Faster Valuation of Financial Derivatives

A promising alternative to Monte Carlo.

Spassimir H. Paskov and Joseph F. Traub

Monte Carlo simulation is widely used to value complex financial instruments. Vast sums are spent annually on these methods.

Monte Carlo methods use random (or, more precisely, pseudo-random) points. If we plot a moderate number of pseudo-random points in two dimensions, we observe regions where there are no points (see, e.g., Traub and Wozniakowski [1994, p. 102]). Rather than use pseudo-random points, it seems attractive to choose points that are as uniformly distributed as possible. There is a notion in number theory called discrepancy, which measures the deviation of a set of points in d dimensions from uniformity. Although the question of which point sets in d dimensions have the lowest discrepancy is open, various *low-discrepancy* point sets are known.

Our study compares the efficacy of low-discrepancy methods with Monte Carlo methods on the valuation of financial derivatives. We use a collateralized mortgage obligation (CMO), provided to us by Goldman Sachs, with ten bond classes (tranches) formulated as the computation of ten integrals of dimension up to 360. We choose this CMO because it has fairly high dimension, and because each integrand evaluation is very expensive it is crucial to sample the integrand as few times as possible. We believe that our conclusions regarding this CMO will hold for many other financial derivatives.

The low-discrepancy sample points chosen for

SPASSIMIR H. PASKOV is a risk analyst at the Union Bank of Switzerland in New York (NY 10171).

JOSEPH F. TRAUB is the Edwin Howard Armstrong professor of computer science at Columbia University in New York (NY 10027).

our tests are Sobol and Halton points. We compare methods based on these points with the classical Monte Carlo method and also with the classical Monte Carlo method combined with antithetic variables.

An explanation of terminology is required. *Low-discrepancy points* are sometimes referred to as *quasi-random* points. Although this latter term is in widespread use, we believe the expression is misleading as there is nothing random about these deterministic points. We prefer to use the terminology *low-discrepancy* or *deterministic.*

We assume the finance problem has been formulated as an integral over the unit cube in d dimensions. We have built a software system called FINDER for computing high-dimensional integrals. It runs on a heterogeneous network of workstations under PVM 3.2 (parallel virtual machine). Because workstations are ubiquitous, this is a cost-effective way to perform large numbers of computations quickly. Of course, FINDER can also be used to compute high-dimensional integrals on a single workstation.

A routine for generating Sobol points is given by Press et al. [1992]. FINDER, however, incorporates major improvements, and the results reported here were obtained using it. One improvement is developing the table of primitive polynomials and initial direction numbers for dimensions up to 360.

This article is based on two years of software construction and testing. Preliminary results were presented to a number of New York City financial institutions in the fall of 1993 and the spring of 1994. A January 1994 article by Traub and Wozniakowski [1994, p. 102] discusses the theoretical issues and reports that "preliminary results obtained by testing certain finance problems suggest the superiority of the deterministic methods in practice." Further results were reported at a number of conferences in the summer and fall of 1994. A June 1994 article in *Business Week* indicates the possible superiority of low-discrepancy sequences.

Details on the CMO, the numerical methods, and the test results are presented by Paskov [1994]. Here we limit ourselves to stating our main findings and indicating typical results. For the sake of brevity, we shall refer to the method that uses Sobol points as the Sobol method.

Our main conclusions regarding the evaluation of this CMO fall into three groups.

Deterministic and Monte Carlo Methods

The Sobol method consistently outperforms the Monte Carlo method. The Sobol method consistently outperforms the Halton method. In particular:

- The Sobol method converges significantly faster than. the Monte Carlo method.
- The convergence of the Sobol method is smoother than the convergence of the Monte Carlo method. This makes automatic termination easier for the Sobol method.
- Using our standard termination criterion, the Sobol method terminates two to five times faster than the Monte Carlo method, often with less error.
- The Monte Carlo method is sensitive to the initial seed.

Sobol, Monte Carlo, and Antithetic Variables Methods

The Sobol method consistently outperforms the antithetic variables method, which in turn consistently outperforms the Monte Carlo method. In particular:

- These conclusions also hold when a rather small number of sample points are used, an important case in practice. For example, for 4,000 sample points, the Sobol method running on a single Sun-4 workstation achieves accuracies within range from one part in a thousand to one part in a million, depending on the tranche, within a couple of minutes.
- Statistical analysis on the small sample case further strengthens the case for the Sobol method over the antithetic variables method. For example, to achieve similar performances at a confidence level of 95%, the antithetic variables method needs from seven to seventy-nine times more sample points than the Sobol method, depending on the tranche.
- The antithetic variables method is sensitive to the initial seed, but convergence of the antithetic variables method is less jagged than convergence of the Monte Carlo method.

