Measuring Portfolio Risk Using Quasi Monte Carlo Methods

Alexander Kreinin, Leonid Merkoulovitch, Dan Rosen and Michael Zerbs

We consider the application of Quasi Monte Carlo methods to risk measurement. A scenario generation technique based on low discrepancy Sobol sequences is compared to one based on pseudo-random number generation. The comparison methodology includes the calculation of the number of scenarios necessary to guarantee the required accuracy level of the Value-at-Risk estimate and the calculation of the relative performance of the two methods. The methodology is illustrated by a case study of a multi-currency portfolio. We demonstrate that Quasi Monte Carlo methods significantly improve the performance of the portfolio simulation and, therefore, reduce the time to obtain reliable risk measurements.

Accurate and efficient risk measurement is at the core of an enterprise-wide risk management strategy. This strategy requires that financial institutions implement a risk management engine to compute market risk and credit risk of their full portfolio. For example, as a regulatory requirement banks report the Value-at-Risk (VaR) of their portfolios daily. However, the need for a robust risk engine extends beyond regulation, and is becoming a central component in the way financial institutions manage their business today.

It is not uncommon for the portfolio of a bank or an insurance company to contain several hundred thousand positions, including substantial volumes of derivative products such as swaps, caps and floors, swaptions and mortgage-backed securities. Various analytical shortcuts can be used to estimate the risks of these portfolios but, in general, simulation methods are necessary to calculate risk accurately when the portfolios contain substantial positions in instruments with optionality, or when the distributions of the underlying risk factors are not normal. However, a full simulation of a large and complex portfolio is computationally expensive, and may not even be achievable within a reasonable time using top-of-the-line computers.

In this paper, we contrast the use of a Quasi Monte Carlo (QMC) method based on low discrepancy Sobol sequences with the standard Monte Carlo (MC) method based on pseudo-random sampling to compute the distribution of future portfolio value and to measure VaR.

QMC methods were first introduced in the finance literature by Boyle (1977), and during the past decade they have received considerable attention for derivatives pricing. Recent papers by Boyle (1996) and by Acworth et al. (1997) have reported results of computational comparisons of QMC versus MC methods for options pricing. Studies indicate that for pricing problems QMC methods outperform MC simulation, and that of the low discrepancy sequences (LDS) tested, the Sobol sequences perform best.

However, QMC methods have not been applied to measure portfolio risk, where the emphasis is

on estimating the tails of distributions and not their averages. To the best of our knowledge, there are no published articles on their application to VaR estimation. However, QMC methods have a strong intuitive appeal for risk management problems. Low discrepancy sequences specifically attempt to cover the space of risk factors "evenly," thus avoiding the clustering usually associated with pseudo-random sampling. Sampling evenly seems to be a desirable property not only when estimating the average of a distribution, but also when searching for exceptional cases in the tails.

Calfisch et al. (1997) note that QMC methods converge slowly when applied to problems in high dimension. It is well-known that the profit and loss distribution of an institution's portfolio is a function of a large number of risk factors, commonly hundreds. Although the application of QMC methods to portfolio VaR calculation seems to be problematic, the dimensionality limitation has been overstated in two ways. It is common to project the numerous risk factors onto a reduced set of risk factors depending on the calculation method used and data availability, and of this reduced set, only a subset will be independent risk factors. The dimensionality of problems of the size of this independent risk factor space is not a limitation to the application of QMC methods (Kreinin et al. 1998).

This paper is organized as follows. We begin with a brief introduction of simulation methods used to measure VaR. This is followed by an explanation of some basic properties of Monte Carlo methods and the ideas behind low discrepancy sequences. The presentation is informal, since our intention is to present in simple, intuitive terms, the basic properties of these techniques and the benefit of applying them to problems in risk management. A formal mathematical treatment can be found in Niederreiter (1992). The next section presents the case study. Finally, we offer some concluding remarks and comment on directions for future development.

Simulation methods for VaR estimation

The Value-at-Risk of a portfolio is generally defined as the maximum loss that is expected to occur at some level of confidence, α, over a specified period of time:

$$Pr\{-\Delta V(t,\mathbf{r}(t)) \le VaR(\alpha)\} = \alpha \qquad (1)$$

where $\Delta V(t,\mathbf{r}(t))$ is the unrealized profit and loss (UP&L) of the portfolio, between today and time t:

$$\Delta V(t,\mathbf{r}(t)) = V(t,\mathbf{r}(t)) - V_0$$

$\mathbf{r}(t)$ is the vector of risk factors, and V_0 is the mark-to-market value of the portfolio today. Usually t is one to ten days, and α is 95% to 99%.

