BEATING MONTE CARLO

Simulation methods using low-discrepancy point sets beat Monte Carlo hands down when valuing complex financial derivatives, report

Anargyros Papageorgiou **and** Joseph Traub

Monte Carlo simulation is widely used to price complex financial instruments, and much time and money have been invested in the hope of improving its performance. However, recent theoretical results and extensive computer testing indicate that deterministic methods, such as simulations using Sobol or Faure points, may be superior in both speed and accuracy.

In this paper, we refer to a deterministic method by the name of the sequence of points it uses, eg, the Sobol method. We tested the generalised Faure sequence due to Tezuka (1995) and a modified Sobol method which includes additional improvements to those documented in Paskov & Traub (1995). We compared these two low-discrepancy deterministic methods with basic Monte Carlo in the valuation of a collateralised mortgage obligation (CMO).

We found that deterministic methods beat Monte Carlo:

☐ **by a wide margin**. In particular:
(i) Both the generalised Faure and modified Sobol methods converge significantly faster than Monte Carlo.
(ii) The generalised Faure method always converges at least as fast as the modified Sobol method and often faster.
(iii) The Monte Carlo method is sensitive to the initial seed.

☐ **for a small number of sample points:**
(i) Deterministic methods achieve a low error level with a small number of points.
(ii) For the most difficult CMO tranche, generalised Faure achieves accuracy of 10^{-2} with 170 points, while modified Sobol uses 600 points. The Monte Carlo method, on the other hand, requires 2,700 points for the same accuracy.
(iii) Monte Carlo tends to waste points due to clustering, which severely compromises its performance when the sample size is small.

☐ **as the sample size and the accuracy demands grow.** In particular:
(i) Deterministic methods are 20 to 50 times faster than Monte Carlo (the speed-up factor) even with moderate sample sizes (2,000 deterministic points or more).
(ii) When high accuracy is desired, deterministic methods can be as much as 1,000 times faster than Monte Carlo.

1. Distribution of 512 pseudo-random points

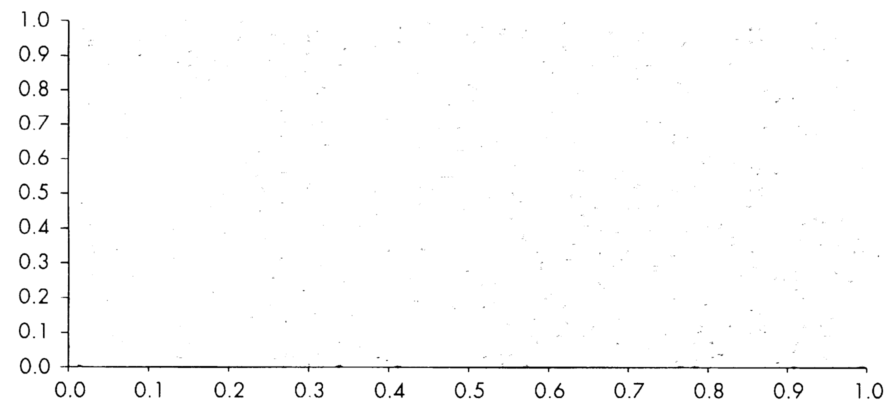

2. Distribution of 512 low-discrepancy points

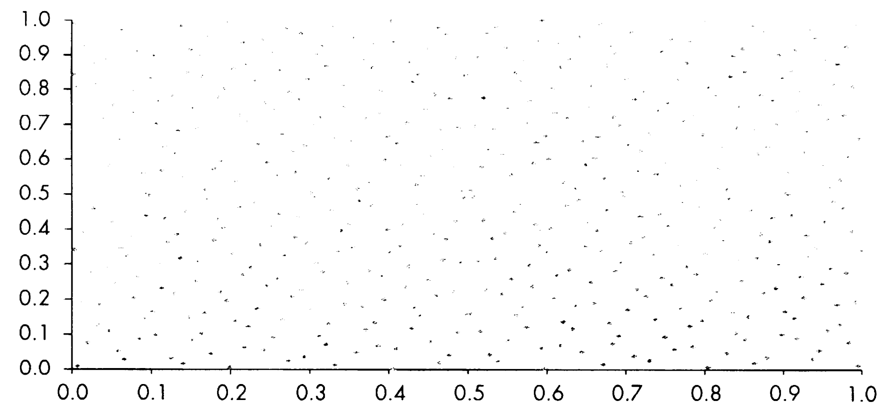

There are two ways of valuing financial derivatives: via paths or as a high-dimensional integral. For simplicity, we will restrict ourselves to a discussion of the integral formulation and, without loss of generality (Paskov, 1996), the integral over the unit cube in d dimensions. For most finance problems, this integral cannot be analytically computed; we have to settle for a numerical approximation.

The basic Monte Carlo method obtains this approximation by computing the arithmetic mean of the integrand evaluated at randomly chosen points. More precisely, only pseudo-random points can be generated on a digital computer and these are used in lieu of random points. There are sophisticated variations of this method but we refer only to the basic version in this paper.

