

LANDMARK PAPER REPRISE – RECENT ADVANCES IN SIMULATION FOR SECURITY PRICING

Proceedings of the 1995 Winter Simulation Conference
ed. C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman

RECENT ADVANCES IN SIMULATION FOR SECURITY PRICING

Phelim Boyle

School of Accountancy
University of Waterloo
Waterloo, Ontario, N2L 3G1

Mark Broadie

Columbia Business School
415 Uris Hall
New York, NY 10027

Paul Glasserman

Columbia Business School
403 Uris Hall
New York, NY 10027

ABSTRACT

Computational methods play an important role in modern finance. Through the theory of arbitrage-free pricing, the price of a derivative security can be expressed as the expected value of its payouts under a particular probability measure. The resulting integral becomes quite complicated if there are several state variables or if payouts are path-dependent. Simulation has proved to be a valuable tool for these calculations. This paper summarizes some of the recent applications and developments of the Monte Carlo method to security pricing problems.

1 INTRODUCTION

The increase in complexity of financial models and securities in recent years has led to greater attention to computational methods in the financial industry. Numerical methods are routinely used for a variety of applications, including the valuation of securities, the estimation of their sensitivities, risk analysis and stress testing of portfolios. Simulation is a useful tool for many of these calculations, evidenced in part by the voluminous literature of successful applications. Examples include the stochastic volatility applications in Duan (1995) and Hull and White (1987); the valuation of mortgage-backed securities in Schwartz and Torous (1989); the valuation of exotic options in Kemna and Vorst (1990); and the valuation of interest-rate derivative claims in Carverhill and Pang (1995) and Ritchken and Sankarsubramanian (1995).

We focus on the use of simulation in pricing *derivative securities*, also called *contingent claims*. These are securities, such as options or futures, whose payouts are determined by the value of certain underlying assets. The prices of derivative securities can be represented as expectations with respect to an appropriate probability measure involving the underlying

assets. If the number of underlying assets is large, or if the rule by which the derivative security derives its value is sufficiently complex, simulation becomes an attractive means for computing the price.

The representation of derivative security prices as expectations is a consequence of a deep result of financial theory; see, e.g., Duffie (1992) for background. Briefly, under a condition called *market completeness*, a derivative security can be replicated through trading in the underlying assets. The absence of arbitrage thus entails a relation between the price of the derivative security and those of the underlying assets. It turns out that, to preclude arbitrage, the price of the derivative security must be the expected value of its discounted payouts with respect to an *equivalent martingale measure*, also called a *risk-neutral probability*. This is the probability measure under which the discounted underlying assets become martingales; i.e., all assets have the same expected rate of return, which must then be the riskless rate.

To make this more concrete, we consider the Black-Scholes option pricing model. A typical model in continuous-time finance of the evolution of the price S_t of a stock or other asset is the stochastic differential equation

$$dS = \mu S dt + \sigma S dW, \quad (1)$$

in which μ is the rate of return, σ is the volatility, and W is a standard Brownian motion process. Under the risk-neutral measure, the drift μ is replaced by the risk-free rate, r , thus making $e^{-rt}S_t$ a martingale. An option to buy the stock at time T at price K (called the *strike price*) will pay $(S_T - K)^+$ at time T . The current price of the option is the expected present value of this payout with respect to the risk-neutral measure; i.e., it is

$$C = E[e^{-rT}(S_T - K)^+],$$

the expectation taken with $\mu = r$ in (1).

This particular expectation can be evaluated in closed-form, resulting in the celebrated Black-Scholes formula (see, e.g., Hull 1993). For purposes of illustration, we nevertheless point out how simulation would be used to compute the expectation. It follows from (1) that S_T has a lognormal distribution. Specifically, under the risk-neutral measure, it admits the representation

$$S_T = S_0 e^{(r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}Z}, \quad (2)$$

where Z is a standard normal random variable. Substituting independent samples Z_1, \dots, Z_n from the standard normal distribution into (2) yields independent samples $S_T^{(i)}$, $i = 1, \dots, n$, of the terminal stock price. An unbiased estimator of the option price is then given by

$$\hat{C} = \frac{1}{n} \sum_{i=1}^n C_i \equiv \frac{1}{n} \sum_{i=1}^n e^{-rT} \max\{0, S_T^{(i)} - K\}. \quad (3)$$

From this simple example, we may abstract the following general steps in pricing by simulation:

- Simulate sample paths of the underlying state variables (e.g., underlying asset prices and interest rates) over the relevant time horizon. Simulate these according to the risk-neutral measure.
- Evaluate the discounted cash flows of a security on each sample path, as determined by the structure of the security in question.
- Average the discounted cash flows over sample paths.