In general, simulation approaches to estimate VaR first generate a large number of joint scenarios, N, on the vector of risk factors $\mathbf{r}(t)$ at the time horizon t. The portfolio is then revalued under each of these scenarios. To obtain an accurate estimate of the portfolio distribution the portfolio is evaluated under each of the N scenarios. A non-parametric estimate of VaR is obtained by ordering the changes in value under each scenario in ascending order, calculating the cumulative probability and identifying that scenario with cumulative probability lower than $1 - \alpha$. Parametric approximations of VaR can also be obtained by fitting a given parametric distribution to the simulation outcomes.

Simulation techniques differ in

- the choice of the underlying distribution of the risk factors
- the way in which samples are drawn from this distribution, and
- the manner in which the valuation function $V(t,r(t))$ is implemented.

Ideally, full valuation of each instrument is performed under each scenario, although this may be rather costly. To accelerate the

computation, approximations such as Taylor series expansions, cashflow bucketing techniques or more general portfolio compression techniques may be used.

Historical simulation draws scenarios from the observable discrete historical changes in the risk factors during a specified period of time. The Monte Carlo simulation approach, on the other hand, starts from a specified joint distribution of these changes (usually estimated from historical data) and applies statistical techniques to draw random samples from this distribution. The most common assumption is that the distribution of the changes in risk factors is joint log-normal, but this need not be the case.

In this paper, we focus on improving the performance of the Monte Carlo simulation by introducing an advanced sampling technique. Standard Monte Carlo methods draw random samples from the distribution using a pseudo-random number generator. By contrast, Quasi Monte Carlo methods use deterministic points generated from a type of mathematical vector-sequences called low discrepancy sequences. The idea behind QMC techniques is that by choosing points in the risk factor space more evenly, we reduce the number of scenarios necessary to achieve a desired level of accuracy in a VaR calculation. In the next sections we present a brief description of LDS and examine the characteristics that result in accurate VaR estimation with a reduced number of scenarios. Thereafter, we present an example comparing the performance of MC methods to that of a QMC method based on the low discrepancy Sobol sequence.

Monte Carlo methods

Let us outline several basic properties of the Monte Carlo method that are necessary to contrast the results to those obtained by QMC methods.

Since the scenarios are drawn randomly in Monte Carlo methods, simulations with a finite number of points, N, starting from different initial seeds, yield different results. Thus, we can obtain probabilistic errors for the estimates. These errors are measured by the standard deviation of the estimate, σ, over many similar simulations. Clearly, as N grows, we expect the difference between the outcomes of any two simulations, and hence σ, to decrease. A well known result from probability theory is that the error varies inversely to the square root of the number of draws:

$$\sigma \sim O\left(\frac{1}{\sqrt{N}}\right)$$

Thus, if we increase the number of draws by a factor of four, we can expect the errors to be reduced by half.

Applying MC methods to calculate the VaR also yields probabilistic bounds on the VaR estimates. If no parametric assumptions are made on the distribution of the UP&L, these bounds are computed using rank statistics as shown in Morokoff et al. (1998).

The main advantages of standard Monte Carlo methods are

- they are generally applicable to all problems
- their rate of convergence is independent of the dimensionality of the risk factor space
- they are very popular and their properties are well known
- they yield probabilistic errors and *a priori* bounds on VaR estimates.

Their main disadvantages are

- pseudo-random number generators tend to generate clusters of points
- they do not explicitly exploit particular features of the problems
- the rate of convergence is slow.

Quasi Monte Carlo methods and low discrepancy sequences

If d is the dimension of the space of independent risk factors, then sampling methods generally sample from a unit d-dimensional hypercube. Whereas Monte Carlo methods are based on points randomly generated from the hypercube,

Quasi Monte Carlo methods are based on deterministic sequences of points that satisfy the property that their elements evenly cover the hypercube. The measure of how evenly a sequence of points covers the region is called **discrepancy**; the more evenly the points are distributed in the region, the lower the discrepancy. Hence, these sequences have been termed low discrepancy sequences.

Although the idea behind LDS is simple, the mathematical theory and the algorithms for generating the sequences are far from trivial. A detailed explanation of the solutions and algorithms is found in Sobol (1967) and Niederreiter (1992).

