QMC techniques for CAT bond pricing

Hansjörg Albrecher, Jürgen Hartinger, and Robert F. Tichy

albrecher@tugraz.at,hartinger@finanz.math.tugraz.at,tichy@weyl.math.tugraz.at
Graz University of Technology, Department of Mathematics
Steyrergasse 30, A-8010 Graz, Austria
This research was supported by the Austrian Science Fund Project S-8308-MAT. The first author is a postdoctoral research fellow at K.U. Leuven (Fellowship F/03/035).

Summary. Pricing of catastrophe bonds leads to integrals with discontinuous and formally infinite-dimensional integrands. We investigate the suitability of Quasi-Monte Carlo methods for the numerical evaluation of these integrals and develop several variance-reduction algorithms. Furthermore, the performance of Quasi-Monte Carlo sequences for asymptotically efficient rare event simulation is examined. Various numerical illustrations are given.

Key words: Quasi-Monte Carlo, insurance linked securities, rare events, importance sampling, variation reduction

1 Introduction

In the last decade securitization of risk gained considerable importance in catastrophe risk management of insurance companies. The high level of worldwide insurance losses both in severity and frequency due to natural catastrophes and the under-capacity of traditional reinsurance encouraged the development of so-called insurance linked securities. These are financial products that transfer some of the insurance risk to the financial market. A survey on existing products and on the history of this alternative to classical reinsurance may be found e. g. in Doherty [8], Gorvett [10] or more recently Muermann [17]. One of the key issues to popularize these products among capital investors is to develop fast algorithms for the correct valuation of their fair prices. Since typically the underlying models are too complex to allow for analytical pricing formulae, efficient numerical methods are of major importance.
In this paper we will focus on the pricing of catastrophe-linked bonds (Cat bonds for short), for which the issuer has the right to default on parts of the principal and/or coupon if a predefined index of insurance losses due to natural catastrophes exceeds a given threshold. This loss index is usually

constructed as a weighted average of the losses of individual insurance companies. For the determination of fair Cat bond prices, the step of modeling the stochastic behavior of this loss index is crucial. In [1], the performance of Quasi-Monte Carlo methods for the pricing of Cat bonds in a model with a compound Poisson loss index process in the presence of stochastic interest rates and default risk of the issuing company was analyzed. Here, we will consider a more complex loss index model recently proposed by Dassios and Jang [6] and investigate the potential of Quasi-Monte Carlo techniques to evaluate the resulting high-dimensional integrals.

In contrast to Monte Carlo methods, the error of Quasi-Monte Carlo (QMC) techniques is deterministically bounded above by the Koskma-Hlawka inequality: Let $\{x_N\}_{1\le n\le N}$ be a point sequence in $[0,1)^s$, $V(f)<\infty$ be the variation of f in the sense of Hardy and Krause and $D_N^*(x_n)$ denote the star discrepancy of $(x_1,\ldots,x_N)$, then

$$\left| \int_{[0,1]^s} f(x)dx - \frac{1}{N}\sum_{i=1}^{N} f(x_i) \right| \le V(f)\; D_N^*(x_n), \tag{1}$$

separating the integration error into a factor only depending on the integrand f and a factor depending on the point sequence $(x_1,\ldots,x_N)$ only. Since the best known sequences (so-called low discrepancy sequences) have a discrepancy of order $\mathcal{O}(\log^s(N)/N)$, QMC techniques are at least asymptotically superior to Monte Carlo simulation, the probabilistic error of which is known to be of order $\mathcal{O}(1/\sqrt{N})$. For an introduction to QMC integration and the construction of low discrepancy sequences we refer to Niederreiter [18]. Although there exist functions for which the bound (1) is attained, it turns out to be a rather conservative bound in most cases. In a seminal paper, Paskov and Traub [19] showed empirically that there are extremely high dimensional problems ($s=360$ and more) occurring in mathematical finance, where QMC methods outperform Monte Carlo algorithms by far. Since then there has been intensive empirical and theoretical research to classify types of integrands that are particularly well suited for QMC integration. The concept of tractability (cf. Sloan and Wozniakowski [20]) and the effective-dimension approach (cf. Wang and Fang [23]) provide partial explanations to these aspects, but the problem is still not perfectly understood.

In this paper we investigate the performance of QMC techniques for integrands occurring in the price determination of Cat bonds. From the view-point of Quasi-Monte Carlo methodology, the valuation of these financial products leads to several interesting phenomena: At least formally, the involved integrands are infinite-dimensional, but have moderate effective truncation dimension (a rigorous definition of the concept of truncation dimension can be found in [23]). One is also faced with integrands with jumps (discontinuities), which are in general not parallel to the axes, so that the variation of the integrands

is unbounded and the Koksma-Hlawka bound becomes useless. We will use a transformation of the integrands to smooth the jumps and compare the convergence behavior with and without smoothing. Moreover, since in practical situations the threshold for the forgiveness trigger of the Cat bond is usually rather high, the QMC evaluation of the corresponding integrals amounts to a sampling of rare events. Rare event sampling is well established using Monte Carlo sequences (see e.g. Asmussen and Glynn [3] and the references therein). This paper aims to analyze the performance of QMC algorithms combined with an asymptotically efficient rare event technique, which applies for light-tailed claim sizes. In case of heavy-tailed distributions, efficient rare event simulation is still not sufficiently developed from a theoretical point of view (cf. [2]). However, for our integrands, we illustrate how solutions of light-tailed problems might be used as control variates for heavy-tailed problems to improve the convergence significantly.
Section 2 introduces the loss-index model and the valuation methodology, Section 3 specifies the details of the proposed simulation techniques and numerical illustrations of its performance are given Section 4. Section 5 concludes.