In the example above, the "paths" consisted simply of the terminal values $S_T^{(i)}$. More generally, (as we will see shortly) a pricing problem may require simulating a discrete-time approximation to the continuous-time process modeled by (1).

The rest of this paper is organized as follows. Section 2 discusses the use of some variance reduction techniques. Section 3 examines the application of low discrepancy sequences (quasi-Monte Carlo methods). Section 4 discusses the estimation of risk measures. Section 5 touches on further topics of current interest.

2 VARIANCE REDUCTION

In this section, we discuss the implementation of three specific variance reduction techniques in security pricing problems. The methods we discuss are antithetic variates, control variates, and moment matching.

2.1 Antithetic Variates

This method is more notable for its widespread familiarity among finance professionals than for its efficacy. Its popularity is no doubt due to its simplicity.

Consider, again, the problem of computing the Black-Scholes price of a call option, as discussed around (3). Though there is no need to use simulation in this case, it serves as a useful illustration. In this setting, the implementation of antithetic variates is particularly simple; for if Z_i has a standard normal distribution, then so does $-Z_i$. The price $\tilde{S}_T^{(i)}$ obtained from (2) with Z_i replaced by $-Z_i$ is thus a valid sample from the terminal stock price distribution. Similarly, each

$$\tilde{C}_i = e^{-rT} \max\{0, \tilde{S}_T^{(i)} - K\}$$

is an unbiased estimator of the option price, and therefore so is

$$\hat{C}_{AV} = \frac{1}{n} \sum_{i=1}^n \frac{C_i + \tilde{C}_i}{2}.$$

Because \hat{C}_{AV} uses twice as many samples as \hat{C} , it is preferred only if

$$2\text{Var}[\hat{C}_{AV}] \leq \text{Var}[\hat{C}],$$

which simplifies to $\text{Cov}[C_i, \tilde{C}_i] \leq 0$. That this condition is in fact met is a simple consequence of the monotonicity of the mapping from Z_i to C_i .

More elaborate options may depend on the entire *path* of stock prices rather than just the terminal value S_T . In this case, it becomes necessary to simulate a discrete-time approximation $\{S_{t_j}, j = 0, \dots, m\}$ of the path $\{S_t, 0 \leq t \leq T\}$. Each $S_{t_{j+1}}$ can be generated from the preceding price S_{t_j} and a normal variate Z_{j+1} according to (2), with T replaced by $t_{j+1} - t_j$ in the exponent. An antithetic path can then be generated using $-Z_1, \dots, -Z_m$ in place of Z_1, \dots, Z_m .

For some further examples of the application of antithetics in finance, see Boyle (1977), Clewlow and Carverhill (1994), and Hull and White (1987).

2.2 Control Variates

The use of control variates to reduce variance is well-known in simulation and has attracted some interest in financial applications. We describe two particularly effective examples specific to the financial setting.

Our first example is an application to *Asian* options proposed by Kemna and Vorst (1990). The payoff on an Asian option depends on the (arithmetic)

average price of the underlying asset. An example of an Asian option pricing problem is the computation of

$$P_A = E[e^{-rT}(\bar{S}_A - K)^+],$$

where

$$\bar{S}_A = \frac{1}{T} \int_0^T S_t dt.$$

As always, the expectation is taken with the rate of return on S equal to the risk-free rate r ; i.e., the expectation is with respect to the equivalent martingale measure.

There is no closed-form expression for P_A because there is no simple characterization of the distribution of \bar{S}_A . Even in discrete time, the distribution of an average of lognormal random variables does not admit a simple expression. In contrast, the *geometric* mean

$$\bar{S}_G = \exp \left\{ T^{-1} \int_0^T \ln S_t dt \right\}$$

is itself lognormally distributed, resulting in a tractable expression for

$$P_G = E[e^{-rT}(\bar{S}_G - K)^+].$$

Based on this observation, Kemna and Vorst (1990) replaced the straightforward estimator $\hat{P}_A = e^{-rT}(\bar{S}_A - T)^+$ with the control-adjusted estimator

$$\hat{P}_A + (P_G - \hat{P}_G),$$

where $\hat{P}_G = e^{-rT}(\bar{S}_G - K)^+$. Because of the strong correlation between \hat{P}_G and \hat{P}_A , they achieved substantial variance reduction using this approach. A further reduction in variance could presumably be obtained by optimizing the coefficient on $(P_G - \hat{P}_G)$, implicitly taken to be 1 in their implementation.