A sequence would cover the hypercube uniformly if the number of points in any possible subset of the hypercube were proportional to the volume of the subset. Discrepancy measures the worst deviation between the volume of the subset and the fraction of the number of points in the subset over the total number of points, over all possible subsets. In practice, however, the discrepancy of a sequence is not easy to determine. Hence, a simpler measure, called star discrepancy**,** is generally used. **Star discrepancy** considers only those subsets having the form of a hypercube, instead of all possible subsets.

An important mathematical result is that the asymptotic form of the star discrepancy, D_N^*, is described by the relation:

$$D_N^* \sim O\left(\frac{(\log N)^d}{N}\right) \qquad (2)$$

Note that star discrepancy decreases asymptotically as *1/N*, but increases with the dimensionality, *d*, of the problem.

Why does QMC simulation seem like a promising tool in risk management? In risk management, we are particularly concerned with the behaviour of the tails of the distribution of changes in value. Intuitively, it seems that a method that samples the space more evenly will better describe behaviour in the tails. We are disposed to think that approximations of multivariate distributions with better coverage of the whole space will be superior to those based on pseudo-random number generation and that QMC methods will allow us to generate a smaller number of scenarios without loss of accuracy in the final approximation.

Figure 1 provides a comparison of two-dimensional pseudo-random vectors and points obtained from a two-dimensional Sobol sequence. Note that the Sobol points cover the unit square more uniformly than do the randomly generated points. Thus, the discrepancy of the Sobol points is lower than that of the pseudo-random points. The star discrepancy of a Sobol sequence is optimal when $N=2^n$, $n=1, 2, \ldots$

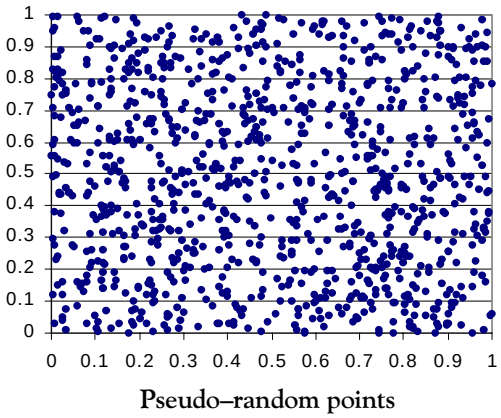

Figure 1: Two-dimensional pseudo-random points vs. two-dimensional Sobol sequences

Since points generated by LDS are deterministic, QMC methods do not yield probabilistic error estimates. (Recent papers use modifications of

LDS that lead to probabilistic errors, e.g., Boyle (1996)). Generally, errors in QMC estimates are measured against the true value of the derivative price which is assumed to be known. In structured experiments, one must choose simple problems for which the true value can be obtained analytically or computed with a large number of draws. In practice this value is not known for every portfolio or new circumstance that arises.

In summary, the advantages of QMC methods are

- they are based on sampling techniques that generate points evenly within the region and avoid the clustering generally associated with MC methods
- they have been well-tested in the non-financial and financial literature (for derivatives pricing).

Their main disadvantages are

- their lack of generality when compared to MC methods means that their effectiveness may be largely dependent on the problem and extensive testing is required
- they do not yield probabilistic errors or *a priori* bounds on VaR estimates
- their rate of convergence depends on the dimensionality of the risk factor space, d
- they may be inefficient for problems in very large dimensions.

Case study

We compare the performance of the MC and QMC methods to measure VaR for a simple multi-currency test portfolio. The accuracy of the estimations is measured with respect to the "true" UP&L distribution of the portfolio, which is computed with negligible error using a very large number of MC scenarios. Our main objectives are to measure the speed-up obtained from QMC methods, and to assess whether it depends on the accuracy desired in the VaR calculation or the confidence level used.

In what follows, the comparison methodology and the test portfolio are described, followed by a presentation of the results.

Comparison methodology

Our first step is to establish the "true" portfolio distribution. The portfolio distribution is computed using a MC simulation with a number of scenarios large enough to guarantee that errors are negligible. Fixing a confidence level for VaR, we can then compare the results of the simulations. Then, for a given simulation, the error, ε, is the absolute difference between the VaR estimate and the "true" VaR.

