Quasi-Monte Carlo Approaches to Option Pricing

John R. Birge*
Department of Industrial and Operations Engineering
University of Michigan
Ann Arbor, MI 48109
(313)764-9422

Abstract: Complications such as varying interest rates and complex contingencies can prohibit analytical computation of option and other derivative prices. A typical approach is to rely on Monte Carlo simulation in these cases and to use statistical properties of assumed random sequences. This approach has two drawbacks. First, the validity of a pseudo-random sequence for statistical randomness is somewhat questionable. Second, even assuming randomness, the error from central limit arguments decreases slowly, inversely proportional to the square root of the number of observations. Quasi-Monte Carlo sequences, on the other hand, do not rely on a randomness assumption and have order of magnitude better asymptotic error rates. We show how quasi-Monte Carlo sequences can be used in option pricing, give some analytical justfication for their preference over crude or pseudo-Monte Carlo methods, and present some empirical evidence based on simple call option models.

Keywords: simulation, option pricing, quasi-random

1. Introduction

The landmark paper of Black and Scholes [1973] spawned many other approaches to option pricing and to the evaluation of general contingent claims (see, e.g., Jarrow and Rudd [1983]). Many of these approaches try to find simple analytical calculations as in the Black-Scholes formula. When the contingencies become complex, interest rates vary, or other simplifying assumptions are invalid, analytical formulas may either contain excessive errors or may not be available. In these cases, Monte Carlo simulation can be used to represent many random phenomena that can influence the option price (see, e.g., Boyle [1977]).

* Based on work supported by the National Science Foundation under Grant ECS 88-15101 and 92-15921. The author thanks Sanjeev Acharya and Jim Wartinbee for assistance in the computational implementations.

The Monte Carlo method (see Hammersley and Handscomb [1964]) has its basis in statistical theory, in particular, the central limit theorem that allows the calculation of confidence intervals about the mean value of a random variable based on a sequence of independent random observations. Practically, the validity of this approach depends on the ability to generate sequences of random numbers. Thus, random number generation is a critical element in using the Monte Carlo method (see Deák [1990]).

The general practice is to use *pseudo-random sequences* that share properties with random sequences, such as frequency of increasing or decreasing runs of various lengths. These pseudo-random sequences are constructed from a variety of techniques involving simple operations on previous numbers in the sequence. Their use and statistical justification in the Monte Carlo method has, however, been questioned (see, e.g., Zaremba [1968]). Even assuming the statistical basis, the error in Monte Carlo estimates decreases at an asymptotic rate that is inversely proportional to the square root of the number of observations, $O(\frac{1}{\sqrt{N}})$.

An alternative is to rely on numerical properties of the sequence and on how closely a sequence fills in each region in the support of the random variables. This measure, called the *discrepancy*, can be reduced by constructing sequences that more deliberately fill in the random support. These sequences are often called *quasi-random* numbers or *low-discrepancy* sequences. Their use in approximating integrals or expectations is called the *quasi-Monte Carlo* method (see Niederreiter [1978] for a survey and comparison with pseudo-random numbers).

The great advantage of quasi-random sequences is that the asymptotic error generally reduces to $O(\frac{1}{N})$. The difficulty is, however, that the constants that can be calculated for this order represent a worst case and may be quite large. In practice, however, the error may be much less. We give some indication of this below.

Calculating the error for any individual example may be difficult but another recent discovery gives even greater potential for the use of quasi-Monte Carlo sequences. Woźniakowski [1991] has shown that the expected error of using particular quasi-Monte Carlo sequences is approximately $O(\frac{1}{N})$ for integrating over a general class of continuous functions in many dimensions. This observation effectively suggests, on average, crude Monte Carlo requires the square of the number of observations required by quasi-Monte Carlo to achieve the same error. This conquering of the *curse of dimensionality* (at least in expectation) has led to renewed interest in improved numerical integration and simulation (see Yam [1991], Guterl [1994]).