Network of Workstations

All the methods benefit by being run on a network of workstations. In particular:

- For N workstations, the measured speedup is at least 0.9N, where $N \leq 25$.
- A substantial computation that took seven hours on a Sun-4 workstation took twenty minutes on the network of twenty-five workstations.

We do not claim that the Sobol method is always superior to the Monte Carlo method. We do not even claim that it is always superior for financial derivatives. After all, the test results reported here are only for one particular CMO. We do believe, however, that it will be advantageous to use the Sobol method for many other types of financial derivatives.

NUMERICAL METHODS

The idea underlying the Monte Carlo method is to replace the integral of f (x), which is a continuous average, by a discrete average over randomly chosen points. More precisely, let D denote the d-dimensional unit cube. We approximate

$$\int_D f(x)\,dx$$

by

$$\frac{1}{n}\sum_{i=1}^{n} f(t_i)$$

It is well-known that, if one chooses n points from a flat distribution, the expected error is

$$E_n(f) = \frac{\sigma(f)}{\sqrt{n}}$$

where $\sigma^2(f)$ denotes the variance of f.

The Monte Carlo method has the advantage that the expected error is independent of dimension, but it suffers from the disadvantage that the rate of convergence is only proportional to $n^{-1/2}$. This motivates the search for methods that converge faster. Low-discrepancy methods also approximate the integral of f (x) by a discrete average, although this time the average is taken over low-discrepancy points. A number of low-discrepancy point sets are known. Here we confine ourselves to Sobol or Halton points. Roughly speaking, both have the property that the rate of convergence is proportional to $(\log n)^d/n$. See Niederreiter [1992] for the theory of low-discrepancy points and references to the literature.

The n^{-1} factor in the convergence formula for low-discrepancy points may be contrasted with the $n^{-1/2}$ convergence of Monte Carlo, suggesting that low-discrepancy methods are sometimes superior to Monte Carlo methods. A number of researchers report that this advantage decreases with increasing dimension. Furthermore, they report that the theoretical advantage of low-discrepancy methods disappears for rather modest values of the dimension, say, $d \leq 30$.

These conclusions are based on mathematical problems specifically constructed for testing purposes or for certain problems arising in physics. As we shall see, tests on 360 dimensional integrals arising from a CMO lead to very different conclusions.

THE FINANCE PROBLEM

We test a collateralized mortgage obligation (CMO) provided to us by Goldman Sachs. This CMO consists of ten tranches that derive their cash flows from an underlying pool of mortgages. Cash flows (interest and principal) received from the pool of mortgages are divided and distributed to each of the tranches according to a set of prespecified rules. The technique of distributing the cash flows transfers the prepayment risk among different tranches. While the actual amount of cash flows obtained will depend upon the future level of interest rates, our problem is to estimate the expected value of the sum of present values of future cash flows for each of the tranches.

The underlying pool of mortgages has a thirty-year maturity, and cash flows are obtained monthly. This leads to 360 cash flows and hence to integration in 360 dimensions. The precise mathematical formulation for this CMO may be found in Section 5 of Paskov [1994].

SOFTWARE SYSTEM FOR COMPUTING HIGH-DIMENSIONAL INTEGRALS

Theory suggests that the low-discrepancy deterministic methods provide an interesting alternative to the Monte Carlo method for computing high-dimensional integrals. We have developed and tested a distributed software system for computing multivariate integrals on a network of workstations. The software also runs on a single workstation.

The software uses a sequence of sample points as follows:

- Halton points.
- Sobol points.