The martingales inevitably present in a security pricing simulation provide another source of control variates. The simplest of these in an option pricing simulation is the underlying asset itself. Because $\{e^{-rt}S_t\}$ is a martingale, we have $E[e^{-rT}S_T] = S_0$, so $(e^{-rT}S_T - S_0)$ provides a simple control. If there are multiple underlying assets, then multiple controls are readily available.

Clewlow and Carverhill (1994) have taken this observation a step further. They simulate a discrete-time approximation $\{S_t, j = 0, \dots, m\}$ of the asset price and build a control variate from the increments $\Delta S_j = S_{t_{j+1}} - S_{t_j}$. They choose the coefficients on these increments to approximate the change in the option price resulting from ΔS_j . Specifically, they use the derivative of the option price with respect to the underlying asset. This approach mimics the trading strategy used to replicate (or *hedge*) the option.

As the time increment used in the discrete-time approximation decreases to zero, the hedging strategy exactly replicates the option, ultimately resulting in a zero-variance simulation estimate. A possible weakness of this approach is the difficulty of computing the derivative of the option price at each time step.

For further applications of the control variate technique in financial simulation, see Boyle (1977), Broadie and Glasserman (1993), Carverhill and Pang (1995), Duan (1995), and Ritchken and Sankarsubramanian (1995).

2.3 Moment Matching

Next we describe a variance reduction technique proposed by Barraquand (1994), who termed it *quadratic resampling*. His technique is based on *moment matching*. As before, we introduce it with the simple example of estimating the call option price on a single asset and then generalize.

As before, let $Z_i, i = 1, \dots, n$, denote independent standard normals used to drive a simulation. The sample moments of the n Z 's will not exactly match those of the standard normal. The idea of moment matching is to transform the Z 's to match a finite number of the moments of the underlying population. For example, the first moment of the standard normal can be matched by defining

$$\tilde{Z}_i = Z_i - \bar{Z}, \quad i = 1, \dots, n, \quad (4)$$

where $\bar{Z} = \sum_{i=1}^n Z_i/n$ is the sample mean of the Z 's. Note that the \tilde{Z}_i 's are normally distributed if the Z_i 's are normal. However, the \tilde{Z}_i 's are not independent. As before, terminal stock prices are generated from the formula

$$\tilde{S}_T(i) = S_0 e^{(r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}\tilde{Z}_i}, \quad i = 1, \dots, n.$$

An unbiased estimator of the call option price is the average of the n values $\tilde{C}_i = e^{-rT} \max(\tilde{S}_T(i) - K, 0)$.

In the standard Monte Carlo method, confidence intervals for the true value C could be estimated from the sample mean and variance of the estimator. This cannot be done here since the n values of \tilde{Z} are no longer independent, and hence the values \tilde{C}_i are not independent. This points out one drawback of the moment matching method: confidence intervals are not as easy to obtain. Indeed, the sample variance of the \tilde{C}_i 's is usually a poor estimate of $\text{Var}[\tilde{C}_i]$.

Equation (4) shows one way to match the first moment of a distribution with mean zero. If the underlying population has mean μ_Z , transformed Z 's can be generated using $\tilde{Z}_i = Z_i - \bar{Z} + \mu_Z$. The idea can

easily be extended to match two moments of a distribution. In this case, an appropriate transformation is

$$\tilde{Z}_i = (Z_i - \bar{Z}) \frac{\sigma_Z}{s_Z} + \mu_Z, \quad i = 1, \dots, n, \quad (5)$$

where s_Z is the sample standard deviation of the Z_i 's and σ_Z is the population standard deviation. Of course, for a standard normal, $\mu_Z = 0$ and $\sigma_Z = 1$. An estimator of the call option price is the average of the n values \tilde{C}_i .

Using the transformation (5), the \tilde{Z}_i 's are not normally distributed even if the Z_i 's are normal. Hence, the corresponding \tilde{C}_i are biased estimators of the true option value. For most financial problems of practical interest, this bias is likely to be small. However, the bias can be arbitrarily large in extreme circumstances (even when only the first moment of the distribution is matched). This can happen, for example, if the transformation (4) changes the support of the Z 's.

