

A Generalized Linear Transformation Method for Simulating Meixner Lévy Processes

Junichi Imai* and Ken Seng Tan†

Abstract—In this paper, we consider an enhanced quasi-Monte Carlo (QMC) method for pricing derivative securities when the underlying asset price follows an exponential Lévy process. In particular, we focus on a special family of the Lévy process known as the Meixner process. The enhanced QMC is based on a generalization of the linear transformation (LT) method of Imai and Tan (2006). The generalized LT method can be used to simulate general stochastic processes and hence has a wider range of application than the original LT which only applies to the Gaussian process. Using some option examples with dimensions ranging from 4 to 250 as test cases, the numerical results suggest that the generalized LT-based QMC substantially outperforms the standard applications of quasi-Monte Carlo and Monte Carlo methods.

Keywords: *Quasi-Monte Carlo, computational finance, derivative securities, dimension reduction*

1 Introduction

In the last few decades, we have observed significant advances in the field of financial mathematics. Sophisticated financial models (such as for modeling the dynamics of the asset prices, interest rates, currencies, etc.) have been proposed. New mathematical tools and innovative numerical methods have also been discovered. The new era of financial mathematics has, in part, been spurred by the celebrated Black-Scholes models [3] introduced in 1973 and, in part, due to the subsequent tremendous expansion and sophistication of the financial markets.

The basic assumption of the Black-Scholes model is the normality of the log-returns of the underlying asset price. Numerous empirical studies, on the other hand, have clearly pointed the inadequacy of the normality assumption since empirically the log-returns of the underlying

typically have higher kurtosis than that of the normal distribution. For this reason, a number of more elaborate models including GARCH models (e.g. see [6]) and models with stochastic volatility (e.g. see [11]) has been proposed. More recently, the Lévy process as an alternate process for modeling the dynamic of the log-returns of the underlying and in derivative pricing has been gaining popularity (e.g. see [23] and [15]).

The sophistication of the models and the complexity of the financial products also imply that only in rare cases there exists tractable pricing formulae for these products. In pricing most exotic derivative securities, we typically resort to numerical methods such as the binomial models, finite difference methods, Monte Carlo (MC), or quasi-Monte Carlo (QMC) methods. In the past decade, QMC has become a popular tool in computational finance. Early finance applications of QMC mainly focus on the Black-Scholes type models; see [14]. More recently, this method has been extended to other more exotic models, particularly the Lévy models (see [8], [19], [2], [16], [13]). Many numerical studies seem to suggest that the success of QMC is intricately related to the notion of *effective dimension*. Dimension reduction techniques such as the Brownian bridge construction ([17] and [4]), the principal component construction [1], and the linear transformation (LT) method [12] have been proposed to further enhance QMC.

In this paper, we provide further numerical evidence on the efficiency of the generalized LT method that has recently been proposed by [13]. The generalization has the advantage that it no longer confines to the Gaussian process, as in the original LT method. This implies that the generalized LT has a wider range of applications, particularly effective for simulating Lévy processes. The rest of the paper is organized as follows. Section 2 describe the original LT method as well as our proposed generalization. Section 3 presents numerical evidences on the effectiveness of the proposed dimension reduction technique. A special family of Lévy process is used in our illustration. Section 4 concludes the paper.

2 An extension of the LT method

In [12], Imai and Tan demonstrated that the efficiency of the QMC can be increased by applying the linear trans-

*J. Imai is Associate Professor in the Faculty of Science and Technology, Keio University, Japan (email: jimai@ae.keio.ac.jp). He is grateful for the research support from a Grant-in-Aid for Scientific Research from the Japan Society for the Promotion of Science (21710157).