EXHIBIT 1
SOBOL AND HALTON RUNS FOR TRANCHE A AND TWO MONTE CARLO RUNS USING RAN2

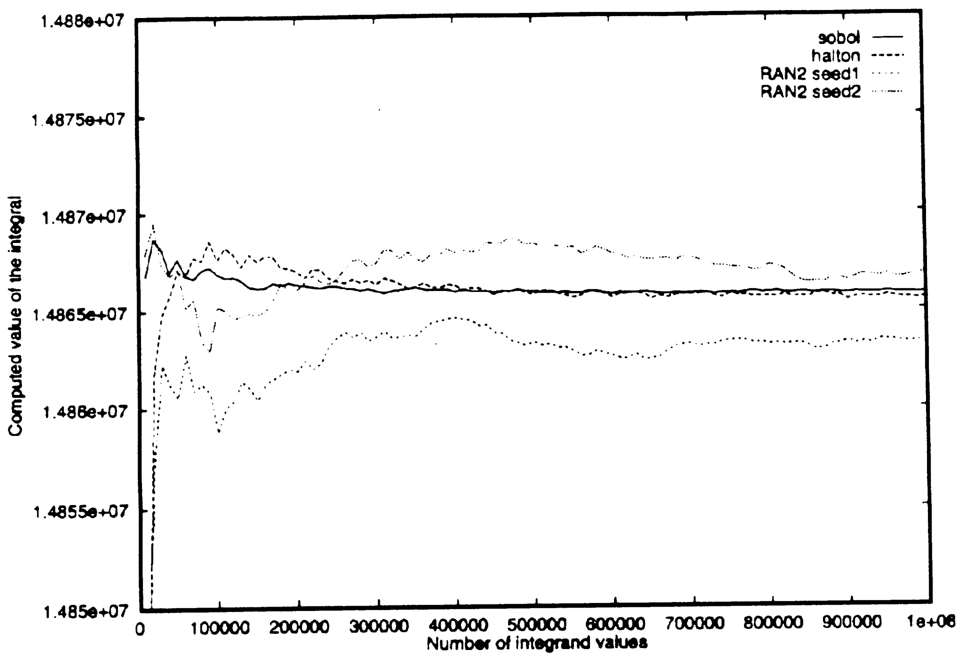

- Uniformly distributed random points.

The user can choose the sequence of sample points from a menu. The software is written in a modular way so other kinds of deterministic and random number generators can be easily added. One or several multivariate functions defined over the unit cube of up to 360 variables can be integrated simultaneously.

A routine for generating Sobol points is given in Press et al. [1992]. Our FINDER system makes major improvements on this routine, and the results reported here were obtained using FINDER One of the improvements is development of the table of primitive polynomials and initial direction numbers for dimensions up to 360.

The software permits the use of various random number generators. RAN1 and RAN2 from Press et al. [1992] are used because of their wide availability and popularity.

COMPARISON OF DETERMINISTIC AND MONTE CARLO METHODS

Selected results of extensive testing of the deterministic and Monte Carlo methods for the CMO are summarized in a number of graphs.

Exhibit 1 shows the results for one of the ten tranches (tranche A) of Sobol, Halton, and Monte Carlo runs with two randomly chosen initial seeds. Throughout this section, we describe results on tranche A. Results for other tranches are similar unless stated otherwise. The pseudorandom generator RAN2 from Press et al. [1992] is used to generate random sample points for the Monte Carlo runs.

It is striking how typical Exhibit 1 is of the vast amount of data we collected. We can conclude:

- The Monte Carlo method is sensitive to the initial seed.
- The deterministic methods, especially the Sobol method, converge significantly faster than the Monte Carlo method.
- The convergence of the deterministic methods, especially of the Sobol method, is smoother than the convergence of the Monte Carlo method. This makes automatic termination easier for the Sobol method.
- The Sobol method outperforms the Halton method.

Exhibit 2 plots the same Sobol and Halton runs versus the arithmetic mean of twenty Monte Carlo runs. The twenty Monte Carlo runs use twenty different randomly chosen initial seeds. We stress that the number of sample points on the x-axis is correct only for the deterministic methods. The actual number of sample points for the averaged Monte Carlo graph is twenty times the number of sample points on the x-axis. The results of the deterministic methods and the averaged Monte Carlo results are approximately the same. After roughly the first 50,000 integrand evaluations, the behavior of the deterministic methods and average Monte Carlo is roughly the same, even though we are using twenty times more random than deterministic points.

In Exhibit 3, an automatic termination criterion is applied to Sobol, Halton, and three Monte Carlo runs. We choose a standard automatic termination criterion. Namely, when two consecutive differences between consecutive approximations using 10,000 i, i = 1, 2, ..., 100, sample points drop below some threshold value for all of the tranches of the CMO, the computational process is terminated.

With the threshold value set at 250, the Sobol run terminates at 160,000 sample points; the Halton run terminates at 700,000 sample points; and the three Monte Carlo runs terminate at 410,000, 430,000, and 780,000 sample points. Hence, the Sobol run terminates two to five times faster than the Monte Carlo runs.

