Proceedings of the 2000 Winter Simulation Conference
J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, eds.

QUASI-MONTE CARLO METHODS IN CASH FLOW TESTING SIMULATIONS

Michael G. Hilgers

Computer Science Department
325 Computer Science
1870 Miner Circle
Rolla, MO 65409-0350, U.S.A.

ABSTRACT

What actuaries call *cash flow testing* is a large-scale simulation pitting a company's current policy obligation against future earnings based on interest rates. While life contingency issues associated with contract payoff are a mainstay of the actuarial sciences, modeling the random fluctuations of US Treasury rates is less studied. Furthermore, applying standard simulation techniques, such as the Monte Carlo method, to actual multi-billion dollar companies produce a simulation that can be computationally prohibitive. In practice, only hundreds of sample paths can be considered, not the usual hundreds of thousands one might expect for a simulation of this complexity. Hence, insurance companies have a desire to accelerate the convergence of the estimation procedure. This paper reports the results of cash flow testing simulations performed for Conseco L.L.C. using so-called *quasi-Monte Carlo* techniques. In these, pseudo-random number generation is replaced with deterministic *low discrepancy sequences.* It was found that by judicious choice of subsequences, the quasi-Monte Carlo method provided a consistently tighter estimate, than the traditional methods, for a fixed, small number of sample paths. The techniques used to select these subsequences are discussed.

1 INTRODUCTION

The insurance industry faces a number of challenging simulation problems. Of course, the actuarial sciences are largely devoted to modeling human life expectancy and the occurrence of disease. But the financial health of the insurance providers themselves heavily depends on making the proper response to the changing economic landscape. Their ability to payoff insurance contracts 30 years from now is strongly tied to investment decisions made today.

Being aware of this, federal regulation requires insurance providers to prove they will be solvent 30 years from now. This area of actuarial science is called *cash flow testing.* Each year, insurance companies must run simulations of their investment earnings pitted against the multiplicity of obligations they incur. In principle, their investment strategies should be market independent in the sense that money will remain to pay off all contracts regardless of the behavior of the marketplace.

Of course, investment strategies of an insurance company are also regulated and are mostly limited to investing in bonds with various levels of limited risk. The value of these bonds is derived from the prevailing interest rates. Since insurance companies are allowed to move their investments from short-term to long-term bonds as the market changes, simulating long-term solvency ultimately focuses on modeling the 90-day and 10-year treasury rates. (Intermediate rates are interpolated from these base rates.)

Federal regulation of cash flow testing requires insurance providers to consider a set of possible interest rate scenarios called the *New York Seven*. These represent seven simple market behaviors from boom to bust. It is important to note the actuarial software required to demonstrate solvency in these situations does not depend directly on the nature of the *New York Seven*, allowing a more ambitious simulation. By modeling the short-term and long-term interest rates as stochastic processes, it is possible to use this existing software to perform a Monte Carlo estimation of the company's future profit, allowing management to fine-tune investment procedures.

There are two major challenges to this simulation process. First, a good model of the interest rates is needed. In this paper, we used an interest rate model commonly accepted in the industry, and our examination of the stochastic scenarios produced using this model found they reasonably replicated the distribution of the historical rates. But even with this foundation, the run-time of the resulting Monte Carlo simulation is stifling. Sometimes days are required to get even a crude estimate.

A possible means of overcoming these challenges involves the use of deterministic *low discrepancy sequences* instead of pseudo-random number in the simulation. Using number theoretic methods, a low discrepancy sequence attempts to fill space uniformly, not mimic randomness. Monte Carlo simulations using these sequences, dubbed *quasi*-Monte Carlo, have outperformed traditional methods in a number of complex financial simulations. While this technique was not expected to shorten the run-time length, it was hoped the *quasi-Monte Carlo* approach would tighten the estimate, given the available computation time for the simulation.

The main difficulty encountered in apply quasi-Monte Carlo methods to cash flow testing simulations is the limited number of sample paths. The performance gains obtained from the use of low discrepancy sequences are observed for a large number of sample paths, maybe even hundreds of thousands. Limited attention has been given to the small sample path case. Motivated by a few existing observations, we used various techniques to encourage rapid convergence in this small sample path setting. We were able to consistently improve the quality of the estimates, as compared to traditional Monte Carlo methods, for a variety of actual insurance products and companies. While some of the techniques we utilized have been suggested in the literature, others are original and based on our empirical investigations. We believe this work strongly justifies further analytical investigation into our construction techniques and are currently pursuing it.

Section 2 examines interest rate models and discusses the one we used. Section 3 introduces low discrepancy sequences and examines how we accelerated convergence of the simulation by using subsequences of the originally posed constructions. In Section 4, we present the results of several simulations using actual corporate models.

2 INTEREST RATE MODELS

The 90-day and 10-year interest rates are modeled using a coupled, two-factor, mean reverting random walk. Of course, *two-factor* refers to the number of rates being modeled. The rates are *coupled* in the sense a change in one affects the other. As an example of this coupling, we would anticipate the 10-year rate to be higher than the 90-day rate, since a longer commitment warrants a larger return. When this is not the case, it is called an *inversion.* There have been several instances of inversions since the 1970's. (See Figure 1.) The rates are *mean reverting* in that they do not grow without bound, unlike traditional stock market models. Instead, they tend to oscillate around an average level for prolong periods of time. Of course, this is due to strong governmental and market forces designed to keep interest rates in check.