†K.S. Tan, corresponding author, holds the Canada Research Chair in Quantitative Risk Management in the Department of Statistics and Actuarial Science, University of Waterloo, Canada, and the Cheung Kong Scholar in the China Institute for Actuarial Science, Central University of Finance and Economics, Beijing, China (email: kstan@uwaterloo.ca). Tan acknowledges the research funding from the Canada Research Chairs Program and NSERC.

formation (LT) dimension reduction technique. This method exploits explicitly the payoff structure of derivative securities and works as long as the underlying assets follow Gaussian process. Suppose we are interested in estimating $E[g(\mathbf{Z})]$, where $g(\mathbf{Z})$ corresponds to the payoff function of a derivative security and \mathbf{Z} is a d -dimensional standardized normal random vector. Problems of this kind are of interest to us as $E[g(\mathbf{Z})]$ can be interpreted as the price of a derivative security when g corresponds to its discounted payoff function. The standard MC estimate of $E[g(\mathbf{Z})]$ is given by the sample average over all simulated g with \mathbf{Z} randomly sampled. The standard QMC application is similar to MC except that \mathbf{Z} is now generated via inverse transforms from a (randomized) low discrepancy sequence.

The essence of the LT-based QMC is to exploit a simple fact that $E[g(\mathbf{Z})] = E[g(\mathbf{AZ})]$ for any orthogonal matrix \mathbf{A} ; i.e. $\mathbf{A}^T \mathbf{A} = \mathbf{I}$ and \mathbf{I} is the identity matrix. This implies that for an arbitrary orthogonal matrix \mathbf{A} , another consistent estimate of $E[g(\mathbf{Z})]$ can be obtained by taking the sample average based on $g(\mathbf{A}\boldsymbol{\varepsilon}_i), i = 1, \dots, N$. To contrast the difference between MC and QMC, it is important to point out that if we were to use MC to estimate $E[g(\mathbf{Z})]$, both $g(\boldsymbol{\varepsilon}_i)$ -based and $g(\mathbf{A}\boldsymbol{\varepsilon}_i)$ -based estimators are equivalent and that there is no difference in terms of numerical efficiency. On the other hand if QMC were used, the numerical accuracy of these estimators can be very different. In fact, [12] demonstrates that a careful choice of \mathbf{A} leads to a much greater precision of $g(\mathbf{A}\boldsymbol{\varepsilon}_i)$ -based QMC estimator relative to the crude $g(\boldsymbol{\varepsilon}_i)$ -based QMC estimator. The numerical examples in [12], [20] and [21] even suggest that the LT-based QMC can be more efficient relative to the Brownian bridge construction-based QMC [17] and the principal component construction-based [1]. One possible explanation accounts for the difference between MC and QMC can be attributed to the impact due to effective dimension.

Motivated by the dimension reduction and its intricate connection to variance decomposition, [12] proposes an algorithm that optimally determining the orthogonal matrix column by column. To describe the algorithm, it is essential to introduce some additional notation. Let \mathbf{A}_k denote the k -th column of \mathbf{A} , $\langle \mathbf{a}, \mathbf{b} \rangle$ denote the inner product between vectors \mathbf{a} and \mathbf{b} , and $\hat{\boldsymbol{\varepsilon}}_k = (v_1, \dots, v_{k-1}, 0, \dots, 0)^T$ denote the d -dimensional vector with arbitrary chosen random variables v_1, \dots, v_{k-1} . Then the optimal \mathbf{A}_k can be obtained by solving the following optimization problem:

$$\begin{aligned} & \max_{\mathbf{A}_k \in \mathbb{R}^d} \left(\frac{\partial g(\mathbf{A}\boldsymbol{\varepsilon})}{\partial \varepsilon_k} \Big|_{\boldsymbol{\varepsilon}=\hat{\boldsymbol{\varepsilon}}_k} \right)^2 \\ & \text{subject to } \|\mathbf{A}_k\| = 1 \text{ and} \\ & \quad \langle \mathbf{A}_j^*, \mathbf{A}_k \rangle = 0, \quad j = 1, \dots, k-1. \end{aligned} \tag{1}$$

We emphasize that the above algorithm is carried out iteratively for $k = 1, 2, \dots, d$. This implies that in the k -th iteration, $\mathbf{A}_j, j = 1, \dots, k-1$ are already optimally determined in the earlier iteration steps. When QMC is combined with the optimally determined \mathbf{A} using the above algorithm to estimate $E[g(\mathbf{Z})]$ (via $E[g(\mathbf{AZ})]$), [12] refers the resulting method as the QMC-LT.