2 A shot noise process for the loss dynamics

In the model of Dassios and Jang [6], the claim index is governed by a doubly stochastic compound Poisson process. A shot noise process defines frequency and impact of catastrophic events and determines the corresponding claim intensity of the Poisson process generating the claims relevant for the index in a given time interval. Models of this type allow to incorporate reporting lags of the occured claims (see also [11]).
Let M_t be the number of catastrophic events in $[0, t]$ induced by a homogeneous Poisson process with intensity ρ. Denote the time of the i-th catastrophe by t_i and let Y_i with $i = 1, \ldots, M_t$ be independent random variables describing the impact of each catastrophe. The shot noise process is then defined as

$$\lambda_t = \lambda_0 e^{-\delta t} + \sum_{i=1}^{M_t} Y_i e^{-\delta(t-t_i)}, \quad (t > 0), \tag{2}$$

where δ denotes the decay rate of the reporting intensity, i.e. low δ indicates long reporting lags. Finally, the loss index is specified by

$$C_t = \sum_{j=1}^{N_t} X_j,$$

where the number of claims N_t is Poisson distributed with parameter $\int_0^t \lambda_u du$. The random variables X_i may be interpreted as single claims or groups of claims and are assumed to be iid with distribution function H. Dassios and

Jang [6] provide closed-form solutions for Γ-distributed claims and exponentially distributed impact variables Y_i. We will use this specific case as a performance check of our algorithms (but in principle our algorithm can also be used for heavy-tailed Y_i). Note that it is assumed that the process λ_t started sufficiently far in the past, so that at time 0, the asymptotic distribution of λ_0 may be used, from which it follows for exponentially distributed Y_i (with parameter α) that λ_0 is $\Gamma(\rho/\delta, 1/\alpha)$-distributed.
The Cat bond is now assumed to have the payoff at maturity time T

$$P_{Cat}(T) = I_{\{C_T \le K\}} + p\, I_{\{C_T > K\}},$$

where I denotes the indicator function, K is the trigger value and p denotes the fraction of the principal that has to be paid in case the loss index is above the trigger at time T. For simplicity, we will always normalize the principal to 1 and set the coupon rate equal to 0, so that the price of the Cat bond at time $t < T$ will always be less than 1. According to general asset pricing theory, the fair price $P_{Cat}(t)$ at time t of such an asset is given as the discounted expectation of its payoff under a suitable risk-neutral measure Q. The correct choice of this risk-neutral measure Q has been debated in the literature, here we follow Merton's approach [15] identifying the forward measure with the physical measure P (see also [5, 12]) so that

$$P_{Cat}(t) = P_{t,T}\, \mathbb{E}\left[1 + (p-1)\, I_{\{C_T > K\}}\right],$$

where $P_{t,T}$ is the price of a risk-free bond with maturity T at time t. In this case the probability generating function of N_T can be determined (cf. [6]):

$$\mathbb{E}\left(\theta^{N_T}\right) = \sum_{j=0}^{\infty} \theta^j\, \mathbb{P}(N_T = j) = \left(\frac{\alpha e^{-\delta T}}{\alpha + \frac{1-\theta}{\delta}(1 - e^{-\delta T})}\right)^{\frac{(1-\theta)\rho}{\alpha\delta^2 + \delta(1-\theta)}}. \tag{3}$$

Then the price of the Cat bond at time 0 with maturity T is given by

$$\begin{aligned} P_{Cat}(0) &= P_{0,T}\, \mathbb{E}\left[1 + (p-1)\, I_{\{C_T > K\}}\right] \\ &= P_{0,T}\left[p + (1-p) \sum_{j=0}^{\infty} \mathbb{P}(N_T = j) F^{*j}(K)\right], \end{aligned} \tag{4}$$

where $F^{*j}(K) = Pr(X_1 + \ldots + X_j \le K)$ denotes the j-th convolution of the claim size distribution F.

3 Simulation methodology

Since in general there are no closed-form expressions for (4) available, one has to develop efficient simulation techniques. Let us for convenience restrict to the case of exponentially distributed Y_i and first consider the crude Monte Carlo algorithm:

Crude Monte Carlo Algorithm

1. For $j = 1, \ldots, n$ do
   a) Generate a $\Gamma(\rho/\delta, 1/\alpha)$-distributed variable λ_0.
   b) Generate catastrophe times t_i by adding up Exp(ρ)-distributed random variables until $t_i > T$. Set $M_T = i - 1$.
   c) For $i = 1, \ldots, M_T$ generate Exp(α)-distributed random variables Y_i.
   d) Calculate $Z_T = \int_0^T \lambda_u du$ according to (2).
   e) Generate a Pois(Z_T)-distributed random variable determining the number N_T of claims.
   f) For $k = 1, \ldots, N_T$ generate independent H-distributed random variables X_k.
   g) Set $B^{(j)} = I_{\{C_T^{(j)} \leq K\}} + p\, I_{\{C_T^{(j)} > K\}}$, where $C_T^{(j)} = \sum_{k=1}^{N_T} X_k$.
2. The final estimate is given by

$$\hat{P}_{Cat}(0) = \frac{P_{0,T}}{n} \sum_{j=1}^{n} B^{(j)}. \tag{5}$$

Here, all non-uniform random variables are generated by standard inversion techniques (possibly by using approximations of the inverse distribution function, see e.g. [7]) and the uniform pseudo-random numbers are obtained using the Mersenne twister [14].

Formally, the involved integrands are not well suited for QMC integration. The dimension of the problem as well as the variation of the integrands is infinite, so that the Koksma-Hlawka bound is useless. In the following, we discuss the non-standard obstacles faced when designing a corresponding QMC algorithm:

Effective dimension

The two (Poisson) variables M_T and N_T may become arbitrarily large, which formally leads to an infinite dimension of the integration problem. However, the probability of large M_T, resp. N_T, is very small and thus their impact is negligible, so that one can safely limit these variables at some large value without causing relevant errors (i.e. the effective dimension of the problem is moderate). Wang and Fang [23] proposed an algorithm to estimate the effective truncation dimension of integration problems based on the ANOVA decomposition of the (square-integrable) integrand. We will first order the variables due to their importance for the integration on a heuristic basis and then apply the algorithm of Wang and Fang to assess the number of relevant variables for the integration for this particular ranking. This procedure helps to understand the performance of QMC methods for our integrands (cf. Section 4).