From (2), the mean and variance of the terminal stock price S_T are also known, so the moment matching idea could be applied to the simulated terminal stock values $S_T(i)$. Our (limited) numerical experience does not provide consistent evidence for choosing between these moment matching methods. Both methods tend to outperform antithetics, but control variates seem most effective of all. This is not surprising, since it can be shown that moment matching is asymptotically equivalent to using moments as controls without optimizing the coefficients.

3 LOW DISCREPANCY SEQUENCES

Many studies have found that for the computation of integrals in moderate dimensions, quasi-Monte Carlo based on *low discrepancy* sequences outperforms ordinary Monte Carlo based on random inputs. Typically, the error in Monte Carlo decreases at rate $O(1/\sqrt{n})$, whereas the error using low discrepancy sequences decreases at rate $O((\log n)^d/n)$, where d is the dimension. For background on low discrepancy methods, see Niederreiter (1992) and Spanier and Maize (1994).


These methods have recently been applied to financial simulation problems with considerable empirical success. Specific applications are more fully described in papers by Birge (1995), Joy, Boyle, and Tan (1995) and Paskov (1994). These suggest that the integrals involved in security pricing may be well suited to quasi-Monte Carlo. There are two reasons why security prices might lend themselves better to quasi-Monte Carlo than do other classes of integrals. The first is that the integrands tend to be fairly smooth, and it is generally recognized that smoothness helps

in using quasi-Monte Carlo. The second reason is that the dimension of pricing problems is often linked to the length of the time horizon over which a security has payouts, and payouts in the distant future are significantly discounted in computing a price. Thus, many of the higher dimensions may contribute little to the value of the integral which potentially reduces the effective dimension of the problem.

Next we test standard Monte Carlo versus the low discrepancy sequences of Faure, Sobol', and Halton. We estimate the price of a discretely sampled geometric average Asian option, which is given by

$$C = E[e^{-rT}(\bar{S} - K)^+],$$

where $\bar{S} = (\prod_{j=1}^d S_j)^{1/d}$ and S_j is the asset price at time jT/d . Since \bar{S} is lognormally distributed, an exact formula is available for C . We generate 500 test problems by selecting random problem parameters. For each test problem, we compute price estimates based on $n = 50,000$ sample paths of the asset price using the four methods. Root-mean-squared (RMS) relative error results are shown in the next figure for problems of dimension $d = 10, 50,$ and 100 .



Results for the Halton sequence were not competitive and are suppressed. RMS error for standard Monte Carlo is nearly independent of the problem dimension. The error with the Sobol' method grows smoothly with the problem dimension and grows erratically for the Faure method (though this could be an artifact of the value of n).

These results are broadly consistent with existing literature. Bratley, Fox, and Niederreiter (1992) conclude that quasi-Monte Carlo is unlikely to outperform ordinary Monte Carlo in dimensions higher than about 12, but their test problems could be more difficult than those that typically arise in a financial context. Paskov (1994) reports successful application of Sobol' sequences in the evaluation of a 360 dimen-

sional integral arising from the pricing of a collateralized mortgage obligation. He uses much larger values of n and this could explain the better performance for such a high dimensional problem.

The application of low discrepancy methods in finance is not immune from other shortcomings associated with these methods: there are no simple, reliable error bounds or termination criteria.

4 ESTIMATING RISK MEASURES

Most of the discussion in this paper centers on the use of Monte Carlo for pricing securities. In practice, the evaluation of price *sensitivities* is often as important as the evaluation of the prices themselves. Indeed, whereas prices for some securities can be observed in the market, their sensitivities to parameter changes typically cannot and must therefore be computed. Since price sensitivities are important measures of risk, the growing emphasis on risk management systems suggests a greater need for their efficient computation.

The derivatives of a derivative security's price with respect to various model parameters are collectively referred to as *Greeks*, because several of these are commonly referred to with the names of Greek letters. (See, e.g., Chapter 13 of Hull (1993) for background.) Perhaps the most important of these — and the one to which we give primary attention — is *delta*: the derivative of the price of a contingent claim with respect to the current price of an underlying asset. The delta of a stock option, for example, is the derivative of the option price with respect to the current stock price. An option involving multiple underlying assets has multiple deltas, one for each underlying asset. In the rest of this section, we discuss various approaches to estimating price sensitivities, especially delta.