There are several advantages associated with the LT-based QMC. First, it exploits explicitly the payoff structure of the derivative securities. Second, and more importantly, the numerical studies conducted in [12], [21] and [20] have concluded the competitive advantage of LT-based QMC relative to other QMC-based methods, even for dimensions of several hundred. On the other hand, one severe limitation of QMC-LT is that it is restricted to a class of function which depends on a vector of normal random variables. This poses no problem for derivative pricing under Black-Scholes framework as illustrated in [12], [20] and [21]. For other models, such as those involving Lévy process, this requirement is no longer met. It is therefore of significant interest to address the general problem of estimating $E[g(\mathbf{X})]$, where $\mathbf{X} = (X_1, \dots, X_d)^T$ is a vector of d iid random variables with arbitrary probability density function (pdf) $f(x)$ and cumulative distribution function (cdf) $F(x)$. We emphasize here that the distribution of \mathbf{X} needs not be normally distributed. Motivated by this, [13] proposes an extension of LT which is based on the following series of transformations. First note that

$$E[g(\mathbf{X})] = \int_{\Omega} g(\mathbf{x}) f(x_1) \cdots f(x_d) dx_1 \cdots dx_d,$$

where Ω is the domain of \mathbf{X} . By substituting $y_i = F(x_i), i = 1, \dots, d$, the above integration reduces to an integration problem over $[0, 1]^d$:

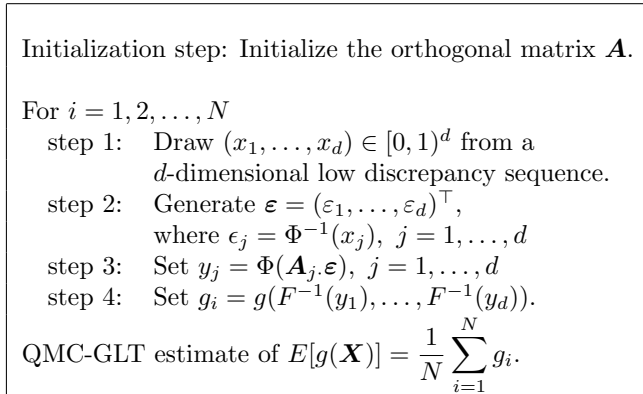
$$E[g(\mathbf{X})] = \int_{[0,1]^d} g(F^{-1}(y_1), \dots, F^{-1}(y_d)) dy_1 \cdots dy_d.$$

Now consider the transformation $Z = \Phi^{-1}(Y)$ where Φ represents the cdf of the standard normal distribution. Then $E[g(\mathbf{X})]$ can be expressed as follows:

$$\begin{aligned} & \int \cdots \int_{-\infty}^{\infty} g(F^{-1}(\Phi(z_1)), \dots, F^{-1}(\Phi(z_d))) \times \\ & \quad \phi(z_1) \cdots \phi(z_d) dz_1 \cdots dz_d \\ & = E[g(F^{-1}(\Phi(Z_1)), \dots, F^{-1}(\Phi(Z_d)))], \end{aligned} \tag{2}$$

where ϕ is the pdf of the standard normal, and $\mathbf{Z} = (Z_1, \dots, Z_d)^T$ is a vector of independent standard normal random variable. The significance of (2) is that after some trivial transformations, the expectation is now taken with respect to the normal distribution. This implies that another consistent estimator of $E[g(\mathbf{X})]$ can be obtained via $E[g(F^{-1}(\Phi(\mathbf{A}_1 \mathbf{Z})), \dots, F^{-1}(\Phi(\mathbf{A}_d \mathbf{Z})))]$, for any orthogonal matrix \mathbf{A} where \mathbf{A}_j corresponds to the j -th row of \mathbf{A} . We refer this approach as the generalized LT method (GLT) and when this method is combined with

Figure 1: QMC-GLT Algorithm for Estimating $E[g(\mathbf{X})]$



QMC, we denote as QMC-GLT. Figure 1 describes the QMC-GLT method in algorithmic form. We now make the following remarks with regards to the GLT-based QMC.