2.1 Review of Interest Rate Models

Unlike the lognormal random walk used in Black & Scholes models of the stock market, interest rate models rarely offer an explicit solution of the underlying stochastic differential equation. Various approximate solutions and special cases have been explored in the literature. Vasicek (1977) proposed a basic mean-reverting model

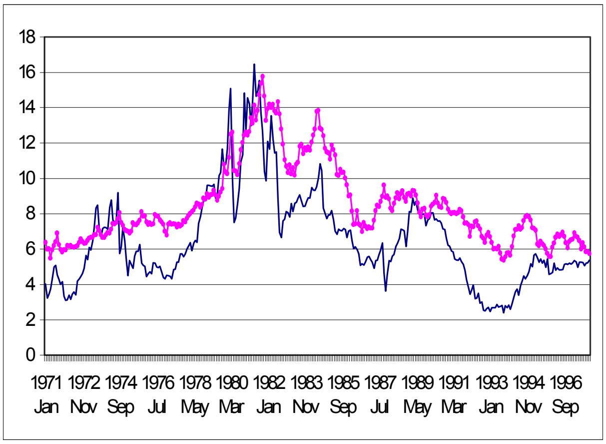

Figure 1: Historical 90-Day (Darker Line) and 10-Year Treasury Rates (Lighter Line).

$$dr_t = (a - br_t)dt + \sigma dW_t^*$$

where r_t is the short term rate, W_t^* is a one dimensional Brownian motion, and a, b, and σ are positive constants. This model is considered *mean-reverting* because it admits a stationary Gaussian distribution with mean a/b and variance $\sigma^2/(2b)$, which is viewed as the mean interest rate distribution. Various other single-factor, mean reverting models have been proposed, such as those of Cox, Ingersoll, and Ross (1985) and Longstaff (1989). Most of the variations have features improving implementation in a particular form of financial derivative.

In simultaniously modeling short-term and long-term interest rates, one must account for their natural coupling. This gives rise to so-called *two-factor* models. The one described by Ho (1995) closely resembles our implementation. (Note that he gives credit to Brennan and Schwartz (1979) and Longstaff and Schwartz (1992).) Consider rates r_t and s_t satifying

$$dr_t = (a - br_t)dt + \sigma_1 r_t dW_t^*$$
$$ds_t = \sigma_2 s_t dV_t^*$$

where (W_t^*, V_t^*) is a two-dimensional Brownian motion with dependent components satisfying $dW_t^* dV_t^* = \rho dt$. σ_1 and σ_2 are the standard deviations of the short-term and long-term rates and ρ is a measure of their correlation. Notice the short-term rate, r_t, is much like the Vasicek model. The long-term rate, s_t, is basically derived from the short-term rate via the correlation of the two Brownian motions. Ang and Sherris (1997) have examined specific application of two-factor models to the cash flow testing problem.

2.2 Discussion of TAS Model

For purposes of the experiments performed at Conseco, we used the interest rate model employed by the Tillinghast Actuarial Systems (TAS), which is the software system used by Conseco to perform cash flow testing, among many other tasks. The TAS algorithm for generating an interest rate scenario is based upon the lognormal distribution. For each quarter over the duration of the simulation (up to 31 years), two uniformly distributed pseudo-random numbers are generated. These are transformed into lognormal deviates, using volatility parameters specified by the actuaries. The 90-day rate is transformed first, and then this value is coupled to the 10-year rate using the correlation factor. These two random walks, as they evolve in time, are subject to a series of somewhat empirical transformations to force a desired qualitative behavior. For example, one transformation involves forcing the return of the interest rates to a specified mean, at a specified rate, in a user-controlled time period. Another empirical transformation controls the duration of an inversion. Finally, artificial boundary conditions are imposed on the interest rates. The rates are given an elastic response, as specified by the actuaries, to upper and lower boundaries. This causes the rates to literally “bounce back” into the desired range. For more detailed information on this interest model, see the formulae section of the TAS documentation

This model includes a number of financial and non-financial parameters. Before proceeding with the simulations, the actuaries study the distribution of the generated scenarios in comparison with historical rates. Non-financial parameters are then adjusted improve the replication of historical behavior.

3 QUASI-MONTE CARLO SIMULATIONS

A traditional Monte Carlo simulation using the TAS generated interest rates calculates the desired quantity, such as profit, for each interest rate scenario. A number of these scenarios are generated, and the desired estimation is the arithmetic average of the profits corresponding to each individual scenario. This amounts to an estimation of the profit’s expected value, which is an integration over the sample space. In this case, the sample space has dimension 248 (8 random numbers per year times 31 years). Traditionally, the Monte Carlo method has been the preferred technique for such high dimensional numerical integration. It can be shown that the errors in the estimates converge to zero like $O(1/\sqrt{N})$ where N is the number of sample paths, Niederreiter (1992).

Recently, attention has been given to a sister technique, known aptly as the quasi-Monte Carlo method. Instead of using pseudo-random numbers, a *low discrepancy sequence* is generated in the unit hypercube. Low discrepancy sequences do not attempt to mimic randomness, but instead attempt to fill the hypercube as uniformly as possible. Numerical integration using low discrepancy sequences in a Monte Carlo-like fashion often has an observed order of convergence of $O(1/N)$. (See Ninomiya and Tezuka (1995) for solid examples of this behavior. Note the actual proven bound is less generous, Niederreiter (1992).)

The claims of the convergence rates are particularly attractive to the actuaries performing cash flow testing simulations since the number of sample paths is limited to hundreds by practical considerations, instead of tens of thousands to millions as in most academic investigations. As an example of the practical limitations, a single scenario for the American Life Company, a multi-billion dollar insurance provider owned by Conseco, takes about 20 minutes on a Pentium II 450MZ. To do a Monte Carlo simulation involving hundreds of scenarios takes days and generates many gigabytes of data and challenges system resouces. Therefore, the simulations are limited to 100 to 200 scenarios in practice, which creates their desire to have a consistently tight estimate for this number of sample paths.