If pseudo-random points from a flat distribution are plotted on the unit square in two dimensions (see figure 1), there are some regions where no sample points occur and others where the points are more concentrated. This is clearly undesirable. Random point samples are wasted due to clustering. Indeed, Monte Carlo simulations with very small sample sizes cannot be trusted. It would be better to place our sample points as uniformly as possible, which is the idea behind low-discrepancy sequences. Discrepancy is a measure of deviation from uniformity; hence low-discrepancy points are desirable. Figure 2 shows a plot of certain low-discrepancy points on the unit square in two dimensions.

A low-discrepancy method approximates the integral by computing the arithmetic mean of the integrand evaluated at low-discrepancy points. Low-discrepancy sequences have been extensively studied.[1] In contrast to the Monte Carlo method, these use deterministic points. They are sometimes said to be quasi-random.

In 1992, the conventional wisdom was that although theory suggested that low-discrepancy methods were sometimes superior to Monte Carlo, this theoretical advantage was not seen for high-dimensional problems. Joseph Traub and a PhD student, Spassimir Paskov, decided to compare the efficacy of low-discrepancy and Monte Carlo methods in valuing financial derivatives.

They used a CMO (Fannie Mae REMIC Trust 1989–2023) provided by Goldman Sachs, with 10 tranches requiring the evaluation of 10 integrals, each over 360 dimensions, and a particular low-discrepancy sequence due to Sobol. The values of the tranches depended on the interest rate and prepayment models used. Paskov & Traub made major improvements in the Sobol points, which led to a more uniform distribution. The improved Sobol method consistently outperformed Monte Carlo (Paskov & Traub, 1995, and Paskov, 1996).

Software construction and testing of low-discrepancy deterministic methods for pricing financial derivatives began at Columbia University in autumn 1992. Preliminary results were shared with a number of New York City financial houses in autumn 1993 and spring 1994. The first published announcement was Traub and Wozniakowski's January 1994 article in *Scientific American*. A more detailed history is given in Paskov & Traub (1995).

In September 1995, IBM announced a product called the Deterministic Simulation Blaster, which uses a low-discrepancy deterministic method (*Risk* Technology Supplement, August 1995, pages 23–24; *Risk* November 1995, page 47). The company claimed a very large improvement over Monte Carlo. However, IBM has not revealed the method for choosing the sample points and the methodology for calculating the speed-up. IBM acknowledges that Columbia University pioneered the use of low-discrepancy methods to price financial derivatives.

Results

We have built a software system called Finder[2] for computing high-dimensional integrals which has modules for generating generalised Faure points and modified Sobol points. It includes major improvements. Indeed, a number of financial institutions have informed us that they could not replicate our results using, for example, the Sobol point generator found in Press *et al* (1992).

We used Finder to price the CMO and to compare low-discrepancy methods with Monte Carlo simulation. As both types of method compute the arithmetic mean of the integrand evaluated at a number of points, the difference in performance depends on the number of points that each method uses to achieve the same accuracy. We observe the least number of points a method needs to achieve and maintain a relative error below a specified level, say, 10^{-2}. This is gauged by the speed-up of one method relative to another, which we define as the ratio of the least number of points required by one method, divided by the least number of points required by the other method. This ensures that both methods maintain the same level of accuracy.

This definition of speed-up is new: we study the convergence and the error of a method throughout a simulation. We believe that this has advantages over speed-up calculations which are based only on the error values at the end of a simulation. Note that our definition of speed-up is a more rigorous requirement than only computing the confidence level of Monte Carlo.

Our extensive testing has shown that fixed accuracy requires different tranches to be treated with different numbers of points. Again, we emphasise that deterministic methods beat Monte Carlo for every tranche. Our results are reported using the residual tranche of the CMO, which depends on all 360 monthly interest rates, as the reference point, since it is the most difficult to price. If this tranche can be priced with a given accuracy using a certain number of samples, the same number of samples will yield at least the same accuracy for the rest of the tranches.

Since pricing models for complicated derivatives are subject to uncertainty, financial houses are often content with relative errors of one part in a hundred. Furthermore, if they

[1] See Paskov (1996) for the formal definition of discrepancy and an extensive bibliography

[2] Finder may be obtained from Columbia University

wish to price a book of instruments, it is critical to use a small number of samples. Deterministic methods achieve a relative error of one part in a hundred using a small number of points. In figures 3 and 4 we see that some 170 generalised Faure points or 600 Sobol points are sufficient for a relative error equal to 10^{-2}, whereas Monte Carlo requires 2,700 points for the same level of accuracy. A very small number of generalised Faure points thus yields an accurate price 16 times faster than Monte Carlo.

A further reduction of the error by a factor of 20 (equal to $10^{-3}/2$) requires about 16,000 generalised Faure points while Monte Carlo may require up to 800,000 random points, which yields a speed-up factor of up to 50. In general, samples using as few as 2,000 generalised Faure points can price the CMO 20 to 50 times faster than Monte Carlo. As far as convergence rates are concerned, for $n \leq 10^4$ generalised Faure points, the error is proportional to $n^{-0.82}$. This error estimate is conservative, since a much higher convergence rate is frequently attained but, in any case, it is very much superior to the $n^{0.5}$ expected Monte Carlo error.