Remark 1. The GLT-based QMC assumes that F is invertible. For applications where F is a complicated function and cannot be inverted analytically, one can still apply GLT by resorting to some high precision numerical inversion techniques for inverting F . For example, we have employed the numerical inversion method of [10] in our numerical illustrations in Section 3.

Remark 2. Both LT and GLT require pre-computation of the orthogonal matrix \mathbf{A} . Initializing all columns of \mathbf{A} can be quite time consuming, particularly for large dimensional application. One way of reducing the computational burden is to exploit the iterative design of the optimization problem. Instead of optimizing all d columns of \mathbf{A} , one can use a sub-optimal \mathbf{A} by only optimizing its first d^* columns with the remaining columns randomly assigned (but subject to the orthogonality conditions). When $d^* \ll d$, this translates into a significant reduction in the pre-computation effort. The numerical examples to be presented later indicate that GLT is so effective at dimension reduction that the loss of efficiency induced by the sub-optimal \mathbf{A} is negligible and more than compensated by the saving in computational burden.

Remark 3. When \mathbf{X} is a vector of d iid normal variates, then the proposed GLT reduces to the original LT method. If we further assume that the orthogonal matrix \mathbf{A} is the identity matrix, then we recover the standard application of QMC.

3 Numerical Illustrations

In this section, we offer some numerical evidences on the effectiveness of the QMC-GLT relative to standard MC

and standard QMC. Subsection 3.1 describes the setup of our numerical examples. Subsection 3.2 evaluates the efficiency of the various simulation methods by comparing the simulated option prices. Subsection 3.3 examines the effectiveness of GLT on dimension reduction.

3.1 Model Setup

For our numerical illustration we use the plain-vanilla European call options and the Asian call options as test cases. The payoffs at the maturity of these options are given, respectively, by

$$h(S_{t_1}, \dots, S_{t_d}) = (S_{t_d} - K)^+ \quad (3)$$

and

$$h(S_{t_1}, \dots, S_{t_d}) = \left(\frac{1}{d} \sum_{i=1}^d S_{t_i} - K \right)^+, \quad (4)$$

where $(x)^+ = \max(x, 0)$, K is a pre-specified strike price of the option, S_t is the price of the underlying asset at time t , $t_i, i = 1, \dots, d$ denotes the discretized set of time points for which the prices are simulated and $t_d = T$. Note that the plain-vanilla option depends on the prices at maturity T while the Asian option is a classical example of a path-dependent option in that its payoff depends on the historical prices.

The fundamental theorem of asset pricing (see [5]) asserts that its time-0 no-arbitrage price is given by

$$E_Q[e^{-rT} h(S_{t_1}, \dots, S_{t_d})], \quad (5)$$

where r is the risk-free rate of return and the expectation is taken with respect to an equivalent martingale measure Q . Even if we were to impose the simplifying assumption that the log-returns of the underlying asset follows a Gaussian process (i.e. Black-Scholes model), simple closed-form solution to the Asian option (5) does not exist. Consequently this calls for numerical methods and this, in turn, has stimulated some ingenious approximation algorithms. For example, dimension reduction techniques that have been applied to pricing Asian option in the Black-Scholes model can be found in [1], [26] and [12].