3.1 Low Discrepancy Sequences

Discrepancy is a set theoretic measure of the distribution of a point set in the unit hypercube of some space. Intuitively, a sequence is declared as having low discrepancy if it uniformly fills space as the number of points in the sequence goes to infinity. The original constructions of low discrepancy sequences belong to Faure (1982), Halton (1962), and Sobol’ (1967), to name a few. Niederreiter (1992) has developed a unifying construction technique that encompasses much of the earlier work. Tezuka (1993) has produced a generalization of Niederreiter’s construction, uniting it with Halton sequences.

The basic elements in constructing a low discrepancy sequence have been the same for most of the past decade and are succinctly described in Niederreiter (1992). The construction process involves sub-dividing the unit hypercube into boxes of fixed volume that have faces parallel to the cube’s faces. The goal is to put a point in each of these boxes before proceeding to a finer scale.

To generate the n^{th} element of a low discrepancy sequence in s dimensional space, a prime integer b is chosen and n is expanded in base b. That is,

$$n = \sum_{r=0}^{\infty} a_r(n) b^r$$

where $a_r(n) \in \{0,1,2,\ldots,b-1\}$. The i^{th} component of the n^{th} element of the sequence is

$$x_n^{(i)} = \sum_{j=1}^{\infty} y_{nj}^{(i)} b^{-j}$$

with

$$y_{nj}^{(i)} = \sum_{r=0}^{\infty} c_{jr}^{(i)} a_r(n) \cdot$$

All of the arithmetic is done modulo b to ensure that $y_{nj}^{(i)} \in \{0,1,2,\ldots,b-1\}$. The collection of coefficients $c_{jr}^{(i)}$ is known as the i^{th} generator matrix. How these matrices are constructed distinguishes the various methods.

3.2 Choice of Generator Matrices

While most of these sequences have been explored in the setting of computational finance, the work of Ninomiya and Tezuka (1995) suggested two types seem to have the best performance for a variety of high-dimensional finance problems. Following their lead, we implemented the Nedereitter construction for a base 2 sequence, (Bratley, Fox, and Niederreiter (1992)), and Tezuka's (1994) so-called Generalized Faure sequence based upon polynomial Halton sequences.

3.2.1 The Niederreiter Generator Matrices

Let $p_i(x)$ be an irreducible polynomial over the finite field F_2. It is used to generate the coefficients of the i^{th} generator matrix by choosing certain coefficients in the Laurent series expansion

$$\frac{x^k}{p_i(x)^j} = \sum_{r=-k}^{\infty} a^{(i)}(j,k,r) x^{-r-1} \cdot$$

The coefficients $a^{(i)}(j,k,r)$ are determined by multiplying both sides of the equality by $p_i(x)^j$, expanding the right-hand-side, grouping together powers of x, and equating coefficients of the x terms. The $c_{jr}^{(i)}$ are chosen from among these in a manner detailed by Niederreiter (1992). Working in base 2 is particularly attractive since the finite field operations reduce to the standard binary XOR and AND. (See Niederreiter (1992) and Bratley et al (1992) for details).

The uniformity of the sequence being generated can be shown to depend on the quantity

$$T_b(s) = \sum_{i=1}^{s} (\deg(p_i) - 1) \cdot$$

This dependence is such that small $T_b(s)$ implies low discrepancy. Therefore, the irreducible polynomials of a specific degree should be exhausted before moving to a higher degree in order to achieve the lowest discrepancy.

3.2.2 The Generalized Faure Generator Matrices

Unlike the recursion formula described above, Tezuka and Tokuyama (1994) produced a closed-form expression for the generator matrix. Their algorithm requires a prime number $b \geq s$. Then the (i,j) component of the h^{th} generator matrix is

$$\sum_{q=0}^{\min(i-1,j-1)} \binom{i-1}{q} \binom{j-1}{q} b_h^{i+j-2q-2} \cdot$$

The numbers $b_1, b_2, \ldots, b_s$ are distinct elements of the finite field of order b. Since the elements of Pascal's triangle can be computed with a simple recursion relation, these generator matrices are a fairly simple computation.

Sequences produced using this type of generator matrix are called "Generalized Faure" because of the similarity with the sequences constructed by Faure. The generator matrices in Faure's construction are powers of the triangular matrix associated with Pascal's triangle. Both Niederreiter and Tezuka have produced general construction algorithms that ultimately include some form of binomial coefficients, suggesting an interesting tie between low discrepancy sequences and combinatorial theory.

3.3 Selected Subsequences

Since low discrepancy sequences are not attempting to mimic randomness, the first elements of the sequence often appear in a very predictable fashion. In fact, a large number of elements may be required before the unit hypercube has been covered. For example, in Figure 2, the first 500 points of the Generalized Faure sequence in 250-dimensional space have been calculated, and the first and second coordinates have been plotted.

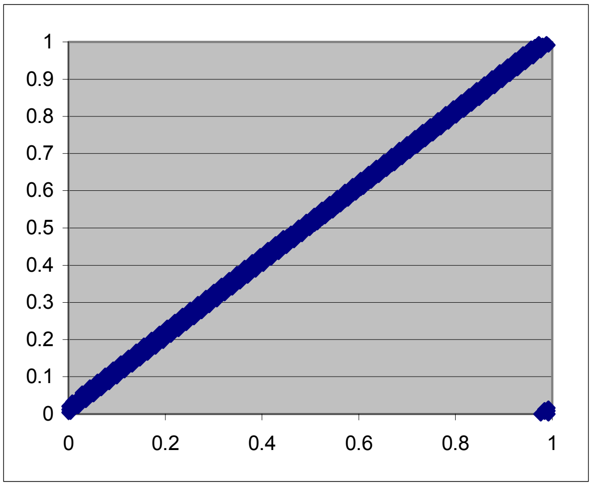

Figure 2: Components 1 and 2 of the First 500 Elements of the Generalized Faure Sequence in 250-dimensional Space.

