

Least-squares Importance Sampling for Monte Carlo security pricing

LUCA CAPRIOTTI*†

Global Modelling and Analytics Group, Investment Banking Division, Credit Suisse Group,
One Cabot Square, London E14 4QJ, UK

(Received 13 August 2006; in final form 22 October 2007)

We describe a simple Importance Sampling strategy for Monte Carlo simulations based on a least-squares optimization procedure. With several numerical examples, we show that such Least-squares Importance Sampling (LSIS) provides efficiency gains comparable to the state-of-the-art techniques, for problems that can be formulated in terms of the determination of the optimal mean of a multivariate Gaussian distribution. In addition, LSIS can be naturally applied to more general Importance Sampling densities and is particularly effective when the ability to adjust higher moments of the sampling distribution, or to deal with non-Gaussian or multi-modal densities, is critical to achieve variance reductions.

Keywords: Monte Carlo methods; Derivatives pricing; Financial derivatives; Financial engineering

1. Introduction

The impressive development of the securities markets has generated in the last few years a steady demand for more and more structured financial products. At the same time, the level of sophistication and complexity of the pricing models employed by investment firms has dramatically increased, in a continuous search for a possible edge against competitors. As a result, an increasing fraction of the models employed in practice is too complex to be treated by analytic or deterministic numerical methods (trees or partial differential equations), and Monte Carlo simulation becomes more often than ever the only computationally feasible means of pricing and hedging.

Although generally easy to implement, Monte Carlo simulations are infamous for being slow. In fact, being stochastic in nature, their outcome is always affected by a statistical error, that can generally be reduced to the desired level of accuracy by iterating the calculation for a long enough time. This comes with a high computational cost since such statistical uncertainties, all things being equal, are inversely proportional to the square root of the number of statistically independent samples. Hence,

in order to reduce the error by a factor of 10 one has to spend 100 times as much computer time. For this reason, to be used on a trading floor, Monte Carlo simulations often require to be run on large parallel computers with a high financial cost in terms of hardware, infrastructure, and software development.

Motivated by this very practical necessity, several approaches to speed up Monte Carlo calculations, such as antithetic variables, control variates, and Importance Sampling, have been proposed over the years (Glasserman 2004). These techniques aim to reduce the variance per Monte Carlo observation so that a given level of accuracy can be obtained with a smaller number of iterations. In general, this can be done by exploiting some information known *a priori* on the structure of the problem at hand, like a symmetry property of the Brownian paths (antithetic variables), the value of a closely related security (control variates), or the form of the statistical distribution of the random samples (Importance Sampling). Antithetic variables and control variates are the most commonly used variance reduction techniques, mainly because of the simplicity of their implementation, and the fact that they can be

*Email: luca.capriotti@credit-suisse.com

†Present Address: 11 Madison Avenue, New York, NY 10010, USA.

accommodated in an existing Monte Carlo calculator with a small effort. However, their effectiveness varies largely across applications, and is sometimes rather limited (Glasserman 2004).

On the other hand, Importance Sampling techniques, although potentially more powerful, have not been employed much in professional contexts until recently. This is mainly because such techniques generally involve a bigger implementation effort, and they are also less straightforward to include in a general Monte Carlo framework. Moreover, when used improperly, Importance Sampling can increase the variance of the Monte Carlo estimators, thus making its integration in an automated environment more problematic. Nonetheless, the potential efficiency gains at stake are so large that the interest in finding efficient Importance Sampling schemes is still very high.

The idea behind Importance Sampling is to reduce the statistical uncertainty of a Monte Carlo calculation by focusing on the most important sectors of the space from which the random samples are drawn. Such regions critically depend on both the random process simulated, and the structure of the security priced. For instance, for a deep out-of-the money call option (Hull 2002), the payoff sampled is zero for most of the iterations of a Monte Carlo simulation. Hence, simulating more samples with positive payoff reduces the variance. This can be done by changing the probability distribution from which the samples are drawn, and reweighing the payout function by the appropriate likelihood ratio (Radon-Nikodym derivative) in order to produce an unbiased result of the original problem (Glasserman 2004).

Most of the work in Importance Sampling methods for security pricing has been done in a Gaussian setting (Reider 1993, Boyle *et al.* 1997, Vazquez-Abad and Dufresne 1998, Glasserman *et al.* 1999a, b, Su and Fu 2000, 2002, Arouna 2003, 2004) such as that arising from the simulation of a diffusion process. In this framework, Importance Sampling is achieved by modifying the drift term of the simulated process in order to drive the Brownian paths towards the regions that are the most important for the evaluation of the security. For instance, for the call option above, this can be obtained by increasing the drift term up to a certain optimal level (Reider 1993, Boyle *et al.* 1997). The different approaches proposed in the literature essentially differ in the way in which such a change of drift is found, and can be roughly divided into two families depending on the strategy adopted. The first strategy, common to the so-called adaptive Monte Carlo methods (Vazquez-Abad and Dufresne 1998, Su and Fu 2000, 2002, Arouna 2003), aims to determine the optimal drift through stochastic optimization techniques that typically involve an iterative algorithm. On the other hand, the second strategy, proposed in a remarkable paper by Glasserman,

Heidelberger, and Shahabuddin (GHS)(Glasserman *et al.* 1999a), relies on a deterministic optimization procedure that can be applied for a specific class of payouts. This approach, although approximate, turns out to be very effective for several pricing problems, including the simulation of a single factor Heath-Jarrow-Morton model (Glasserman *et al.* 1999b and portfolio credit scenarios (Glasserman and Li 2005).

In this paper, we propose the Least-squares Importance Sampling (LSIS) as an effective and flexible variance reduction method for Monte Carlo security pricing. This approach, originally proposed in physics for the optimization of quantum mechanical wave functions of correlated electrons (Umrigar *et al.* 1998), is applied here in a financial setting. In LSIS the determination of the optimal drift—or, more generally, of the most important regions of the sample space—is formulated in terms of a least-squares minimization. This technique can be easily implemented and included in an existing Monte Carlo code, and simply relies on a standard least-square algorithm for which several optimized libraries are available. We will show that LSIS provides efficiency gains analogous to those of previously proposed methods for problems that can be formulated in terms of the determination of the optimal mean vector of a multivariate Gaussian density. In addition, LSIS can be naturally applied to more general Importance Sampling densities such as non-Gaussian or multi-modal distributions.

In the following section, we begin by reviewing the main ideas behind Importance Sampling in a generic Monte Carlo framework. Then in section 3 we specialize the discussion to the Gaussian setting, and we briefly review the recent adaptive strategies, and the GHS approach of Glasserman *et al.* 1999a. The rationale of LSIS is introduced in section 4 together with the essential implementation details, and in section 5 we present the results of our numerical experiments. Here we perform a systematic study of the variance reductions obtained by means of LSIS for a variety of test cases, including a comparison with recent Importance Sampling techniques. Finally, we draw our conclusions in section 6.