4.1 Finite-Difference Approximations

Consider the problem of computing the delta of the Black-Scholes price of a call option; i.e., computing

$$\Delta = \frac{dC}{dS_0},$$

where C is the option price and S_0 is the current stock price. There is an explicit expression for delta, so simulation is not required, but the example is useful for purposes of illustration. A crude estimate of delta is obtained by independently generating two discounted option payoffs, $\hat{C}(S_0)$ and $\hat{C}(S_0 + \epsilon)$, from initial prices S_0 and $S_0 + \epsilon$ (according to (2)-(3)) and computing the finite-difference ratio

$$\tilde{\Delta} = \epsilon^{-1}[\hat{C}(S_0 + \epsilon) - \hat{C}(S_0)]. \quad (6)$$

Repeating this many times and averaging we obtain an estimator converging to

$$\epsilon^{-1}[C(S_0 + \epsilon) - C(S_0)], \quad (7)$$

where $C(\cdot)$ is the option price as a function of the current stock price.

This discussion suggests that to get an accurate estimate of Δ we should make ϵ small. However, because we generated S_T and $S_T(\epsilon)$ independently of each other, we have

$$\text{Var}[\tilde{\Delta}] = \epsilon^{-2}(\text{Var}[\hat{C}(S_0 + \epsilon)] + \text{Var}[C(S_0)]) = O(\epsilon^{-2}),$$

so the variance of $\tilde{\Delta}$ becomes very large if we make ϵ small. To get an estimator that converges to Δ we must let ϵ decrease slowly as n increases, resulting in slow overall convergence. A general result of Glynn (1989) shows that the best possible convergence rate using this approach is typically $n^{-1/4}$. Replacing the forward difference estimator in (6) with the central difference $(2\epsilon)^{-1}[\hat{C}(S_0 + \epsilon) - \hat{C}(S_0 - \epsilon)]$ typically improves the optimal convergence rate to $n^{-1/3}$. These rates should be compared with $n^{-1/2}$, the rate ordinarily expected from Monte Carlo.

Better estimators can generally be obtained using the method of common random numbers, which, in this context, simply uses the same Z for $\hat{C}(S_0)$ and $\hat{C}(S_0 + \epsilon)$. Denote by $\hat{\Delta}$ the finite-difference approximation thus obtained. For fixed ϵ , the sample mean of independent replications of $\hat{\Delta}$ also converges to (7). But if $\hat{C}(S_0)$ and $\hat{C}(S_0 + \epsilon)$ are positively correlated when simulated with common random numbers, then $\text{Var}[\hat{\Delta}] \leq \text{Var}[\tilde{\Delta}]$. That they are in fact positively correlated is easily verified in this example.

The impact of this variance reduction is most dramatic when ϵ is small. A simple calculation shows that, using common random numbers,

$$E[|\hat{C}(S_0 + \epsilon) - \hat{C}(S_0)|^2] = O(\epsilon^2), \quad (8)$$

and therefore that

$$\text{Var}[\epsilon^{-1}\{\hat{C}(S_0 + \epsilon) - \hat{C}(S_0)\}] = O(1);$$

i.e., the variance of $\hat{\Delta}$ remains *bounded* as $\epsilon \rightarrow 0$, whereas we saw previously that the variance of $\tilde{\Delta}$ increases at rate ϵ^{-2} . Thus, the more precisely we try to estimate Δ (by making ϵ small) the greater the benefit of common random numbers. Moreover, this indicates that to get an estimator that converges to Δ we may let ϵ decrease faster as n increases than was possible with $\tilde{\Delta}$, resulting in faster overall convergence. An application of Proposition 2 of L'Ecuyer and Perron (1994) shows that a convergence rate of $n^{-1/2}$ can be achieved.

The dramatic success of common random numbers in this example relies on the fast rate of mean-square convergence of $\hat{C}(S_0 + \epsilon)$ to $\hat{C}(S_0)$ evidenced by (8). This rate does not apply in all cases. It fails to hold, for example, in the case of a *digital option* paying a fixed amount B if $S_T > K$ and 0 otherwise. The price of this option is $C = e^{-rT}BP(S_T > K)$; the obvious simulation estimator is

$$\hat{C}(S_0) = e^{-rT}B\mathbf{1}_{\{S_T > K\}},$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator of the event $\{\cdot\}$. Because $\hat{C}(S_0)$ and $\hat{C}(S_0 + \epsilon)$ differ only when $S_T \leq K < S_T(\epsilon)$, we have

$$E[|\hat{C}(S_0 + \epsilon) - \hat{C}(S_0)|^2] = O(\epsilon),$$

compared with $O(\epsilon^2)$ for a standard call. As a result, delta estimation is more difficult for the digital option, and a similar argument applies to barrier options generally. Even in these cases, the use of common random numbers can result in substantial improvement compared with differences based on independent runs.