Another aspect of the importance of knowing at least approximately the ranking of importance of the variables shows the following refined version of the Koksma-Hlawka inequality:

$$\left| \frac{1}{N} \sum_{n=1}^{N} f(y_n) - \int_{[0,1]^s} f(u)du \right| \leq \sum_{l=0}^{s-1} \sum_{F_l} D_N^* \left(y_n^{(F_l)} \right) V^{(s-l)}(f^{(F_l)}), \tag{6}$$

where the second sum is extended over all $(s-l)$-dimensional faces F_l of the form $y_{i_1} = \cdots = y_{i_l} = 1$, the discrepancy $D_N^*(y_n^{(F_l)})$ is computed in the face of $[0,1]^s$ containing $\left(y_n^{(F_l)}\right)$, and $V^{(s-l)}(f^{(F_l)})$ is the variation (in the sense of Vitali) of the restriction $f^{(F_l)}$ of f to F_l (see e.g. [9]).
Thus, the integration error can be kept low by assigning the first dimensions of the QMC sequences (which in most constructions have a lower discrepancy in the lower-dimensional faces, see e.g. [18]) to those variables which are responsible for most of the integrand's variation.

Variation and jumps

Due to the involved characteristic functions, our integrands have discontinuities that are not parallel to the axes. This implies infinite variation of the integrands in the sense of Hardy and Krause (see [18]), even when we fix the maximum number of claims. Literature on QMC algorithms for integrands with jumps is scarce. Due to a result of Zaremba (cf. [18]), for any point set $\omega_N = \{x_1\}_{i=1,\ldots,N}$, $x_i \in U^s$ and convex set $C \subset U^s$, we have

$$\left| \frac{1}{N} \sum_{i=1, x_i \in C}^{N} f(x_i) - \int_{U^s} f(x) 1_C \, dx \right| \leq (V(f) + |f(1,\ldots,1)|) \, J_N(\omega),$$

where $f(x)$ is a function of bounded variation $V(f)$ on U^s and $J_N(\omega)$ denotes the isotropic discrepancy of ω. A similar bound can be given for the set of Jordan measurable sets, but the best known sequences in terms of the isotropic discrepancy lead to an upper bound $\mathcal{O}\left(N^{-1/s} \log N\right)$ for the convergence order, which is much worse than Monte Carlo. Based on computational experiments and heuristical arguments, Berblinger et. al [4] and Morokoff and Caflish [16] conjecture a convergence rate of $\mathcal{O}\left(N^{-1/2-1/2s}\right)$ for integrands with discontinuities.

Wang [22] proposes several techniques to "smooth away" jumps. We will use conditional Monte Carlo to smooth the jumps. For fixed N_T, write

$$\begin{aligned} \mathbb{E}\left[I_{\{C_T > K\}}\right] &= \mathbb{P}(X_1 + \ldots + X_{N_T} > K) \\ &= \mathbb{E}(\mathbb{P}(X_1 + \ldots + X_{N_T} > K | X_1, \ldots, X_{N_T - 1})) \\ &= \mathbb{E}\left[1 - H\left(K - \sum_{i=1}^{N_T - 1} X_i\right)\right]. \end{aligned} \tag{7}$$

Note that this step smoothes the integrand and saves one dimension. Moreover, as a conditional Monte Carlo estimator, the variance of (7) is reduced.

Rare events

If, as in realistic situations, the trigger K is very large, the event $\{C_T > K\}$ is rare. Crude Monte Carlo algorithms then have poor performance (e.g. the relative error defined by standard deviation of the estimate over the required quantity is not bounded). Thus we adopt an asymptotically efficient importance sampling algorithm of [2] and analyze its performance in combination with QMC methods. The combination of QMC and rare event techniques seems to be new and deserves an empirical analysis. We will use exponential tilting of our distributions: for fixed N_T, replace

$$\mathbb{P}(C_T > K) = \int 1_{\{x_1+\ldots+x_{N_T}>K\}} \prod_{i=1}^{N_T} dH(x_i) \tag{8}$$

$$= \int 1_{\{x_1+\ldots+x_{N_T}>K\}} \frac{e^{\theta(x_1+\ldots+x_{N_T})}}{\hat{H}(\theta)^{N_T}} \prod_{i=1}^{N_T} d\tilde{H}(x_i), \tag{9}$$

where $\hat{H}(t)$ is the moment generating function of H and $d\tilde{H}(x_i)/dH(x_i) = e^{-\theta x_i}\hat{H}(\theta)$. So we now sample from the new distribution $\tilde{H}$. Asymptotic efficiency can then be achieved by choosing θ in such a way that $\mathbb{E}_{\tilde{H}}[C_T] \approx K$. This approach requires the existence of the moment-generating function of H.

An additional way to reduce the variation for problems with general claim distribution H is to use exact solutions for some simple cases (provided in our case by Dassios and Wang [6]) as control variables. In particular when the correlation between the control variate and the required estimate is high, this leads to a significant additional variance reduction of the problem (cf. Example 3 in Section 4).