2. Importance sampling

Let us consider the general problem of estimating the expectation value of a scalar function, $G(Z)$, depending on a d -dimensional real random vector $Z=(Z_1, \dots, Z_d)$ with joint probability density $P(Z)$,

$$V = E_P[G(Z)] = \int_D dZ G(Z)P(Z), \quad (1)$$

where D is the domain of possible values of the state variables Z .[†] In a financial derivatives context, $G(Z)$

[†]In the present discussion we will treat the Z_i 's as continuous variables, however all the results also apply if some or all of them can assume only a discrete set of values. For any such variable, the symbol $\int dZ_i$ is to be interpreted as a sum over the possible outcomes. Here and in the following, with a slight abuse of notation, we indicate by Z both the random variables and the dummy variables of integration.

would typically represent the discounted payout of a certain security, and $P(Z)$ would be the risk-neutral probability density of an arbitrage free market (Harrison and Kreps 1979, Baxter and Rennie 1996, Musiela and Rutkowski 2002). For instance, for the familiar call option in the Black-Scholes framework (Hull 2002) one has $d = 1$, $P(Z) = (2\pi)^{-1/2} \exp(-Z^2/2)$ and

$$G(Z) = e^{-rT} \left(X_0 \exp \left[\left(r - \frac{\sigma^2}{2} \right) T + \sigma Z \right] - K \right)^+, \quad (2)$$

where r is the risk-free interest rate, σ is the volatility, X_0 and K are, respectively, the spot and strike price, and T the maturity of the option.

Whenever the dimension d of the state variable Z is large (say $d \geq 4$) standard numerical quadrature approaches become highly inefficient, and Monte Carlo methods are the only feasible route for estimating expectation values of the form (1). To do so, one interprets Equation (1) as a weighted average of the payout function $G(Z)$ over the possible configurations Z with weights given by the probability density $P(Z)$. This immediately leads to the simplest (and crudest) Monte Carlo estimator which is obtained by averaging the payout function over a sample of N_p independent values of the random variable Z generated according to the probability density $P(Z)$,

$$V \simeq \bar{V} = \frac{1}{N_p} \sum_{i=1}^{N_p} G(Z_i), \quad Z_i \sim P(Z). \quad (3)$$

In particular, the central limit theorem (Kallenberg 1997) ensures that, for big enough samples, the values of the estimator \bar{V} are normally distributed around the true value, and converge for $N_p \rightarrow \infty$ towards V , namely

$$V \simeq \frac{1}{N_p} \sum_{i=1}^{N_p} G(Z_i) \pm \frac{\Sigma}{\sqrt{N_p}}, \quad (4)$$

where $\Sigma^2 = E_P[G(x)^2] - E_P[G(x)]^2$ is the variance of the estimator and can be similarly approximated by

$$\Sigma^2 \simeq \frac{1}{N_p} \sum_{i=1}^{N_p} (G(Z_i) - \bar{V})^2. \quad (5)$$

Although equation (4) ensures the convergence of the Monte Carlo estimator to the expectation value (1), its practical utility depends on the magnitude of the variance, Σ^2 . Indeed, the square root convergence in (4) implies that the number of replications N_p that are (asymptotically) necessary to achieve a given level of accuracy is proportional to the variance of the estimator.† Roughly speaking, such a quantity is relatively small whenever the function $G(Z)$ is approximately constant over the region of values of Z that is

represented the most among the random samples, i.e. the region that contains most of the probability mass of $P(Z)$. This is generally not the case for most of the pricing problems encountered in practice, and the calculation of accurate estimates of the expectation value (1) may require large sample sizes N_p , thus becoming computationally demanding.

However, the choice of extracting the random variable Z according to the probability density $P(Z)$, although natural, is by no means the only possible one. Indeed, the Monte Carlo integration can be performed by sampling an arbitrary probability density $\tilde{P}(Z)$ provided that the integral is suitably reweighed. In fact, using the identity

$$\int_D dZ G(Z) P(Z) = \int_D dZ \frac{G(Z) P(Z)}{\tilde{P}(Z)} \tilde{P}(Z), \quad (6)$$

an alternative estimator of the expectation value (1) is readily found as

$$V \simeq \tilde{V} = \frac{1}{N_p} \sum_{i=1}^{N_p} W(Z_i) G(Z_i) \quad Z_i \sim \tilde{P}(Z), \quad (7)$$

with the weight function given by $W(Z) = P(Z)/\tilde{P}(Z)$. The variance of the new Monte Carlo estimator reads

$$\tilde{\Sigma}^2 = \int_D dZ (W(Z)G(Z) - V)^2 \tilde{P}(Z), \quad (8)$$

and critically depends on the choice of the sampling probability density $\tilde{P}(Z)$. For non-negative functions $G(Z)$, the optimal choice of $\tilde{P}(Z)$ is the one for which $\tilde{\Sigma}$ vanishes, namely:

$$P_{\text{opt}}(Z) = \frac{1}{V} G(Z) P(Z). \quad (9)$$

In fact, the Monte Carlo estimator corresponding to such an *optimal sampling density* reads

$$\tilde{V} \simeq \frac{1}{N_p} \sum_{i=1}^{N_p} W(Z_i) G(Z_i) = \frac{1}{N_p} \sum_{i=1}^{N_p} V, \quad (10)$$

leading to a constant value V on each Monte Carlo replication, and resulting therefore in zero variance.‡ Unfortunately, such a choice is not really viable as the normalization constant, V , is the expectation value (1) we want to calculate in the first place.

3. Gaussian framework

The Importance Sampling approaches proposed in the literature usually apply or are generally formulated within a Gaussian setting, i.e. in a context where the density $P(Z)$ defining the expectation value (1) is a d -dimensional standard normal distribution. For example, this is the

†In particular, the Monte Carlo integration becomes unfeasible if the variance of the estimator diverges, giving rise to the so-called *sign-problem* instability. Although this problem is the crux of Monte Carlo simulations in several branches of the physical sciences [see, e.g., Sorella, S. and Capriotti, L., *Phys. Rev. B*, 2000, 61, 2599], this issue does not usually affect financial contexts.

‡It is possible to show (Press *et al.* 2002) that, when $G(Z)$ does not have a definite sign, the optimal sampling density has the similar form $P_{\text{opt}} = |G(Z)|P(Z)/V$, although in this case the optimal variance is not zero.

case for a Monte Carlo simulation of a vector diffusive process of the form (Glasserman 2004)

$$dX(t) = \mu(X(t), t) dt + \sigma(X(t), t) dW_t, \tag{11}$$

e.g. for the calculation of the price of an option depending on the path followed by $X(t)$ within a certain time interval $[0, T]$. Here the process $X(t)$ and the drift $\mu(X, t)$ are both L -dimensional real vectors, W_t is an N -dimensional standard Brownian motion, and the volatility, $\sigma(X, t)$, is an $L \times N$ real matrix.

Continuous-time processes of the form (11) are typically simulated by sampling $X(t)$ on a discrete grid of points, $0 = t_0 < t_1 < \dots < t_M = T$, by means, for instance, of a Euler scheme†

$$X_{i+1} = X_i + \mu(X_i, t) \Delta t_i + \sigma(X_i, t) \sqrt{\Delta t_i} Z_{i+1}, \tag{12}$$

where $X_i = X(t_i)$, $\Delta t_i = t_{i+1} - t_i$, and Z_{i+1} is an N -dimensional vector of independent standard normal variates. In this representation, each discretized path for the vector process $X(t)$ can be put into a one-to-one correspondence with a set of $d = N \times M$ independent standard normal variables Z . Hence, the original problem of evaluating the expectation value of a functional of the realized path of the process $X(t)$ can be formulated as in (1). More precisely, $G(Z)$ is given by the discretized payout functional, and the probability density is given by a d -dimensional standard normal distribution

$$P(Z) = (2\pi)^{-d/2} e^{-Z^2/2}, \tag{13}$$

where $Z^2 = Z \cdot Z$.