Obviously, the components were nearly identical. Increasing to 3000 points, Figure 3 demonstrates the banded nature of the filling process. Apparently, a large number of points would be required before the hypercube is covered. With a moments reflection, it is seen an individual coordinate of the initial elements of the sequence is of the from

$$a(251)^{-1} + b(251)^{-2},$$

since 251 is the first prime larger than 250. For this choice of generator matrix, it is also observed that for all coordinates, we have

$$x_{(n)}^{(i)} = a_0(n)(251)^{-1} + \ldots .$$

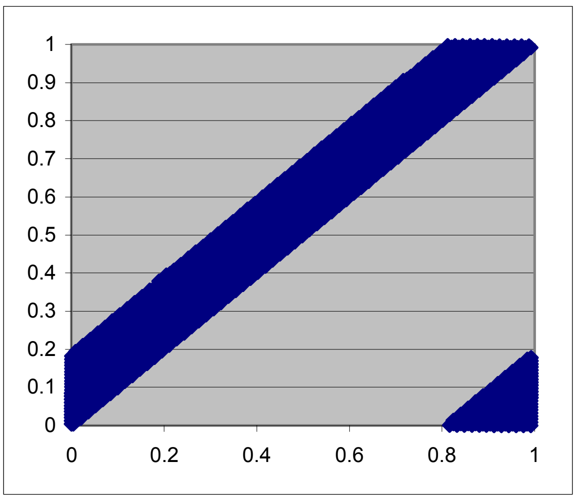

Figure 3: Components 1 and 2 of the First 3000 Elements of the Generalized Faure Sequence in 250-dimensional Space.

Hence the work of covering the cube is being done by higher order terms, suggesting that over 250^2 elements are needed to "cover" the cube. From this, one might deduce Generalized Faure sequences are not appropriate for cash flow testing, since this number is many orders of magnitude too high for their simulations. However, using the subsequence constructions to be discussed, some success was achieved.

On surface appearance, the Niederreiter sequences have more promise because the generator matrices are not identical for the first term in the expansion. Furthermore, the method does not require working in a base higher than the dimension of the space. In fact, base 2 is standard due to the simplicity of the finite field operations.

Figure 4 seems to support our hopes. The distribution of the first 500 points has few gaps and offers a uniform covering. However Figure 5 is less encouraging. As can be seen, coordinates 47 and 48 of the sequence were quite similar. After careful examination, it is revealed the generator matrices for these coordinates are nearly identical, and this is traceable to the irreducible polynomials used to produce the matrices. The polynomials differ only slightly in the highest powers. In fact, many of the irreducible polynomials of the finite field of order 2 differ only in two coefficients and that is for a term of degree 8 or higher. Said in another way, the generator matrices for these coordinates will create elements that differ only for larger values of n. Once again, this seems to preclude cash flow testing simulations.

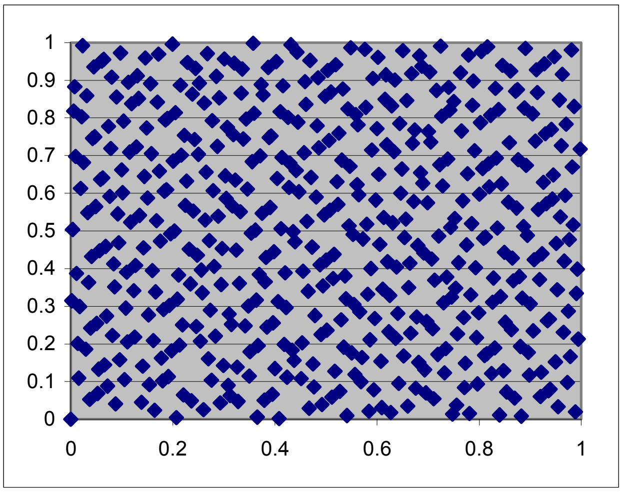

Figure 4: Coordinates 1 and 2 of the First 500 Points in a Niederreiter Sequence Base 2.

3.3.1 Subsequence Selection

Our motivation for using subsequences to accelerate the simulation process originated with a paper by Kocis and Whiten (1997). To provide a historical perspective of their work, it had been suggested earlier that $s = 40$ was the practical limit for quasi-Monte Carlo simulations (Bratley et al (1992)). In the time since that paper, a number of applied simulations in high dimension have observed excellent performance. For instance, Paskov and Traub (1995) saw excellent convergence results for a collateralized mortgage obligation with $s = 360$. As an attempt to formally challenge this limitation, Kocis and Whiten (1997) performed a thorough investigation of numerical integration over domains of dimension up to 400. They concluded low discrepancy sequences performed well, thereby further dispelling the concerns of Bratley et al (1992). Kocis and Whiten also recognized the problem with the initial elements systematically filling space and introduced *leaping* over some elements of the sequence to accelerate the filling of space. By *leaping* they mean, selecting a periodic subsequence, say every p^{th} term. Their analysis was focused primarily on Halton sequences, though they did offer some comment on Sobol' sequences. They did not examine either of the sequences in this paper. Yet their empirical success encouraged our similar investigation using the Niederreiter and Generalized Faure sequences.