Since in situations where the number of function evaluations is not too high, Quasi-Monte Carlo methods are especially competitive for low dimensions, the best numerical performance can be obtained by implementing hybrid Monte Carlo techniques, i.e. QMC sequences for the initial dimensions and the remaining dimensions are then simulated by crude Monte Carlo (cf. [21]). The ranking of the dimensions is done on a heuristic basis with (6) in view. In our analysis we used 50 dimensions of QMC sequences.
As an alternative to the deterministic Koksma-Hlawka bound (which as mentioned earlier is not helpful for our application) for the integration error, one can obtain probabilistic error bounds by randomizing quasi-Monte Carlo methods, usually at the expense of some of the convergence speed (see e.g. [13]). For the purpose of comparison, we will include a random-start Halton algorithm proposed by Hickernell and Wang [24] in our set of QMC techniques applied in Section 4.

4 Numerical analysis

In this part we present numerical illustrations for the proposed algorithms, comparing importance sampling, smoothing and control variate effects for Monte Carlo and QMC techniques in three examples. In the first two examples exact solutions are available for comparison. For simplicity, we assume that the interest rate is zero during the lifetime of the Cat bond (a positive interest rate would only effect the term $P_{0,T}$ and is thus not of interest for the present study). First, we try to analyze the effect of discontinuities of the integrand on the convergence speed of the algorithm.

Example 1: For the model in Section 2, let $p = 0.5$, $\delta = 1.5$, $\rho = 2$, $\alpha = 2$. Thus the mean number of claims N within a year is $8/3$. Furthermore, we assume that the maturity of the Cat bond is one year ($T = 1$) and the claims X_i are exponentially distributed with mean 5. In this case exact solutions for the Cat bond price can be approximated by using a series expansion of (3) and the fact that the nth convolution F^{*n} is $\Gamma(n, 1/5)$-distributed. The series in (4) is then truncated after sufficiently many terms so that the significant digits are exact (see Table 1 below).

Note that the crude Monte Carlo algorithm presented in Section 3 induces two types of non axes-parallel jumps in the corresponding integrand, namely during the evaluation of M_T and N_T and those due to the characteristic functions. In this example we will circumvent the former by using the probability distribution of N_T (determined by a series expansion of (3)) and obtaining N_T by inversion from a uniform random variate. This simplifies the algorithm and allows to solely study the effect of jumps due to the characteristic functions. Utilizing (7) we then construct a smoothed algorithm with bounded variation and compare its converge speed with the original algorithm. Since here it is easily possible to rank the variables according to their contribution to the variance, we are able to compute the effective truncation dimension of the problem. Table 2 shows the results for both crude and smoothed algorithm as a function of K.

K	5	15	20	25	35	50
$P_{Cat}(0)$	0.688000	0.828120	0.875373	0.910710	0.955471	0.985153

Table 1. Exact Cat bond prices for some triggers K in Example 1

We tested the crude and the smoothed algorithm with Halton, Sobol, Faure, randomized Halton and crude Monte Carlo sequences and compared the results with the available exact solution. Figure 1 shows least-squares fits of

	Crude Algorithm						Smoothed Algorithm					
	L						L					
K	1	2	3	5	8	10	1	2	3	5	8	10
5	60	85	96	99.9	100	100	91	98	99.7	100	100	100
15	52	66	80	96	99.9	100	70	84	93	99.2	100	100
20	49	60	73	93	99.7	100	65	78	88	98	100	100
25	46	55	66	87	99.7	100	62	71	83	96	99.8	100
35	40	47	57	76	96	99.9	53	61	71	88	98.7	100
50	33	38	42	59	87	95	40	49	56	74	93	98

Table 2. Cumulated variance of the L most important variables in Example 1 (in % of total variance, 100 denotes values larger than 99.94)

the absolute simulation errors of the various methods in a double-logarithmic diagram for four values of K (for randomized Halton and Monte Carlo, we used the arithmetic average of the absolute errors of 25 independent runs). The slopes of these lines reflecting the convergence rate of the algorithm are given in Table 3.

	Crude Algorithm					Smoothed Algorithm				
K	MC	S	H	RH	F	MC	S	H	RH	F
5	−0.49	−0.69	−0.70	−0.70	−0.74	−0.51	−0.98	−0.98	−0.84	−0.80
15	−0.52	−0.47	−0.44	−0.63	−0.69	−0.54	−0.80	−0.98	−0.76	−0.85
20	−0.52	−0.68	−0.72	−0.62	−0.59	−0.53	−0.73	−0.93	−0.72	−0.67
25	−0.50	−0.66	−0.56	−0.57	−0.75	−0.52	−0.77	−0.74	−0.71	−0.71
35	−0.49	−0.60	−0.60	−0.55	−0.74	−0.50	−0.82	−0.61	−0.61	−0.75
50	−0.47	−0.52	−0.61	−0.56	−0.75	−0.45	−0.67	−0.64	−0.57	−0.78

Table 3. Exponents x of the empirical convergence rate $\mathcal{O}(N^x)$ of the various simulation algorithms in Example 1.

As expected, for low-effective dimension problems, QMC algorithms clearly outperform Monte Carlo in terms of convergence speed. This effect becomes weaker for larger values of K (and thus larger effective dimension). Whereas the improvement of smoothing is less pronounced for the Monte Carlo algorithm, one observes that smoothing has a considerable effect on the convergence rate for QMC sequences. This is due to the fact that in addition to bounding the variation of the integrand, smoothing significantly lowers the effective truncation dimension of the problem (cf. Table 2). Both effects are more pronounced for (randomized) Halton sequences than for the Sobol case.

