

Quasi-Monte Carlo Methods in Numerical Finance

Corwin Joy • Phelim P. Boyle • Ken Seng Tan

Enron Corporation, 1400 Smith Street, Houston, Texas 77002-7361

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

Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario N2L 3G1 Canada

This paper introduces and illustrates a new version of the Monte Carlo method that has attractive properties for the numerical valuation of derivatives. The traditional Monte Carlo method has proven to be a powerful and flexible tool for many types of derivatives calculations. Under the conventional approach pseudo-random numbers are used to evaluate the expression of interest. Unfortunately, the use of pseudo-random numbers yields an error bound that is probabilistic which can be a disadvantage. Another drawback of the standard approach is that many simulations may be required to obtain a high level of accuracy. There are several ways to improve the convergence of the standard method. This paper suggests a new approach which promises to be very useful for applications in finance. Quasi-Monte Carlo methods use sequences that are deterministic instead of random. These sequences improve convergence and give rise to deterministic error bounds. The method is explained and illustrated with several examples. These examples include complex derivatives such as basket options, Asian options, and energy swaps.

(Monte Carlo Simulation; Quasi-random Sequences; Faure Sequences; Numerical Finance; Derivative Valuation)

1. Introduction

Numerical methods have come to play an increasingly important role in modern quantitative finance. This is due to the fact that most finance models have analytical solutions in only a few special cases. In many cases the price of a given financial instrument can be assumed to evolve according to a second order partial differential equation. Different techniques are used to solve this equation and compute the sensitivities of the price to different variables. One approach is to discretize the differential equation and solve the resulting set of difference equations. Another popular technique is to use a lattice or tree approximation. Monte Carlo methods are often used when these methods are difficult to implement due to the complexity of the problem.

In recent years there has been renewed interest in Monte Carlo methods. This is due to the flexibility of

the method in handling increasingly complex¹ financial instruments and the advent of the powerful workstations has significantly reduced the execution time required for the simulation. For example, Monte Carlo algorithms are normally more efficient than competing procedures if there are several state variables or other complexities. In addition the relevant algorithms are often fairly straightforward to design. Boyle (1977) provides an early discussion of the use of Monte Carlo methods in pricing options. The disadvantages² of Monte Carlo methods are that the error term is proba-

¹ Many of the new "exotic" securities involve several underlying assets, different currencies and path dependency.

² Another disadvantage is that the method still has difficulties in handling the early exercise feature of American options despite notable progress by Tilley (1993) and Grand, Vora and Weeks (1993).

bilistic and that it can be computationally burdensome to achieve a high level of accuracy. To reduce the computational burden of standard (or crude) Monte Carlo methods a variety of variance reduction methods have been proposed. There are many such techniques and they include control variates, antithetic variates, stratified sampling and importance sampling.

The purpose of this paper is to introduce another technique for improving the efficiency of the Monte Carlo method. This method known as the quasi-Monte Carlo method and was pioneered in computational physics where large scale Monte Carlo calculations are common. The basic idea is to use numbers that are deterministic rather than random. These deterministic sequences can be used in place of the usual random (or pseudo-random) numbers to obtain faster convergence with known error bounds. The use of quasi-Monte Carlo methods leads to more efficient numerical procedures with the added benefit of deterministic error bounds. The monograph by Niederreiter (1992) gives a good survey of quasi-Monte Carlo methods.

In this paper we introduce quasi-Monte Carlo methods and discuss different applications in numerical finance. In §2 we introduce the problem and motivate the method by analyzing a simple European call. The price of this call can be expressed in terms of an integral and we discuss the connection between numerical integration and standard Monte Carlo approaches. In §3 we discuss the procedure for valuing complex derivative securities such as Asian options using the Monte Carlo approach and we examine ways to improve the convergence. In §4 we introduce a particular type of quasi random sequences known as Faure sequences. These sequences have attractive properties for many applications in finance. We show that they provide swifter convergence and better error bounds than standard Monte Carlo methods in several practical settings. In later sections we use Faure sequences to value a range of complex derivative securities. We demonstrate that quasi-Monte Carlo methods are more efficient for obtaining prices and price sensitivities than standard Monte Carlo methods.

2. Use of Deterministic Sequences to Value Standard Options

We begin by illustrating the use of deterministic sequences to value a standard European option. The op-

tion price is written as a discounted expectation over the terminal price distribution. We discuss the numerical valuation of this integral and its relationship to the standard Monte Carlo procedure. Then we show how a judicious choice of the "random" numbers can improve the accuracy of the procedure. The main idea is that since the use of random numbers introduces error we can do better by preselecting the numbers in a deterministic fashion.

The current price of a European derivative security can be written as its discounted expectation under the equivalent martingale (or Q measure)

$$\text{Price} = e^{-rT} E_Q[g(A(T))] \quad (1)$$

where T is the time until maturity, $A(T)$ is the asset price at time T , $g(\cdot)$ is the payoff function, and r is the riskless rate.

Assume that the asset price dynamics follow geometric Brownian motion with constant drift, μ and volatility, σ . The current price of the derivative security is obtained by integrating the terminal payoff under the risk neutral measure. Hence we have

$$\begin{aligned} \text{Price} &= e^{-rT} \int_{-\infty}^{\infty} g[A(0)e^{((r-(\sigma^2/2))T+\sigma\epsilon\sqrt{T})}] \frac{1}{\sqrt{2\pi}} e^{-\epsilon^2/2} d\epsilon \\ &= \int_{-\infty}^{\infty} h(\epsilon) \frac{1}{\sqrt{2\pi}} e^{-\epsilon^2/2} d\epsilon \\ &= \int_0^1 h(\Phi^{-1}(x)) dx \\ &= \int_0^1 f(x) dx. \end{aligned} \quad (2)$$

In the first line of the above equations the argument of g (the term in square brackets) denotes the terminal distribution of the underlying asset. Under our assumptions this asset has a lognormal distribution which we can write in terms of the exponential of an appropriate normal distribution. Thus the variable ϵ has a normal distribution and so the final term is the standard normal density function. In the second line we have used the function h to make the expression more compact. The critical step is in the third line where we map the normal variate ϵ into the interval $[0, 1]$ through the inverse transformation.

We transform the integral of h from minus infinity to plus infinity to the integral of $f(x)$ over the uniform distribution with range $[0, 1]$. This simplification permits us to obtain efficient approximations to the integral by carefully selecting sample points in $[0, 1]$. This technique leads to more efficient approximations of the original integral.

We illustrate the standard Monte Carlo method using the same example. First we generate n unit normal variates; $\epsilon_1, \epsilon_2, \dots, \epsilon_N$. The crude Monte Carlo estimate of the option price is:

$$\frac{1}{N} \sum_{i=1}^N h(\epsilon_i) \quad \text{where } \epsilon_i \sim N(0, 1). \quad (3)$$

This estimate has an error of $O(N^{-1/2})$. However, if we were to use a more uniformly spaced grid by choosing the N evenly spaced quantiles of $N(0, 1)$, we would obtain the N panel method (or midpoint rule) for evaluating

$$\int_0^1 f(x) dx. \quad (4)$$

This has an error of $O(N^{-1})$. For example, if $N = 4$ and we wish to use the four panel method we split the interval $[0, 1]$ into four equal subintervals

$$[0, \frac{1}{4}], [\frac{1}{4}, \frac{2}{4}], [\frac{2}{4}, \frac{3}{4}], [\frac{3}{4}, 1]$$

and then use the midpoint of each subinterval to evaluate the integral. This would be equivalent to using the following values of ϵ in the Monte Carlo routine

$$\begin{aligned} \bar{\epsilon} &= \{\Phi^{-1}[\frac{1}{8}], \Phi^{-1}[\frac{1}{4} + \frac{1}{8}], \Phi^{-1}[\frac{2}{4} + \frac{1}{8}], \Phi^{-1}[\frac{3}{4} + \frac{1}{8}]\} \\ &= [-1.15035, -0.31864, 0.31864, 1.15035]. \end{aligned}$$

One difficulty with the panel method is that we may not know in advance how many simulations will be needed to obtain the desired level of accuracy. In fact, in actual applications we may want to continue a simulation until we reach a given level of precision. In this situation we need a sequence of numbers that is evenly spaced and that somehow "knows how" to fill in the gaps that have been left by previous elements in the sequence. Sequences with this property are known as low discrepancy³ sequences. We can think of the discrepancy as a

quantitative measure of the deviation of the set of numbers from a uniform distribution. As a crude example of such a sequence, suppose we bisect the initial interval $[0, 1]$, and then go on to bisect recursively the resulting subintervals. We would then obtain the following sequence of nested N panel methods for any $N = 2^k - 1$, $k \geq 1$, by taking the midpoints of these intervals:

$$\begin{aligned} \bar{\epsilon} &= \{\underbrace{\Phi^{-1}[\frac{1}{2}], \Phi^{-1}[\frac{1}{4}], \Phi^{-1}[\frac{3}{4}]}_{N=1}, \underbrace{\Phi^{-1}[\frac{1}{8}], \Phi^{-1}[\frac{3}{8}], \Phi^{-1}[\frac{5}{8}], \Phi^{-1}[\frac{7}{8}]}_{N=3}\} \\ &\quad \underbrace{\hspace{15em}}_{N=7} \end{aligned} \quad (5)$$

3. Valuation of Complex Derivatives

Our discussion of European options on a single asset shows that by a judicious choice of the "random" points we can improve the accuracy of the estimate. We saw that the trick is not to have any randomness at all but to select deterministic sequences. In the case of path dependent options or options involving several assets where there are no analytical solutions, we are more likely to require Monte Carlo approaches. For these cases we require s dimensional sequences of "random" numbers. We discuss this below and later show how to generate such sequences.