4.2 Direct Estimates

Even with the improvements in performance obtained from common random numbers, derivative estimates based on finite differences still suffer from two shortcomings. They are biased (since they compute difference ratios rather than derivatives) and they require multiple resimulations: estimating sensitivities to d parameter changes requires repeatedly running one simulation with all parameters at their base values and d additional simulations with each of the parameters perturbed. The computation of 10–50 Greeks for a single security is not unheard of, and this represents a significant computational burden when multiple resimulations are required.

Over the last decade, a variety of *direct* methods have been developed for estimating derivatives by simulation. Direct methods compute a derivative estimate from a single simulation, and thus do not require resimulation at a perturbed parameter value. Under appropriate conditions, they result in unbiased estimates of the derivatives themselves, rather than of a finite-difference ratio. Our discussion focuses on the use of *pathwise* derivatives as direct estimates, based on *infinitesimal perturbation analysis*.

The pathwise estimate of the true delta dC/dS_0 is the derivative of the sample price \hat{C} with respect to S_0 . More precisely, it is

$$\frac{d\hat{C}}{dS_0} = \lim_{\epsilon \rightarrow 0} \epsilon^{-1}[\hat{C}(S_0 + \epsilon) - \hat{C}(S_0)],$$

provided the limit exists with probability 1. If $\hat{C}(S_0)$ and $\hat{C}(S_0 + \epsilon)$ are computed from the same Z , then provided $S_T \neq K$, we have

$$\frac{d\hat{C}}{dS_0} = \frac{d\hat{C}}{dS_T} \frac{dS_T}{dS_0} = e^{-rT} \mathbf{1}_{\{S_T > K\}} \frac{S_T}{S_0}. \quad (9)$$

At $S_T = K$, C fails to be differentiable; however, since this occurs with probability zero, the random variable $d\hat{C}/dS_0$ is almost surely well defined.

The pathwise derivative $d\hat{C}/dS_0$ can be thought of as a limiting case of the common random numbers finite-difference estimator in which we evaluate the limit analytically rather than numerically. It is a direct estimator of the option delta because it can be computed directly from a simulation starting at S_0 without the need for a separate simulation at a perturbed value S_0 . This is evident from the expression in (9). The question remains whether this estimator is unbiased; that is, whether

$$E \left[\frac{d\hat{C}}{dS_0} \right] = \frac{dC}{dS_0} \equiv \frac{d}{dS_0} E[\hat{C}].$$

The unbiasedness of the pathwise estimate thus reduces to the interchangeability of derivative and expectation. The interchange is easily justified in this case; see Broadie and Glasserman (1993) and Fu and Hu (1993) for this example and others.

The utility of this technique rests on its applicability to more general models. In Broadie and Glasserman (1993), pathwise estimates are derived and studied (both theoretically and numerically) for Asian options and a model with stochastic volatility. For example, the Asian option delta estimate is simply

$$e^{-rT} \frac{\bar{S}}{S_0} \mathbf{1}_{\{\bar{S} > K\}},$$

where \bar{S} is the average asset price used to determine the option payoff. Evaluating this expression takes *negligible* time compared with resimulating to estimate the option price from a perturbed initial stock price. The pathwise estimate is thus both more accurate and faster to compute than the finite-difference approximation. These advantages extend to a wide class of problems.

As already noted, the unbiasedness of pathwise derivative estimates depends on an interchange of derivative and expectation. In practice, this generally means that the security payoff should be a pathwise continuous function of the parameter in question. The standard call option payoff $e^{-rT} \max\{0, S_T - K\}$ is continuous in each of its parameters. An example where continuity fails is a digital option with payoff

$e^{-rT} \mathbf{1}_{\{S_T > K\}} B$, with B the amount received if the stock finishes in the money. Because of the discontinuity at $S_T = K$, the pathwise method (in its simplest form) cannot be applied to this type of option.