We suggest some direct uses of this observation below, although it does not directly apply to simple options which correspond only to a small subset of all functions considered in Woźniakowski's result. The general asymptotic error bounds in quasi-Monte Carlo methods do, however, indicate that they can improve on option pricing with pseudo-Monte Carlo. This paper presents both analytical and empirical evidence for this claim.

In Section 2, we present our basic approaches for option pricing, which we interpret as integration over a region with dimension equal to the number of distinct periods of price changes. Section 3 gives some analytical comparisons of quasi- and pseudo-random sequence approaches. Section 4 presents our empirical observations. Section 5 gives conclusions.

2. Monte Carlo Option Pricing

The Monte Carlo method generally refers to procedures for approximating the integral, $I = \int_R f(x)dx$, of a function, f, over a region, R, where $f : \Re^n \to \Re$, and $R \subset \Re^n$. The basic idea is to choose a sequence of points, $\{x^k\}$, $i = 1, \ldots, N$, uniformly distributed in R and to estimate I by

$$\hat{I}(N) = \frac{1}{N} \sum_{k=1}^{N} f(x^k). \tag{1}$$

If the sequence $\{x^k\}$ indeed represents independent random observations, then the $\{f(x^k)\}$ values can also be considered observations of independent random variables with mean, I. To distinguish the random variables from their observations we use bold-face notation, e.g., $\mathbf{x}^i$.

Assuming I and a finite second moment, $N\hat{\mathbf{I}}(N)$ can be interpreted as the sum of independent random variables. The central limit theorem can be invoked to determine the asymptotic normal distribution of $\hat{\mathbf{I}}(N)$ and to obtain a confidence interval on I using an observation of $\hat{\mathbf{I}}(N)$ (for derivations, see, for example, Feller [1971]).

The fundamental result is that $\hat{\mathbf{I}}(N)$ is asymptotically normally distributed with mean I and variance $\frac{1}{N}V_f$, where $V_f = \int_R \|f(x)\|^2 dx - I^2$, the variance of $f(\mathbf{x})$. Thus, $\sqrt{N}\frac{(\hat{\mathbf{I}}(N)-I)}{\sqrt{V_f}}$ is asymptotically a standard normal with zero mean and unit variance. If Φ denotes the distribution function of the standard normal, then a level α confidence interval can be obtained from an observation $\hat{I}(N)$ of $\hat{\mathbf{I}}(N)$ by noting, $\alpha = P[\Phi^{-1}(\frac{1-\alpha}{2}) \leq \sqrt{N}\frac{\hat{\mathbf{I}}(N)-I}{\sqrt{V_f}} \leq \Phi^{-1}(1-\frac{1-\alpha}{2})]$, so that with probability α,

$$I \in [\hat{I}(N) - \frac{\sqrt{V_f}}{\sqrt{N}}\Phi^{-1}(\frac{1-\alpha}{2}), \hat{I}(N) + \frac{\sqrt{V_f}}{\sqrt{N}}\Phi^{-1}(\frac{1-\alpha}{2})]. \tag{2}$$

The two primary difficulties in using this method are that deterministic sequences generally replace the random variables in $\hat{\mathbf{I}}(N)$ and that the interval in (2) only changes proportionally to $\frac{1}{\sqrt{N}}$. The deterministic sequences are pseudo-random sequences that pass tests for randomness, but then the applicability of (2) is rather tenuous.

A further implementation difficulty is that single pseudo-random streams are frequently used to obtain each component of the n-dimensional variable x. The result may be that a sequence that appears quite uniform in one dimension may be distributed with substantial collinearity (or lattice structure) when applied sequentially for higher dimensions. In these cases, generation of more points simply fills in each line (or hyperplane in higher dimensions). The error can never improve beyond the error in integrating along the lines. This error may be especially high if the integrand concerns the sum of the components as in option pricing.