Here we are interested in the effectiveness of GLT on non-Gaussian applications. We assume the dynamic of the log-returns of the underlying asset follows a particular family of Lévy process called Meixner process. A Meixner Lévy process $\{X_t, t \geq 0\}$ is a stochastic process which starts at zero, i.e. $X_0 = 0$, has independent and stationary increments, and X_t has the Meixner distribution with pdf given by

$$f^{\text{Meixner}}(x; a, b, d, m) = \frac{(2 \cos(\frac{b}{2}))^{2d}}{2a\pi\Gamma(2d)} \exp\left(\frac{b(x-m)}{a}\right) \left| \Gamma\left(d + \frac{i(x-m)}{a}\right) \right|^2,$$

where $a > 0, -\pi < b < \pi, d > 0, m \in \mathfrak{R}$, and $\Gamma(\cdot)$ is the Euler gamma function. Option pricing associated with this model can be found in [7]. See [22] for additional studies, particularly with emphasis on finance applications and fitting of the financial data. For example, by fitting Meixner Lévy process to the daily log-returns of the Nikkei-225 Index for the period January 1, 1997 to December 31, 1999, [22] obtains the following set of parameter values: $a = 0.02982825, b = 0.12716244, d = 0.57295483, m = -0.00112426$. These parameter values are used to simulate the asset prices as reported in the following two subsections.

3.2 Efficiency in terms of accuracy and computational effort

For option specifications, we set the initial asset price $S_0 = 100$, strike price $K = 100$, interest rate $r = 4\%$, maturity $T = 1$ year, and asset prices are sampled at quarterly, monthly, weekly and daily time intervals so that $d = 4, 12, 50$, and 250 , respectively.¹ Note that d also corresponds to the nominal dimension of the options. This allows us to assess the impact of the nominal dimensions on the various simulation techniques by merely increasing the frequency of the monitoring time points.

For each set of the option contract specification, we estimate its price using three simulation techniques: the standard MC, the standard QMC, and the QMC-GLT. In order to simulate the trajectories of the asset prices, we need a fast and efficient Meixner variate generator. Here we use the numerical inversion method [10] while [13] uses [9]. The newer inversion algorithm [10] is more efficient than [9] for two reasons. One is that we can avoid inverting the function numerically with the Newton method by adopting the Newton interpolation and hence it is more efficient especially when the computation of the pdf is time-consuming. The other reason is that we can use the same subintervals for the Gauss-Lobatto integration and the Newton interpolation, which makes the algorithm simpler and also increases its efficiency.

The method of GLT requires pre-computation of the orthogonal matrix \mathbf{A} , which in turn depends problem of interest. This implies that we need to solve for the optimal \mathbf{A} for each plain-vanilla option and Asian option. As pointed out in Remark 2 that a significant saving in computation time can be achieved by partially optimizing \mathbf{A} . In our numerical results, we only optimize the first four columns of \mathbf{A} , regardless of the nominal dimensions of the problems. Hence only the option examples with $d = 4$ are simulated optimally under the QMC-GLT. To further isolate the effect of the dimension reduction as induced by the GLT method, the same set of scrambled Sobol' low discrepancy sequence is used for both standard QMC and QMC-GLT.

¹Assume a year has 50 trading weeks or 250 trading days.

Table 1: Simulated plain-vanilla call options and Asian call option prices under the Meixner Lévy process with $d = 4, 12, 50, 250$. Efficiency ratios of QMC and QMC-GLT, relative to MC, are reported in parentheses.

| | MC | QMC | QMC-GLT |
|----------------------------|--------|-------------|--------------|
| Plain-vanilla call options | | | |
| $d = 4$ | 11.918 | 11.951(122) | 11.952(1147) |
| $d = 12$ | 11.898 | 11.953(9) | 11.952(731) |
| $d = 50$ | 11.975 | 11.986(1.5) | 11.954(515) |
| $d = 250$ | 12.027 | 11.959(0.4) | 11.950(5) |
| Asian call options | | | |
| $d = 4$ | 7.985 | 8.015(340) | 8.016(1826) |
| $d = 12$ | 7.112 | 7.133(11) | 7.133(680) |
| $d = 50$ | 6.802 | 6.810(3) | 6.795(106) |
| $d = 250$ | 6.743 | 6.711(0.4) | 6.709(6) |