Consider an Asian option where the payoff is based on the arithmetic average of the asset prices taken at certain dates known as reset points. We restrict attention to European contracts. Standard or (crude) Monte Carlo methods can be used to value such contracts. The continuous distribution is sampled at discrete time steps. Duffie (1992, p. 200) points out that if we replace the continuous time process with a discrete time process using s evenly spaced time steps then the error for the usual Euler discretization is $O(s^{-1})$. This error should be added to whatever error is incurred in evaluating the discrete time version of the integral. So, unless the security has a set of natural discrete times on which its value depends, the number of time steps should be chosen with care. Ideally one ought to choose a time step that has an error which matches the error introduced by estimating the discrete integral.

³ We give a more technical definition of discrepancy in Appendix B.

Suppose that the time step has already been chosen to obtain the discrete time version of the continuous time model. Assume that the path dependent security's terminal value at time T depends only on the prices A_1, A_2, \dots, A_s at s intermediate times

$$0 \leq t_1, t_2, \dots, t_s \leq T.$$

We assume the price increments are independent, take the expected discounted payoff under the risk neutral measure, and transform the integral as before to obtain

$$\text{Price} = \int_{I^s} f(x) dx \quad (6)$$

where $I^s = [0, 1]^s$.

We might be tempted to use the multidimensional version of the trapezoidal⁴ rule to evaluate this integral. Unfortunately the trapezoidal rule does not work well for a large number of time steps because the error bound is $O(N^{-2/s})$. Niederreiter (1992) gives a clear discussion of the issue of error bounds. If we use standard Monte Carlo methods using pseudo-random numbers the error bound is $O(N^{-1/2})$. Quasi-Monte Carlo methods which use deterministic sequences provide deterministic error bounds $O(N^{-1}(\log N)^s)$ for suitably chosen points. Such sequences are known as quasi-random sequences. There are several types of quasi-random sequences such as Halton sequences, Sobol' sequences and Faure sequences. Fox (1986, p. 374) explains the advantages of Faure sequences. In the next section we show how Faure sequences are generated.

4. Quasi-random Numbers: Faure Sequences

The basic idea of quasi-Monte Carlo is to use a set of points that are carefully selected in a deterministic fashion. We call these numbers quasi-random even though they are perfectly deterministic and have no random component. In this section we describe how to generate quasi-random sequences and describe a particular class of quasi-random sequences known as Faure sequences. Faure sequences have certain advantages for the valuation of high dimensional integrals. This leads to effi-

cient algorithms for the computation of prices and price sensitivities for complex derivative securities.

We first consider the case of a one-dimensional sequence of quasi-random numbers. Let r be any prime number (≥ 2). Any integer n has a unique expansion in terms of base r . We can generate a number in the interval $[0, 1)$ by reflecting the expansion in base r about the "decimal point." An example will clarify the procedure. Suppose $r = 3$ and $n = 7$. We can write 7 in base 3 as

$$7 = 2(3^1) + 1(3^0) = 21.$$

When we reflect 21 (in base 3) about the "decimal point" we obtain

$$\phi_3(7) = \frac{1}{3} + \frac{2}{9} = \frac{5}{9}.$$

This is clearly a number in the interval $[0, 1)$. The next number in the sequence is $\phi_3(8) = \frac{8}{9}$ and the first 9 numbers in this sequence are

$$\frac{9}{27}, \frac{18}{27}, \frac{3}{27}, \frac{12}{27}, \frac{21}{27}, \frac{6}{27}, \frac{15}{27}, \frac{24}{27}, \frac{1}{27}.$$

Notice how the new numbers that are added tend to fill in the gaps in the existing sequence. The general expression for n in terms of the base r is

$$n = \sum_{j=0}^m a_j(n)r^j. \quad (7)$$

Only a finite number of these $a_j(n)$ will be non-zero. The corresponding quasi-random number according to this procedure is

$$\phi_r(n) = \sum_{j=0}^m a_j(n)r^{-j-1}. \quad (8)$$

Note that if we have a sequence of integers running from r^k to $r^{k+1} - 1$ then the maximum distance between any two points in the corresponding quasi-random sequence must be given by r^{-k} . This happens because the most significant digit in the representation of n in base r becomes the least significant digit when we reflect about the decimal point. This means that as points are added they "know" how to fill the gaps evenly. One-dimensional quasi-random sequences of this type are known as van der Corput sequences and are described by Niederreiter (1992) and Faure (1982).

There are different ways to generalize the one-dimensional quasi-random numbers described above to

⁴ See Press et al. (1992) for a discussion of the trapezoidal rule.

Table 1 Three Dimensional Faure Sequences

n	$a_0(n)$	$a_1(n)$	$a_2(n)$	ϕ_n^1	ϕ_n^2	ϕ_n^3
1	1	0	0	1/3	1/3	1/3
2	2	0	0	2/3	2/3	2/3
3	0	1	0	1/9	4/9	7/9
4	1	1	0	4/9	7/9	1/9
5	2	1	0	7/9	1/9	4/9
6	0	2	0	2/9	8/9	5/9
7	1	2	0	5/9	2/9	8/9
8	2	2	0	8/9	5/9	2/9
9	0	0	1	1/27	16/27	13/27
10	1	0	1	10/27	25/27	22/27
11	2	0	1	19/27	7/27	4/27

s -dimensional quasi-random sequences. For our purposes Faure sequences have considerable advantages. The original discussion of these sequences by Faure (1982) is somewhat terse. Fox (1986) and Niederreiter (1992) give more detailed discussions.

We now describe how to generate an s -dimensional Faure sequence. Let r be the smallest prime number that is $\geq s$ and ≥ 2 . We start as before by representing any integer n in terms of base r as

$$n = \sum_{j=0}^m a_j^1(n)r^j. \quad (9)$$

The first Faure number is given by reflecting about the "decimal point" as before.

$$\phi_r^1(n) = \sum_{j=0}^m a_j^1(n)r^{-j-1}. \quad (10)$$

The remaining elements of the sequence are found recursively. First assume we know all the $a_j^{k-1}(n)$. The $a_j^k(n)$ are generated as follows:

$$a_j^k(n) = \sum_{i \geq j}^m {}^iC_j a_i^{k-1}(n) \bmod r, \quad (11)$$

where⁵ ${}^iC_j = i!/j!(i-j)!$. Thus the next level of coefficients is obtained by multiplying by an upper triangular matrix with elements

⁵ Note that for numerical accuracy, ${}^iC_j \bmod r$ is usually calculated via ${}^iC_j \bmod r = ({}^{i-1}C_j \bmod r + {}^{i-1}C_{j-1} \bmod r) \bmod r$.

$$\begin{matrix} {}^0C_0 & {}^1C_0 & {}^2C_0 & {}^3C_0 \cdots \\ 0 & {}^1C_1 & {}^2C_1 & {}^3C_1 \\ 0 & 0 & {}^2C_2 & {}^3C_2 \\ 0 & 0 & 0 & {}^3C_3 \end{matrix}$$

The successive points in the Faure sequence are obtained via

$$\phi_r^k(n) = \sum_{j=0}^m a_j^k(n)r^{-j-1}, \quad 2 \leq k \leq s, \quad (12)$$

that is, we reflect $a^k(n)$ about the decimal point to obtain $\phi^k(n)$.

This recursive procedure permits us to generate the s points corresponding to the integer n in the Faure sequence based on r ($\geq s$).

