathbf{x}_i$ account for most of the variability of the estimator. The rationale is that the QMC point sets P_n are usually constructed so that the projections $P_n(I)$ on subspaces indexed by the dimensions in $I \subseteq \{1, 2, \ldots, s\}$ are very well distributed when $|I|$ is small or contains indices that are close to each other (i.e., $I = \{i_1, \ldots, i_t\}$ where $1 \le i_1 < \ldots < i_t \le s$ and $i_t - i_1$ is small). Thus, if the function f in (3) can be rewritten so that the corresponding subsets of variables $\{x_{i_1}, \ldots, x_{i_t}\}$ explain most of the variability of f, then the error should be reduced because these important subsets will be integrated by well-distributed point sets $P_n(I)$. This argument does not hold for MC because no projection $P_n(I)$ is more regularly distributed than any other when the points of P_n are i.i.d. uniformly distributed over $[0, 1)^s$.

In certain situations, several ways of defining f are available which all give the same distribution for $f(\mathbf{U})$ when $\mathbf{U}$ is uniform over $[0, 1)^s$, but for which the importance of the first few variables differs significantly. One of these situations is the generation of a (discretized) Brownian motion $W(\cdot)$ over a time interval $[0, T]$. A standard way of generating $W(\cdot)$ is to discretize $[0, T]$ by defining $t_j = jT/d$ for $j = 0, \ldots, d$, for some constant d, and generate $W(t_j)$ for $j = 0, 1, \ldots, d$ in succession by exploiting the fact that the $W(t_{j+1}) - W(t_j)$ are i.i.d. normal. An alternative is the *Brownian bridge* technique, also called *Ciesielski discretization*, which exploits the fact that for a Brownian motion $W(\cdot)$, the distribution of $W(t)$ conditional on $(W(u) = a,\ W(v) = b)$ is known for $u < t < v$. Assuming that d is a power of 2, one can then generate the path $\{W(t_i),\ 1 \le i \le d\}$ in a non-sequential way as follows: First generate $W(t_d)$, then $W(t_{d/2})$, then $W(t_{d/4})$ and

$W(t_{3d/4})$, and so on. In this way the effective dimension can be reduced because the first few $U(0,1)$ variates explain most of the variability of the path. Using this technique to reduce the effective dimension in the context of QMC has first been suggested by Caflisch and Moskowitz (1995).

5 EXAMPLES AND NUMERICAL RESULTS

We now experiment with lattice rules combined with other variance reduction techniques, for two classes of problems. In our implementation, the processes $\{\sigma(t), 0 \le t \le T\}$ and $\{S(t), 0 \le t \le T\}$ have been discretized using the Euler scheme with $t_j = jT/d$, $j = 0, \ldots, d$. For all the experiments, we take $r = 0.05$ and $T = 0.5$.

In terms of computation times, using lattice rules (LR) is faster than plain MC (because the points $\mathbf{x}_i$ are faster to generate). Using the Brownian bridge (BB) method increases the work compared with the standard sequential scheme, but its combination with LR is approximately as fast as the standard scheme with MC (depending on which computer/compiler is used).

5.1 Lookback Option

The first type of lookback option that we consider gives the right to sell the primitive asset at the high:

$$Y_{\text{high}} = M^S(T) - S(T),$$

where $M^S(T) = \max_{0 \le t \le T} \{S(t)\}$. The second type is a *call option on the maximum*:

$$Y_{\max} = \left(M^S(T) - K\right)^+, \tag{4}$$

where K is a positive constant called the *strike price*. No analytic formula is available for these two types of lookback options when the volatility is stochastic. Under the BS model (with constant volatility) closed-form analytic formulas for the exact prices $E(Y_{\text{high}})$ and $E(Y_{\max})$ have been established by Goldman, Sosin, and Gatto (1979) and Conze and Viswanathan (1991), respectively.