To overcome the difficulty of collinearity, pseudo-random vector generation schemes are possible (see, for example, Niederreiter [1991]). Another option is to use quasi-random numbers that explicitly consider the full dimension of the problem. These sequences avoid the randomness assumption and attempt to achieve low errors based purely on numerical (not statistical) analysis. In the next section, we give some of the analytical error properties of these sequences.

Assuming streams of pseudo- or quasi-random vectors, $\hat{I}(N)$ can be constructed by simply evaluating f at each point. We will assume that R is the unit hypercube, $[0,1]^n$, in $\Re^n$, and that f is constructed so the components of $\mathbf{x}$ are independently distributed. In this way, we need only generate uniform points in the unit hypercube and then compute f.

To illustrate the approach, we assume a simple model of a call option (without early exercise) and find the appropriate f. We suppose that n is the number of periods into which the time to maturity is divided. In practice, the size of periods might vary and their endpoints might correspond to potential interest rate changes, dividend dates, or possible exercise points in an American option. We could include correlations among periods, but, for our illustrations, we assume independence of price changes across periods.

Now, assuming that stock returns have a lognormal distribution, we can simply invert the uniform components of x to obtain n normally distributed random values, sum these values, and exponentiate to obtain the price appreciation after n periods. Using the standard risk neutrality argument (see, e.g., Cox and Ross [1976]), the call price can then be discounted by a riskless bond rate to achieve the call option price.

To state this more formally, assume that $S_T(\mathbf{x})$ is the random stock price at maturity date, T, C_t is the call price at time t, K is the strike price, and r is a riskless return (assumed constant for this discussion). Making the additional assumptions of frictionless markets and no early exercise, we can write C_t as

$$C_t = e^{-r(T-t)}(\int_{[0,1]^n} (S_T(x) - K)^+ dx). \tag{3}$$

Hence, in (1), we have $f(x) = e^{-r(T-t)}(S_T(x) - K)^+$. We just need to define S_T. In general, we might consider many independent factors that all combine to effect the price at time T. In the simplest case, each x_i corresponds to the quantile of the proportional price change in period i where price changes, $(\mathbf{S}_{t_i}/\mathbf{S}_{t_{i-1}})$, $i = 1, \ldots, n$, among periods are independent. In this way, we let $\delta_i(\mathbf{x}_i) = (\mathbf{S}_{t_i}/\mathbf{S}_{t_{i-1}})$ and then

$$S_T(\mathbf{x}_1, \ldots, \mathbf{x}_n) = S_t \Pi_{i=1}^n \delta_i(\mathbf{x}_i). \tag{4}$$

We also assume for ease of exposition that the proportional price changes are lognormally distributed with mean return $(r - \frac{\sigma^2}{2})(t_{i+1} - t_i)$ and standard deviation $\sigma\sqrt{t_{i+1} - t_i}$, so that

$$\frac{log(\delta_i(x_i)) - (r - \frac{\sigma^2}{2})(t_{i+1} - t_i)}{\sigma\sqrt{t_{i+1} - t_i}} = \Phi^{-1}(x_i). \tag{5}$$

Combining (3), (4) and (5), we have

$$C_t = e^{-r(T-t)}(\int_{[0,1]^n} (S_t(\Pi_{i=1}^n e^{\Phi^{-1}(x_i)\sigma\sqrt{t_{i+1}-t_i} + (r - \frac{\sigma^2}{2})(t_{i+1}-t_i)}) - K)^+ dx). \tag{6}$$

Now, Monte Carlo or quasi-Monte Carlo approaches to evaluating (6) only require finding sequences of points, $x^k, k = 1, \ldots, N$, as in (1) and using the error results that accompany these sequences. We could use more complicated integrands than those in (6) but we choose the simple form above because analytical results are so readily available. Since we assume the price changes in each period are independent and lognormally distributed, their product, S_T, is also lognormally distributed with mean return, $(r - \frac{\sigma^2}{2})(T-t)$, and standard deviation, $\sigma\sqrt{T-t}$. This observation leads to the Black-Scholes' formula for

$$C_t = S_t \Phi\left(\frac{log(S_t/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right)$$

$$-Ke^{-r(T-t)}\Phi\left(\left(\frac{log(S_t/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right) - \sigma\sqrt{T-t}\right). \tag{7}$$