As an example, suppose that $s = 3$. In this case we take $r = 3$. The first few elements of the three-dimensional Faure sequence in this case are given in Table 1. To get a visual feel for how Faure sequences fill up the unit square, we have prepared Figures 1 and 2. In this case we selected $r = 2 = s$. We denote the n th element of the series by

$$\phi_n = (\phi_n^1, \dots, \phi_n^s).$$

To apply these results to a Monte Carlo simulation based on normal random variables, we map the point in the s -dimensional cube by applying the inverse normal function to the sequence⁶

$$\epsilon_n = \Phi^{-1}(\phi_n) = \{\Phi^{-1}(\phi_n^1), \dots, \Phi^{-1}(\phi_n^s)\}. \quad (13)$$

These Faure values may be substituted for the normal vectors that would be returned by a pseudo-random number generator. This makes the implementation simple since very little effort is needed to upgrade from pseudo-random numbers to quasi-Monte Carlo numbers. Code to implement the above algorithm is available in Fox (1986). We can also use this procedure to generate correlated random variables using appropriate

⁶ Note that it is incorrect to use the standard Box-Muller transform (Press (1992), p. 289) to map ϵ_n to the unit interval. This is because the Box-Muller transform does not preserve the matrix on \bar{F} and thus fails to preserve the low discrepancy of the original Faure sequence. In other words, if we use the Box-Muller transform, then the even spacing of the Faure sequence will be scrambled, resulting in the loss of our low error bound.

Figure 1 Points 1 to 128

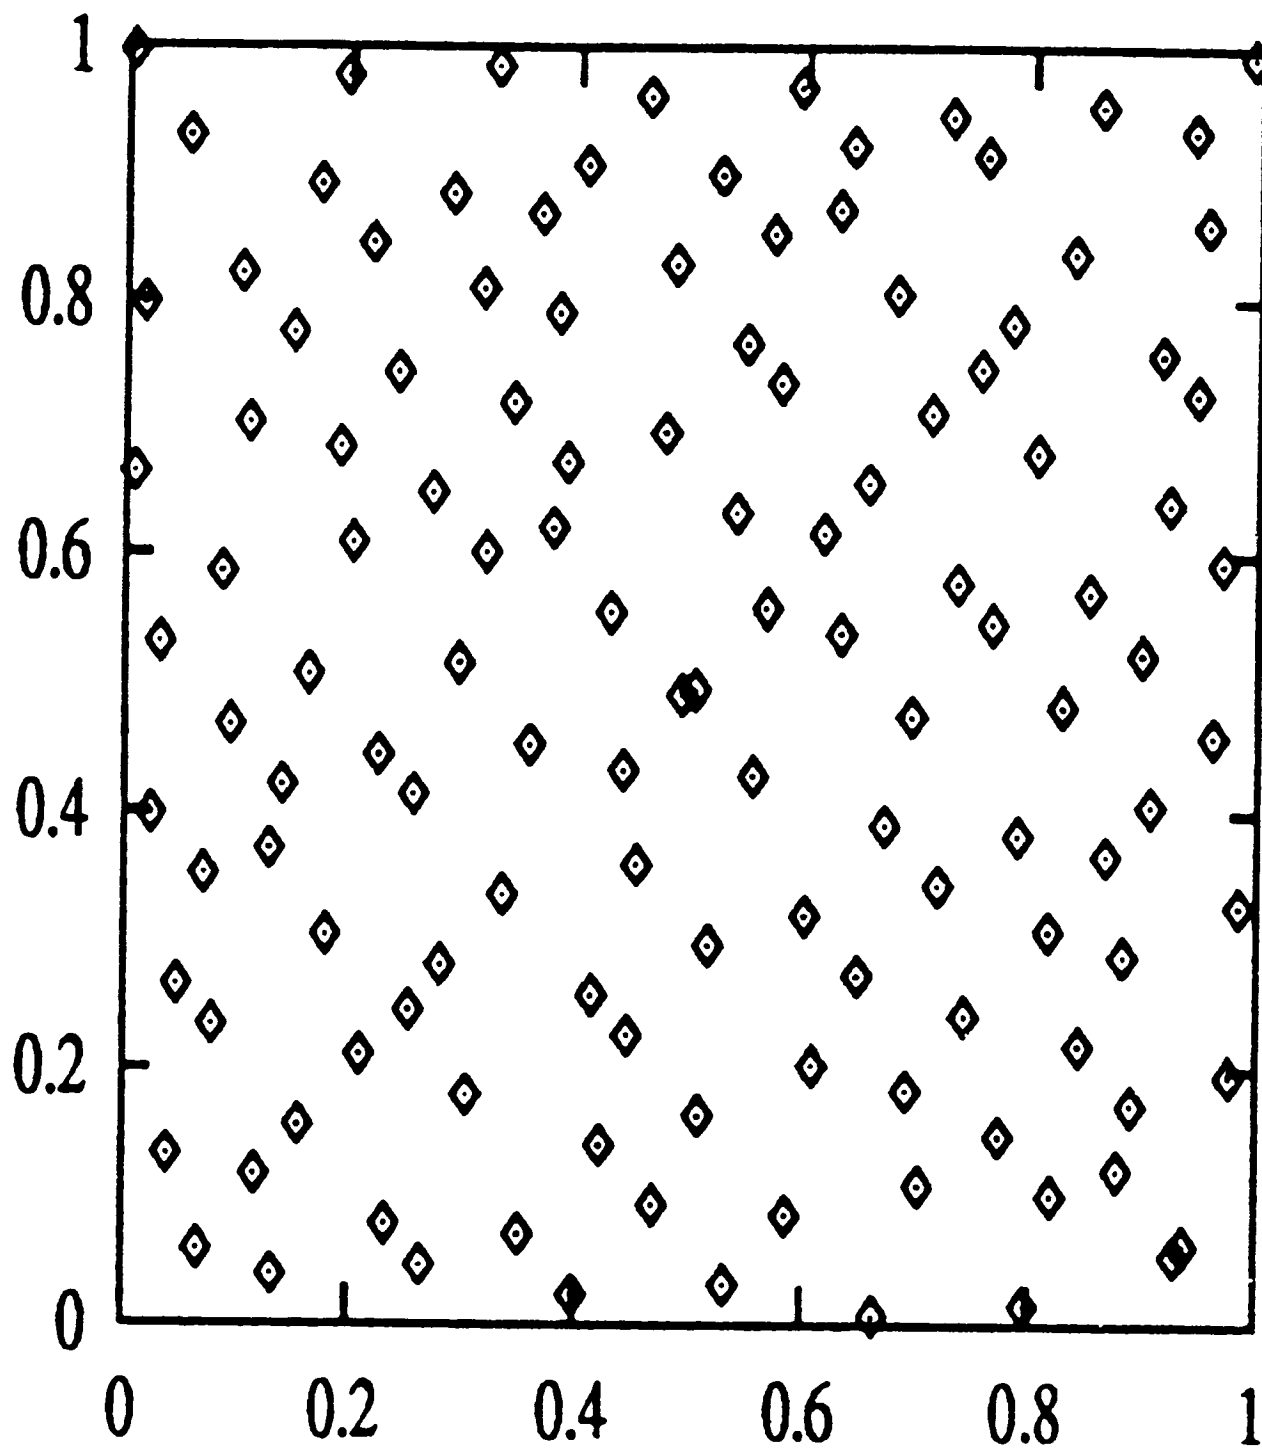
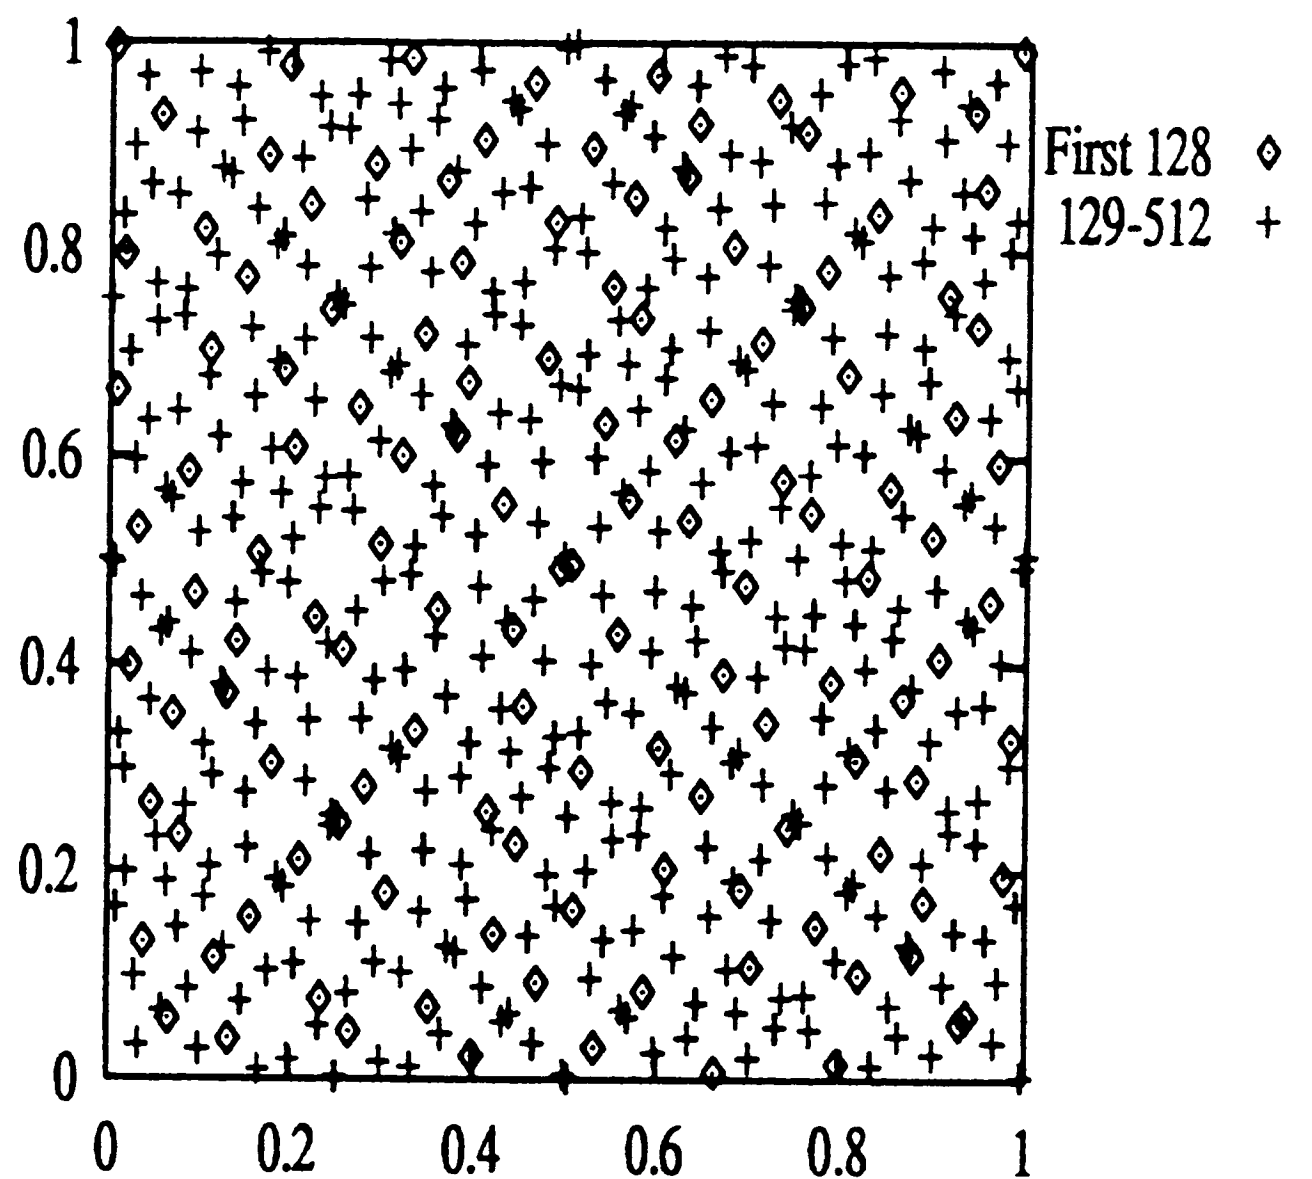