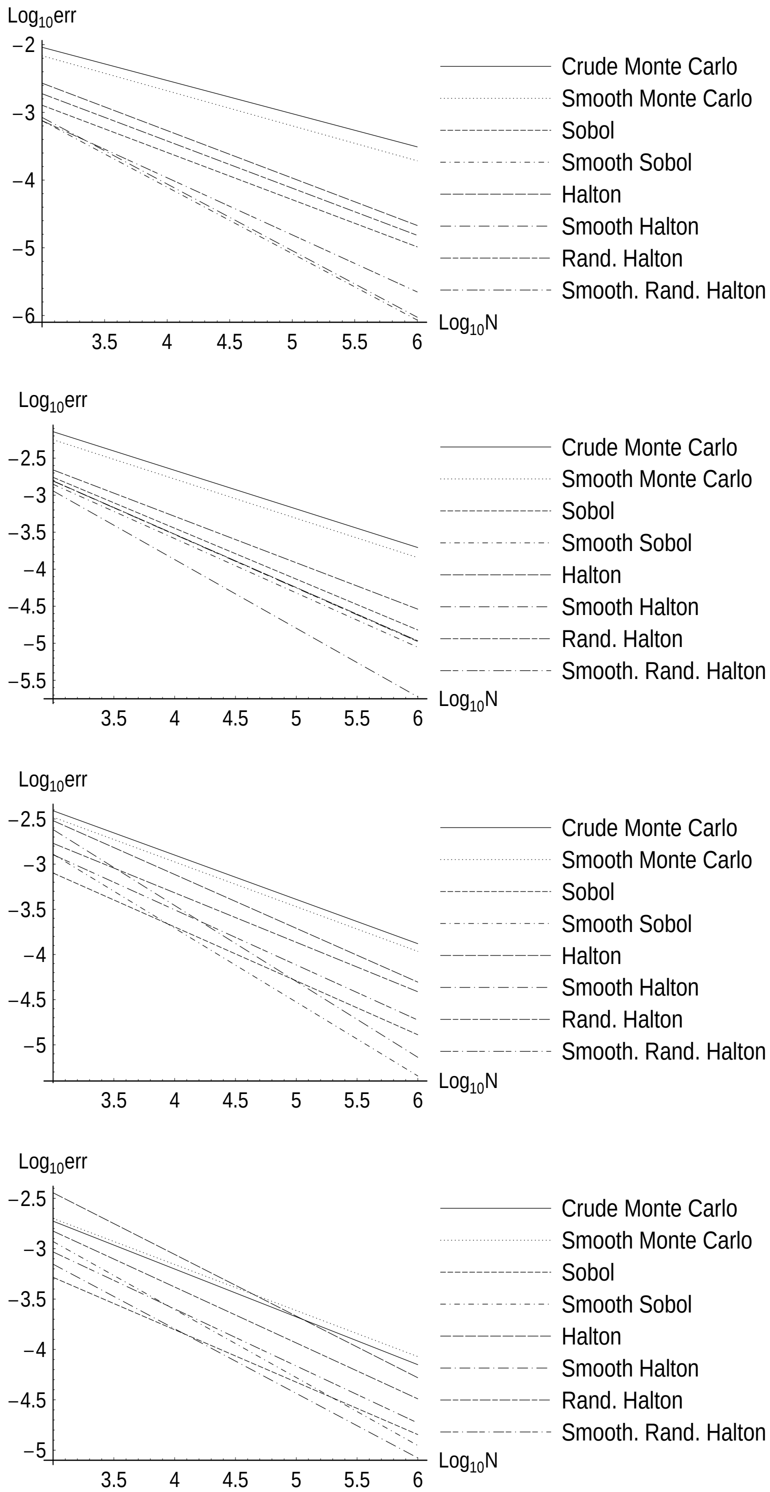

Fig. 1. Least square fits of the simulation error as a function of N in Example 1 for various values of K (topdown: $K = 5, 20, 35, 50$)

Example 2: Let us now consider a set of parameters also used in [6], namely $p = 0.5$, $\alpha = 1, \delta = 0.3, \rho = 4$ and $T = 1$ and $X_i \sim \text{Exp}(1)$. Thus $\mathbb{E}N_1 = 13.33$. Analogously to Example 1, it is possible to derive the exact price of the Cat Bond in this case. However, in contrast to the preceding example, for the simulation here we do not calculate the distribution of N_T explicitly, but use the original algorithm proposed in Section 3 (since the distribution of N_T will in general not be available). The purpose of this example is to illustrate the efficiency of importance sampling. Table 5 compares the performance of the crude algorithm with the importance sampling algorithm based on (9). Absolute and relative errors (defined as the error of crude Monte Carlo over the error of the alternative method) for various values of N and K are given. Figure 2 depicts the corresponding log-log plots and also includes the performance of the corresponding smoothed algorithms.
The question of assigning the individual dimensions of the QMC variates to the random variables of the problem turns out to be crucial in this case. We have chosen the ordering on a heuristic basis (including the implementation of several good candidates and comparisons) to keep the effective truncation dimension as low as possible.
Due to the larger number of claims and the additional requirement of (quasi-) random variates in the evaluation of N_T, one faces a higher effective truncation dimension than in Example 1 (see Table 4). Again, for small K, QMC sequences clearly outperform Monte Carlo sequences due to low effective truncation dimension, and this effects diminishes as K and thus the effective dimension of the problem increases. For large K, the randomized Halton turns out to be the best crude algorithm. Smoothing is not really beneficial in this example, but importance sampling, originally designed for Monte Carlo applications, turns out to work very well for QMC sequences (note the improvement of the convergence speed for large K, which is the rare event situation). One also observes that the importance of the first few variables is dramatically increased. The combination of Sobol sequences and importance sampling seems to be by far the best for this example.