Recall that N is the number of samples drawn in a simulation. For a given MC simulation, we define the **effective number of scenarios**, $N(\alpha,\varepsilon)$, as the minimum number of scenarios required to estimate $VaR(\alpha)$ with an error that does not exceed ε. To obtain $N(\alpha,\varepsilon)$, for both MC and QMC simulations we add scenarios to the simulation until the error in the VaR remains below the desired level when more scenarios are added. Since the convergence of the VaR estimates with both techniques is not necessarily monotonic, N does not necessarily represent the first time that the error in the simulation falls below ε.

Since MC simulations initiated with different seeds generate different VaR estimates, the effective number of scenarios, $N(\alpha,\varepsilon)$, is itself a random variable. Its probability distribution can be estimated by performing a large number of simulations with different initial seeds. We denote the β quantile of this distribution by $N_\beta(\alpha,\varepsilon)$. Thus, for example, $N_{0.95}(0.99,0.02)$ (β=95%, α=99%, ε=2%) is the number of scenarios required to be 95% certain that the estimation of VaR (99%) has an error of at most 2%. Note the distribution of N cannot be determined *a priori* with a single simulation, but only after a number of similar simulations with different initial seeds. In contrast to MC, the effective number of scenarios for QMC is deterministic, not a random variable.

We define the speed-up of the QMC relative to the MC simulation as:

$$S_{\beta}(\alpha,\varepsilon) = \frac{N_{\beta}(\alpha,\varepsilon)}{N^*(\alpha,\varepsilon)} \qquad (3)$$

where $N^*(\alpha,\varepsilon)$ is the effective number of scenarios for the QMC method. Note that we have explicitly made this speed-up a function of the VaR level, the desired accuracy in the VaR estimation, and the confidence level for the MC simulation to achieve a given accuracy.

Test portfolio and data

The test portfolio is part of a suite of test data used for benchmarking in the industry (Marshall and Seigel 1996). Its base currency is USD and it contains fourteen positions, both long and short, in fixed-rate government bonds, in maturities ranging from 182 days to 10 years, in five currencies: USD, DEM, FRF, ITL and JPY. The mark-to-market of the portfolio is 357.3 million USD.

The zero curves in each currency are modeled using 16 node points, except for the JPY curve which has 15 nodes. Market data for the zero curves is as of September 26, 1997. We assume that the vector of 83 risk factors (16*4 + 15 + 4 FX) is log-normal with zero mean and covariance matrix as published by RiskMetrics™ on that day (J.P. Morgan 1996).

An important property of the covariance matrix is that although its dimensionality is 83, its *rank* is less than 40 because the maximum number of independent risk factors is less than 40. (Mathematically, this means that the matrix contains fewer than 40 non-zero (numerically) eigenvalues and, hence, independent eigen-vectors.) This property is an intrinsic characteristic of covariance matrices computed using exponentially weighted moving averages, such as those provided by RiskMetrics ™ (J.P. Morgan 1995), regardless of the dimensionality of the original risk factor space.

Although the detailed explanation of this point is beyond the scope of this paper, it is important to note that the dimensionality of the MC and QMC simulation arises from the number of independent risk factors (the rank of the matrix), not from the number of underlying risk factors.

Results

Our first step is to compute the "true" 1-day UP&L distribution of the portfolio and the VaR profile. Figure 2 depicts the portfolio VaR against the level of confidence, α. This graph is calculated using a very large number of MC scenarios (over 800,000). Hence, for all intents, the errors are negligible. From these results, the one-day VaR(95%) is 4.81 million USD.

Figure 2: True VaR vs. level of confidence, α

In what follows, we compare the performance of QMC and MC for VaR(α), with α=95%.

Monte Carlo simulation

Figure 3 shows the error in VaR(95%) versus the number of scenarios, for three sample MC runs with different seeds. Note that the convergence for each run is not monotonic. Also, note the dispersion of the effective number of scenarios, $N(0.95,*)$. For example, at a 2% error, the effective number of scenarios, $N(0.95,0.02)$, varies between 4000 and 9700.

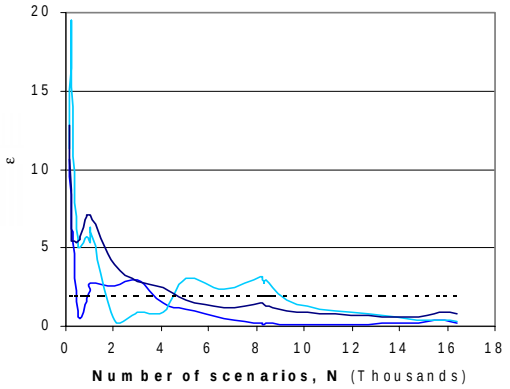

Figure 3: Error vs. number of scenarios for three sample runs

To estimate the distribution of the effective number of scenarios, $N(\alpha,\varepsilon)$, we perform 50 simulations initiated with different seeds, each with 16,384 scenarios.