Figure 2 Points 1 to 512



combinations of uncorrelated random variables. This is useful when we are dealing with an instrument whose payoff depends on the values of several assets.

5. Applications: Use of Quasi-Monte Carlo for Standard European Options

In this section we illustrate the implementation of quasi-Monte Carlo methods in the case of plain vanilla European options. The exact values for the prices and sensitivities are available in these cases and we use these as benchmarks when examining convergence. We use graphs to compare the convergence properties of crude Monte Carlo with quasi-Monte Carlo. The graphs illustrate vividly the improvement in convergence when quasi-Monte Carlo methods are used.

Our base case contracts are standard European calls and puts with the following parameters:

Current price, S	100
Strike price, K	100
Riskless rate, r	10%
Time to maturity, T	1 year
Volatility, σ	30%

The Black Scholes option values for the standard call and put are

Call Value 16.734 Put Value 7.218

In Figure 3 we illustrate the convergence of the call option values under both the standard Monte Carlo

Figure 3 Standard European Call Option

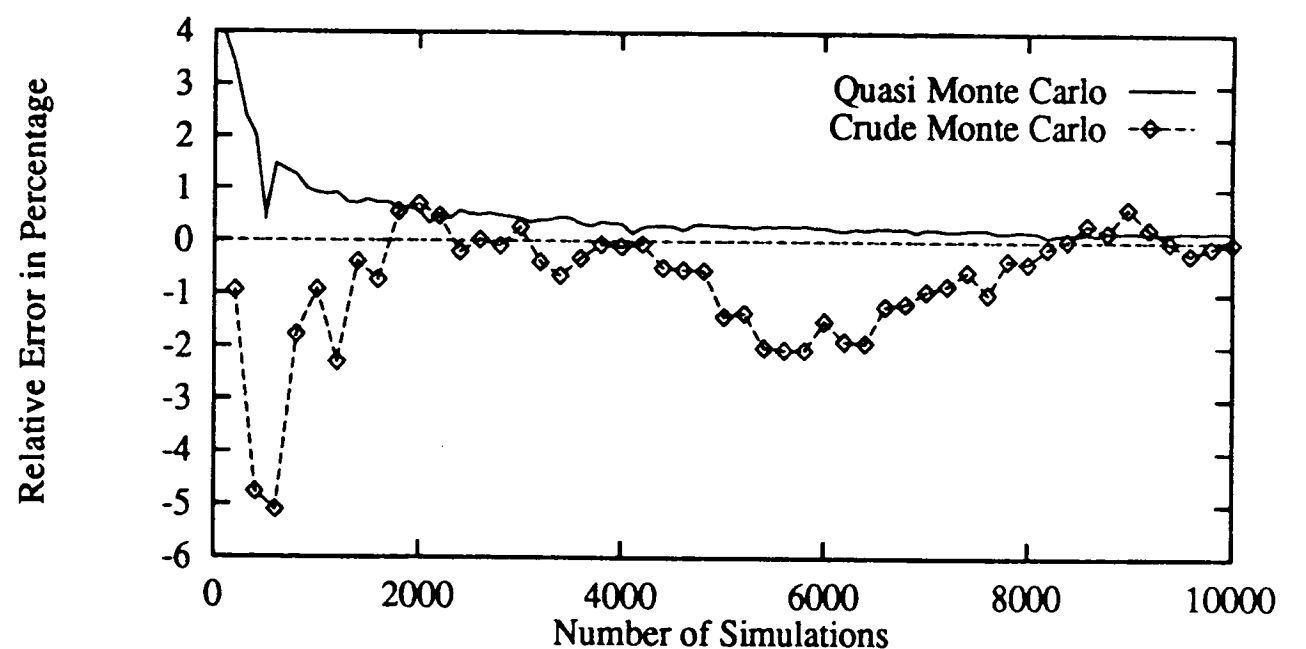
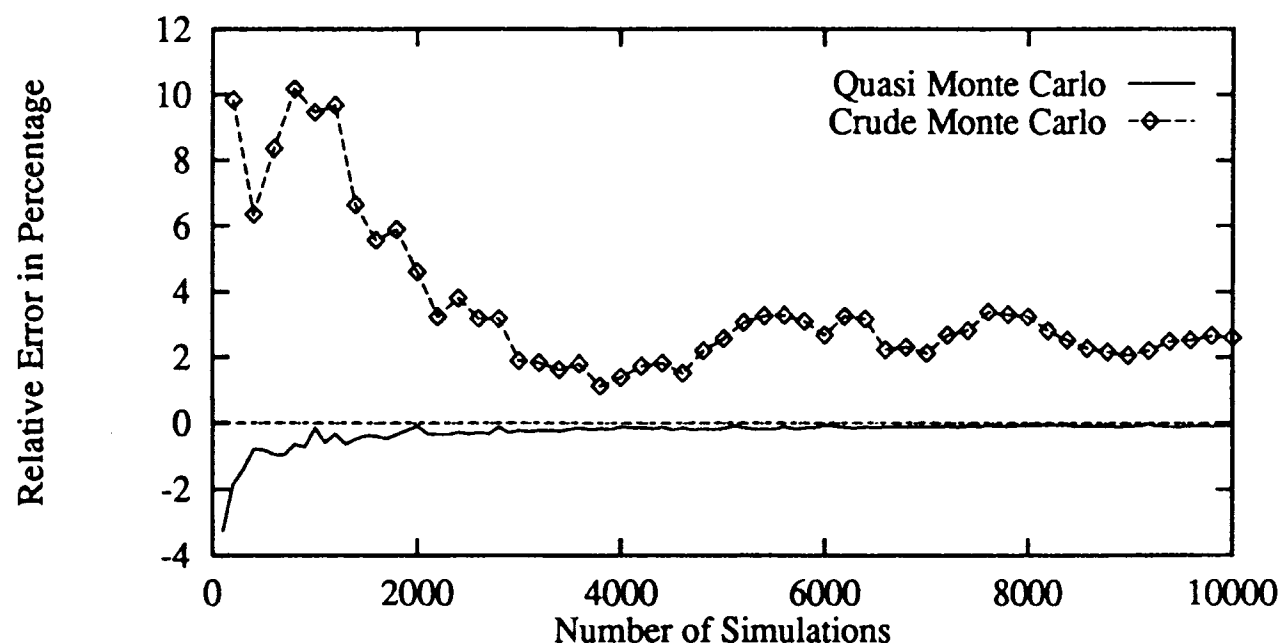


Figure 4 Standard European Put Option



approach and the quasi-Monte Carlo approach. The quasi-Monte Carlo numbers have been computed using Faure sequences (with $r = 2$). The numbers plotted represent the relative error in percentage between the estimated value and the true Black Scholes value. Observe that the quasi-Monte Carlo estimates display better convergence and that the deviations from the true prices are much less erratic. In these applications we generally start the simulations at $n = r^4$ as recommended by Fox (1986).