	Crude Algorithm					**IS Algorithm**				
	L					L				
K	5	10	15	20	25	5	10	15	20	25
5	46	65	86	92	99	45	65	87	92	99
25	31	37	45	56	60	49	53	57	60	64
35	16	25	26	27	35	34	41	47	50	58

Table 4. Cumulated variance of the L most important variables in Example 2 (in % of total variance)

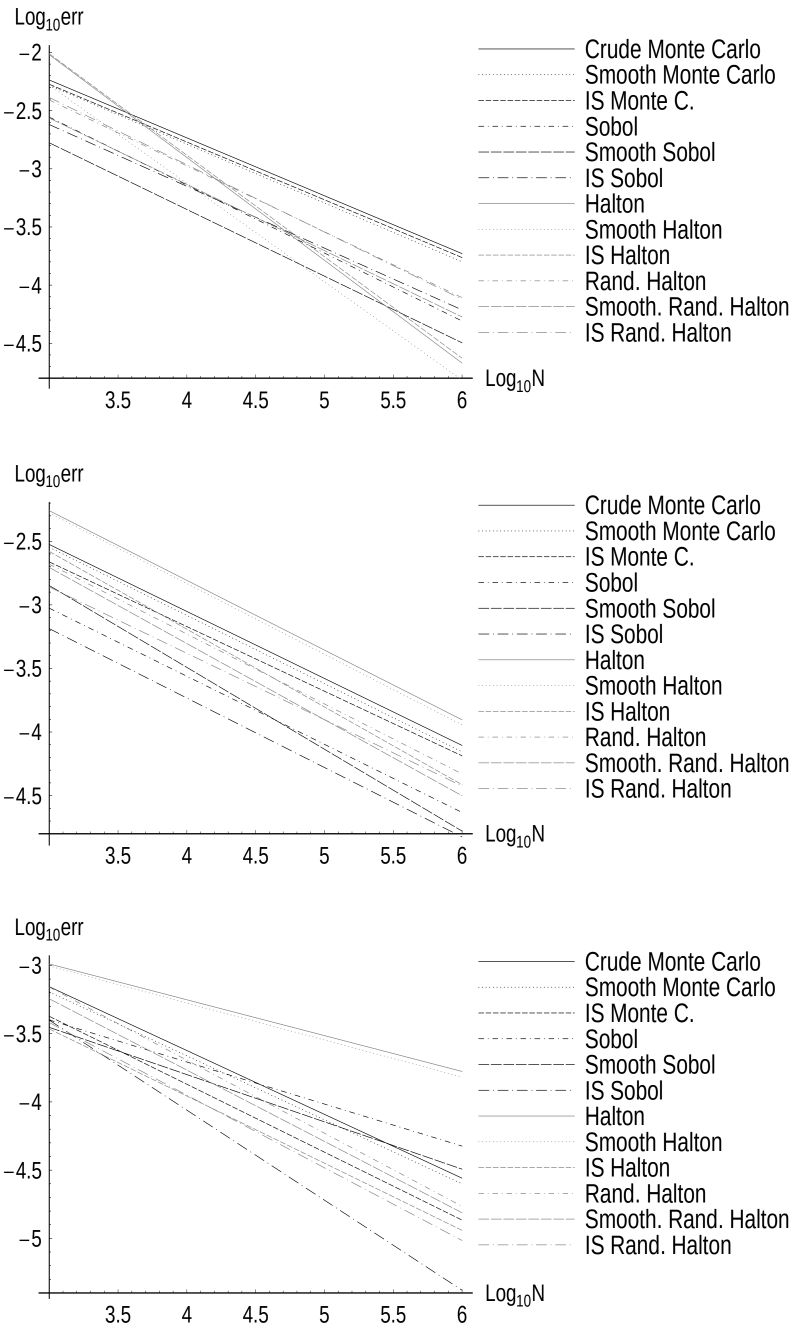

Fig. 2. Least square fits of the simulation error as a function of N in Example 2 for various values of K (topdown: $K = 5, 25, 35$.)

	Crude Algorithm					**Importance Sampling Algorithm**				
N	2^{12}	2^{14}	2^{16}	2^{18}	2^{20}	2^{12}	2^{14}	2^{16}	2^{18}	2^{20}
$K = 5$	**Exact value** = 0.53214									
Monte Carlo	1.6E-3	7.2E-4	4.8E-4	1.8E-4	1.0E-4	2.0E-3	7.6E-4	3.8E-4	1.5E-4	9.1E-5
	1	1	1	1	1	0.81	0.94	1.26	1.14	1.22
Sobol	4.6E-5	4.2E-4	5.8E-4	1.0E-4	2.0E-5	2.9E-4	3.4E-4	6.8E-4	2.6E-5	7.3E-6
	34.06	1.71	0.83	1.68	5.29	5.33	2.07	0.72	6.68	13.88
Halton	2.5E-3	9.7E-4	1.1E-4	1.0E-4	1.7E-5	2.6E-3	1.2E-3	1.2E-4	6.9E-5	3.4E-5
	0.62	0.74	4.33	1.78	5.90	0.60	0.61	4.05	2.58	2.97
Rand. Halton	7.7E-4	4.7E-4	2.2E-4	8.9E-5	4.2E-5	7.6E-4	5.0E-4	1.9E-4	8.8E-5	4.0E-5
	2.00	1.54	2.22	2.00	2.50	2.02	1.44	2.51	2.04	2.50
Faure	3.0E-4	4.8E-4	1.4E-3	3.2E-3	7.5E-5	8.7E-4	7.6E-5	1.5E-3	3.6E-4	1.0E-4
	5.21	1.49	0.34	0.55	1.36	1.77	9.48	0.31	0.50	0.96
$K = 25$	**Exact value** = 0.97671									
Monte Carlo	1.4E-3	8.4E-4	3.8E-4	1.3E-4	8.3E-4	1.2E-3	4.7E-4	2.6e-4	1.1E-4	6.7E-5
	1	1	1	1	1	1.21	1.80	1.43	1.11	1.23
Sobol	5.5E-4	1.6E-4	1.3E-4	8.7E-5	2.6E-5	6.0E-4	1.9E-4	1.4E-4	1.4E-5	3.7E-5
	2.63	5.19	2.92	1.48	3.17	2.38	4.40	2.56	9.52	2.22
Halton	4.6E-3	6.8E-4	2.6E−-4	4.7E-4	2.3E-4	1.2E-3	8.2E-4	4.2E−-4	1.5E-4	1.3E-5
	0.31	1.26	1.48	0.27	0.35	1.15	1.04	0.90	0.87	6.56
Rand. Halton	1.0E-3	4.0E-4	2.0E-4	8.4E-5	5.1E-5	5.4E-4	3.2E-4	1.3E-4	1.1E-4	4.0E-5
	1.43	2.10	1.80	1.53	1.61	2.64	2.58	2.88	1.21	2.08
Faure	9.1E-3	3.3E-3	1.2E-3	1.4E-4	5.3E-7	2.3E-3	1.3E-3	2.0E-4	1.1E-4	5.6E-5
	0.15	0.24	0.31	0.93	1.55	0.61	0.66	1.87	1.22	1.47
$K = 35$	**Exact value** = 0.99842									
Monte Carlo	4.3E-4	1.8E-4	7.5E-5	6.3E-5	3.8E-5	2.1E-4	1.2E-4	6.3E-5	3.0E-5	1.6E-5
	1	1	1	1	1	1.98	1.42	1.19	2.10	2.41
Sobol	1.9E-4	4.8E-4	1.1E-4	9.0E-5	2.5E-5	2.8E-4	1.2E-4	8.9E-6	1.4E-5	2.1E-5
	2.24	0.38	0.67	0.70	1.49	1.55	1.52	8.39	4.44	1.81
Halton	8.1E-4	7.2E-4	4.3E-4	2.1E-4	9.8E-5	1.2E-4	7.9E-5	9.4E-5	1.8E-5	2.7E-5
	0.53	0.25	0.17	0.29	0.40	3.45	2.29	0.80	3.41	1.42
Rand. Halton	3.2E-4	1.7E-4	7.3E-5	3.5E-5	2.0E-5	1.7E-5	7.7E-5	4.0E-5	2.2E-5	8.9E-6
	1.35	1.06	1.02	1.78	1.89	2.49	2.37	1.86	2.82	4.33
Faure	3.1E-4	2.8E-4	4.8E-4	1.6E-4	1.9E-5	3.7E-5	2.3E-4	1.9E-4	3.9E-5	3.2E-5
	1.38	0.65	0.15	0.39	2.09	1.71	0.79	0.39	1.59	1.20