Even though the Sobol method terminates faster, it is often more accurate than the Monte Carlo method. Details may be found in Paskov [1994].

ANTITHETIC VARIABLES

An important advantage of Monte Carlo and deterministic methods

EXHIBIT 2
SOBOL AND HALTON RUNS FOR TRANCHE A AND AN AVERAGE OF TWENTY MONTE CARLO RUNS USING RAN2

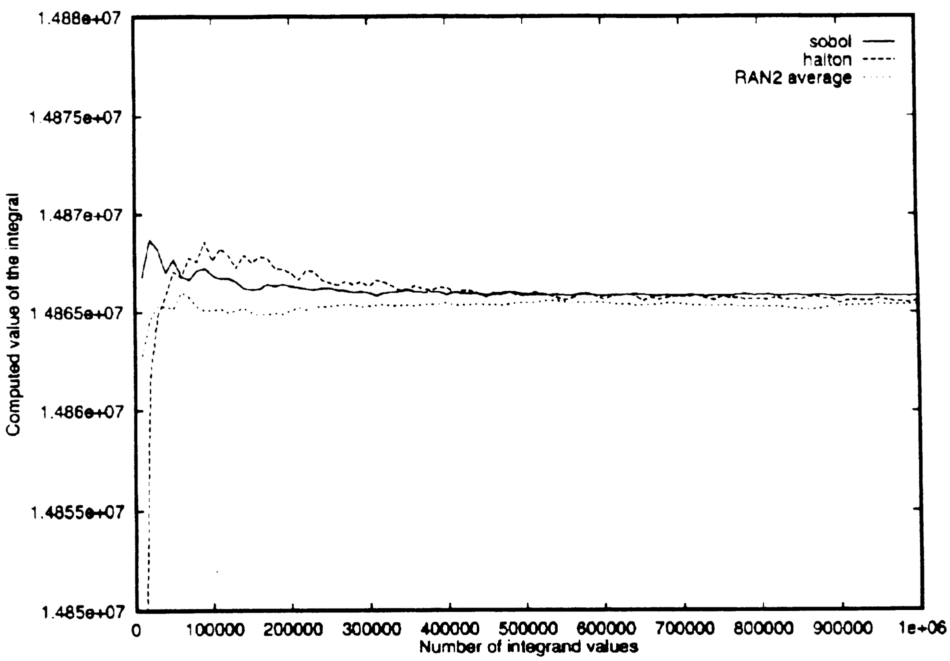

EXHIBIT 3
AUTOMATIC TERMINATION CRITERION APPLIED TO SOBOL, HALTON, AND THREE MONTE CARLO RUNS USING RAN2 FOR TRANCHE A

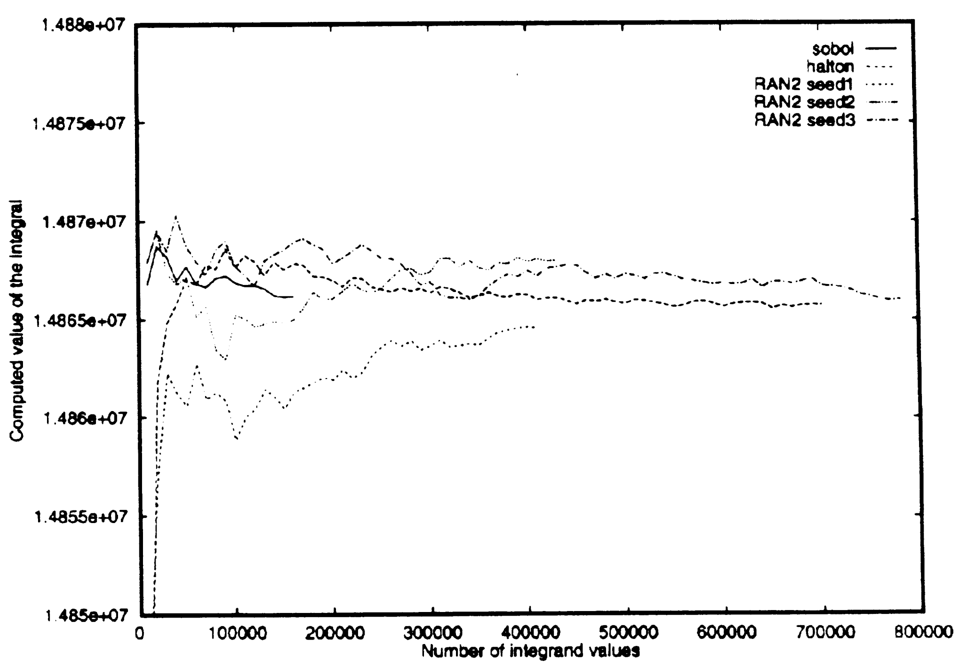

is that they can be utilized very generally. This is important in a number of situations:

- If a financial house has a book with a wide variety of derivatives, it is advantageous to use methods that do not need to be tuned to a particular derivative.
- If a new derivative has to be priced, there is no immediate opportunity to tailor a variance reduction technique to a particular integrand.

Variance reduction techniques are commonly used in conjunction with Monte Carlo methods. Although variance reduction techniques can be very powerful, they can require considerable analysis before being applied. We therefore limit ourselves here to just one variance reduction technique: antithetic variables. The advantage of antithetic variables is its ease of use. Tests reveal that it is superior to the Monte Carlo method for our CMO problem. We emphasize that antithetic variables is not a palliative; it can be inferior to the Monte Carlo method.

Exhibit 4 is analogous to Exhibit 1. It compares the results of Sobol, Halton, and antithetic variables runs with two randomly chosen initial seeds. The data graphed in Exhibit 4 are typical of our results.

From these results we conclude that for this CMO:

- The Sobol method consistently outperforms the antithetic variables method.
- Convergence of the antithetic variables method is [illegible] jagged than convergence of the Monte Car[illegible] method.
- The antithetic variables method consistently outperforms the Monte Carlo method.

Further results regarding antithetic variables may be found in Paskov [1994].

SMALL NUMBER OF SAMPLE POINTS

Results for a small number of points are sometimes of special importance for people who evaluate CMOs and other derivative products. They need methods that can evaluate a derivative in a matter of minutes. Rather low accuracy, on the order of 10^{-2} to 10^{-4}, is often sufficient. The integrands are complicated and computationally expensive. Furthermore, many may have to be evaluated on a daily basis with limited computational resources, such as workstations.

We therefore compare the performance of the Sobol method with Monte Carlo and antithetic variables for 4,000 sample points. This leads to reasonable results and takes less than a couple of minutes of workstation CPU time. We believe that comparable results may hold for other mortgage-backed securities and interest rate derivatives. We drop the Halton method from consideration in this section because it is outperformed by both the Monte Carlo and antithetic variables methods

EXHIBIT 4
SOBOL AND HALTON RUNS FOR TRANCHE A AND TWO ANTITHETIC VARIABLES RUNS USING RAN2

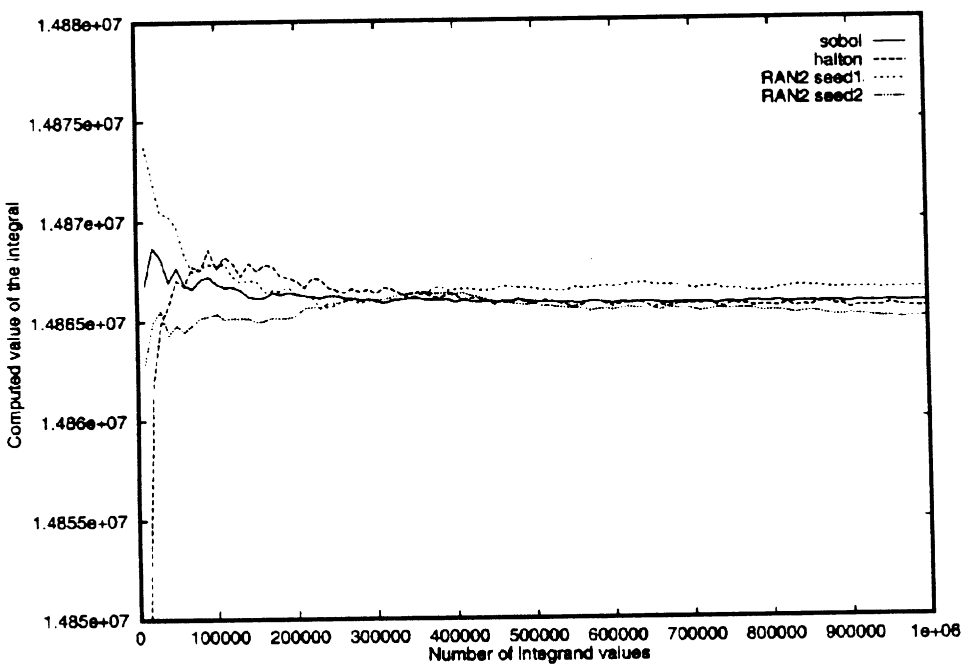

EXHIBIT 5
Number of "Wins" of the Monte Carlo Method and the Sobol Method

Tranche	Monte Carlo	Sobol
A	3	17
B	0	20
C	3	17
D	3	17
E	2	18
G	0	20
H	0	20
J	0	20
R	8	12
Z	4	16