Figure 4 shows the same information for a standard put option. The quasi-Monte Carlo approach is clearly superior here as well. In many applications one is interested in the computation of the Greeks as well as the computation of the price. Broadie and Glasserman (1993) have recently discussed the different approaches to the computation of price sensitivities in the context of standard Monte Carlo methods. These

Figure 5 Standard European Call Option (Delta)

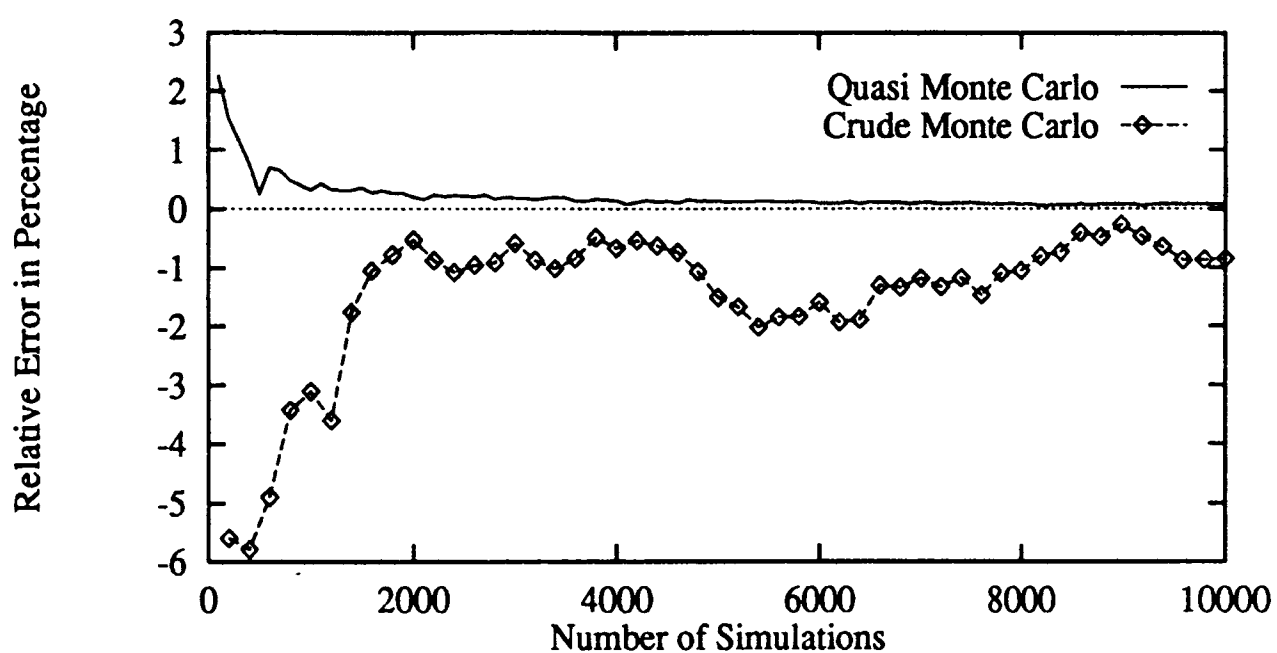
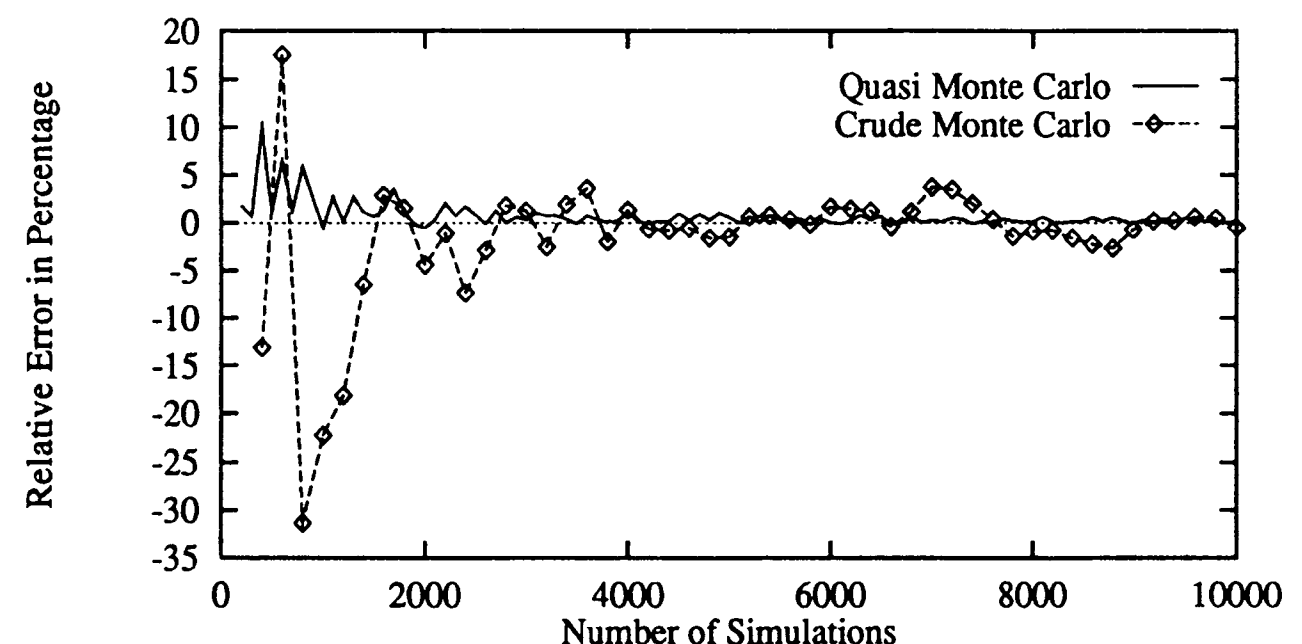


Figure 6 Standard European Call Option (Gamma)



different approaches can also be used with quasi-Monte Carlo methods but in each case the quasi-Monte Carlo approach displays better convergence than the standard Monte Carlo approach. This is illustrated in Figures 5 and 6. Figure 5 shows the convergence of delta for the European call option and Figure 6 provides the same information for the gamma of the call option. These graphs demonstrate that the quasi-Monte Carlo approach outperforms the standard method.

6. Applications of Quasi-Monte Carlo to Value Complex Securities

In practice the most likely applications of the Monte Carlo approach concern situations where there is no known analytic solution or when the analytical solution involves prohibitive computation to produce the answer. Many contracts of current interest fall into this category. In this section we analyze several types of complex derivatives. These include basket options, Asian options, and options that require an HJM simulation framework. We start with a complex option for which there is a simple closed-form solution.

6.1. Options on the Geometric Means of a Portfolio
 Our first example involves a European option on the geometric mean of several assets. Under the Black Scholes assumptions there is a simple expression for the value of this contract. The detailed formula is given in Boyle (1993). The parameters used for this example are

Number of assets	3
Initial asset prices, $S_i(0)$	100, $i = 1, 2, 3$
Volatilities, σ_i	0.3, $i = 1, 2, 3$
Correlations, ρ_{ij}	0.5, $i = 1, 2, j > i$
Strike price, K	100
Riskless rate, r	10%
Time to maturity, T	1 year

With these parameters the value of a European call on the geometric mean of the three assets is 13.771. We can also estimate the value of this option using Monte Carlo simulation methods. Figure 7 illustrates the convergence properties of standard Monte Carlo and quasi-Monte Carlo in this case. The quasi-Monte Carlo numbers were produced using a three-dimensional Faure sequence. The graph shows the superiority of the quasi-Monte Carlo approach in terms of convergence and error bounds. We now turn to an analysis of other types of complex derivatives.