As a prototypical example of an exotic option pricing problem, treated as a test case for the most recently proposed Importance Sampling strategies (Glasserman *et al.* 1999a, Su and Fu 2000, 2002, Arouna 2003), in the following we will consider different arithmetic Asian-style options under a Black-Scholes log-normal model (Baxter and Rennie 1996). In this case, the underlying asset can be simulated on the time grid relevant for the payout function by means of the exact recursion

$$X_{i+1} = X_i \exp \left[(r - \sigma^2/2) \Delta t_i + \sigma \sqrt{\Delta t_i} Z_{i+1} \right], \tag{14}$$

where r is the risk-free interest rate, σ is the constant $N \times L$ real volatility matrix, and σ^2 is an L -dimensional vector of components $\sigma_k^2 = \sum_{j=1}^N \sigma_{kj}$. For an arithmetic Asian-style claim, the discounted payout function is of the form

$$G(Z) = e^{-rT} \Phi(\bar{X}), \tag{15}$$

where $\bar{X} = (1/M) \sum_{i=1}^M X_i$, and $\Phi(\bar{X})$ is some function of L variables. Clearly, European-style options are recovered in the special case $M = 1$.

3.1. Gaussian trial distributions

The simplest possible strategy for Importance Sampling in a Gaussian framework is to try to guide the sampled

paths towards the most important regions of the configuration space (i.e. where the contribution of the integrand is the largest) by means of a change of the drift terms of the process (11) or (12). The corresponding trial probability density reads

$$\tilde{P}_{\tilde{\mu}}(Z) = (2\pi)^{-d/2} e^{-(Z-\tilde{\mu})^2/2}, \tag{16}$$

where $\tilde{\mu}$ is a d -dimensional vector, and the weight function, as also expected from the Girsanov theorem (Musielka and Rutkowski 2002), is

$$W_{\tilde{\mu}}(Z) = \exp[-\tilde{\mu} \cdot Z + \tilde{\mu}^2/2]. \tag{17}$$

A variety of approaches for the determination of the drift vector $\tilde{\mu}$ minimizing the variance of the estimator (8) has recently been proposed in the literature (Vazquez-Abad and Dufresne 1998, Glasserman *et al.* 1999a, b, Su and Fu 2000, 2002, Arouna 2003). These can be roughly classified into two families depending on the strategy adopted. The first strategy, common to the so-called adaptive Monte Carlo methods, aims to determine the optimal drift vector through a stochastic minimization of the variance. Such minimization differs in detail in the various methods but always involves an iterative procedure, to be performed in a preliminary Monte Carlo simulation.

In particular, Su and Fu (2000, 2002), building upon previous work of Vazquez-Abad and Dufresne (1998), used a gradient-based stochastic approximation, dubbed infinitesimal perturbation analysis, in order to estimate the optimal *uniform shift* of the drift for the diffusion (12), minimizing the variance of the estimator (8). In the notation of this section, this translates into working with a trial density of the form (16) where the drift vector $\tilde{\mu}$ has components all equal to a single optimization parameter. The improvement of this method with respect to that of Vazquez-Abad and Dufresne (1998) is that the minimization is carried out under the original probability measure (as we also do for LSIS), while in the latter the minimization was formulated under the trial probability measure. As a result, the stochastic minimization applies also for non-differentiable payout, thus making the approach more general. The application of this technique to partial average Asian options in a Black-Scholes market, and to caplets under the Cox-Ingersoll-Ross model, provides significant variance reductions (Su and Fu 2000, 2002).

Along similar ideas, Arouna (2003) has recently proposed a different stochastic optimization method for the determination of the optimal sampling density (16). Here, in contrast to the previous approach, all the components of the drift vector are independently optimized. The method relies on a truncated version of the Robbins-Monro algorithm that is shown to converge asymptotically to the optimal drift, and to provide an effective variance reduction in a variety of cases. However, as remarked by the same author, a critical aspect of the practical implementation of the

†The use of other discretization schemes does not alter the present discussion.

Robbins-Monro algorithm is that it depends on the size of the iterative step. Hence, particular care needs to be taken in order for the algorithm to be efficient.†

On the other hand, the second strategy, proposed by Glasserman, Heidelberger, and Shahabuddin (GHS) (Glasserman *et al.* 1999a), relies on a saddle point approximation to minimize the variance of the estimator (8), or, equivalently, of its second moment (in the original measure)

$$m_2(\tilde{\mu}) = \int_D dZ W_{\tilde{\mu}}(Z) G(Z)^2 P(Z). \quad (18)$$

In fact, if the payout function $G(Z)$ is positive definite, by defining $F(Z) = \log G(Z)$ one can approximate equation (18) with the zero-order saddle point expansion

$$(2\pi)^{-d/2} \int_D dZ \exp[2F(Z) - \tilde{\mu} \cdot Z + \tilde{\mu}^2/2 - Z^2/2] \\ \simeq C \exp \left[\max_{Z \in D} (2F(Z) - \tilde{\mu} \cdot Z + \tilde{\mu}^2/2 - Z^2/2) \right],$$

where C is a constant. As a result, within this approximation, the problem of determining the optimal change of drift boils down to finding the vector μ such that

$$\max_{Z \in D} (2F(Z) - \tilde{\mu} \cdot Z + \tilde{\mu}^2/2 - Z^2/2) \quad (19)$$

is minimum. It is easy to show that this is obtained by choosing $\tilde{\mu}^* = Z^*$ where Z^* is the point that solves the optimization problem

$$\max(F(Z) - Z^2/2), \quad (20)$$

or, equivalently, for which the payout times the original density, $G(Z) P(Z)$, is maximum, i.e. Z^* corresponds to the maximum of the optimal sampling density (equation (9)). The simplest interpretation of the saddle point approach is therefore that it approximates the zero variance density by means of a normal distribution with the same mode and unit (co)variance.

The saddle point approach can be expected to be particularly effective in reducing the variance of the Monte Carlo estimator whenever the logpayout function $F(Z)$ is close to linear in the portion of the configuration space where most of the probability mass of $P(Z)$ lays. Moreover, GHS have also shown that stratifying (Glasserman 2004) the d -dimensional random vector Z along the optimal μ provides a spectacular variance reduction in certain cases (Glasserman *et al.* 1999a, b). However, whenever the optimal sampling probability density (9) cannot be accurately represented by a single Gaussian with the same mode and unit (co)variance, the saddle point approximation is less beneficial. In particular, as will also be illustrated in section 5, this approach turns out to be less effective whenever the structure of the payout function $G(Z)$ is such that the optimal sampling density (9) has a width

which is very different from that of the original distribution, or is multi-modal.

In the following section we describe an alternative least-squares strategy that is straightforward to implement and flexible enough to be applied in a generic Monte Carlo setting. Indeed, Least-squares Importance Sampling (LSIS) is not limited to the determination of the optimal change of drift in a Gaussian model. Instead, it can be applied to any Monte Carlo simulation provided that a reasonable guess of the optimal sampling density is available. For this reason, in the next section we will momentarily leave the Gaussian framework, and we will describe the rationale of LSIS in a more general setting.