Of course, we do not need to use Monte Carlo approaches to evaluate (6) if (7) or other analytical representations are available. However, when more complicated types of integrals are included, then such straightforward analytical solutions may not be available. We use the simple form here to allow us to evaluate the accuracy of pseudo- and quasi-random schemes against a known analytical value. In the next section, we present some analytical results about the accuracy of these methods.

3. Analytical Error Analysis of Quasi-Monte Carlo Schemes

The main basis for error analysis of pseudo-random schemes is the use of the confidence interval in (2). In the option pricing case presented here, the error in this estimate is proportional to the standard deviation of the integrand divided by $\sqrt{N}$. If this quantity is unknown, a sample variance can be used to estimate it with a t-distribution in place of the normal. The result allows asymptotic error bounds with confidence level α (assuming true randomness). The $\frac{1}{\sqrt{N}}$ proportionality in the error decrease means, for example, that achieving a confidence interval that is $1/10$ of a standard deviation requires 100 observations but that achieving an error of $1/100$ of a standard deviation requires $10,000$ observations (assuming the asymptotics apply).

In contrast to this slow error decrease, quasi-random numbers may provide better asymptotic error bounds that are proportional to $\frac{1}{N}$ instead of $\frac{1}{\sqrt{N}}$. We will consider two such sequences here. The basic motivation for them is to achieve low *discrepancy*, the difference between the fraction of sequence points within any region of the unit hypercube and the volume of that region. The discrepancy should decrease to zero.

One example of a quasi-random sequence is constructed from irrationals, such as the roots of the first n primes. We will denote this sequence as $\{\alpha^k\}$, where component i of the kth sample is

$$\alpha^k(i) = k\sqrt{P(i)} - \lfloor k\sqrt{P(i)} \rfloor, \tag{8}$$

where $P(i)$ is the ith prime number, $i = 1, \ldots, n$, and $\lfloor k\sqrt{P(i)} \rfloor$ denotes the greatest integer less than or equal $k\sqrt{P(i)}$. Hammersley and Handscomb [1968] show the following (which we present without proof).

Theorem 1. *If the Fourier series of f is absolutely convergent, then there exists a constant c_1 such that:*

$$(\int_{[0,1]^n} f(x)dx - \frac{1}{N}\sum_{k=1}^{N} f(\alpha^k)) \leq c_1 \frac{1 + logN}{N}. \tag{9}$$

Thus, there exists a constant c_1 for any f with these properties that achieves an error bound with probability one that is proportional to $\frac{1+logN}{N}$. In contrast to the straight Monte Carlo approach, the growth in required iterations to achieve given accuracy is much less than quadratic. For example, to reduce the error bound in (9) to $1/10$ of c_1 requires 49 points while reducing the error bound to $1/100$ of c_1 requires 647 points.

While these analytical results for the quasi-random sequence appear quite good, they depend on the constant c_1 which is not easily estimated. They do, however, indicate that asymptotic results may be much better with quasi-random than pseudo-random sequences.

Another sequence we shall use is called the *Halton sequence* (Halton [1960]) based on representing integers in terms of primes. We suppose that k has a representation in base $P(i)$ as $k = a_1 P(i)^{k-1} + a_2 P(i)^{k-2} + \cdots + a_k$, and let $\phi^k(i) = \sum_{i=1}^{k} a_{k-i+1} P(i)^{-i}$. The sequence is then $\phi^k = (\phi^k(1), \ldots, \phi^k(n))$. Halton shows that this sequence also has low discrepancy. Using a standard result on functions of bounded variation, we can state this as the following.