6.2. Basket Options

Table 2 shows the results of applying Faure sequences to basket options, which are options on a basket of different assets in pre-specified quantities (see Gentle (1993) for a description). The spot price of a basket is the sum of these assets' prices multiplied by their respective quantities. The resulting distribution is thus a convolution of the individual assets' distribution. Because the resulting convolution is not easy to work with, these options have proved analytically intractable and hence numerical methods are used to value them. For

Figure 7 European Call Option on Geometric Mean of 3 Stock Prices

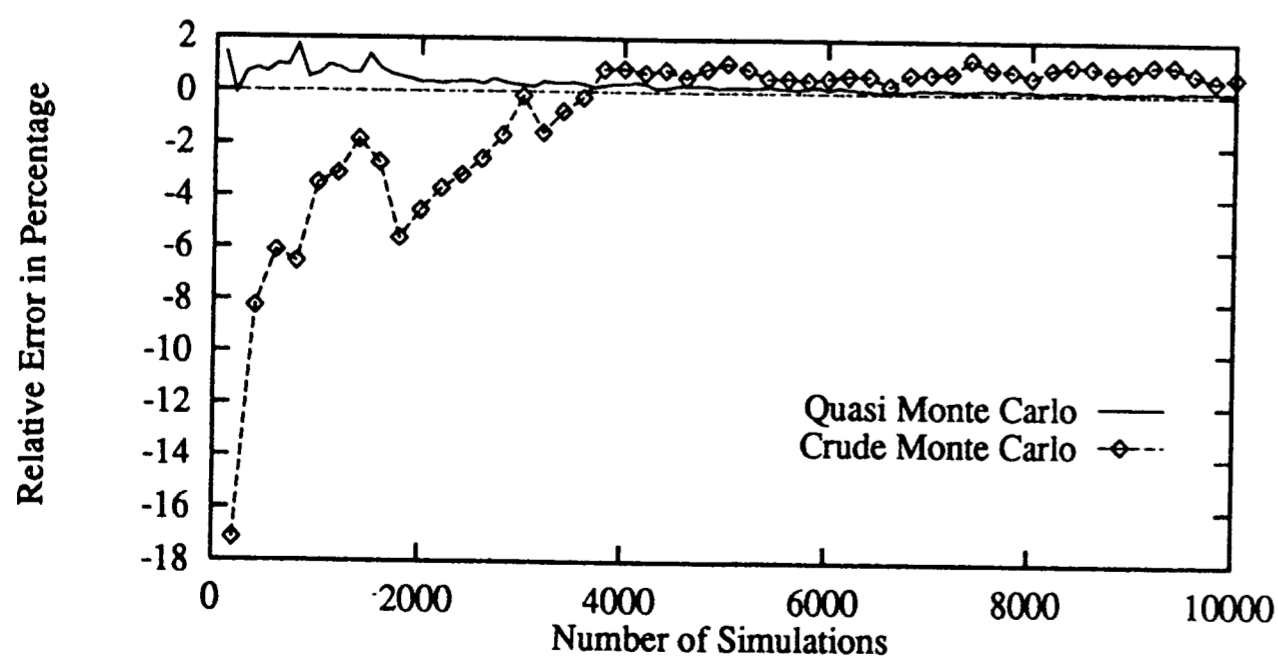


Table 2 Basket Call Option Valuation

Initial Price for LSC: 3.09/MMBTU		
Initial Price for NG: 2.20/MMBTU		
Volatility of LSC: 20% NG: 30%		
Correlation of both Commodities: 20%		
Strike Price: 2.80/MMBTU Riskless Rate: 9%		
Iterations	Crude Monte Carlo	Quasi-Monte Carlo
100	0.1386	0.1445
200	0.1562	0.1489
500	0.1621	0.1524
1,000	0.1455	0.1526
5,000	0.1545	0.1543
10,000	0.1493	0.1545
25,000	0.1553	0.1547
50,000	0.1569	0.1548
100,000	0.1552	0.1548
500,000	0.1544	0.1548

the purposes of this study, a basket of 60% Light Sweet Crude and 40% Natural Gas was chosen. In other words, the spot price of the basket was given by

Basket Spot

$$= 60\% * \text{Price of NYMEX Light Sweet Crude} / \text{MMBTU}^7$$

$$+ 40\% * \text{Price of NYMEX Natural Gas} / \text{MMBTU}.$$

It was further assumed that the correlation between these two commodities was 20%. A call option on this basket was then priced by using both Monte Carlo simulation and Faure sequences. As expected, Faure sequences proved markedly superior in both speed and accuracy for pricing this option. This also provides a demonstration of the applicability of Faure sequences to correlated random variables.

⁷ MMBTU stands for million British Thermal Units and is the standard unit of measurement on the NYMEX for natural gas. Light Sweet Crude and Natural Gas, however, are usually quoted in barrels and hence need to be converted to MMBTUs to make the two prices compatible in terms of units.

Table 3 Pricing of 52-Week Asian Options

52-Week Arithmetic Average Option with 53 Reset Points Spot Price: 100 Riskless Rate: 9% Volatility: 50%						
Iterations	Strike Price = 110		Strike Price = 100		Strike Price = 90	
	Crude	Quasi	Crude	Quasi	Crude	Quasi
100	9.207	13.131	10.206	17.043	19.493	21.942
200	11.778	10.224	15.405	14.219	18.902	19.298
400	10.156	9.180	11.818	13.327	16.854	18.642
800	8.761	10.965	12.593	14.991	17.140	20.118
1,600	9.706	8.953	13.206	12.928	17.873	18.171
3,200	9.280	8.820	12.934	12.801	18.135	18.074
4,800	9.191	8.747	12.965	12.673	18.326	17.919
9,600	9.009	8.894	12.684	12.804	18.121	17.990
19,200	9.195	9.101	12.754	12.993	18.018	18.164
38,400	9.152	9.057	13.033	12.949	18.147	18.116
76,800	8.987	9.052	13.039	12.948	18.214	18.114
153,600	8.999	9.062	12.948	12.962	18.110	18.128
307,200	9.026	9.063	12.914	12.963	18.106	18.128
614,400	9.002	9.065	12.922	12.965	18.085	18.131
1,228,800	9.030	9.063	12.885	12.963	18.064	18.128
2,457,600	9.040	9.062	12.906	12.962	18.100	18.128
Turnbull ^a	9.10 (0.03)		12.98 (0.03)		18.14 (0.03)	

^a These results are given by Turnbull (1992), with the "standard error" bounds reported in brackets.

Table 4 Pricing of 5-Week Asian Options

5-Week Asian Option with 6 Reset Points Spot Price: 100 Riskless Rate: 9% Strike Price: 100 Volatility: 50%		
Iterations	Crude Monte Carlo	Quasi-Monte Carlo
100	2.869	3.876
200	3.899	3.878
400	3.706	3.590
800	3.435	3.612
1,600	3.512	3.618
3,200	3.559	3.605
4,800	3.484	3.607
9,600	3.584	3.605
19,200	3.625	3.612
38,400	3.645	3.611
76,800	3.590	3.611
153,600	3.592	3.612
307,200	3.618	3.612
614,400	3.617	3.612

6.3. Asian Options

The next type of option to be examined was a set of Asian options drawn from a comparative study by Stuart Turnbull in *Risk* magazine (Turnbull (1992)). The averaging period was quite long for these options (52 weekly prices) and as a result the dimension of the problem was much higher than in the previous example (53 time steps versus 2 random variables for the basket option). This caused slower convergence for both methods but Faure sequences were still an improvement as shown in Table 3. To show that this slow convergence is due to high dimensionality, an example is given in Table 4 of a similar Asian option with an averaging period of 5 weeks. For the 5-week option, convergence of the Faure sequence is very fast.

6.4. HJM Factor Models

The final type of option to be priced was the option to enter into a natural gas swap. Because swap prices are highly sensitive to the shape of the forward price curve, it

is extremely important to use a model which can correctly capture price curve reshaping. One way to do this is through the use of a three factor HJM model along the lines of Cortazar and Schwartz (1992). The factor loadings are estimated through historical data and then scaled to match historical volatilities (if this were being used for an actual trade, the factors would be scaled to match implied volatilities). In Table 5 we price a six month option to enter into a one year swap with a constant daily volume of 10,000 MMBTUs per day. Once again, Faure sequences are more accurate than plain Monte Carlo even though the dimensionality of the problem is fairly high ((5 months + 1 partial month) \times 3 factors = 18). As a benchmark the last row compares the results to the price obtained by using Black's model for options on commodities. The volatility of 14% that is used here represents historical swap volatility for the last year and is well within the range of recent implied volatilities.