The process we used to select the subsequence involved two parameters: the starting integer N_0 and the leaping factor p. The subsequence used in the simulation can be expressed as

$$\chi_n^{(i)} = x_{N_0+np}^{(i)}.$$

The governing premise in the selection of N_0 and p is the quasi-Monte Carlo simulation will be improved, when only a limited number of sample paths is available, if the subsequence χ_n covers the 250-dimensional hypercube as uniformly as possible. (It should be noted there is no additional computational cost in this implementing this premise since the intermediate p elements do not need to be calculated.) One approach to implementing this premise would be to generate N points for a given N_0 and p then calculate the discrepancy of the resulting point set. This could be repeated for a large set of choices of these parameters. The values selected for implementation would be the ones corresponding to the lowest discrepancy. Unfortunately, the mathematical definition of discrepancy of a point set in the unit hypercube involves a supermum over all Lebesgue-measurable subsets. This, of course, would require an uncountable number of operations.

An alternative approach would be to use the L^2-discrepancy examined by Hickernell (1996). Calculating the L^2-discrepancy involves a double summation over the point set of size N. Using the process described above, it should be possible to estimate optimal choices of N_0 and p for fixed N. This approach is currently under investigation.

What is offered in this paper is a softer, more empirical version of the above stated procedure, which can be (and was) accomplished in a corporate setting. While this method will probably not find the optimal parameters in a strict mathematical sense, it was easy to do and has performed well in a variety of applications.

Basically, a qualitative study was done. Some thought was given to properties p should process, and then cross sectional pictures, such as those in Figures 2-5, were examined. The goal was to find values of N_0 and p that would make the cross section in Figure 5 look more like that in Figure 4, for example, uniformly across all 250 coordinates for the chosen number of points N.

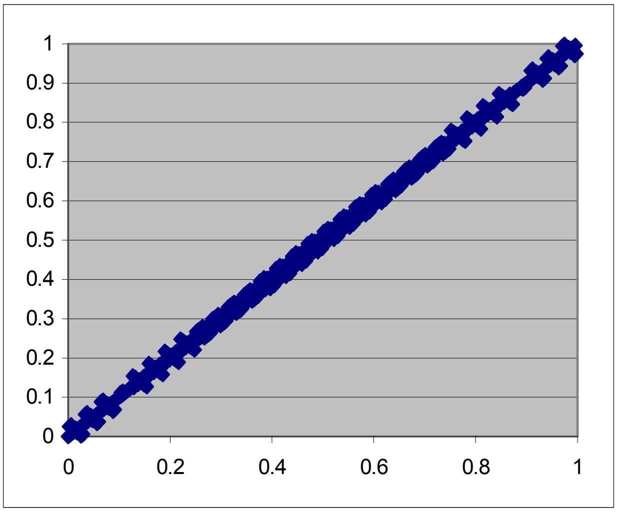

Figure 5: Same as Figure 4 Except Coordinates 47 & 48.

3.3.1.1 Niederreiter Sequences

Beginning with the Niederreiter sequence, the choice of p focused on the irreducible polynomials for base 2. The algorithm used to produce the generator matrices, as described above, utilizes a recursive relationship involving the coefficients of the polynomials. To produce a low discrepancy sequence in a 250-dimensional hypercube, the first 250 irreducible polynomials are required. Grouping the polynomials by degree and exhausting the list of polynomials of one degree before going to the next higher, means that degree 11 is reached. Furthermore, it was observed that often only two coefficients of the polynomials differ. For example,

$$11111011011_2$$
$$11111101011_2.$$

(Note that the coefficients are just 0 or 1 in the finite field of order 2, so polynomials can be succinctly expressed as a string of bits.) Intuitively speaking, there will be a similarity between the generator matices of adjacent coordinates for similar polynomials. Furthermore, the base 2 expansion of n will need some of the higher powers of 2 in order to captialize on the difference between the generator matrices. As a starting point, we examined $p = 2^r$ where r is larger than the degree of the term for which the polynomials differ. It was soon observed that $p = 2^{10} = 1024$ produced a nice dispersion for 500 points, see Figure 6.

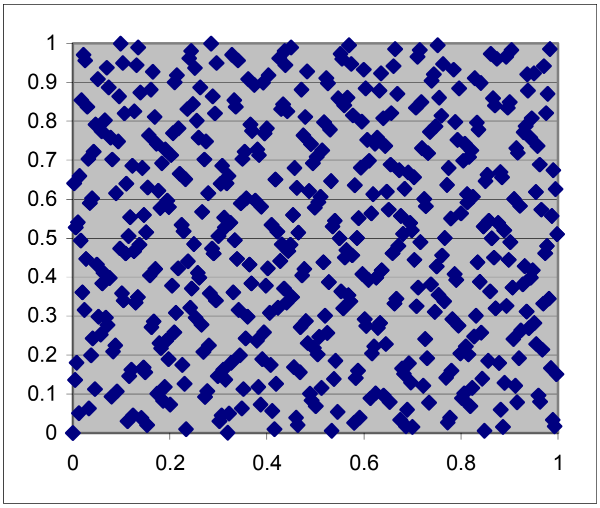

Figure 6: Same as Figure 5 with a Leap of 1024.

One other observation resulting from this qualitative study of the discrepancy of subsequences of the Niederreiter construction was that the coordinates produced by irreducible and non-primitive polynomials seemed problematic. Niederreiter required irreducible polynomials for his construction so he could perform a partial fraction decomposition of a rational function. Further restricting to primitive polynomials would raise the value he labeled as $T_2(250)$, which describes the fineness of the distribution of the sequence thereby increasing the discrepancy of the sequence. Ultimately, for our cash flow testing simulation, all non-primitive polynomials were omitted. In fact, some primitive polynomials were omitted, because they did not cooperate with a choice of p that seemed to work well with most of the other 250 coordinates.