Monte Carlo also exhibits a great sensitivity to the seed of the pseudo-random number generator. So, unless we are dealing with the result of a fairly long simulation, we cannot have much confidence. The long simulations needed yield a deterministic method speed-up of about 1,000.

3. Generalised Faure and Monte Carlo errors

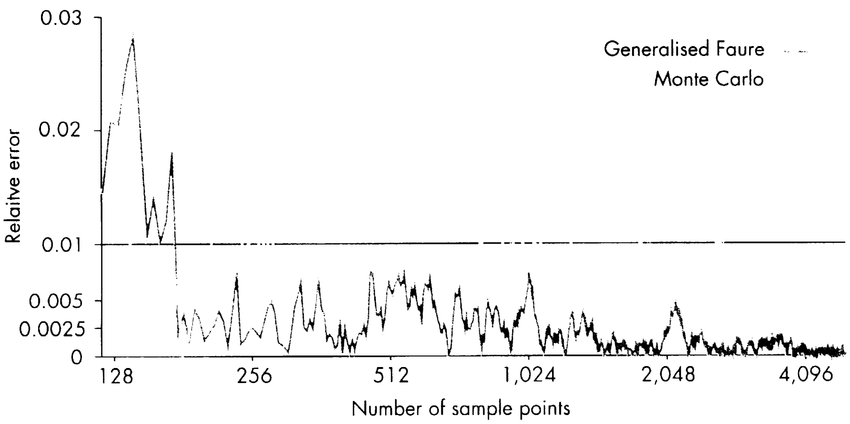

4. Generalised Faure and Sobol errors

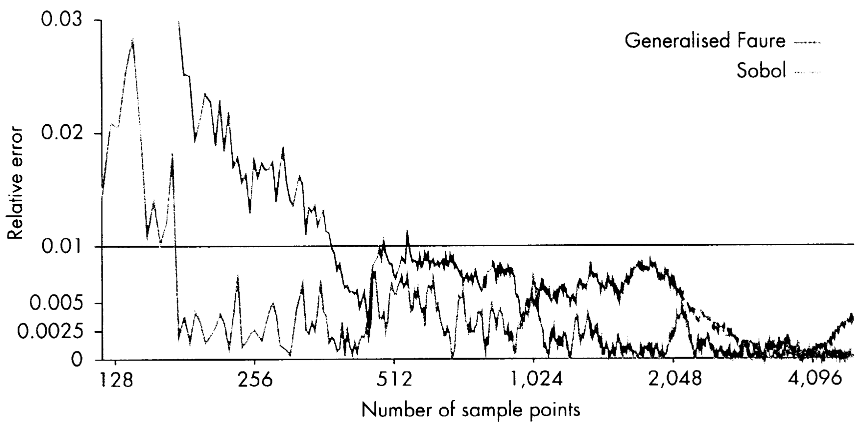

Conclusion

The best of the deterministic methods we have tested is the one based on generalised Faure points. These usually achieve the same accuracy as the Sobol points 2.5 to 6.5 times faster. They can also be produced efficiently, at a cost similar to that for random points, and a only a small number of points is needed to price the CMO. In contrast to some other deterministic sequences, generalised Faure points can be easily produced in very high dimensions. It is much more complicated to obtain the improved Sobol points that we have been using in very high dimensions.

Finder contains features that further improve the quality of the approximation obtained by the generalised Faure method without any additional computational costs. Finally, preliminary but very encouraging results indicate that generalised Faure points can efficiently price financial derivatives modelled in more than 1,500 dimensions. Future work will include: further improvements in Finder; a comparison of the performance of low-discrepancy and Monte Carlo methods on other financial derivatives; studying the possibilities that low-discrepancy methods offer for risk management; and designing new low-discrepancy methods tailored for financial computations. ■

The authors both work at Columbia University. Anargyros Papageorgiou is a researcher e-mail: ap@cs.columbia.edu

Joseph Traub is the Edwin Howard Armstrong professor of computer science e-mail: traub@cs.columbia.edu

Note: a version of this paper was presented on April 15, 1996 at Mathematical Problems in Finance, *organised by the Institute for Advanced Study, Princeton, New Jersey. This work was supported in part by the National Science Foundation*

References

Paskov, SH, and JF Traub, 1995, *Faster valuation of financial derivatives*, The Journal of Portfolio Management, Fall 1995, pages 113–120

Paskov, SH, 1996, *New methodologies for valuing derivatives*, in Mathematics of Derivative Securities, edited by S Pliska and M Dempster, Isaac Newton Institute, Cambridge University Press, Cambridge, UK, forthcoming

Press, W, S Teukolsky, W Wetterling and B Flannery, 1992, *Numerical recipes in C*, second edition, Cambridge University Press, Cambridge

Tezuka, S, 1995, *Uniform random numbers: theory and practice*, Kluwer Academic Publishers, Boston

Traub, JF, and H Wozniakowski, 1994, *Breaking intractability*, Scientific American, January 1994, pages 102–107