Let μ_1 and μ_2 denote respectively $E(\gamma(T)Y_{\text{high}})$ and $E(\gamma(T)Y_{\max})$. Since each of these two products require the computation of $M^S(T)$, we need to generate a path for the 3-dimensional process $(\{S(t_i)\}, \{\sigma(t_i)\}, \{M^S(i)\},\ 1 \le i \le d)$ in each simulation, where

$$M^S(i) = \max_{t_{i-1} \le t \le t_i} S(t)$$

is generated under the conditional distribution given $(S(t_{i-1}), S(t_i), \sigma(t_{i-1}))$, following the technique of Beaglehole, Dybvig, and Zhou (1997). The dimension s is thus equal to $3d$, which is large since we need to take d large enough so that the discretization error is small.

For this problem, we reduce the variance via control variates and antithetic variates, in addition to the use of lattice rules. Our first control variate, Y_1, is taken as the price of the option when the volatility is fixed at $\sigma(0)$, i.e., under the BS model. This Y_1 should control the variability of the paths $\{S(t_i), 1 \le i \le d\}$ and $\{M^S(i), 1 \le i \le d\}$, given the trajectory of the volatility process. To control the variability of the volatility process, we take $Y_2 = \sigma(T)$ as a second control variate. We consider the GBMP and SRMRP models for $\sigma(\cdot)$ (we can compute $E(Y_2)$ for these models).

We also use *antithetic variates* (AV), but only for $S(\cdot)$ and $M^S(\cdot)$ (and not for $\sigma(\cdot)$, because the estimator is not monotone as a function of the uniforms used to generate $\sigma(\cdot)$; see Bratley, Fox, and Schrage 1987 for a discussion of why monotonicity is relevant). We must also be careful about how we assign the $U(0,1)$ variates to the three different paths. Since the BB technique is only used for $S(\cdot)$ and $\sigma(\cdot)$, we decided to use the following assignment: $M^S(i)$ uses U_i, $S(t_i)$ uses U_{d+2i-1} and $\sigma(t_i)$ uses U_{d+2i}. This should put more importance on the first $d + w$ variables, for some small $w > 0$.

In Tables 1, 2 and 3, the parameters for the GBMP are $S(0) = 100$, $\sigma(0) = 0.15$, $\alpha = 0.05$, $\theta = 0.08$. For the SRMRP in Tables 4 and 5, we take $\sigma(0) = 0.15$, $\bar{\sigma} = 0.15$, $\theta = 0.08$ and $\kappa = 1.5$. For the option on $Y_{\max}$, we use $K = 100$. CV1 stands for the estimator that uses the control variate Y_1 only, while CV12 is for the one that uses both control variables. The multiplier a used for the lattice rules, throughout Table 1 to 5 is $a = 178$.

The ratios (variance reduction factors) given in the following tables for MC are precise at least to their first

Table 1: Estimated variance reduction factors for Y_{high}, GBMP and $\rho = 0$. Estimated price = 7.60

	MC	LR	LR + BB
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	8.9	19
AV	5.6	13	34
CV1	190	280	1100
CV1 + AV	560	320	2000
CV12	330	280	1200
CV12 + AV	560	320	2000
$(m, n, d) = (100, 1021, 128)$			
naive	1.0	6.2	16
AV	5.6	8.9	27
CV1	190	540	1100
CV1 + AV	560	590	1800
CV12	330	530	1100
CV12 + AV	560	590	1800

Table 2: Estimated variance reduction factors for Y_{high}, GBMP and $\rho = 0.2$. Estimated price = 10.11

	MC	LR	LR + BB
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	7.9	16
AV	5.5	13	30
CV1	18	39	70
CV1 + AV	49	73	120
CV12	23	39	71
CV12 + AV	49	73	120

Table 3: Estimated variance reduction factors for Y_{max}, GBMP and $\rho = 0$. Estimated price = 10.07