The cumulative distribution function of $N(0.95,0.02)$ is illustrated in Figure 4. The figure also shows the fit of a gamma distribution. The 5th percentile of the distribution gives the number of scenarios required to ensure the VaR estimate is within 2% accuracy, 95% of the time. This value is $N_{0.95}(0.95, 0.02) \approx 14000$.

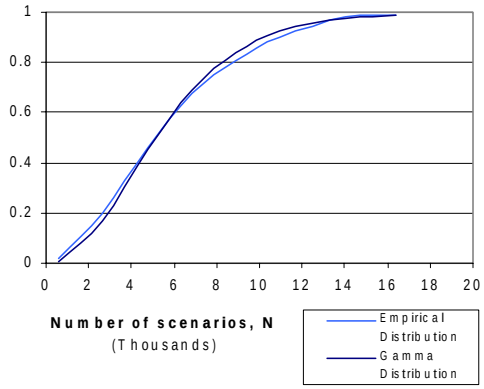

Figure 4: Cumulative distribution function of N

Quasi Monte Carlo simulation with Sobol sequences

For comparison, we determine the number of scenarios required to achieve 2% accuracy with the QMC method. Since the QMC method is deterministic, there is no notion of confidence levels. Figure 5 illustrates the convergence behaviour of QMC simulation, as well as that of the average and of the 5th percentile for the MC simulation.

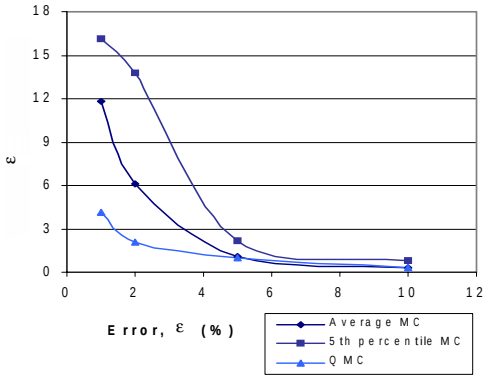

Figure 5: Sobol vs. mean N and 95%N for 2% error

From Figure 5, the effective number of scenarios for QMC at a 2% error is $N^*(0.95, 0.02) = 2048$. Hence, the speed-up at a 2% error is $S_{0.95}(0.95, 0.02) = 6.7$. Compared to the MC method, the QMC method requires almost seven times fewer scenarios to compute VaR(99%) with a 2% error at a 95% confidence level. Five out of 100 times, a MC simulation with roughly seven times more scenarios will yield a worse VaR estimate than can be achieved by the QMC method. However, we do not know *a priori* when this will happen.

Note from Figure 5 that the QMC method becomes less effective as a lower accuracy is required. Within a 5% error, the speed-up is only $S_{0.95}(0.95, 0.05) = 2.1$.

For MC simulation we have plotted the average of the distribution of $N(0.95,\varepsilon)$ in addition to the 5th percentile in order to provide a more general picture of its dispersion. However, it is more meaningful to measure the speed-ups derived from the QMC method with respect to a high confidence level of the MC simulation since we are interested in calculating accurate VaR estimates consistently, not on average. This is not an issue in deterministic QMC simulation.

Table 1 summarizes the speed-ups obtained with the QMC method for VaR(95%) and VaR(99%) at 2% and 5% error levels in the estimation. For VaR(99%), more than 1% of the runs did not converge in over 16,000 scenarios, hence the speed-up was estimated from the gamma approximation to the distribution of N (Figure 4).

	Error, ε	
Confidence Level, α	2%	5%
95%	6.7	2.1
99%	9.4	2.7

Table 1: Speed-ups for QMC compared to MC (β=95%)

As this example shows, the QMC method provides substantial speed-ups for VaR

calculations over MC methods. Although the QMC method generally outperforms the MC method, it appears to provide higher benefits when higher accuracy and a higher confidence level are required. Thus, for example, the speed-up for a VaR calculation with 2% accuracy increases from 6.7 to 9.4 if we want a 99% confidence level as opposed a 95% confidence level in the tails ($S_{0.95}(0.99, 0.02) = 9.4$).