This analysis was attempted for $N = 500, 250, 100$. The smaller values of N increased the values of p. For $N = 500$, the value of N_0 didn't seem to matter and was taken to be 1. The value $p = 2^{10} + 2^6 = 1088$ seemed to perform best. Likewise, for $N = 250$, the values of $N_0 = 2^7 + 1 = 129$ with $p = 2^{17} - 1 = 131071$ were identified as reasonable choices. Little success was achieved for $N = 100$.

3.3.1.2 Generalized Faure Subsequences

It proved to be more difficult to produce appropriate choices of N_0 and p for Generalized Faure subsequences. The diagonal bands of Figures 2 & 3 seemed present in most subsequences. As with the Niederreiter subsequences, powers of the base b seemed a reasonable place to begin. For 250-dimensional sequences, the next larger prime number is $b = 251$. Examining Figure 7, we see that for $N = 500$, $N_0 = 1$, and $p = b = 251$ a series of diagonal strips was produced.

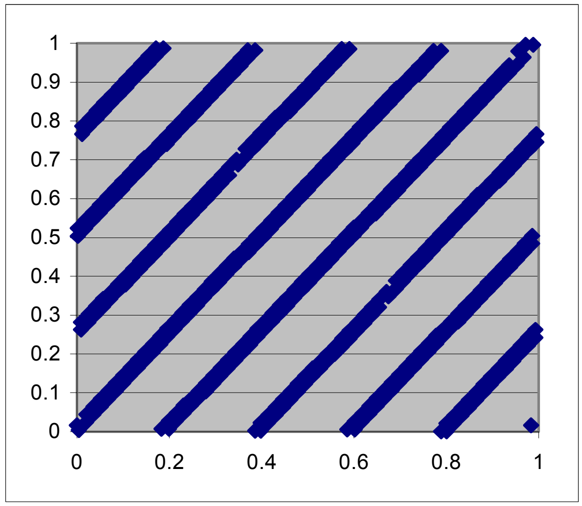

Figure 7: Same as Figure 2 with Leap of 251.

Many other values of the leap parameter were tried. Most involved powers of the base shifted by a small amount. Inevitably, some coordinates produced cross-sectional graphs like Figure 7 or with even fewer diagonal strips. Examining the generator matrices offered some insight as to why these patterns are produced, but no choice of parameters seemed to uniformly remove it. These matters are under further investigation. For purposes of this simulation, we used the values of the parameters already mentioned.

4 CASH FLOW TESTING SIMULATIONS

Three simulations were performed. The first was a simulation of a single premium deferred annuity. The others involved cash flow simulations of entire companies owned by Conseco. A mid-sized company, Massachusetts General Life, was chosen for the second simulation. The final simulation involved the largest insurance provider owned by Conseco, American Life Company. In all of the studies, the quantity estimated was the present value of the book profit after taxes.

4.1 A Single Premium Deferred Annuity

The numerical experiments began with a quasi-Monte Carlo simulation of a single product an insurance company might carry. We chose a single premium deferred annuity product carried by one of Conseco's subsidiary companies. This product can be viewed as an interest bearing bank account into which a single deposit is made. The amount of interest paid depends on current market conditions. First, we attempted to measure the improvement the use of our subsequences might offer relative to the originally posed algorithms.

Figure 8 shows the estimate produced by the standard algorithm is still increasing after 750 interest rate scenarios. The simulations utilizing the subsequences seem to have converged fairly quickly and are not far from the resulting estimate at the desired value of 250 scenarios.

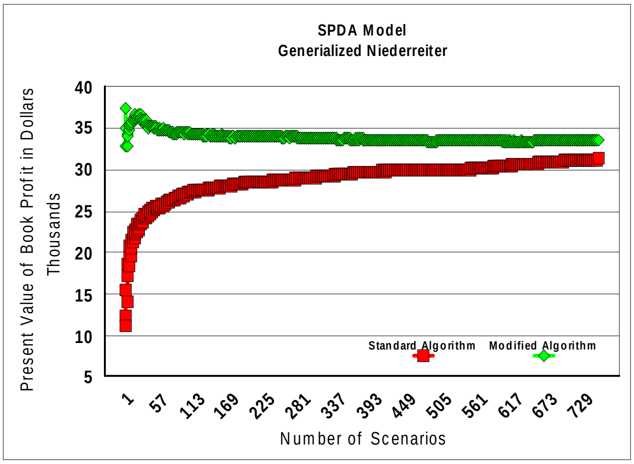

Figure 8: Niederreiter with and without Leaping

The situation is even more dramatic for the Generalized Faure sequences, as shown in Figure 9. As observed earlier, the standard algorithm produces a sequence that fills the cube in a slow, systematic fashion. The estimate is again increasing and has not converged by 750 scenarios. However, the subsequence selection discussed above produced an estimate that seems to have approximately converged within the specified number of scenarios, even with the problems that where identified.

Figure 10 shows the modified quasi-Monte Carlo simulations converge to the same estimate and also agree with a traditional Monte Carlo simulation using the pseudo-random number generator in TAS. Furthermore, the estimates produced by the use of low discrepancy subsequences are much tighter in the 100 to 200 ranges than was the traditional simulation approach. So our use of subsequences in the quasi-Monte Carlo method in this cash flow testing simulation achieved the desired goal.