Table 5. Absolute errors for crude and importance sampling algorithms in Example 2 (including the relative error w.r.t. the crude Monte Carlo estimate).

Example 3: Let us consider Example 2 again, but now we assume that $X_i \sim LN(-\log 2/2, \sqrt{\log 2})$, i.e. the claims have a lognormal distribution with parameters chosen in such a way that the first two moments coincide with the exponential distribution of Example 2. In this case exact solutions are not available any more and are approximated by a Monte-Carlo run with very large N ($N = 50 \times 10^6$). We investigate the suitability of using the exact solutions for the exponential claim sizes X_i as control variates for our heavy tailed case. Proceeding in the usual way, the estimate (5) is thus replaced by

$$\hat{P}_{Cat,CV}(0) = \frac{P_{0,T}}{n} \sum_{j=1}^{n} \left(B_L^{(j)} - \beta(B_E^{(j)} - B_E^*) \right), \tag{10}$$

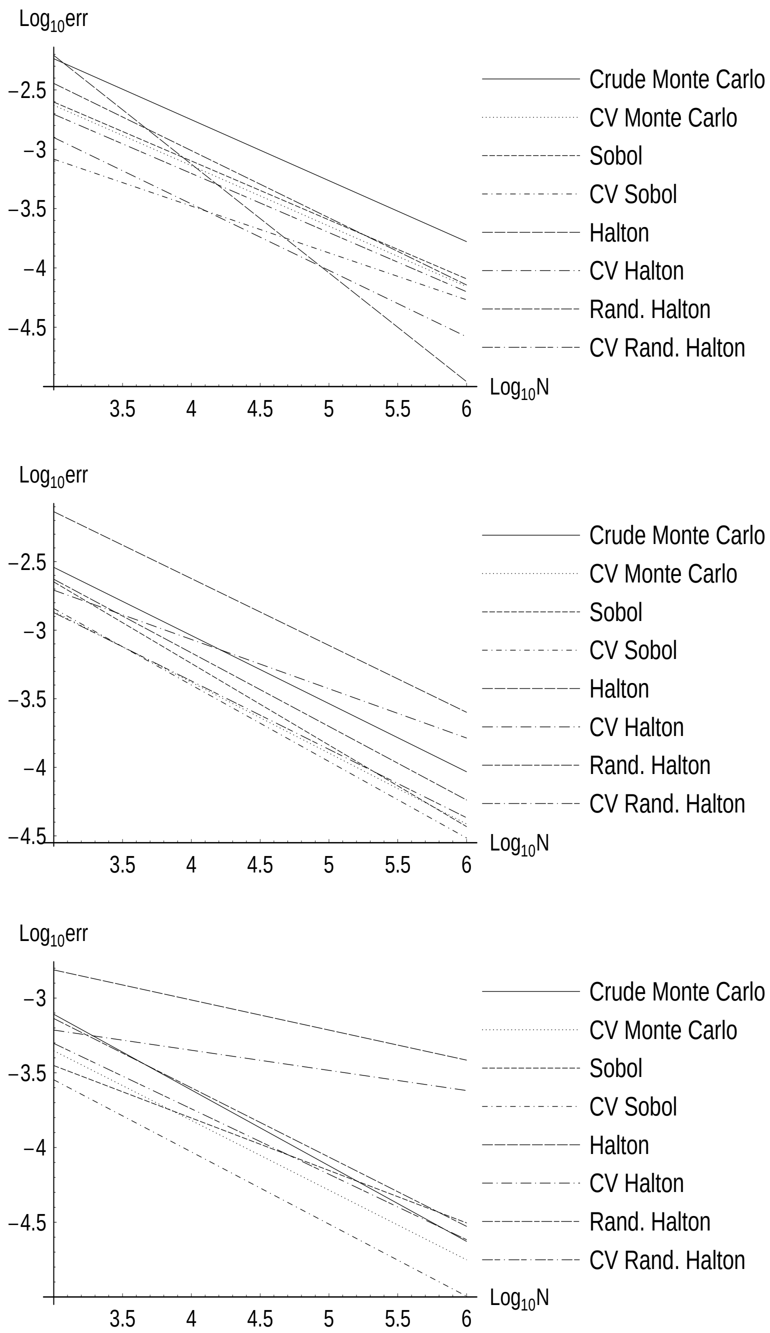

Fig. 3. Least square fits of the simulation error as a function of N in Example 3 for various values of K (topdown: $K = 5, 25, 35$.)