	MC	LR	LR + BB
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	12	39
AV	6.7	15	65
CV1	290	590	2100
CV1 + AV	700	750	2400
CV12	450	590	2100
CV12 + AV	700	750	2400
$(m, n, d) = (100, 1021, 128)$			
naive	1.0	7.1	25
AV	6.6	9.6	52
CV1	280	550	1400
CV1 + AV	700	660	2300
CV12	440	560	1400
CV12 + AV	700	660	2300

2 digits, since the estimators for MC are based on mn replications. When LR is implied, the ratios are less accurate (say $\pm 10\%$) since the LR estimators are based on $m = 100$ replications only. However, these variations do not affect the general conclusions we draw from the results shown in the tables.

From these results, we see that the first control variable works very well in all cases, reducing the variance for MC by factors of at least 180 for GBMP (when $\rho = 0$) and 46 for SRMRP. The control variable Y_2 helps with the MC method (without AV, for GBMP) but is practically useless with LR. Also, LR alone is sometimes worse than MC with GBMP. This is not unexpected since the dimension of this problem is large (192 or 384). However, using BB makes LR better than MC in every case, by a factor of at least 3. The combination of CV1, AV, LR, and BB reduces the variance by a factor of nearly 2000 in Table 1, over 2000 in Table 3, and somewhat less in the other tables. We computed the estimated correlations between Y_1 and Y_{high} (or Y_{max}) for the two models and it turned out to be higher with GBMP than SRMRP. Also, the first control variable works better with Y_{max} than with Y_{high}: The payoff for Y_{high} depends on both $M^S(T)$ and $S(T)$ whereas for Y_{max}, it only depends on $M^S(T)$. Thus, switching from constant to stochastic volatility increases the variance of Y_{max} *less* than that of Y_{high}.

Table 4: Estimated variance reduction factors for Y_{high}, SRMRP and $\rho = 0$. Estimated price = 7.49

	MC	LR	LR + BB
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	8.8	18
AV	5.3	13	33
CV1	47	74	240
CV1 + AV	73	570	720
CV12	72	73	250
CV12 + AV	155	570	770
$(m, n, d) = (100, 1021, 128)$			
naive	1.0	6.5	15
AV	5.3	8.4	29
CV1	47	110	230
CV1 + AV	73	422	630
CV12	72	108	240
CV12 + AV	156	430	630

Table 5: Estimated variance reduction factors for Y_{max}, SRMRP and $\rho = 0$. Estimated price = 9.96

	MC	LR	LR + BB
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	11	36
AV	6.5	17	62
CV1	72	160	440
CV1 + AV	130	700	1100
CV12	100	160	440
CV12 + AV	240	690	1200
$(m, n, d) = (100, 1021, 128)$			
naive	1.0	6.5	23
AV	6.5	10	53
CV1	71	140	270
CV1 + AV	130	540	970
CV12	99	140	280
CV12 + AV	240	540	970

Table 2 reports the same experiment as Table 1, but using $\rho = 0.2$. This means that $S(\cdot)$ depends more on the Brownian motion $W_2(\cdot)$ that drives the volatility process, so that the correlation between the option's payoff and Y_1 decreases. This is confirmed by the smaller reduction factors given by CV1.

To give an idea of the computation time, it took about three minutes, on a Pentium II computer, to compute the LR+BB+CV12+AV estimator of Table 5, with $d = 64$, and the precision is around 0.5 cent at the 95% confidence level.

5.2 Partial Hedging

The basic idea used to evaluate an option in a frictionless, arbitrage-free and complete market is to find the initial cost of a *replication portfolio* that hedges the option. (Of course, in the stochastic volatility models considered here, a second traded asset would have to be introduced to complete the market.) The value of this portfolio matches with certainty the option's payoff at the exercise date. Alternatively, the hedger of an option may opt for *partial replication* (Ben Ameur, Breton, and L'Ecuyer 1999) since it costs less than a full replication, thus allowing an initial gain. However, this agent may fall in default, with a probability that depends on the partial hedging event. Sellers of options who do not perfectly hedge their risk are commonplace in real life.