Table 1 reports the simulated option prices. These values are based on 30 independent batches, with each batch consists of 4096 trajectories. To gauge the efficiency of QMC and QMC-GLT relative to MC, we calculate the efficiency ratio defined as:

$$\text{Efficiency Ratio} = \frac{\sigma_{MC}^2 t_{MC}}{\sigma_x^2 t_x},$$

where σ_{MC}^2 and t_{MC} denote, respectively, the estimated variance and the computation time for the MC method. Similarly, σ_x^2 and t_x are the corresponding estimates for method x , which is either QMC or QMC-GLT in our numerical comparisons. Consequently, the above ratio is a better measurement of efficiency since it takes into the account of both accuracy (as measured by the variance) and computational effort (as measured by the computation time). An efficiency ratio greater than one implies that method x is more efficient relative to MC and vice versa if the ratio is less than one. The efficiency ratios are reported in parentheses in Table 1.

The simulated results demonstrate that the efficiency of QMC deteriorates with dimensions, consistent with the findings in [25]. While the standard QMC offers a substantial improvement relative to MC for low dimensional applications, its effectiveness quickly diminishes with increasing dimensions. In fact when $d = 250$, QMC is less effective than MC for both plain-vanilla and Asian options.

The method of QMC-GLT also exhibits a decline in efficiency with increasing dimensions. The deterioration, however, is less pronounced which in part can be attributed to its effectiveness on dimension reduction (see next subsection). For dimension as low as $d = 4$, QMC-GLT attains a remarkable efficiency gain of more than a thousand for both option examples. When the dimension increases to as high as $d = 250$, the simulated results

still indicate an order of improvement of at least 5 times, even when we take into consideration the additional computational effort required to pre-determine the optimal orthogonal matrix \mathbf{A} .

3.3 Efficiency in terms of dimension reduction

In this subsection, we provide additional evidence on the effectiveness of the GLT-based QMC by examining its effectiveness on dimension reduction. This is motivated by the numerical studies in [17] (see also [4], [1], [26], [27]) which suggests that one way of increasing the efficiency of the underlying QMC is via dimension reduction. This could be attributed to the greater uniformity inherence in the lower dimensional structure of the low-discrepancy sequences (see [28]). By using the analysis of variance (ANOVA) decomposition, [4] formally introduces two notions of effective dimension known as the truncation dimension and the superposition dimension. The effective dimension of function f in the truncation sense is defined as the smallest integer d_T such that $\sum_{u \in \{1, 2, \dots, d_T\}} \sigma^2(f_u) \geq p \sigma^2(f)$ where $\sigma^2(f_u)$ represents the variance of f attributes to the set u , $\sigma^2(f)$ represents the total variance of the function f and p (typically close to 1, say 0.99) corresponds to some percentile. This idea is related to the sensitivity indices [24] or the dimension distribution [18]. When the truncation dimension of a function f is much smaller than its nominal dimension d , then f is said to have low effective dimension and its total variance is sufficiently captured by the first d_T components.

A simple way of gauging the effectiveness of dimension reduction is to compute the following ratio:

$$\text{CER}(d) = \frac{\sum_{u \in \{1, 2, \dots, d\}} \sigma^2(f_u)}{\sigma^2(f)}.$$

Note that the numerator sums up all the variances contributed by the components up to dimension d while the denominator represents the total variance. Hence the ratio is between 0 and 1 and can be interpreted as the cumulative explanatory ratio (CER). By definition, $\text{CER}(d_T) \geq p$.