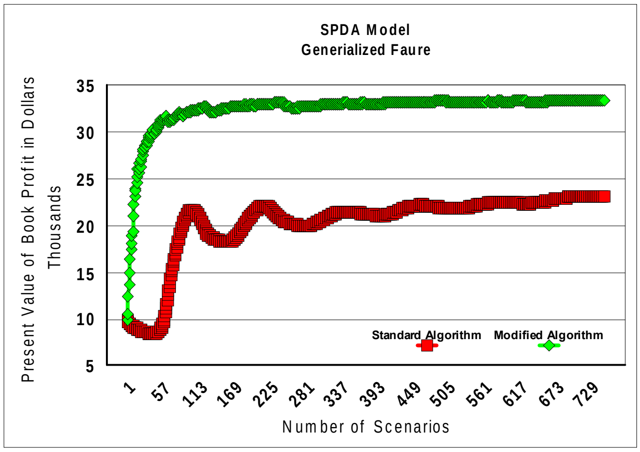

Figure 9: Generalized Faure with and without Leaping.

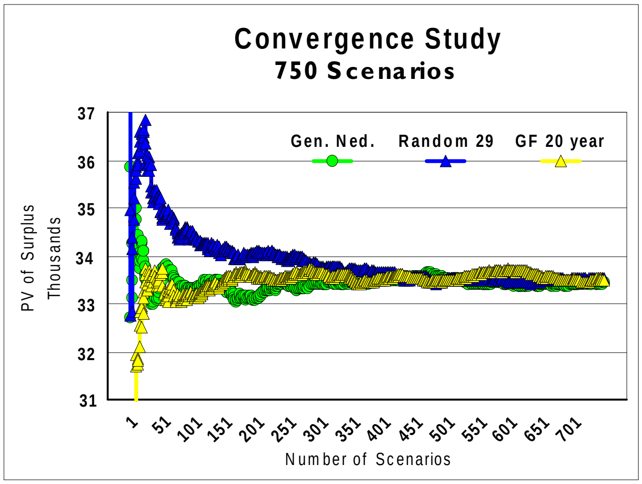

Figure 10: Comparing against Random Number Generation.

4.2 Massachusetts General Life

Encouraged by the results of a cash flow simulation of a single product, the experiment continued with the simulation of an entire company. Now facing the real world constraints of time and space, this simulation was limited to 400 interest rate scenarios. It is difficult to conclude from any of the performances that the simulation had converged. Note that the dollar amounts involved were large. The book profit was on the order of $27,000,000.

In Figure 11, the relative errors of several of the simulations are plotted. The Niederreiter subsequences and TAS generated random numbers were in some agreement at the end of the simulation, so their estimations were used to calculate the relative errors. Notice the Generalized Faure subsequences did not perform as well for this much larger simulation. Also, the estimate they produce depends, in this case, strongly on the seed used to initiate the pseudo-random number generation. This has been observed in other financial simulations and cannot be

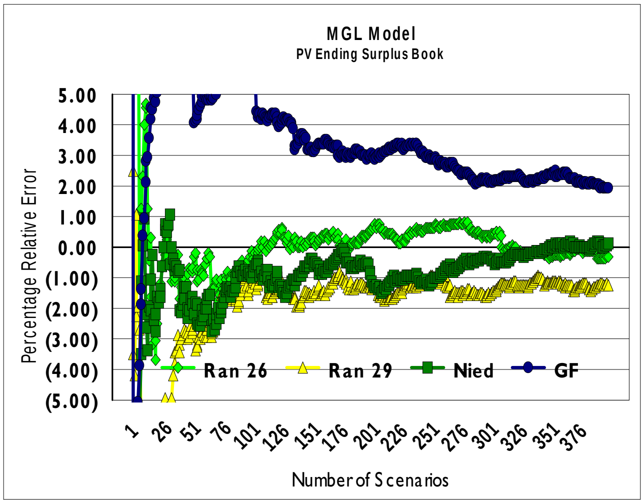

Figure 11: Convergence Study of MGL Including Two Different Random Number Sequences Differing by Choice of Initial Seed.

remedied by increasing the number of sample paths. Since the seed is a non-financial parameter in the simulation, this dependence is particularly disturbing in the arena of cash flow testing.

We observe the Niederreiter subsequences are within 1% to 2% of the final value in the desired range of 100 to 200 scenarios. However, the simulation using the pseudo-random numbers beginning with 26 as initial seed *happened* to provide a better estimate in this range, while the simulation with 29 as the initial seed did not. Perhaps, the most significant fact to conclude from the first two stages of this simulation process is that the Niederreiter subsequences have perform consistently and acceptably in two very different situations. The same cannot be said for the traditional methods.

4.3 American Life Company

American Life Company is a large company whose holdings are on the order of a billion dollars. It carries a wide variety of customers and products. As a result, the cash flow testing simulation is almost prohibitively large. The simulations being reported herein took days to run on a dedicated computer. In practice, the runs are often interrupted by server problems and so on. It is difficult for the team of actuaries with the responsibility of doing the testing to perform the estimations using more than 100 scenarios. For these simulation experiments, 250 scenarios were used and the run time was 72 hours. Nearly a gigabyte of data was produced for each simulation.

The Generalized Faure subsequences continued their poor performance for the larger simulations. In fact, the performance was so poor, it has been omitted from the graph to reduce the visual clutter.