Theorem 2. *If f is of bounded variation over $[0,1]^n$, then there exists a constant c_2 such that:*

$$(\int_{[0,1]^n} f(x)dx - \frac{1}{N}\sum_{k=1}^{N} f(\phi^k)) \leq c_2 \frac{(logN)^n}{N}. \tag{10}$$

Comparisons between Theorems 1 and 2 are difficult again because of the difficulty in estimating the constants. In general, c_2 is much less than c_1. Another advantage of the Halton sequence may be in computation since $\{\alpha^k\}$ constructions require many digits of $\sqrt{P(i)}$ when k is high. With these caveats, we can still find the worst-case reductions of error from c_2 in (10). For three ($n = 3$) dimensions, with the Halton sequence, reducing the error bound to 1/10 of c_2 requires 6910 points while reducing the error bound to 1/100 of c_2 requires $176,000$ points. A somewhat better comparison is that the additional digit of accuracy requires 25 times more samples compared to 13 times as many samples using the $\{\alpha^k\}$ sequence and 100 times with standard Monte Carlo.

We use the Halton sequence given above because it does not a priori require the total number of observations N. Knowing N (or fixing N at certain values) can reduce the error order to $\frac{(logN)^{n-1}}{N}$ using, for example, the Hammersley sequence (see Hammersley and Handscomb [1964]). We use methods for variable N, however, to observe the errors directly as N increases.

Another sequence, proposed by Faure [1982], has demonstrated smaller errors for the same number of observations (Fox [1986]) with the same asymptotic error rate as in (10) but with generally additional computation time. In our computational comparison, we use this sequence as well. We also include a sequence due to Sobol [1967] (see also Antonov and Saleev [1979]) with the same asymptotic rate and different practical computational characteristics.

The Hammersley points have found further justification in a construction by Woźniakowski [1991]. He shows that shifting the Hammersley points leads to a sequence of points that achieves the best average case error for numerical integration over continuous real functions on the unit hypercube equipped with the Wiener sheet measure (the mean over all functions at any point is zero and the covariance between values at two points is the product of minimum coordinate values of the points). Woźniakowski shows that to achieve an error of ϵ requires, on average for functions in this class, $O((log\epsilon)^{-(\frac{n-1}{2})}\epsilon^{-1})$ function evaluations. This contrasts with Monte Carlo's requiring $O(\epsilon^{-2})$ observations and grid points which require $O(\epsilon^{-n})$ evaluations.

The average case asymptotic advantage over Monte Carlo and any grid point scheme has generated much speculation about the potential improvements in integration over Monte Carlo methods. The average case result would not extend directly to the simple option example we consider here but, for path determined derivatives (e.g, options on weighted average values of securities over a fixed time horizon or CMOs (see Paskov and Traub [1995] for an implementation)), the class of functions may fit quite closely (each component in determining the price over time is a Wiener process). This theoretical basis remains of interest, but the results are only for average cases and, again, asymptotic constants may be large. Our goal in the next section is to show that the error of quasi-Monte Carlo may indeed be much better than these analytical results for functions that arise in option pricing.

4. Empirical Observations

The results in Section 3 provide some motivation for the use of quasi-random points over pseudo-random points in evaluating integrals such as the option price in (3). A difficulty in using the analytical results with this model is, however, the estimation of the constants in the error analyses. Indeed, the function with the inverse normal integrand is not even of bounded variation. Nevertheless, the low discrepancy of the points in $\{\alpha^k\}$ and $\{\phi^k\}$ leads to low errors in calculating these integrals.

All computer runs in this section were performed using the FORTRAN compiler on an IBM RS6000 workstation at the University of Michigan. For a pseudo-random number generator, we used the portable FORTRAN generator by Schrage [1979] because of its general acceptance (see, e.g., Law and Kelton [1991]), ease of use, and efficiency. For the option pricing model, we tested ten, thirty, and one hundred eighty period models. We used identical distributions in each period to make the analytical calculations easier. The initial price (S_t) and strike price (K) were first both set to 40 and then K was varied to 35 and 45. The annual riskfree interest rate was $r = 0.1$ with annual volatilities of $\sigma = 0.1, 0.3$ and 0.6.