6.5. Portfolio Risk and Scientific Stress Tests

One final application that Faure sequences are particularly well suited for is the simulation of risks in large

Table 5 Pricing of Swaptions Under a 3 Factor HJM Model

Cost to Enter into an At The Money Natural Gas Swap Option
 Valuation Date: Sept. 20, 1994
 Inception of the Swap: March 1, 1995
 Maturity of the Swap: Feb 28, 1996
 Maturity of the Swaption: Feb 28, 1995
 Volume: 10,000 MMBTUs per day
 Fixed Price of the Swap: Current Mid-Market Price for a Swap
 of the above Duration and Volume

Iterations	Crude Monte Carlo (in \$)	Quasi-Monte Carlo (in \$)
100	231,593	268,027
200	236,071	288,062
400	271,236	300,659
800	267,513	270,850
1,600	271,635	268,295
3,200	258,667	268,025
4,800	265,364	266,580
9,600	267,193	267,198
19,200	265,433	267,094
38,400	266,137	266,910
76,800	263,069	267,538
153,600	265,279	267,302
Black-Scholes at 15% Volatility	260,930	

portfolios. The first reason for this is that, since each simulation pass can take quite a long time, using Faure sequences to simulate prices is an excellent way to decrease the time needed to obtain good accuracy. The second reason for using Faure sequences is that they can be directly compared to a systematic set of stress tests based upon the simulation model. This analogy comes from the fact that Faure sequences are generated so as to form very evenly spaced points in the probability space. The result is that if one uses Faure sequences in conjunction with one's simulation module, one can be sure that every possibility has been "tried out" and that there is not some kind of disaster scenario that has not been simulated due to an uneven set of random numbers. In other words, what Faure sequences do in this context is to systematically try out different combinations of prices under the model being used and test each and every one of these combinations to determine the effect on the portfolio. This ensures that all reasonable combinations will be tried, and eliminates the possibility that a poor set of random numbers might fail to uncover a crucial weakness or a potential portfolio collapse. This realization ends the dichotomy between ordinary stress tests and Monte Carlo simulation and helps one to understand how they relate to each other. Naturally, even Faure sequences cannot eliminate the risk that the model may fail, so ordinary stress tests are still very important.

7. Summary and Conclusion

Quasi-Monte Carlo methods provide a way to improve the accuracy and reliability of Monte Carlo simulation. The key idea is to use deterministic sequences known as quasi-random sequences. This results in better convergence and deterministic error bounds. In this paper we used the quasi-Monte Carlo approach to value a range of complex derivative securities that are of practical interest. This was done through the use of Faure sequences which have in general good convergence properties even when a large number of time steps are needed.

We illustrated the properties of the quasi-Monte Carlo approach in our numerical examples. It is evident from our graphs that the quasi-Monte Carlo approach leads to improved convergence and lower error bounds.

The theoretical error bounds for the quasi-Monte Carlo approach are given in Appendix B. It turns out that these bounds are loose. In practice any option that has a reasonably regular payoff function will give far higher convergence than these loose bounds indicate.

Wozniakowski (1991) has derived average case error bounds for real continuous functions on the unit cube. He shows that the average difference between the exact value of a multivariate integral and its Monte Carlo estimate is closely related to the discrepancy of the sample points that are used in the Monte Carlo approximation. This means that low discrepancy sequences lead to low average case error as well as low worst case error. Pas-kov (1994) gives further discussion of these bounds.

We can also improve the performance of quasi-Monte Carlo methods by using traditional variance reduction techniques.⁸ In addition we can apply Richardson extrapolation in conjunction with quasi-Monte Carlo methods to further enhance the accuracy. The above Faure sequences can be generated about as quickly as normally distributed random numbers. Furthermore, the theory shows that these sequences are robust in the sense that any subsequence of Faure numbers will give good convergence. Thus these sequences can be used much like pseudo-random number generators. Our conclusion is that this technique will prove to be a very powerful tool for many types of problems in computational finance.⁹

Although the present paper has concentrated on Faure sequences other low discrepancy sequences may have advantages over Faure sequences. In particular Boyle, Broadie, and Glasserman (1996) show that Sobol'

sequences outperform Faure sequences in many instances. Sobol' sequences also have advantages in that they tend to be more evenly dispersed¹⁰ throughout the unit hypercube than Faure sequences. The application of low discrepancy series to problems in finance is a topic of current interest. Some of the issues include the selection of appropriate termination criteria, how to exploit the regularity properties of the integrand and the performance of various low dimensional sequences with each other and with standard Monte Carlo at very high dimensions for a range of situations.

¹⁰ As we move to higher dimensions the Faure points are not evenly dispersed throughout the hypercube. If we project the 50th and 51st dimension we find strong patterns and large gaps. This point is discussed in Boyle, Broadie, and Glasserman (1996). We are also grateful to Alan Jung for discussions on this issue.

Appendix A: Computing the Normal Inverse Function

The authors are indebted to B. Moro for his generosity in providing the following algorithm.

This Appendix explains one procedure for inverting the cumulative normal density function. The inverse, $\Phi^{-1}(x)$, of cumulative normal density function

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-(t^2/2)} dt \quad (14)$$

can generally be computed using the Newton-Raphson iterative method. However, this is too slow to be useful in practice. The procedure illustrated below has high accuracy for all values of Φx in the interval $[10^{-10}, 1 - 10^{-10}]$ while retaining speed of simple rational approximation.

If $|y| \leq 0.42$, where $y = \Phi x - 0.5$, $\Phi^{-1}(x)$ is computed using a rational approximation as given by Beasley and Springer (1977),

$$\Phi^{-1}(x) = y \frac{\sum_{n=0}^3 a_n y^{2n}}{\sum_{n=0}^4 b_n y^{2n}} \quad (15)$$

where a_n and b_n are suitably chosen coefficients. For $|y| > 0.42$, we approximate $\Phi^{-1}(x)$ by a truncated Chebyshev series as

$$\Phi^{-1}(x) = \begin{cases} \sum_{n=0}^8 c_n T_n(z) - \frac{c_0}{2} & \text{if } y > 0 \\ \frac{c_0}{2} - \sum_{n=0}^8 c_n T_n(z) & \text{otherwise,} \end{cases} \quad (16)$$

Here $z = k_1 * [2 * \log(-\log(0.5 - |y|)) - k_2]$ and the constants k_1 and k_2 are chosen such that $z = -1$ when $\Phi(x) = 0.92$ and $z = 1$ when $\Phi(x) = 1 - 10^{-12}$. The upper limit was set rather arbitrarily. One can easily derive Chebyshev approximations which will be valid for $\Phi(x)$ up to

⁸ Quasi-Monte Carlo methods perform even better when combined with variance reduction techniques. For example, to apply the control variate method the original integral $\int f(x)dx$ is replaced with an integral of the form $\int [f(x) - g(x)]dx$ where $g(x)$ has the same type of behavior as $f(x)$. Clearly if crude Monte Carlo can be used to value $f(x) - g(x)$ then quasi-Monte Carlo can value this difference even more accurately. Two common variance reduction techniques that do not improve upon standard quasi-Monte Carlo methods are antithetic variables and stratified sampling. Quasi-Monte Carlo methods are intrinsically designed to keep on subdividing the integration region into uniformly spaced parts with every additional simulation point. Quasi-Monte Carlo methods already incorporate the advantages associated with these two variance reduction techniques.

⁹ The authors are grateful to the referees and Robert Heinkel for constructive comments.

Table 6 Constants for Inverting Normal Variates

n	a_n	b_n	n	c_n
0	2.50662823884	1.00	0	7.7108870705487895
1	-18.61500062529	-8.47351093090	1	2.7772013533685169
2	41.39119773534	23.08336743743	2	0.3614964129261002
3	-25.44106049637	-21.06224101826	3	0.0373418233434554
4		3.13082909833	4	0.0028297143036967
			5	0.0001625716917922
	k_1	k_2	6	0.0000080173304740
	0.4179886424926431	4.2454686881376569	7	0.0000003840919865
			8	0.0000000129707170

$1 - 10^{-15}$, but that would come with increased computational cost. Once the coefficients c_n are computed (see Press and Teukolsky (1992), Eq. (5.8.7)), Eq. (16) can be efficiently evaluated using Clenshaw's formula (see Press and Teukolsky (1992), Eq. (5.8.11)). The required constants a_n, b_n, c_n, k_1 and k_2 are reproduced in Table 6.

The rational approximation using Eq. (15) has the largest absolute error of 3.0×10^{-9} in the interval $[0.5, 0.92]$. In the tails of distribution, Eq. (16) retains this accuracy for up to 6 standard deviations.

Appendix B: Error Bounds for Faure Sequences

This Appendix discusses the error bounds for Faure sequences. We compare the error bounds for quasi-Monte Carlo based on Faure sequences with the error bounds from crude Monte Carlo. We begin with three definitions and an important theorem. First we define the counting function of a subset of a set of points.