Again, it is difficult to argue that any of the methods had converged. To produce this relative error plot, the average of the final estimates of all three methods was used. Two traditional Monte Carlo simulations were included to demonstrate the variance in the estimate due to changes in the initial seed. One of the traditional simulations produced an estimate that was reasonable in the 100 to 200 scenario range, while the other was off several percent. Meanwhile, the Niederreiter subsequences were in the 1% to 2% error range over the critical number of interest rate scenarios. Hence, the method of subsequence construction examined in this paper produced a quasi-Monte Carlo simulation of cash flow testing that *consistently* out-performed the traditional approach in this limited number of sample paths environment.

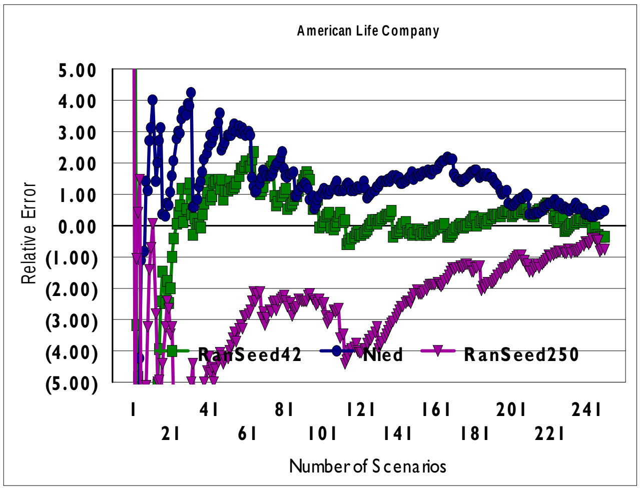

Figure 12: ALC with Two Different Random Number Sequences and the Niederreiter Subsequence

5 CONCLUDING REMARKS

On the basis of these results, Conseco has chosen to use the quasi-Monte Carlo method utilizing the Niederreiter subsequences as examined in this paper. But this does not mean the problem being addressed has reached a conclusion. A less qualitative approach to subsequence selection is desirable. Perhaps the L^2-discrepancy method previously mentioned holds the answer. We are investigating. Furthermore, a thorough examination of the Generalized Faure method is in order. Since it offers the simplest generator matrix construction, it is desirable to use these sequences; yet, its performance in this small number of sample paths environment was less than satisfactory.

It is believed that the empirical results of this study show that further study of the use of low discrepancy sequences in high dimensional, low number of sample paths problems is warranted.

ACKNOWLEDGMENTS

The author thanks Conseco for their financial support through this study. In particular, he is deeply indebted to Randy Gantt F.S.A. for the resources and technical assistance he provided. He would also like to acknowledge the technical and physical support from Mary Avila F.S.A. and Steve Avila Ph.D. Thank you.

REFERENCES

Ang, A., and M. Sherris. 1997. Interest rate risk management: Developments in interest rate term structure modeling for risk management and valuation of interest-rate-dependent cash flows. *North American Actuarial Journal* 1(3):1-26.

Bratley, P., B. L. Fox, and H. Niederreiter. 1992. Implementation and tests of low-discrepancy sequences. *ACM Transactions on Modeling and Computer Simulation* 2(3):195-213.

Brennan, M. J., and E. S. Schwartz. 1979. A continuous time approach to the pricing of bonds. *Journal of Banking and Finance* 3:133-155.

Cox, J. C., J. E. Ingersoll, and S. A. Ross. 1985. A theory of term structure of interest rates. *Econometrica* 53:385-407.

Faure, H. 1982. Discrepance de suites associees a un systeme de numeration (en dimension s). *Acta Arithematia* 41:337-351.

Halton, J. H. 1960. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numerical Mathematics* 2:84-90.

Hickernell, F. J. 1996. The mean square discrepancy of randomized nets. *ACM Transactions on Modeling and Computer Simulation* 6(4):274-296.

Ho, T. S. Y. 1995. Evolution of interest rate models: A comparison. *The Journal of Derivatives* Summer:9-19.

Kocis, L. and W. J. Whiten. 1997. Computational investigations of low discrepancy sequences. *ACM Transactions on mathematical Software* 23(2):266-294.

Longstaff, F. A., and E. S. Scwartz. 1992. Interest rate volatility and the term structure: A two-factor general equilibrium model. *Journal of Finance* 47:1259-1282.

Longstaff, F. A. 1989. A nonlinear general equilibrium model of the term structure of interest rates. *Journal of Financial Economics* 23:195-224.

Ninomiya, S., and S. Tezuka. 1995. Toward real-time pricing of complex financial derivatives. *Applied Mathematical Finance* 3:1-20.

Niederreiter, H. 1992. *Random number generation and quasi-Monte Carlo methods.* Philadelphia, PA: SIAM.

Paskov, S. and J. F. Traub. 1995. Faster valuations of financial derivatives. *The Journal of Portfolio Management* 22(1):113-120.

Sobol', I. M. 1967. The distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* 7(4):86-112.

Tezuka, S. 1993. Polynomial arithmetic analogue of Halton sequences. *ACM Transactions of Modeling and Computer Simulation* 3(2):99-107.

Tezuka, S. and T. Tokuyama. 1994. A note on polynomial arithmetic analogue of Halton sequences. *ACM Transactions of Modeling and Computer Simulation* 4(3):279-284.

Vasicek, O. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5:177-188.

AUTHOR BIOGRAPHY

MICHAEL G. HILGERS is an Associate Professor in the Computer Science Department at UMR. He received his BS in Computer Science, MS in Applied Mathematics, MS in Engineering Mechanics from UMR and his Ph.D. in Applied Mathematics from Brown University. He is a member of ACM and SIAM. His interests include computational finance, computational mechanics, and matrix analysis. He has acted as consultant developing software for the financial services industry for several years. His email and web addresses are `<hilgers@umr.edu>` and `<www.umr.edu/~hilgers>`.