The three horizons, $T - t = 10$, $T - t = 30$ and $T - t = 180$ days, were chosen to give some indication of changes from relatively low to high dimensions. The analytical values from (7) range from $C_t(T - t = 10, \sigma = 0.1) = 0.322$ to $C_t(T - t = 180, \sigma = 0.6) = 7.52$. Figures 1-13 give indications of the relative accuracy merit of each sequence scheme. We then discuss the computational burden of each method and its relevance in practice.

Figures 1 to 3 give the relative error in percentage for the ten, thirty, and one hundred eighty period models, $S_t = K = 40$, and $\sigma = 0.3$ for each of the pseudo-random sequence (*pseudo*), the sequence using prime roots (*alpha*, $\{\alpha_k\}$) plus the Halton, Faure and Sobol sequences, as the number of samples, N, increases. In Figure 1 ($T - t$=10), Sobol has consistently the lowest error until Faure reaches the same level at around 28,000 iterations. The *alpha* sequence error is also below 0.5% after 5,000 iterations. Halton shows steady improvement, but *pseudo* displays significant oscillation.

In Figure 2, *alpha* performs best up to 28,500 iterations, where Faure then becomes consistently the most accurate. Sobol's error is quite low but shows some oscillation. Halton is again consistent but not very accurate, while *pseudo* suffers once more from oscillation particularly at low iteration numbers.

In Figure 3, Halton and Faure are not present because their errors never reached 2% or less. The *alpha* sequence generally has the lowest errors. The *pseudo* sequence has significant oscillation while Sobol has less pronounced oscillation.

Some general observations are that the *alpha* and Sobol sequences achieve relatively low errors quite quickly. The pseudo-random sequence shows significant oscillation in errors. The Halton and Faure sequences are the most consistent in error improvement but they are not competitive in the higher dimensions in these trials.

Figures 4 and 5 give the option values obtained for $T - t = 180, \sigma = 0.3$ under each sequence to illustrate the performance of Halton and Faure. Note that the values from these sequences are converging at relatively slow rates from values far from the actual. Figure 5 shows that *pseudo* and Sobol are biased somewhat above the true value while *alpha* obtains values below the true value. (This tendency is not, however, uniformly true but does indicate that other variance reduction techniques, such as antithetic variates, may be useful.)

The effect of varying volatilities is explored in Figures 6 ($\sigma = 0.1$) and 7 ($\sigma = 0.6$) for $T - t = 180$. In general, the results here are quite similar to those for $\sigma = 0.3$ except that errors are higher for greater volatilities. Again, Halton and Faure never achieve low errors in 100,000 iterations, *pseudo* has significant oscillations, and *alpha* produces slightly more consistent values than Sobol.

Figures 8 and 9 demonstrate the effect of varying the strike price at out-of-the-money or in-the-money values with $K = 35$ and $K = 45$. The results are again quite similar to the previous examples for $T - t = 180$. The Sobol sequence, however, appears slightly more consistent than *alpha* for $K = 45$ and greater than 70,000 iterations.

The advantage of the quasi-Monte Carlo procedures depends on terminating a simulation with an accurate result earlier than the pseudo-Monte Carlo method. The pseudo-Monte Carlo method's benefit depends on the extent to which one can apply the confidence interval as in (2) to set a priori run lengths for given desired errors. As we saw above, it is not clear how (2) is interpreted with this deterministic sequence. In our case, (2) is often violated for the sample chosen. The meaning of the α-level confidence is not clear with the fixed sequence. The use of sample variances may also lead to significant errors when variances are large but not well detected in the standard Monte Carlo scheme (see, e.g., Fox [1986])).