4. Least-squares Importance Sampling

A practical approach to the search for an effective Importance Sampling density can be formulated in terms of a nonlinear optimization problem. For this purpose, let us consider a family of trial probability densities, $\tilde{P}_\theta(Z)$, depending on a set of N_θ real parameters $\theta = (\theta_1, \theta_2, \dots, \theta_{N_\theta})$. The variance of the estimator corresponding to $\tilde{P}_\theta(Z)$ (equation (8)) can be written in terms of the *original* probability density $P(Z)$ as

$$\tilde{\Sigma}_\theta^2 = E_P[W_\theta(Z)G^2(Z)] - E_P[G(Z)]^2, \quad (21)$$

with $W_\theta(Z) = P(Z)/\tilde{P}_\theta(Z)$. Hence, the optimal Importance Sampling density within the family $\tilde{P}_\theta(Z)$ is that for which the latter quantity, or, equivalently, the second moment (18) or

$$E_P[W_\theta(Z)G^2(Z)] \quad (22)$$

is minimum. The crucial observation is that the Monte Carlo estimator of this quantity,

$$m_2(\theta) \simeq \frac{1}{N'_p} \sum_{i=1}^{N'_p} (W_\theta(Z_i)^{1/2} G(Z_i))^2, \quad Z_i \sim P(Z), \quad (23)$$

can be interpreted as a nonlinear least-squares fit of a set of N'_p data points (x_i, y_i) with a function $y = f_\theta(x)$ parameterized by θ , with the correspondence $y_i \rightarrow 0$, $x_i \rightarrow Z_i$, and $f_\theta(x) \rightarrow W_\theta(Z)^{1/2} G(Z)$. The latter is a standard problem of statistical analysis that can be tackled with a variety of robust and easily accessible numerical algorithms, such as the so-called Levenberg-Marquardt method (Press *et al.* 2002).

Alternatively, to improve the numerical stability of the least-squares procedure, it is convenient in some situations to minimize, instead of (22), the pseudo-variance

$$S_2(\theta) = E_P[(W_\theta(Z)^{1/2} G(Z) - V_T)^2] \\ \simeq \frac{1}{N'_p} \sum_{i=1}^{N'_p} (W_\theta(Z_i)^{1/2} G(Z_i) - V_T)^2, \quad (24)$$

†B. Arouna, Private communication.

where the constant V_T is a guess of the option value. Indeed, the minimization of (24) is equivalent to that of the real variance of the estimator (21) as

$$S_2(\theta) = \tilde{\Sigma}_\theta^2 + (E_P[G(Z)] - V_T)^2. \tag{25}$$

The algorithm for the determination of the optimal sampling density within a certain trial family can therefore be summarized as it follows:

- generate a suitable number N'_p of replications of the state variables Z according to the *original* probability density $P(Z)$;
- choose a trial probability density $\tilde{P}_\theta(Z)$, and an initial value of the vector of parameters θ ;
- set $x_i \rightarrow Z_i$, $f_\theta(x) \rightarrow W_\theta(Z)^{1/2} G(Z)$ and $y_i \rightarrow 0$ (respectively $y_i \rightarrow V_T$) and call a least-squares fitter, say $\text{LSQ}[x, y, f_\theta(X), \theta]$, providing the optimal $\theta = \theta^*$ by minimizing the second moment of the estimator $m_2(\theta)$ (equation (23)) (respectively $S_2(\theta)$ (equation (24))).

Once the optimal parameters θ^* have been determined through the least-squares algorithm, one can perform an ordinary Monte Carlo simulation by sampling the probability density $\tilde{P}_{\theta^*}(Z)$, and calculating expectation values according to equation (7).

As will be illustrated with the numerical examples of section 5, this procedure does not add a significant overhead to the simulation, because just a relatively small number of replications $N'_p \ll N_p$ is usually required to determine the optimal parameter θ^* . This is due to the fact that the configurations over which the optimization is performed are fixed. As a result of this form of *correlated sampling* (Umrigar *et al.* 1988), the difference in the $m_2(\theta)$'s for two sets of values of the parameters being optimized is much more accurately determined than the

values of the $m_2(\theta)$'s themselves. This rather surprising feature is rooted in the fact that the minimization of equation (23) as a means to optimize the trial density, $\tilde{P}_\theta(Z)$, can be justified in terms of a genuine maximum likelihood criteria (Bressanini *et al.* 2002), and it is therefore independent on how accurately $m_2(\theta)$ approximates the quantity (22). In the stochastic optimization literature, replacing the minimization of the expectation value (22) with its deterministic counterpart is also known as sample path optimization or sample average approximation (Shapiro 1996, Kleywegt *et al.* 2001).

Clearly, the fact that only a limited number of Monte Carlo samples is required for the optimization of the trial density limits the overhead introduced by the LSIS algorithm. This makes LSIS a practical strategy for variance reduction.

In the following we will demonstrate the effectiveness of this simple strategy by applying it to a variety of test cases, and we will compare the present approach with those most recently proposed in the literature (Glasserman *et al.* 1999a, Su and Fu 2000, 2002, Arouna 2003).

5. Numerical examples

5.1. European options

We will start by considering European-style options. These are the simplest possible examples of financial relevance that allow one to illustrate the LSIS strategy. In particular, we will consider first standard call and put options, written on a single underlying asset, $X(t)$, following the log-normal process (14). Here $N=M=L=1$ and $\Delta t_0=T$. The payout function reads as in equation (15) with $\bar{X} \rightarrow X_1 \equiv X(T)$, $\Phi(X) = (X - K)^+$ and $\Phi(X) = (K - X)^+$ for the call and the put, respectively, and K is the strike price.

In these cases, the probability density $P(Z)$ is a univariate Gaussian distribution, and the option value (1) involves the integration of a function of a single variable. As discussed in section 2, Importance Sampling techniques seek a sampling probability density $\tilde{P}_\theta(Z)$ as close as possible to the optimal one (equation (9)) (see figure 1). The simplest choice for $\tilde{P}_\theta(Z)$, in this setting, is a Gaussian distribution of the form (16) (with $d=1$), so that the only parameter θ to optimize is the drift $\tilde{\mu}$.

The LSIS algorithm described in section 4 is particularly straightforward to implement, and involves a pre-simulation stage in which a set of N'_p standard normal random numbers are generated and stored, and a least-squares optimization routine is invoked. As anticipated, we found that the least-squares fitter was able to determine successfully the optimal $\tilde{\mu}$ with as little as $N'_p \simeq 50$ Monte Carlo replications, thus making the approach useful in practice.

In tables 1 and 2 we compare the results obtained with LSIS with those obtained by means of the Robbins–Monro (RM) adaptive Monte Carlo (as quoted by Arouna (2003)), and the saddle point approach of GHS




Figure 1. Sampling probability density functions for a European call option with $T=1$, $r=0.05$, $\sigma=0.3$, $X_0=K=50$ as obtained with LSIS (optimizing just the drift, $\text{LSIS}(\tilde{\mu})$, and both the drift and the volatility, $\text{LSIS}(\tilde{\mu}, \tilde{\sigma})$) and the saddle point approximation (GHS). On this scale the results for $\text{LSIS}(\tilde{\mu})$ and GHS are indistinguishable. The original (13) and the optimal (9) sampling densities are also shown for comparison.

(Glasserman *et al.* 1999a). Here, as an indicator of the efficiency gains introduced by the different strategies of Importance Sampling, we have defined the variance ratio as

$$VR = \left(\frac{\sigma(\text{Crude MC})}{\sigma(\text{IS})} \right)^2, \quad (26)$$

where the numerator and denominator are the statistical errors (for the same number of Monte Carlo paths) of the Crude and the Importance Sampling estimators, respectively. In particular, for this and the following examples, we have estimated the uncertainties on the statistical errors and on the VR's by performing a standard error analysis of the outcome of several Monte Carlo simulations for different random number seeds.

We found that the different methods produce a significant and comparable variance reduction. In particular, intuitively, the change of drift is more effective for low volatility, and deep in and out of the money options. In this case, LSIS and GHS optimized trial densities $\tilde{P}_{\tilde{\mu}}(Z)$ are very similar, as shown in figure 1. This could be expected as, in this case, the optimal Importance Sampling density (9) can be effectively approximated by a Gaussian with the same mode and unit variance, so that the GHS approach produces accurate results.