DEFINITION 1. Let P be a point set consisting of points x_1, x_2, \dots, x_N . For any subset B of \bar{I}^s define

$$A(B; P) = \sum_{n=1}^N c_B(x_n),$$

where c_B is the characteristic function of B . Thus $A(B; P)$ is the counting function that indicates the number of n with $1 \leq n \leq N$ for which $x_n \in B$.

Next we define discrepancy.

DEFINITION 2. Let \mathcal{B} be a family of Lebesgue measurable subsets of \bar{I}^s . The discrepancy $D_N(\mathcal{B}; P)$ of a point set P is defined as

$$D_N(\mathcal{B}; P) = \sup_{B \in \mathcal{B}} \left| \frac{A(B; P)}{N} - \lambda_s(B) \right|,$$

where $\lambda_s(B)$ is the Lebesgue measure of B .

DEFINITION 3: STAR DISCREPANCY. The star discrepancy $D_N^*(P) = D_N^*(x_1, \dots, x_N)$ of the point set P is defined by $D_N^*(P) = D_N(\mathcal{J}^*; P)$ where \mathcal{J}^* is the family of all subintervals of I^s of the form $\prod_{i=1}^s [0, u_i]$.

Now comes the key theorem.

THEOREM (KOKSMA-HLAWKA). Let $I^s = [0, 1]^s$ and let f have bounded variation $V(f)$ on I^s in the sense of Hardy and Krause. Then, for any $x_1, \dots, x_N \in I^s$ we have

$$\left| \frac{1}{N} \sum_{n=1}^N f(x_n) - \int_{I^s} f(u) du \right| \leq V(f) D_N^*(x_1, \dots, x_N).$$

Intuitively, the D star discrepancy $D_N^*(x_1, \dots, x_N)$ represents how "evenly" x_1, \dots, x_N are distributed within \bar{I}^s . It turns out that in the one-dimensional case, the best choice for x_1, \dots, x_N are the evenly spaced points $x_i = (2i - 1)/2N$ for $1 \leq i \leq N \in [0, 1]$. This gives the minimum possible discrepancy of $D_N^*(x_1, \dots, x_N) = 1/2N$ (Niederreiter 1992, p. 23). However, this set has the drawback that in order to obtain this minimal discrepancy N must be specified in advance. For Faure sequences, the resulting has low discrepancy for all N which allows the simulation to be easily extended until the desired accuracy is obtained. Furthermore, the projection of Faure sequences onto lower dimensional hypercubes $I^d, d \leq s$, will still have discrepancy $O(N^{-1}(\log N)^d)$. This means that the first d terms from a higher order Faure sequence of order s can be used for an integral of lower dimension d .

For the Faure sequences under consideration, we have the following result from Faure (1982):

Given the Faure sequence x_1, \dots, x_N of dimension s

$$D_N^*(x_1, \dots, x_N) \leq A_s \frac{(\log N)^s}{N} + O\left(\frac{(\log N)^{s-1}}{N}\right) \text{ where}$$

$$A_s = \frac{1}{s!} \left(\frac{q_s - 1}{2 \log q_s} \right)^s,$$

with q_s as the smallest prime number that is greater than or equal to s .

Table 7 Comparison of the Theoretical Error Bound for Both Crude and Quasi-Monte Carlo Methods

s	N	$N^{-1/2}$	$A_s N^{-1} (\log N)^s$
5	10^4	0.01	0.163
10	10^8	0.0001	19.262

This series has nearly¹¹ the smallest known set of constants, A_s . When working with a fixed number of simulations N , Niederreiter (1992, p. 97) suggests using the set $(n/N, x_n) \in I^s$, where $n = 0, 1, \dots, N - 1$, which has $D^* \leq O((\log N)^{s-1}N^{-1})$. Furthermore, in the s dimensional case it is widely believed¹² that for any N element point set (x_1, \dots, x_N) we must have

$$D_N^*(x_1, \dots, x_N) \geq O(N^{-1}(\log N)^{s-1}).$$

This suggests that $O(N^{-1}(\log N)^{s-1})$ is the best order of convergence that we can expect from any deterministic set.

If we compare the Faure error bound with theoretical Monte Carlo it may not seem that quasi-Monte Carlo methods are a significant improvement (see Table 7).

In practice, this table is deceptive because Faure sequences perform much better than this indicates. This is due to the fact that quasi-random Faure points are very evenly spaced while most crude Monte Carlo simulations are run using pseudo-random number generators whose D^* discrepancies are higher than the D^* discrepancies of Faure sequences (see Niederreiter (1992) for a discussion of the discrepancies of pseudo-random number generators).

References

- Beasley, J. D. and S. G. Springer, "The Percentage Points of the Normal Distribution," *Applied Statistics*, 23 (1977), 118–121.
- Boyle, P., "New Life Forms on the Option Landscape," *J. Financial Engineering*, 2, 3 (1993), 217–252.
- , "Options: A Monte Carlo Approach," *J. Financial Economics*, 4 (1977), 323–338.
- , Broadie M., and P. Glasserman, "Monte Carlo Methods for Security Pricing," forthcoming in *J. Economics Dynamics and Control*.
- Bratley, P. and B. L. Fox, "ALGORITHM 659: Implementing Sobol's Quasi Random Sequence Generator," *ACM Trans. Math. Software*, 14 (1988), 88–100.
- Broadie, M. and P. Glasserman, "Estimating Security Price Derivatives Using Simulation," Working Paper, Columbia University, New York, 1993.
- Cortazar, G. and E. Schwartz, "The Valuation of Commodity Contingent Claims," Working Paper, Anderson Graduate School of Management, UCLA, Los Angeles, CA, 1992.
- Duffie, D., *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ, 1992.
- Faure, H., "Discr pance de suites associ es   un syst me de num ration (en dimension s)," *Acta Arithmetica*, 41 (1982), 337–351.
- Fox, P., "ALGORITHM 647: Implementation and Relative Efficiency of Quasirandom Sequence Generators," *ACM Trans. Math. Software*, 12 (1986), 362–376.
- Gentle, D., "Basket Weaving," *Risk*, June (1993), 51.
- Grant, D., D. Vora, and D. Weeks, "Path Dependent Options: Extending the Monte Carlo Approach," Working Paper, University of New Mexico, Albuquerque, NM, 1993.
- 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. and D. Handscomb, *Monte Carlo Methods*, John Wiley & Sons, New York, 1964.
- Hlawka, E., "Funktionen von beschr nktter Variation in der Theorie der Gleichverteilung," *Ann. Math. Pura Appl.*, 54 (1961), 325–333.
- , F. Firneis, and P. Zinterhof, *Zahlentheoretische Methoden in der Numerischen Mathematik*, Oldenbourg, Vienna, Austria, 1981.
- Hua, L. and Y. Wang, *Applications of Number Theory to Numerical Analysis*, Springer, Berlin, 1981.
- Niederreiter, H., "Point Sets and Sequences with Small Discrepancy," *Monatsh. Math.*, 104 (1987), 273–337.
- , *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Philadelphia, PA, 1992.
- , "Quasi-Monte Carlo Methods for Multidimensional Numerical Integration," in H. Brass and G. Hammerlin (Eds.), *Numerical Integration III, International Series of Numer. Math.*, Birkh user, Basel, 85, 1988, 157–171.
- Paskov, S., "Computing High Dimensional Integrals with Applications to Finance," Technical Report CUCS-023-94, Dept. of Computer Science, Columbia University, New York, 1994.
- Press, W., S. Teukolsky et al., *Numerical Recipes in C, 2nd ed.*, Cambridge University Press, New York, 1992, 309–315.
- Sobol', I. M., "The Distribution of Points in a Cube and the Approximate Evaluation of Integrals," *USSR Computational Math. and Math. Physics*, 19 (1967), 86–112.
- Tilley, J. A., "Valuing American Options in a Path Simulation Model," *Trans. Society of Actuaries*, 45 (1993), 499–520.
- Wozniakowski, H., "Average Case Complexity of Multivariate Integration," *Bulletin (New Series) of the American Mathematical Society*, 24, 1 (January 1991), 185–194.
- Turnbull, S., "Average Intelligence," *Risk*, February (1992).

¹¹ A slight improvement on the Faure sequence has been made by Niederreiter (1987) but we prefer to stay with Faure sequences since their calculation is simpler. Sobol' (1967) sequences also have this same order of convergence but their constants are never better than Faure sequences and are in fact worse for all $s \geq 8$ (Niederreiter (1992), p. 96). However, Sobol' sequences are at least better than Halton (1960) sequences which have the same order of convergence but display super exponential growth in their constants as the dimension s increases.

¹² This can be shown for dimensions $s = 1$ and 2 but the best general result that has been proved for $s \geq 3$ is $D_N^*(x_1, \dots, x_N) \geq O(N^{-1}(\log N)^{(s-1)/2})$ (Niederreiter 1992, p. 32).

Accepted by Robert Heinkel; received November 2, 1994. This paper has been with the authors $1\frac{1}{2}$ months for 1 revision.