We consider various measure of termination efficiency in Figures 10-13. These results give the ratio of the stopping times from the quasi-Monte Carlo procedures to the time for *pseudo*. Each trial's result (given as $T-t$, σ, K (if not 40)) is displayed. Figure 10 shows this result for one alternative, suggested in Fox [1986], to compare values at $2N$ with values at N observations and to stop if these values are within some relative tolerance (0.5% in Figure 10) of each other. Note that *alpha* terminates within half the time of *pseudo* in all but the 180, 0.6 case. Sobol has somewhat less consistency but terminates earlier than *pseudo* in all but the $K = 45$ case. Halton terminates earlier than *pseudo* for $T-t = 10$ and $T-t = 30$ but has difficulties for $T-t = 180$ as mentioned above. Faure also has difficulties for $T-t = 30$ in addition to the 180-day trials in this test.

Another approach, similar to statistical tests based on mean and range, is to consider the range of observations over some interval compared to the overall mean value. We consider this measure for an interval of 5000 observations and an error of 0.5% in Figure 11. In this test, *alpha* and Sobol consistently require fewer observations than *pseudo* with *alpha* approaching termination as much as ten times faster than *pseudo*. Halton and Faure again terminate earlier than *pseudo* for ten and thirty days but not for 180-day models.

To consider other stopping conditions, we consider the last iteration with errors above some tolerance in Figure 12 for an error of 1% and Figure 13 for an error of 0.5%. For the 1% error, *alpha* achieves this result at least twice as fast as *pseudo* in every case and is ten times faster in several cases. Sobol terminates in this approach faster than *pseudo* in each case and is almost 50 times faster in the ten-day models. Halton and Faure again have difficulties for the higher dimensions.

For the 0.5% error in Figure 13, *alpha* again achieves earlier termination than *pseudo* but is not as consistent as with the larger error. Sobol is, however, more consistent in improvements relative to *pseudo* in this case. The results for Halton and Faure are as in the other tests.

A possible advantage of the pseudo-random scheme is in efficiency of calculation time. In general, we would anticipate the use of these techniques only in problems which are too complicated for analytical solutions and, in which, each function evaluation is much more costly than the random number generation time. For this simple example, function evaluation was quite fast so the times vary slightly. The results for ratios of sequence generation times to the *pseudo* generation times appear in Table 1 below.

Method	**Ratio to Pseudo**
Alpha	1.2
Halton	2.5
Faure	$5-10$
Sobol	$9-10$

Table 1. Ratio of computation times over all trials to pseudo-random generation times.

Our results are consistent with Fox [1986] and Bratley and Fox [1988]. In general, they report Halton times approximately twice their *pseudo* times and Faure and Sobol times approximately 5-6 times the values for *pseudo*. Bratley and Fox [1988], however, report much improved (by an order of magnitude) timings for Sobol with a nonstandard XOR operator in the compiler.

Since many derivative evaluations would generally require much greater computation times than number generation times, the results in Table 1 have less relevance. Overall, the results here indicate that the *alpha* sequence may have the most consistency for 1% accuracy. Sobol may have greater consistency for smaller error values, but this accuracy should be weighed against computation time. Faure in these trials is quite accurate in low dimensions but not effective in higher dimensions. Halton does not appear as effective as the other methods.

These results are consistent with Paskov [1994] and Paskov and Traub [1995] who applied the Sobol and Halton sequences to CMO evaluations. In their studies, Sobol achieves greater accuracy than pseudo-random schemes as well as the Halton sequence.

5. Conclusions

We have presented quasi-Monte Carlo approaches to option pricing. We demonstrated some of their analytical advantages over the use of pseudo-random streams in standard

Monte Carlo. We also gave some empirical evidence that quasi-random streams can produce more accurate results in fewer iterations than standard Monte Carlo. Quasi-Monte Carlo also offers the advantage of less fluctuation than standard Monte Carlo and makes on-line error approximations possible.