However, the LSIS method is not limited to Importance Sampling strategies based on a pure change

of drift, and one can easily introduce additional optimization parameters in the trial density. For instance, in this example it makes sense to introduce the sampling volatility, $\tilde{\sigma}$,

$$\tilde{P}_{\tilde{\mu}, \tilde{\sigma}}(Z) = (2\pi\tilde{\sigma}^2)^{-1/2} e^{-(Z-\tilde{\mu})^2/2\tilde{\sigma}^2}. \quad (27)$$

As illustrated in figure 1, by adjusting both $\tilde{\mu}$ and $\tilde{\sigma}$, one obtains a trial density closer to the optimal one. This corresponds to an additional variance reduction of up to over one order of magnitude, as shown in tables 1 and 2.

Another simple example in which the ability to adjust the width of the sampling density proves effective is the Butterfly spread option,

$$\Phi(X) = (X - K_1)^+ - 2(X - K_2)^+ + (X - K_3)^+, \quad (28)$$

with $K_1 < K_2 < K_3$. In this case, a pure change of drift can only force the expected end point of the sampled paths $X(T)$ to fall in the money, $K_1 < X(T) < K_3$. However, for values of the volatility or the maturity large enough, a significant fraction of the sampled paths will still result in a zero payout. Hence, a pure change of drift is not very effective in reducing the variance. This is illustrated in table 3. Instead, by quenching the sampled volatility, one can increase the number of paths falling in the money thus reducing the variance. For instance, in the example of table 3, for $S_0 = 50$ the saddle point change of drift is very small, and does not alter the variance. Instead, by

Table 1. Comparison between LSIS, the adaptive Robbins–Monro (RM) algorithm (as quoted by Arouna (2003)), and the saddle point approach of Glasserman *et al.* (1999a) (GSH): price of a European call option on a log-normal asset (14) for different values of the volatility σ , and of the strike price K . The variance reduction (VR) is defined in Equation (26). The parameters used are $r = 0.05$, $X_0 = 50$, $T = 1.0$, and the number of simulated paths is 1 000 000 for Crude MC, LSIS and GHS, and 50 000 for RM. Results for LSIS obtained by optimizing the drift only (LSIS($\tilde{\mu}$)), and both the drift and the volatility (LSIS($\tilde{\mu}, \tilde{\sigma}$)) are reported. The uncertainties on the least significative digits of the option prices, and variance reductions are reported in parentheses. Note that Arouna (2003) report no error estimate quoted for the RM results. However, from the ratio of the simulated paths, it is sensible to estimate the errors on the RM variance reductions as about $\sqrt{20} \simeq 4$ times larger than those quoted for LSIS($\tilde{\mu}$).

σ	K	Crude MC Price	LSIS($\tilde{\mu}$)		LSIS($\tilde{\mu}, \tilde{\sigma}$)		RM		GHS	
			Price	VR	Price	VR	Price	VR	Price	VR
0.1	30	21.4633(50)	21.46294(49)	104(1)	21.46296(12)	1700(100)	21.47	112(4)	21.46294(50)	100(1)
	50	3.4032(39)	3.4019(14)	7.8(1)	3.4046(10)	15(1)	3.41	7.8(4)	3.4018(14)	7.8(1)
	60	0.2315(11)	0.23112(19)	33.5(5)	0.23132(12)	84(5)	0.23	31(2)	0.23126(19)	33.5(5)
0.3	30	21.598(15)	21.5912(37)	16.4(1)	21.5984(21)	51(1)	21.63	16.8(4)	21.5973(39)	14.8(2)
	50	7.114(11)	7.1169(35)	9.9(5)	7.1159(21)	27(1)	7.12	11(2)	7.1146(35)	9.9(1)
	60	3.4954(83)	3.4514(21)	15.6(1)	3.4523(14)	35(1)	3.45	15.2(4)	3.4508(22)	14.2(1)

Table 2. Same as table 1 for a European put option.

σ	K	Crude MC Price	LSIS($\tilde{\mu}$)		LSIS($\tilde{\mu}, \tilde{\sigma}$)		RM		GHS	
			Price	VR	Price	VR	Price	VR	Price	VR
0.1	40	0.004235(98)	0.0041650(47)	435(6)	0.0041616(41)	571(9)	0.0042	350(24)	0.0041650(47)	435(6)
	50	0.9636(19)	0.96419(64)	8.8(1)	0.96436(38)	25(2)	0.97	9.6(4)	0.96385(63)	9.1(1)
	60	7.3052(46)	7.3059(19)	5.9(1)	7.3056(11)	17(1)	7.31	6.3(4)	7.3047(19)	5.9(1)
0.3	30	0.13397(83)	0.13445(13)	41(1)	0.13448(10)	69(2)	0.13	38(4)	0.013440(13)	40.8(5)
	50	4.6794(65)	4.6761(27)	5.8(1)	4.6767(16)	16.5(5)	4.68	6.2(4)	4.6761(27)	5.8(1)
	60	10.5203(97)	10.5236(44)	4.9(1)	10.5266(26)	13.9(2)	10.54	4.8(4)	10.5223(46)	4.4(1)

Table 3. European butterfly spread option (28) on a log-normal asset (14) for different values of the spot price: comparison between the saddle point approach of Glasserman *et al.* (1999a) (GHS), and LSIS with the optimization of both drift and volatility (LSIS($\tilde{\mu}, \tilde{\sigma}$)). The parameters used are $r=0.1$, $K_1=45$, $K_2=50$, $K_3=55$, $T=1.0$, $\sigma=0.3$, and the number of simulated paths is 1 000 000.

X_0	Crude MC Price	GHS		LSIS($\tilde{\mu}, \tilde{\sigma}$)	
		Price	VR	Price	VR
30	0.15818(69)	0.15732(33)	4.4(8)	0.157659(40)	298(5)
40	0.4871(11)	0.48666(98)	1.26(1)	0.48703(11)	100(5)
50	0.6274(13)	0.6274(13)	1.00(1)	0.62750(11)	140(1)
60	0.5156(12)	0.5157(10)	1.44(1)	0.515636(93)	166(2)
70	0.32961(97)	0.32875(67)	2.10(1)	0.329281(73)	177(2)

allowing the sampling volatility to vary, we obtain $\tilde{\mu}^* \simeq -0.02$ and $\tilde{\sigma}^* \simeq 0.14$, reducing the variance by more than two orders of magnitude.

5.2. Asian call option

As a second example we consider here the single underlying ($L=1$) arithmetic Asian option, equation (15), with a call payout, $\Phi(Z) = e^{-rT}(\bar{X} - K)^+$, and M observation dates.

In this case, the calculation of the expectation value (1) involves a multi-dimensional integral over the probability density (13). However, the implementation of LSIS is no more difficult than for the European payout. Indeed, the fact that the sample points Z_i are in the present case M -dimensional vectors is practically irrelevant for the LSIS procedure. All the least-square fitter needs is a method that, given a configuration Z_i and the vector θ , provides the value of the target function $f_\theta(Z_i) = W_\theta(Z_i)^{1/2}G(Z_i)$.

In order to compare LSIS with the other approaches proposed in the literature we have considered a Gaussian trial density of the form of equation (16). Here such density depends on as many parameters, $\tilde{\mu} = (\tilde{\mu}_0, \dots, \tilde{\mu}_{M-1})$, as there are sampling dates in the Euler discretization of the process (12). A typical outcome of the optimized drift vector produced by the LSIS procedure is shown in figure 2.