The problem of discontinuities often arises in the estimation of γ , the second derivative of an option price with respect to the current price of an underlying asset. Consider, again, the standard call option. We have an expression for $d\hat{C}/dS_0$ in (9) involving the indicator $\mathbf{1}_{\{S_T > K\}}$. This shows that $d\hat{C}/dS_0$ is discontinuous in S_T , preventing us from differentiating pathwise a second time to get a direct estimator of γ .

To address the problem of discontinuities, Broadie and Glasserman (1993) construct *smoothed* estimators. These estimators are unbiased, but not as simple to derive and implement as ordinary pathwise estimators. Broadie and Glasserman also investigate the use of the *likelihood ratio method* for derivative estimation. This method differentiates the probability density of an asset price, rather than the outcome of the asset price itself. The domains of this method and the pathwise method overlap, but neither contains the other. When both apply, the pathwise method generally has lower variance.

Overviews of these methods can be found in Glasserman (1991), Glynn (1987), and Rubinstein and Shapiro (1993). For discussions specific to financial applications see Broadie and Glasserman (1993) and Fu and Hu (1993).

5 FURTHER TOPICS

We conclude this paper with a brief discussion of other recent developments in the application of Monte Carlo methods to security pricing.

There have recently been some advances made on the problem of pricing *American* contingent claims by simulation. These are securities whose cash flows depend on decisions of the owner as well as on the path of the underlying asset or assets. (When no decisions are involved, the security is called *European*.) Pricing an American contingent claim involves determination of an optimal policy and of the present value of a security's payouts under that policy. The optimization involved makes this a difficult problem for simulation. Other numerical methods, however, face difficulties when the dimension of the problem is large. Tilley (1993), Grant, Vora, and Weeks (1994), Barraquand and Martineau (1995), and Broadie and Glasserman (1995), have proposed simulation-based methods to price American-style securities.

Reider (1993) and Nielsen (1994) have explored the possibility of using *importance sampling* to speed up

the computation of option prices. This technique changes the underlying probability measure to give greater weight to paths with otherwise low probability. The resulting estimate is then weighted by a likelihood ratio to eliminate bias resulting from the change of measure. Reider's implementation changes the drift in a process with continuous state space; Nielsen's changes the transition probabilities in a binomial lattice. Both techniques show potential for variance reduction.

An important issue in any security-pricing simulation (but one that we have not addressed) is the approximation of a diffusion by a discrete-time process. Kloeden and Platen (1992) discuss a variety of methods for constructing discrete-time approximations with different orders of convergence. For any such scheme, decreasing the time step can be expected to give more accurate results, but at the expense of greater computational effort. Duffie and Glynn (1993) analyze this trade-off and characterize asymptotically optimal time steps as the overall computational effort grows.

REFERENCES

- Barraquand, J. 1994. Numerical Valuation of High Dimensional Multivariate European Securities, working paper, Digital Equipment Corporation, Paris, France (to appear in *Management Science*).
- Barraquand, J., and D. Martineau. 1995. Numerical Valuation of High Dimensional Multivariate American Securities, working paper, Salomon Brothers International Ltd. (to appear in *Journal of Financial and Quantitative Analysis*).
- Birge, J.R. 1995. Quasi-Monte Carlo Approaches to Option Pricing, working paper, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109.
- Boyle, P. 1977. Options: A Monte Carlo Approach, *Journal of Financial Economics*, Vol. 4, 323-338.
- Bratley, P., B.L. Fox, and H. Niederreiter. 1992. Implementation and tests of low-discrepancy sequences, *ACM Transactions on Modelling and Computer Simulation*, Vol. 2, No. 3, 195-213.
- Broadie, M., and P. Glasserman. 1993. Estimating Security Price Derivatives by Simulation, PaineWebber Working Paper Series in Money, Economics, and Finance #PW-93-11, Columbia Business School (to appear in *Management Science*).
- Broadie, M., and P. Glasserman. 1995. Pricing American-Style Securities Using Simulation, working paper, Columbia University.
- Carverhill, A., and K. Pang. 1995. Efficient and Flexible Bond Option Valuation in the Heath, Jarrow