The analyses in this paper concentrated on crude pseudo-Monte Carlo and quasi-Monte Carlo schemes. In practice, variance reduction techniques are used to lower the error in either crude scheme. Additional variance reduction techniques should have roughly equivalent effects on pseudo- or quasi-Monte Carlo approaches (see, for example, Fox [1986]). In the quasi-Monte Carlo case, however, the reduction in error is caused by reducing discrepancy rather than variance.

These results show that quasi-Monte Carlo may have an advantage in a simple model where analytical results are available. The advantage may be even greater in more complex models (such as the CMO evaluations in Paskov and Traub [1995]) which require substantial computation for each function evaluation. They may also be applicable to evaluating path-based options due to their excellent average case behavior which is an order of magnitude better than standard Monte Carlo.

References

Antoneev, I.A. and V.M. Saleev. "An economic method ofc computing LP_τ-sequences." *USSR Computational Mathematics and Mathematical Physics*, 19 (1979), 252-256.

Black, F., and M. Scholes. "The pricing of options and corporate liabilities." *Journal of Political Economics*, 81 (1973), 637-659.

Boyle, P. "Options: A Monte Carlo approach." *Journal of Finance*, 32 (1977), 323-338.

Bratley, P., and B.L. Fox. "Algorithm 659, Implementing Sobol's quasirandom sequence generator." *ACM Transactions on Mathematical Software*, 14 (1988), 88-100.

Cox, J., and S. Ross. "The valuation of options for altenative stochastic processes." *Journal of Financial Economics*, 3 (1976), 145-166.

I. Deák. *Random Number Generators and Simulation*, Akadémiai Kiadó, Budapest (1990).

Faure, H. "Discrépance de suites associées à un système de numération (en dimension s)." *Acta Arithmetica*, XLI (1982), 337-351.

Feller, W. *An Introduction to Probability Theory and Its Applications, Volume II*, New York:Wiley (1971).

Fox, B. L. "Implementation and relative efficiency of quasirandom sequence generators." *ACM Trans. Math. Software*, 12 (1986), 362-376.

Guterl, F. "Suddenly, number theory makes sense to industry." *Business Week*, June 20 (1994), 172-174.

Halton, J. H. "On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. " *Numer. Math.* , 2 (1960), 84-90.

Hammersley, J. M., and D.C. Handscomb. *Monte Carlo Methods*, London:Methuen (1964).

Jarrow, R. A., and A. Rudd. *Option Pricing*, Homewood, Illinois: Irwin (1983).

Law, A. M., and W. D. Kelton. *Simulation Modeling and Analysis*, New York:McGray-Hill (1991).

Niederreiter, H. "Quasi-Monte Carlo methods and pseudorandom numbers." *Bull. Amer. Math. Soc.*, 84 (1978), 957-1041.

Niederreiter, H. "Recent trends in random number and random vector generation." *Annals of Operations Research*, 30 (1991), 323-346.

Paskov, S.H. "New methodologies for valuing derivatives." Technical Report, Department of Computer Science, Columbia University, New York (October 1994).

Paskov, S.H., and J. Traub, "Faster valuation of financial derivatives." Technical Report, Department of Computer Science, Columbia University, New York (January 1995).

Schrage, L. "A more portable random number generator." *ACM Trans. Math. Software*, 5 (1979), 132-138.

Sobol, I.M. "On the distribution of points in a cube and the approximate evaluation of integrals." *USSR Computational Mathematics and Mathematical Physics*, 7 (1967), 236-242.

Woźniakowski, H. "Average-case complexity of multivariate integration." *Bull. Amer. Math. Soc.* (new series), 24 (1991), 185-194.

Yam, P. "Math Exorcise." *Scientific American*, July 1991, 29.

Zaremba, S.K. "The mathematical basis of Monte Carlo and quasi-Monte Carlo methods." *SIAM review*, 10 (1968), 303-314.