Here we consider the (simple) problem of finding the initial gain when partially hedging an European option on a hedging event that depends on the final primitive asset price. More specifically, we want to estimate $E[X]$, where

$$X = \gamma(T)(S(T) - K)^+ \mathbf{1}_{\{S(T) \le b\}}, \tag{5}$$

$b > K$ is the parameter defining the hedging event (i.e., the hedger speculates that $S(T)$ will be smaller than b), and K is the strike price of the option. The naive estimator for $E[X]$ can be computed by generating the path of $(\{S(t_i)\}, \{\sigma(t_i)\}, 1 \le i \le d)$, so the dimension of this problem is $2d$.

For this problem, we use *conditional Monte Carlo* (CMC) to reduce the variance of the naive estimator. As explained by Hull and White (1987) and Willard (1997), the conditional expectation of X given the trajectory of the volatility process $\{\sigma(t_i),\ 1 \le i \le d\}$ can be computed analytically:

$$\begin{aligned} & E(X \mid \sigma(t_1), \ldots, \sigma(t_d)) \\ = \;& S(0)(\Phi(d_1) - \Phi(d_1')) - K e^{-rT}(\Phi(d_2) - \Phi(d_2'), \end{aligned}$$

where $d_1 = (\ln(S(0)/K) + (r + \bar{\sigma}/2)T)/\sqrt{\bar{\sigma} T}$, $d_1' = (\ln(S(0)/b) + (r + \bar{\sigma}/2)T)/\sqrt{\bar{\sigma} T}$, $d_2 = d_1 - \sqrt{\bar{\sigma} T}$, $d_2' = d_1' - \sqrt{\bar{\sigma} T}$, $\bar{\sigma} = (\sum_{i=1}^{d} \sigma^2(t_i))/d$ and $\Phi(\cdot)$ is the cumulative normal distribution. This follows from the fact that the distribution of $S(t_d)$ conditional on $\sigma(t_1), \ldots, \sigma(t_d)$ is known (see also Ben Ameur, Breton, and L'Ecuyer 1999 for details and other cases).

This CMC estimator has *provably* less variance than the naive MC one (this follows from standard variance decomposition; see Bratley, Fox, and Schrage 1987). It also reduces the work, because the process $S(\cdot)$ needs not be generated explicitly. Moreover, the dimension of the problem is reduced from $2d$ to d, which should help QMC methods (when combined with CMC). The number d of dimensions may still be too large though, and this is where the BB technique comes to the rescue.

In Table 6, we compare MC, LR, and LR + BB, using either the naive estimator, CMC, or CMC + AV. Different values of n and d are used. With the naive estimator, the BB technique is applied to the simulation of both Brownian motions $W_1(\cdot)$ and $W_2(\cdot)$, and the $U(0,1)$ variates are interleaved: $U_1, U_3, \ldots, U_{2d-1}$ go to $W_1(\cdot)$ while $U_2, U_4, \ldots, U_{2d}$ go to $W_2(\cdot)$. The volatility model is the MRPP. There is no analytical solution for $E[X]$ in this case. The parameters of the model are $S(0) = 100$, $\sigma(0) = 0.15$, $\bar{\sigma} = 0.15$, $\theta = 0.08$, $\kappa = 1.5$, $K = 100$ and $b = 120$. The multiplier a used for the lattice rules when $n = 251$ is $a = 46$, and $a = 325$ for $n = 1021$.

Table 6: Estimated variance reduction factors for a partially hedged portfolio. Estimated initial value = 3.96.