5.2.1. Comparison with the saddle point approximation.

As shown in table 4, LSIS provides a very effective variance reduction for the Asian call options considered, similar to those obtained with the approach of GHS. This could be expected on general grounds as in the case of the European call and put options. Indeed, it is not difficult to verify that the single asset Asian call option generates a single-mode optimal sampling density $P_{\text{opt}}(Z) \propto G(Z)P(Z)$, with covariance close to one. Hence, as discussed in the previous section, the optimal sampling density can be accurately represented by a Gaussian distribution, and the saddle point method provides an accurate approximation of the optimal drift vector. Indeed, as shown in figure 2, the optimal drift vectors obtained with the LSIS and the GHS

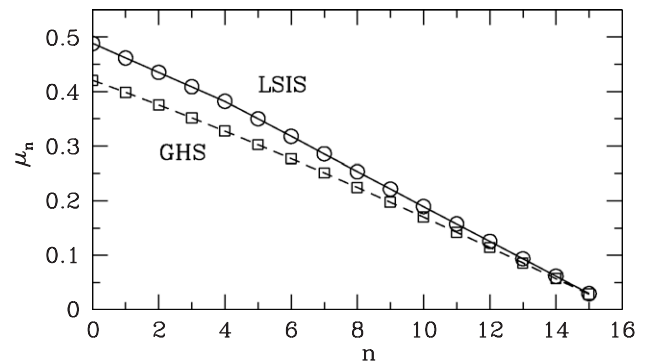


Figure 2. Optimal drift vector as obtained with the LSIS and the GHS procedures for an arithmetic Asian call option (15), $\Phi(\bar{X}) = (\bar{X} - K)^+$ on a log-normal asset (14) for $M=16$, $\sigma=0.3$, $X_0=K=50$, $r=0.05$, and $T=1.0$.

approaches are quite similar, with an overlap of ~ 0.99 when both are normalized to one.

As anticipated, we found that a few hundred Monte Carlo configurations and 10–20 iterations of the least-squares fitter were typically sufficient to determine the optimal drift vector in (16) for $M \simeq 50$, thus making the overhead of the pre-simulation stage rather limited. However, as the number of time steps or, more generally, the number of components of the drift vector increases, the complexity of the optimization problem also increases. Nevertheless, as suggested by Glasserman *et al.*(1999b), one can significantly reduce the computation time associated with the optimization stage by approximating the drift vector with a continuous function parameterized by a small number of parameters. These are, in turn, tuned by the least-square algorithm in order to determine an approximate optimal drift vector. We have found that a particularly effective realization of this approach is to approximate the drift vector by a piecewise linear function, parameterized by its values where it changes slope (the so-called *knot points*). For instance, in the present examples, optimizing over four to eight equally spaced knot points provides variance reductions practically indistinguishable from those obtained by a full optimization. In addition, the drift vector generally changes continuously with the market parameters. As a result, additional computational saving can be obtained

Table 4. Comparison between LSIS and the saddle point approach of Glasserman *et al.* (1999a) (GSH): price of Asian options (15) on a log-normal asset (14) for different values of the volatility σ , the strike K , and the number of observation dates M . The parameters used are $r=0.05$, $X_0=50$, $T=1.0$, and the number of simulated paths is 1 000 000.

M	σ	K	Crude MC Price	LSIS		GHS	
				Price	VR	Price	VR
16	0.1	45	6.0565(29)	6.05522(88)	10.86(5)	6.05537(89)	10.62(5)
		50	1.9198(22)	1.91994(80)	7.56(5)	1.91914(83)	7.03(5)
		55	0.20272(74)	0.20235(16)	21.4(2)	0.20237(16)	21.4(2)
16	0.3	45	7.1545(77)	7.1531(26)	8.8(1)	7.1529(27)	8.13(1)
		50	4.1730(63)	4.1714(20)	9.9(1)	4.1712(21)	9.0(1)
		55	2.2135(48)	2.2115(13)	13.6(7)	2.2116(14)	11.8(7)
64	0.1	45	5.9967(28)	5.99510(87)	10.4(5)	5.99500(85)	10.9(5)
		50	1.8467(21)	1.84522(78)	7.2(3)	1.84525(81)	6.7(3)
		55	0.17519(67)	0.17448(14)	23(2)	0.17443(14)	23(1)
64	0.3	45	7.0257(75)	7.0204(25)	9.0(2)	7.0204(26)	8.3(2)
		50	4.0271(61)	4.0222(19)	10.3(5)	4.0220(20)	9.3(5)
		55	2.0849(46)	2.0794(13)	12.5(5)	2.0794(13)	12.5(5)

Table 5. Importance Sampling plus stratification. Comparison between LSIS and the GHS approach of Glasserman *et al.* (1999a) for the Asian option of table 4.

M	σ	K	LSIS + SS		GHS + SS	
			Price	VR	Price	VR
16	0.3	45	7.15284(25)	950(20)	7.15266(24)	1030(10)
		50	4.17122(18)	1225(15)	4.17118(18)	1225(30)
		55	2.21183(11)	1900(100)	2.21183(11)	1900(50)
64	0.3	45	7.02075(23)	1060(30)	7.02076(23)	1060(30)
		50	4.02251(17)	1290(30)	4.02250(17)	1290(30)
		55	2.07967(13)	1320(100)	2.07965(12)	1470(100)

by starting the pre-simulation from a drift vector optimized for a case with a similar set of parameters.

As anticipated in section 3.1, Importance Sampling can also be naturally combined with Stratified Sampling by choosing the optimal drift vector as the stratification direction (for further details, see Glasserman *et al.* (1999a)). This, as shown in table 5, gives rise to a further reduction of the variance, which can be of several order of magnitude, depending on the option parameters, thus resulting in quite remarkable savings in computational time.

5.2.2. Comparison with adaptive Monte Carlo approaches.

It is also interesting to compare the LSIS results with those obtained by means of the recently proposed adaptive Monte Carlo methods briefly discussed in section 3.1. As illustrated in table 6, we found that our approach achieves variance reductions similar to those obtained with the Robbins–Monro algorithm of Arouna (2003). However, for the example considered, LSIS appears to perform better for small values of the volatility and deep in and out of the money options. This is particularly remarkable when considering that the pre-simulation stage required by the Robbins–Monro algorithm involves sampling a number of Monte Carlo paths much larger than that required by the LSIS approach. For instance, in the present example, the number of iterations

quoted by Arouna (2003) is 400 000, while the LSIS results were obtained by sampling just $N'_p \simeq 400$ configurations. In addition, LSIS does not have any exogenous parameter to be fine-tuned in order to achieve an efficient optimization, thus making the approach easier to be automated.

Finally, in order to make a comparison with the adaptive Monte Carlo approach of Su and Fu (2000, 2002), we have slightly modified the Asian option example and considered a payout depending on a partial average

$$\bar{X} = \frac{1}{M - M_0} \sum_{i=M_0}^M X_i, \quad (29)$$

where M_0 is the time step index corresponding to the first observation. The results we have obtained in this case are reported in table 7, and show that LSIS generally outperforms the approach of Su and Fu (2000, 2002). This is expected as LSIS involves a more general class of sampling densities (16) as the change of drift is not constrained to be uniform along the simulation path.