Finally, it should also be noted that the speed-ups for the QMC method are higher for $\beta > 95\%$. For example, the speed-up for a VaR(95%) calculation with 2% accuracy increases from 6.7 to 8.4 times if we want to be 99% certain of this accuracy ($S_{0.99}(0.95, 0.02) = 8.4$).

Concluding remarks

We have introduced a QMC method based on Sobol sequences for VaR estimation, and compared its performance to standard MC simulation. Based on a case study of a simple multi-currency portfolio, we find that the QMC method provides substantial speed-ups when compared to the standard MC method. In our experiments, the QMC method is between two to nine times faster than the MC simulation. Although we report the results for only one, reproducible, case study, we have obtained similar results in various other applications to real institutional portfolios.

As shown in the paper, this speed-up increases with

- the confidence level for the VaR estimate, α
- the accuracy required for the VaR estimation, ε
- the certainty required in the MC simulation to reach the required accuracy in VaR, β.

This speed-up is most sensitive to the certainty required for the MC simulation. It is also proportional to the confidence level and to the accuracy required on VaR, as would be expected from the convergence properties of both methods.

QMC methods have not been applied generally to measure portfolio risk. We argue, however, that they have a strong intuitive appeal for risk management problems. We have further shown through an example that significant computational gains are obtained with QMC simulation on problems of up to dimension 40. This is sufficient for VaR analyses based on the RiskMetrics™ covariance matrices and, more generally, any covariance matrix estimated with exponentially weighted moving averages, since the ranks of such matrices generally do not exceed this number.

There are still numerous challenges in the application of the QMC methods to risk management. The estimation of VaR in many instances requires simulation in spaces of high dimensionality. For example, in multi-step simulations over time, the dimensionality of the problem is dramatically increased.

Several advanced solutions to these problems include

- devise strategies that can be used in higher dimensions. Examples of such methods include mixed QMC and MC strategies and scrambled sequences.
- attempt to reduce the dimensionality of the problem. Techniques such as Principal Component Analysis are generally used for this purpose.
- introducing information on the portfolio structure *before* the simulation to reduce dimensionality.
- decompose the problem into a set of problems, each of them simpler and of lower dimension. This is the idea behind some stratified sampling techniques such as that proposed by Jamshidian and Zhu (1997).

Acknowledgements

We would like to thank Rustom Barua for his conscientious efforts in testing our numerical results.

References

Acworth, P., M. Broadie and P. Glasserman, 1997, "A comparison of some Monte Carlo and Quasi Monte Carlo techniques for option pricing." in *Monte Carlo and Quasi Monte Carlo in Scientific Computing*, Springer-Verlag.

Boyle, P. P., M. Broadie and P. Glasserman, 1996, "Monte Carlo methods for security pricing," *Journal of Economic Dynamics and Control*.

Boyle, P. P., 1977, "Options: a Monte Carlo approach," *Journal of Financial Economics*, 4(2) 323-338.

Caflisch, R. E., W. Morokoff and A. Owen, 1997, "Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension," *Computational Finance*, 1(1): 27-46.

Jamshidian, F. and V. Zhu, 1997, "Scenario simulation theory and methodology," *Finance and Stochastics*, 4(1): 43-67.

Kreinin, A., L. Merkoulovitch, D. Rosen and M. Zerbs, 1998, "Efficiency analysis of the advanced Monte Carlo methods in RiskWatch™." Research Report, Algorithmics Inc.

Marshall, C. L. and M. Siegel, "Value-at-Risk: in search of a risk measurement standard," Draft Document, Massachusetts Institute of Technology, April 1996.

Morokoff, W., R. Lagnado, A. Owen, 1998, "Tolerance for risk," *Risk*, 11(6): 78-83.

Niederreiter, H., 1992, "Random number generation and Quasi Monte Carlo methods," SIAM, Philadelphia.

RiskMetrics™, J.P. Morgan, 1995.

RiskMetrics™ Technical Document, 4th Edition, J.P. Morgan, December 18, 1996.

Schoenmakers, J.G.M. and A.W. Heemink, 1997, "Fast valuation of financial derivatives." *Computational Finance*, 1(1): 47-62.

Sobol, I.M., 1967, "Multi-dimensional quadrature formulas and Haar functions," *Computational Mathematics and Mathematical Physics*, 86(7).

Sobol, I. M. , B. Shukman, 1993, "Random and quasi random sequences: numerical estimates of uniformity of distribution," *Mathematical and Computational Modelling*, 18(8), 39-45.