	MC	LR	LR + BB
$(m, n, d) = (100, 251, 64)$			
naive	1.0	1.3	35
CMC	5.5e4	2.8e6	4.1e6
CMC + AV	5.4e6	1.4e7	7.4e7
$(m, n, d) = (100, 1021, 64)$			
naive	1.0	1.6	39
CMC	5.6e5	6.1e6	1.0e7
CMC + AV	5.6e6	1.7e7	1.0e8
$(m, n, d) = (100, 251, 128)$			
naive	1.0	1.6	36
CMC	5.6e4	2.3e6	3.5e6
CMC + AV	5.7e6	1.2e7	7.3e7
$(m, n, d) = (100, 1021, 128)$			
naive	1.0	1.7	36
CMC	5.7e4	7.2e6	1.1e7
CMC + AV	5.8e6	1.5e7	1.1e8

In Table 6, we see that LR improves upon MC in all cases, even without BB. This differs from what was observed by Willard (1997), who used different QMC methods for a slightly different problem, with $d = 64$. BB brings a significant variance reduction compared with LR alone. CMC and its combination with AV bring spectacular improvements, by factors ranging from (roughly) 5 millions to 100 millions. In terms of computation time, it means that in only one second, an estimator with very high accuracy (3e-5 at the 95% level) can be obtained when using LR+BB+CMC+AV. Note that the improvement of LR upon MC is more important with the CMC estimators than with the naive method, whose respective dimensions are d and $2d$. For comparison, Willard (1997) observed a variance reduction by (roughly) a factor of 100 with CMC, and an additional factor of 10 when CMC was combined with QMC, for his problem.

ACKNOWLEDGMENTS

This work has been supported by NSERC-Canada grant No. ODGP0110050 and FCAR Grant No. 00ER3218 to the second author, and via an FCAR-Québec scholarship to the third author.

REFERENCES

Acworth, P., M. Broadie, and P. Glasserman. 1997. A comparison of some Monte Carlo and quasi-Monte Carlo techniques for option pricing. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, ed. P. Hellekalek and H. Niederreiter, number 127 in Lecture Notes in Statistics, 1–18. Springer-Verlag.

Beaglehole, D. R., P. H. Dybvig, and G. Zhou. 1997. Going to extremes: Correcting simulation bias in exotic option valuation. *Financial Analysts Journal*, 62–68.

Ben Ameur, H., M. Breton, and P. L'Ecuyer. 1999. Partial hedging for options based on extreme values and passage times. Technical Report G–99–15, GERAD, Ecole des Hautes Études Commerciales, Montréal.

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–654.

Boyle, P., M. Broadie, and P. Glasserman. 1997. Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21:1267–1321.

Bratley, P., B. L. Fox, and L. E. Schrage. 1987. *A Guide to Simulation*. Second ed. New York: Springer-Verlag.

Caflisch, R. E., and B. Moskowitz. 1995. Modified Monte Carlo methods using quasi-random sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, ed. H. Niederreiter and P. J.-S. Shiue, number 106 in Lecture Notes in Statistics, 1–16, New York. Springer-Verlag.

Caflish, R. E., W. Morokoff, and A. Owen. 1997. Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension. *The Journal of Computational Finance*, 1(1):27–46.

Conze, A., and Viswanathan. 1991. Path dependent options: The case of lookback options. *Journal of Finance*, 46:1893–1906.

Cox, J. C., J. E. Ingersoll, and S. A. Ross. 1985. A theory of the term structure of interest rates. *Econometrica*, 53:385–407.

Cranley, R., and T. N. L. Patterson. 1976. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914.

Detemple, J., and C. J. Osakwe. 1997. The valuation of volatility options. Working paper. Faculty of Management, McGill University, Montreal.

Duffie, D. 1996. *Dynamic Asset Pricing Theory*. second ed. Princeton University Press.

Fishman, G. S. 1996. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer Series in Operations Research, New York: Springer-Verlag.

Glasserman, P., P. Heidelberger, and P. Shahabuddin. 1999. Asymptotically optimal importance sampling and stratification for pricing path dependent options. *Journal of Mathematical Finance*, 9(2):117–152.

Goldman, M. B., H. B. Sosin, and M. A. Gatto. 1979. Path dependent options: "buy at the low, sell at the high". *Journal of Finance*, 34:1111–1127.