where $B_L^{(j)}$ ($B_E^{(j)}$) are calculated for lognormal (resp. exponential) claim size distributions H, and B_E^* is the known exact value for the exponential case. The weighting factor β is chosen such that the variance of (10) is minimized, that is

$$\beta = \frac{\mathrm{Cov}(B_L, B_E)}{\mathrm{Var}(B_E)},$$

these quantities being themselves estimated from the same simulation.
Figure 3 gives the double-logarithmic plots comparing the least square fits of the simulation errors with and without control variates as a function of N. As in the preceding examples, the efficiency improvement of QMC over Monte Carlo decreases with increasing K. Control variate based algorithms always outperform their crude counterparts and the QMC-Sobol-Control variate algorithm always performs among the best.

5 Conclusion

We have developed QMC algorithms for the valuation of Cat bonds in a model recently proposed by Dassios and Jang [6]. Formally, the integrands involved in the evaluation are not suited for QMC integration as they are infinite-dimensional and have unbounded variation. Crude QMC algorithms outperform Monte Carlo only for problems with small effective dimension. Nevertheless, suitable modifications as smoothing, importance sampling and control variate techniques lead to QMC algorithms which yield significantly faster convergence rates than their Monte Carlo counterparts. Furthermore, we observed that the size of the effective dimension seems to have much more influence on the simulation performance than the occurrence of discontinuities.

References

1. H. Albrecher, J. Hartinger and R.F. Tichy. Multivariate Approximation Methods for the Pricing of Catastrophe-Linked Bonds, *Internat. Ser. Numer. Math.* 145, 21–39, 2003.
2. S. Asmussen, K. Binswanger and B. Højgaard. Rare events simulation for heavy-tailed distributions, *Bernoulli* 6(2), 303–322, 2000.
3. S. Asmussen, P. Glynn. *Stochastic simulation: with a view towards stochastic processes.* Springer, to appear, 2004.
4. M. Berblinger, Ch. Schlier and T. Weiss. Monte Carlo Integration with quasi-random numbers: experience with discontinuous integrands, *Comput. Phys. Comm.* 99, 151–162, 1997.
5. S.H. Cox and H.W. Pedersen. Catastrophe Risk Bonds, *North American Actuarial Journal* 4, 56–82, 2000.
6. A. Dassios and J.W. Jang. Pricing of catastrophe reinsurance and derivatives using the Cox process with shot noise intensity, *Finance Stoch.* 7(1), 73–95, 2003.

7. L. Devroye. *Non-Uniform Random Variate Generation*, Springer, 1986.
8. N. Doherty. Financial innovation in the management of catastrophe risk. *Journal of Applied Corporate Finance* 10, 134–170, 1997.
9. M. Drmota and R.F. Tichy. *Sequences, Discrepancies and Applications*, volume 1651 of *Lecture Notes in Mathematics*, Springer-Verlag, 1997.
10. R. Gorvett. Insurance securitization: The development of a new asset class. *In: Securitization of Risk*, Casualty Actuarial Society, 133–173, 1999.
11. C. Klüppelberg and T. Mikosch. Explosive Poisson shot noise processes with applications to risk reserves, *Bernoulli* 1(1-2), 125–147, 1995.
12. J. Lee and M. T. Yu. Pricing default-risky cat bonds with moral hazard and basis risk. *Journal of Risk and Insurance* 69, 25–44, 2002.
13. P. L'Ecuyer and Ch. Lemieux. Recent advances in randomized quasi-Monte Carlo methods, *Internat. Ser. Oper. Res. Management Sci.* 46, 419–474, 2002.
14. M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM Trans. on Modeling and Computer Simulation* 8, 3–30, 1998.
15. R.C. Merton. Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144, 1976.
16. W.J. Morokoff and R.E. Caflisch. Quasi-Monte Carlo integration, *Journal of Computational Physics* 122, 218–230, 1995.
17. A. Muermann. Catastrophe derivatives. *In: Encyclopedia of Actuarial Science*, Eds: J. Teugels and B. Sundt, John Wiley and Sons, to appear, 2004.
18. H. Niederreiter. *Random number generation and quasi-Monte Carlo methods*, SIAM, 1992.
19. S.H. Paskov and J. Traub. Faster Valuation of Financial Derivatives, *Journal of Portfolio Management* 22, 113–120, 1995.
20. I.H. Sloan and H. Woźniakowski. When are quasi-Monte Carlo algorithms efficient for high-dimensional integrals?, *Journal of Complexity* 14(1), 1–33, 1998.
21. J. Spanier. Quasi-Monte Carlo methods for particle transport problems, *Lecture Notes in Statist.* 106, 121–148, 1995.
22. X. Wang. Improving the rejection sampling method in quasi-monte carlo methods, *Journal of Computational and Applied Mathematics* 114, 231–246, 2000.
23. X. Wang and K.T. Fang. The effective dimension and quasi-monte carlo integration, *Journal of Complexity* 19, 101–124, 2003.
24. X. Wang and F.J. Hickernell. Randomized Halton sequences, *Math. Comput. Modelling* 32(7-8), 887–899, 2000.