for 4,000 sample points. Sometimes computational speed is paramount. It would therefore also be of interest to study a smaller number of points.

For each of the ten tranches we compute twenty approximate answers using the Monte Carlo method with twenty random initial seeds. For each tranche we also compute an approximation using Sobol points. We compute as well the relative errors of all these approximations. To compute the relative errors we needed estimates of the true answers, which we obtained using antithetic variables with 20 million points.

The results are summarized in Exhibit 5. We say a method wins if it has a smaller relative error. (Recall we are fixing the number of samples at 4,000.) Sobol points win for every tranche. In total, the Sobol method wins 177 times out of 200 cases; that is almost 90% of the time.

Exhibit 6 shows the result of comparing the Sobol method with the antithetic variables method.

EXHIBIT 6
Number of "Wins" of the Antithetic Variables Method and the Sobol Method

Tranche	Antithetic Variables	Sobol
A	9	11
B	1	19
C	6	14
D	10	10
E	11	9
G	2	18
H	3	17
J	2	18
R	8	12
Z	9	11

The Sobol method wins for eight of the tranches, ties for one, and loses for one. In total, Sobol wins almost 70% of the time.

The Sobol method achieves accuracies ranging from one part in a thousand to one part in a million, depending on the tranche. It takes about 103 seconds to compute the Sobol results and about 113 seconds to compute the antithetic variables results for all ten tranches running on a Sun-4 workstation.

CLOSING REMARKS

We performed statistical analysis for a small number of sample points. Methodology and results are reported in Section 9 of Paskov [1994]. Here we confine ourselves to noting just one conclusion.

- Statistical analysis on the small sample case further strengthens the case for the Sobol method over the antithetic variables method. For example, to achieve similar performances at a confidence level of 95%, the antithetic variables method needs from seven to seventy-nine times more sample points than the Sobol method, depending on the tranche.

In closing, we suggest some directions for future work:

- Compare the performance of low-discrepancy and Monte Carlo methods on other financial derivatives.
- Test the performance of other known low-discrepancy sequences on various derivatives.
- Characterize analytic properties of classes of financial derivatives and design new methods tuned to these classes.
- Study error reduction techniques for deterministic methods.
- Because results for a small number of samples are often of special interest in finance, it would be attractive to design new deterministic sequences that are very uniformly distributed for a small number of points.
- There are numerous open theoretical problems concerning high-dimensional integration and low-discrepancy sequences. We believe that their solution will aid in the design of better methods for finance problems.

ENDNOTE

The authors thank the National Science Foundation and the Air Force Office of Scientific Research for partial support. They are grateful to Goldman Sachs for providing the finance problem. The authors benefited also from discussions with S. Baker, A. Belur, B. Bojanov, T. Boult, P. Boyle, P. Cheah, J. Davis, S. Heinrich, P. Karasinski, J. Langsam, N. Marinovich, P. Niculescu, E. Novak, D. Schutzer, B. Shukhman, I. M. Sobol, V. Temlyakov, S. Tezuka, J. Tilley, I. Vanderhoof, T. Warnock, A. Werschulz, R. Wertz, and H. Wozniakowski.

REFERENCES

Niederreiter, H. *Random Number Generation and Quasi-Monte Carlo Methods*. CBMS-NSF, 63, Philadelphia: SIAM, 1992.

Paskov, S.H. "New Methodologies for Valuing Derivatives." Technical Report, Computer Science Department, Columbia University, October 1994.

Press, W., S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C*, Second Edition. Cambridge: Cambridge University Press, 1992.

"Suddenly, Number Theory Makes Sense to Industry." *Business Week*, June 20, 1994, pp. 172-174.

Traub, J.F., and H. Wozniakowski. "Breaking Intractability." *Scientific American*, January 1994, pp. 102-107.