5.3. Multi-modal examples

All the examples considered in the previous sections were characterized by a single mode optimal sampling

Table 6. Comparison between LSIS and the adaptive Robbins–Monro (RM) method (as quoted by Arouna (2003)) for the price of an Asian call option (15) on a log-normal asset (14) for different values of the volatility σ , and of the strike price K . The parameters used are $r = 0.1$, $X_0 = 50$, $T = 0.5$, and the number of simulated paths is 1 000 000 for Crude MC and LSIS, and 800 000 for RM. The uncertainties on the RM variance reductions are rough estimates based on the number of simulated paths.

M	σ	K	Crude MC Price	LSIS		RM	
				Price	VR	Price	VR
20	0.1	40	10.7820(21)	10.78243(22)	91(2)	10.73	58(2)
		50	1.6045(16)	1.60398(61)	6.9(1)	1.53	7.0(1)
		55	0.04919(29)	0.049204(54)	28.8(5)	0.037	6.0(5)
	0.3	40	10.8284(62)	10.8277(17)	13(1)	10.77	14(1)
		50	3.1373(44)	3.1352(15)	8.6(1)	2.99	9.0(1)
		60	0.3790(16)	0.37861(35)	20.9(5)	0.320	28.0(5)

Table 7. Comparison between LSIS and the adaptive approach of Su and Fu (SF) (as quoted by Su and Fu (2000)) for the price of a partial average Asian call option (15) and (29) on a log-normal asset (14) for different values of the volatility σ , and of the strike price K . The parameters used are $r = 0.05$, $X_0 = 100$, $T = 1.0$, $M = 365$, $M_0 = 305$, and the number of simulated paths is 1 000 000 for Crude MC and LSIS, and 50 000 for SF. The uncertainties on the SF variance reductions are rough estimates based on the number of simulated paths.

σ	K	Crude MC Price	LSIS		SF	
			Price	VR	Price	VR
0.2	100	9.792(14)	9.7816(46)	9.3(1)	9.747	6.7(3)
	110	5.413(11)	5.4256(33)	11.1(1)	5.397	8.2(3)
	120	2.7758(77)	2.7669(20)	14.8(2)	2.730	11.0(6)
	130	1.3038(52)	1.3096(12)	18.8(1)	1.284	17.0(4)
	140	0.5861(35)	0.58298(61)	33(1)	0.575	25(3)
	150	0.2421(22)	0.24592(30)	54(4)	0.241	44(12)
	160	0.1003(14)	0.09973(14)	100(8)	0.0980	85(20)
	170	0.03761(84)	0.039000(64)	172(20)	0.0380	173(80)
0.3	100	13.323(21)	13.3527(65)	10.4(5)	13.295	7(1)
	110	9.182(18)	9.1564(54)	11.1(5)	9.103	8(1)
	120	6.1041(15)	6.1258(39)	14.8(5)	6.059	10(1)
	130	4.020(12)	4.0034(31)	15.0(5)	3.985	12(3)
	140	2.5624(98)	2.5769(23)	18.1(7)	2.556	15(2)
	150	1.6410(80)	1.6325(13)	38(1)	1.603	22(2)
	160	1.0143(62)	1.02355(90)	47(2)	1.006	30(5)
	170	0.6408(50)	0.63681(61)	67(3)	0.623	36(6)

density (9). In these cases, an Importance Sampling strategy based on the Gaussian trial density (16) or (27) provides good results. In this section we will consider instead option pricing problems that are characterized by multi-modal optimal sampling densities. These are very common in practice, and typically arise if the payout function $G(Z)$ is not monotonic, or for claims written on multiple assets. We will show that, in these cases, a simple Importance Sampling strategy based on a pure change of drift of a Gaussian density proves ineffective. In contrast, LSIS allows one to work with trial densities tailored to the problem at hand, including multi-modal distributions. As in section 5.1, we begin by illustrating the effectiveness of LSIS with a simple European-style payout.

5.3.1. Straddle. A simple example of a pricing problem characterized by a multi-modal sampling density is the European straddle:

$$\Phi(X) = (X - K)^+ + (K - X)^+. \tag{30}$$

As illustrated in figure 3, the corresponding optimal sampling density, $P_{\text{opt}}(Z) \propto G(Z) \exp(-Z^2/2)$, has two well-separated maxima. As a result, saddle point approaches based on a Gaussian distribution, e.g. centered on the higher of the two modes, provide a poor approximation of the optimal sampling density. Indeed, as shown in table 8, the saddle point method *increases* the variance with respect to the crude Monte Carlo estimator. On the contrary, LSIS based on the simple trial density (16) provides a small shift of the drift towards the higher mode of the optimal density (figure 3). This can be easily understood. In fact, the best way to approximate a bi-modal distribution by means of a Gaussian density is to center the latter around the center of mass of the former, in this case $\tilde{\mu} \simeq 0.24$. Yet, since the two maxima of the optimal density are well separated, its overlap with the trial density is still rather poor. As a result, LSIS based on a trial density of the form (16) produces only a small variance reduction. Better results can be obtained by adjusting the width of the trial density. Indeed, as also illustrated in figure 3,

Downloaded By: [Capriotti, Luca] At: 18:57 19 August 2008




Figure 3. Sampling probability density functions for a straddle option (30) with $r=0.05$, $X_0=K=50$, $\sigma=0.3$, $T=1$, as obtained with LSIS (optimizing just the drift, $LSIS(\tilde{\mu})$, both the drift and the volatility, $LSIS(\tilde{\mu}, \tilde{\sigma})$, and using the bi-modal trial density (31), $LSISbm$) and the saddle point approximation (GHS). The original (13) and the optimal (9) sampling densities are also shown for comparison.

Table 8. Importance Sampling results for the European-style straddle (30) obtained by means of the saddle point approximation (GHS) and LSIS, with the optimization of the drift ($LSIS(\tilde{\mu})$), and the drift and the volatility ($LSIS(\tilde{\mu}, \tilde{\sigma})$). The parameters used are $r=0.05$, $X_0=K=50$, $\sigma=0.3$, $T=1$, and correspond to the probability densities sketched in figure 3. The number of simulated paths is 1 000 000.

Crude MC Price	GHS		LSIS($\tilde{\mu}$)		LSIS($\tilde{\mu}, \tilde{\sigma}$)	
	Price	VR	Price	VR	Price	VR
11.803(10)	11.800(33)	0.10(1)	11.8038(88)	1.30(1)	11.8047(55)	3.00(2)

by increasing the standard deviation of the sampling density one achieves a better approximation of the optimal density. This leads to a sizable reduction of the variance of a factor of 3.

An even better ansatz for the optimal distribution is clearly represented by a bi-modal trial density of the form

$$\tilde{P}(Z) = (2\pi)^{-d/2} [w_a e^{-(Z-\mu_a)^2} + w_b e^{-(Z-\mu_b)^2}], \quad (31)$$

where $w_a + w_b = 1$ (and $d=1$ in this specific case) which can be optimized over μ_a, μ_b , and w_a . The simulation of a density of this form is straightforward as it simply implies choosing on each Monte Carlo step one of the two Gaussian components in (31), and sampling a configuration Z_i according to it. This can be done by extracting an auxiliary uniform random number $\xi \in [0, 1]$, and extracting Z_i according to the first Gaussian component if $\xi < w_a$, and according to the second otherwise. The optimized distribution obtained using this bi-modal trial density is

sketched in figure 3 and corresponds to a further sizable reduction of the variance (table 8). Remarkably, the same form of sampling density (31) can also be used for Asian-style straddles (here $d=M$, μ_a and μ_b are M -dimensional vectors) and produces a very effective variance reduction, as illustrated in table 9.