- and Morton Framework, working paper, School of Business, Indiana University.
- Clelow, L., and A. Carverhill. 1994. On the Simulation of Contingent Claims, *Journal of Derivatives*, Vol. 2, 66–74.
- Duan, J.-C. 1995. The GARCH Option Pricing Model, *Mathematical Finance*, Vol. 5, 13–32.
- Duffie, D. 1992. *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton, New Jersey.
- Duffie, D., and P. Glynn. 1993. Efficient Monte Carlo Simulation of Security Prices, working paper, Stanford University (to appear in *Annals of Applied Probability*).
- Fu, M.C., and J.-Q. Hu. 1993. Sensitivity Analysis for Monte Carlo Simulation of Option Pricing, working paper, University of Maryland (to appear in *Probability in the Information and Engineering Sciences*).
- Glynn, P.W. 1989. Optimization of Stochastic Systems via Simulation, in *Proceedings of the Winter Simulation Conference*, The Society for Computer Simulation, San Diego, California, 90–105.
- Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*, Kluwer Academic Publishers, Norwell, Mass.
- Grant, D., G. Vora, and D. Weeks. 1994. Path-Dependent Options: Extending the Monte Carlo Simulation Approach, working paper, Anderson School of Management, University of New Mexico, Albuquerque, New Mexico 87131–6471.
- Hull, J. 1993. *Options, Futures, and Other Derivative Securities*, 2nd edition, Prentice-Hall, Englewood Cliffs, New Jersey.
- Hull, J., and A. White. 1987. The Pricing of Options on Assets with Stochastic Volatilities, *Journal of Finance*, Vol. 42, 281–300.
- Joy C., P.P. Boyle, and K.S. Tan. 1995. Quasi-Monte Carlo Methods in Numerical Finance, working paper, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.
- Kemna, A.G.Z., and A.C.F. Vorst. 1990. A Pricing Method for Options Based on Average Asset Values, *Journal of Banking and Finance*, Vol. 14, 113–129.
- Kloeden, P., and E. Platen. 1992. *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, New York.
- L'Ecuyer, P., and G. Perron. 1994. On the Convergence Rates of IPA and FDC Derivative Estimators, *Operation Research*, Vol. 42, 643–656.
- Niederreiter, H. 1992. *Random Number Generation and Quasi-Monte Carlo Methods*, CBMS-NSF, Vol. 63, SIAM, Philadelphia, Pa.
- Nielsen, S. 1994. Importance Sampling in Lattice Pricing Models, working paper, Management Science and Information Systems, University of Texas at Austin.
- Paskov, S.P. 1994. New Methodologies for Valuing Derivatives, working paper, Department of Computer Science, Columbia University, New York, NY 10027.
- Reider, R. 1993. An Efficient Monte Carlo Technique for Pricing Options, working paper, Wharton School, University of Pennsylvania.
- Ritchken, P., and L. Sankarsubramanian. 1995. Volatility Structures of Forward Rates and the Dynamics of the Term Structure, *Mathematical Finance*, Vol. 5, No. 1, 55–72.
- Rubinstein, R., and A. Shapiro. 1993. *Discrete Event Systems*, Wiley, New York.
- Schwartz, E.S., and W.N. Torous. 1989. Prepayment and the Valuation of Mortgage-Backed Securities, *Journal of Finance*, Vol. 44, No. 2, 375–392.
- Spanier, J., and E.H. Maize. 1994. Quasi-Random Methods for Estimating Integrals using Relatively Small Samples, *SIAM Review*, Vol. 36, No. 1, 18–44.
- Tilley, J.A. 1993. Valuing American Options in a Path Simulation Model, *Transactions of the Society of Actuaries*, Vol. 45, 83–104.

AUTHOR BIOGRAPHIES

PHELIM BOYLE holds the J. Page R. Wadsworth Chair in Finance at the University of Waterloo in Waterloo Canada. His research interests span topics in computational finance and actuarial science. He has held visiting positions at several universities throughout the world and is on the editorial board of a number of finance journals. He holds a Ph.D. from Trinity College Dublin and is a fellow of the Institute of Actuaries.

MARK BROADIE joined the faculty of the Columbia University Graduate School of Business in 1983. His current research interests include portfolio optimization and the pricing of derivative securities. He holds a Ph.D. from Stanford University and a B.S. from Cornell University.

PAUL GLASSERMAN joined the faculty of the Columbia University Graduate School of Business in 1991. Prior to that, he was a Member of Technical Staff in the Operations Research department of AT&T Bell Laboratories in Holmdel, NJ. He holds a Ph.D. from Harvard University and an A.B. from Princeton University.