Hickernell, F. J., H. S. Hong, P. L'Ecuyer, and C. Lemieux. 1999. Extensible lattice sequences for quasi-monte carlo quadrature. submitted.

Hull, J., and A. White. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42:281–300.

Joy, C., P. P. Boyle, and K. S. Tan. 1996. Quasi-Monte Carlo methods in numerical finance. *Management Science*, 42:926–938.

L'Ecuyer, P. 1994. Efficiency improvement via variance reduction. In *Proceedings of the 1994 Winter Simulation Conference*, 122–132. IEEE Press.

L'Ecuyer, P., and C. Lemieux. 1999. Quasi-monte carlo via linear shift-register sequences. In *Proceedings of the 1999 Winter Simulation Conference*. IEEE Press. to appear.

Lemieux, C., and P. L'Ecuyer. 1998. Efficiency improvement by lattice rules for pricing asian options. In *Proceedings of the 1998 Winter Simulation Conference*, 579–586. IEEE Press.

Lemieux, C., and P. L'Ecuyer. 1999a. A comparison of monte carlo, lattice rules and other low-discrepancy point sets. In *Monte Carlo and Quasi-Monte Carlo Methods 1998*, ed. H. Niederreiter and J. Spanier, Lecture Notes in Computational Science and Engineering, New York. Springer-Verlag. to appear.

Lemieux, C., and P. L'Ecuyer. 1999b. Selection criteria for lattice rules and other low-discrepancy point sets. submitted.

Morokoff, W. J., and R. E. Caflisch. 1994. Quasi-random sequences and their discrepancies. *SIAM Journal on Scientific Computing*, 15:1251–1279.

Morokoff, W. J., and R. E. Caflish. 1995. Quasi-Monte Carlo integration. *Journal of Computational Physics*, 122:218–230.

Paskov, S., and J. Traub. 1995. Faster valuation of financial derivatives. *Journal of Portfolio Management*, 22:113–120.

Sloan, I. H., and S. Joe. 1994. *Lattice Methods for Multiple Integration*. Oxford: Clarendon Press.

Tezuka, S. 1998. Financial applications of Monte Carlo and quasi-Monte Carlo methods. In *Random and Quasi-Random Point Sets*, ed. P. Hellekalek and G. Larcher,

volume 138 of *Lecture Notes in Statistics*, 303–332. New York: Springer.

Willard, G. A. 1997. Calculating prices and sensitivities for path-dependent derivatives securities in multifactor models. *Journal of derivatives*, 5:45–61.

AUTHOR BIOGRAPHIES

HATEM BEN AMEUR is currently completing his Ph.D. in the "Service de l'enseignement des Méthodes Quantitatives de Gestion" at the "École des Hautes Études Commerciales" of Montreal. He is also assistant professor in "L'École des Sciences de Gestion" at the "Université du Québec à Montréal". He works on simulation procedures for option pricing.

PIERRE L'ECUYER is a professor in the "Département d'Informatique et de Recherche Opérationnelle", at the University of Montreal. He received a Ph.D. in operations research in 1983, from the University of Montréal. He obtained the *E. W. R. Steacie* grant from the Natural Sciences and Engineering Research Council of Canada for the period 1995–97. His main research interests are random number generation, efficiency improvement via variance reduction, sensitivity analysis and optimization of discrete-event stochastic systems, and discrete-event simulation in general. He is an Area Editor for the *ACM Transactions on Modeling and Computer Simulation*. More details at: `http://www.iro.umontreal.ca/~lecuyer`, where several of his recent research articles are available on-line.

CHRISTIANE LEMIEUX is currently a Ph.D. student in the "Département d'Informatique et de Recherche Opérationnelle", at the University of Montréal. She works on efficiency improvement by quasi-Monte Carlo methods. She completed her M.Sc. in mathematics (actuarial science) at the same university in 1996.