5.3.2. Basket call options. Basket options are another very common class of contingent claims that can give rise to multi-modal optimal densities. This is illustrated in figure 4 for a very simple European-style call option on the maximum of $L=2$ underlying assets following the process (14)

$$\Phi(X) = (\max(X_1, X_2) - K)^+. \quad (32)$$

Also in this case, LSIS based on a bi-modal trial density of the form (31) provides a significant variance reduction that persists when introducing Asian features in the payout (see table 10). We obtained similar results for

Table 9. Variance reduction obtained with LSIS for a European ($M=1$) and Asian style ($M=16$ and 64) straddle (30) using the bi-modal sampling density (31). The parameters used are $r=0.05$, $X_0=K=50$, $\sigma=0.3$, $T=1$, and correspond for $M=1$ to the probability densities sketched in figure 3. The number of simulated paths is 1 000 000.

M	Crude MC Price	LSIS	
		Price	VR
1	11.803(10)	11.8009(44)	5.17(5)
16	7.0604(58)	7.0559(26)	4.98(5)
64	6.8200(55)	6.8163(26)	4.47(5)




Figure 4. Optimal sampling density (Equation (9)) for the European basket call option (32) for $r=0.05$, $K=100$, $X_0^{(1)}=100$, $X_0^{(2)}=105$, $\sigma_1=\sigma_2=0.3$, $T=1$.

Table 10. Importance Sampling results for the Asian-style basket option (32) obtained by means of LSIS, using the bi-modal sampling distribution (31). The parameters used are $r=0.05$, $T=1$, $X_0^{(1)}=100$, $X_0^{(2)}=105$, $\sigma_1=0.3$, $\sigma_2=0.3$, and correspond (for $K=100$ and $M=1$) to the probability densities sketched in figure 4. The number of simulated paths is 500 000.

M	K	Crude MC Price	LSIS	
			Price	VR
1	90	34.555(41)	34.587(14)	8.58(7)
	100	26.621(39)	26.589(15)	6.76(3)
	110	19.749(36)	19.776(13)	7.7(1)
16	90	24.876(24)	24.8700(87)	7.6(1)
	100	16.428(22)	16.4081(81)	7.4(1)
	110	9.711(19)	9.6918(68)	7.8(1)
32	90	24.537(23)	24.5539(84)	7.50(5)
	100	16.082(22)	16.0734(82)	7.2(1)
	110	9.381(18)	9.3752(68)	7.0(1)

Downloaded By: [Capriotti, Luca] At: 18:57 19 August 2008

other multi-asset options by generalizing the form of the trial density (31).

6. Conclusions

In this paper we have described a simple Importance Sampling technique based on a least-squares minimization of straightforward implementation. The resulting strategy, dubbed Least-square Importance Sampling (LSIS), provides an effective variance reduction technique that lends itself to a variety of applications.

We have presented several numerical examples in a diffusive setting, and we have shown that LSIS, when restricted to the optimization of the drift of a diffusion process, provides variance reductions similar to those obtained with existing approaches (Glasserman *et al.* 1999a, Su and Fu 2000, 2002, Arouna 2003). However, LSIS is not limited to the determination of the optimal mean of a Gaussian sampling distribution. This is important when the ability to adjust the width of such a distribution, or to sample non-Gaussian and multi-modal densities, is critical to mimic the form of the optimal sampling density and to obtain sizable variance reductions.

The LSIS strategy can be applied to any Monte Carlo setting provided that a reasonable ansatz for the optimal sampling density is available. This makes LSIS a flexible Importance Sampling approach, that can be used across a variety of financial applications, ranging from value at risk (VaR) estimation, to portfolio credit risk management. This is currently the object of further investigation.

Acknowledgements

It is a pleasure to acknowledge Gabriele Cipriani, David Shorthouse, and Mark Stedman for stimulating discussions, and Paul Glasserman for an enlightening lecture that inspired this work. Useful suggestions by the anonymous referees are also gratefully acknowledged. The opinion and views expressed in this paper are uniquely those of the author, and do not necessarily represent those of Credit Suisse Group.

References

Arouna, B., Robbins–Monro algorithms and variance reduction in finance. *J. Comput. Finan.*, 2003, 7, 1245–1255.

- Arouna, B., Adaptive Monte Carlo method, a variance reduction technique. *Monte Carlo Meth. Applic.*, 2004, 10, 1–24.
- Baxter, M. and Rennie, A., *Financial Calculus: An Introduction to Derivative Pricing*, 1996 (Cambridge University Press: Cambridge).
- Boyle, P., Broadie, M. and Glasserman, P., Monte Carlo methods for security pricing. *J. Econ. Dyn. Contr.*, 1997, 21, 1267–1321.
- Bressanini, D., Morosi, G. and Mella, M., Robust wave function optimization procedures in quantum Monte Carlo methods. *J. Chem. Phys.*, 2002, 116, 5345–5350.
- Glasserman, P., *Monte Carlo Methods in Financial Engineering*, 2004 (Springer: New York).
- Glasserman, P., Heidelberger, P. and Shahabuddin, P., Asymptotically optimal importance sampling and stratification for pricing path-dependent options. *Math. Finan.*, 1999a, 9, 117–152.
- Glasserman, P., Heidelberger, P. and Shahabuddin, P., Importance sampling in the Heath–Jarrow–Morton framework. *J. Deriv.*, 1999b, 7, 32–50.
- Glasserman, P. and Li, J., Importance sampling for portfolio credit risk. *Mgt. Sci.*, 2005, 51, 1643–1656.
- Harrison, J. and Kreps, D., Martingales and arbitrage in multiperiod security markets. *J. of Econ. Theory*, 1979, 20, 381–408.
- Hull, J.C., *Options, Futures and Other Derivatives*, 2002 (Prentice Hall: Englewood Cliffs, NJ).
- Kallenberg, O., *Foundations of Modern Probability*, 1997 (Springer: New York).
- Kleywegt, A.J., Shapiro, A. and de Mello, T.H., The sample average approximation method for stochastic discrete optimization. *SIAM J. Optimiz.*, 2001, 12, 479–502.
- Musiela, M. and Rutkowski, M., *Martingale Methods in Financial Modelling*, 2002 (Springer: New York).
- Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P., *Numerical Recipes in C++: The Art of Scientific Computing*, 2002 (Cambridge University Press: Cambridge).
- Reider, R., Working Paper (unpublished), 1993.
- Shapiro, A., Simulation-based optimization: convergence analysis and statistical inference. *Stochast. Mod.*, 1996, 12, 425–454.
- Su, Y. and Fu, M.C., Simulation in financial engineering: importance sampling in derivative securities pricing. In *Proceedings of 2000 Winter Simulation Conference*, edited by J. Joines, R. Barton, K. Kang, and P. Fishwick, pp. 587–596, 2000 (IEEE Press: Piscataway, NJ).
- Su, Y. and Fu, M.C., Optimal importance sampling in securities pricing. *J. Comput. Finan.*, 2002, 5, 27–50.
- Umrigar, C.J., Wilson, K.G. and Wilkins, J.W., Optimized trial wave functions for quantum Monte Carlo calculations. *Phys. Rev. Lett.*, 1988, 60, 1719–1722.
- Vazquez-Abad, F. and Dufresne, D., Accelerated simulation for pricing Asian options. In *Proceedings of 1998 Winter Simulation Conference*, edited by D. Medeiros, E. Watson, J. Carson, and M. Manivann, pp. 1493–1500, 1998 (IEEE Press: Washington, DC).