Using the same option examples as in the last subsection, Table 2 produces $\text{CER}(d)$, for $d = 1, \dots, 5$ for both QMC and QMC-GLT. The ratio can be estimated numerically (based on MC with 100,000 sample size) using the procedure described in [26]. It is reassuring to note that the GLT is extremely effective at dimension reduction, even for large nominal dimension. For the option examples with $d = 250$ and the standard application of QMC, the first dimension accounts only 1% of the total variance while under the GLT, it captures at least 97% of the total variance. These results also justify using a sub-optimal orthogonal matrix \mathbf{A} to simulate the Meixner Lévy process. The loss of accuracy induced by the sub-optimal \mathbf{A}

Table 2: Cumulative explanatory ratio (in percentage) up to first five dimensions

| d | | dimension u | | | | |
|-----------------------|---------|---------------|-----|-----|-----|-----|
| | | 1 | 2 | 3 | 4 | 5 |
| Plain-vanilla options | | | | | | |
| 4 | QMC | 19 | 42 | 69 | 100 | - |
| | QMC-GLT | 100 | 100 | 100 | 100 | - |
| 12 | QMC | 6 | 12 | 19 | 26 | 34 |
| | QMC-GLT | 100 | 100 | 100 | 100 | 100 |
| 50 | QMC | 1 | 3 | 5 | 6 | 8 |
| | QMC-GLT | 100 | 100 | 100 | 100 | 100 |
| 250 | 0 | 0 | 0 | 1 | 1 | 1 |
| | QMC-GLT | 98 | 98 | 98 | 98 | 98 |
| Asian options | | | | | | |
| 4 | QMC | 44 | 78 | 95 | 100 | - |
| | QMC-GLT | 99 | 100 | 100 | 100 | - |
| 12 | QMC | 16 | 33 | 47 | 60 | 71 |
| | QMC-GLT | 100 | 100 | 100 | 100 | 100 |
| 50 | QMC | 4 | 8 | 12 | 17 | 21 |
| | QMC-GLT | 100 | 100 | 100 | 100 | 100 |
| 250 | QMC | 1 | 1 | 2 | 3 | 3 |
| | QMC-GLT | 97 | 97 | 98 | 98 | 98 |

tends to be negligible due to the overwhelming success of the dimension reduction.

4 Conclusion

In this paper, we consider the generalized LT method and the numerical inversion method recently proposed by [13] and [10], respectively. The numerical illustrations, involving simulating plain-vanilla options and Asian options in the context of the Meixner Lévy process, demonstrated the competitive advantage of the QMC-GLT relative to the the standard MC and QMC. This conclusion is consistent with that reported in [13] which focuses on the generalized hyperbolic Lévy process. The real advantage of the GLT lies in its generality. It can be used to simulate a wide range of stochastic process, in addition to the Gaussian process.

It should be pointed out that [16] proposes a Brownian-bridge based approach for simulating Lévy process. It will be of interest to compare the relative efficiency of the proposed GLT to that of [16]. We leave this for future research.

References

- [1] P. Acworth, M. Broadie, and M. Glasserman. A comparison of some Monte Carlo and quasi-Monte Carlo methods for option pricing. In P. Hellekalek, G. Larcher, H. Niederreiter, and P. Zinterhof, editors, *Monte and quasi-Monte Carlo Methods 1996*,

- volume 127 of *Lecture Notes in Statistics*, pages 1–18. Springer-Verlag, 1998.
- [2] A.N. Avramidis and P. L'Ecuyer. Efficient Monte Carlo and quasi-Monte Carlo option pricing under the variance gamma model. *Management Science*, 52(12):1930–1944, 2006.
- [3] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [4] R.E. Caflisch, W. Morokoff, and A.B. Owen. Valuation of mortgage-backed securities using brownian bridges to reduce effective dimension. , *Journal of Computational Finance*, 1(1):27–46, 1997.
- [5] F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300:463–520, 1994.
- [6] J-C Duan. The GARCH option pricing model. *Mathematical Finance*, 5:13–32, 1995.
- [7] B. Grigelionis. Process of Meixner type. *Lithuanian Mathematics Journal*, 39(1):33–41, 1999.
- [8] J. Hartinger and M. Predota. Simulation methods for valuing Asian option prices in a hyperbolic asset price model. *IMA Journal of Management Mathematics*, 14:65–81, 2003.
- [9] W. Hörmann and J. Leydold. Continuous random variate generation by fast numerical inversion. *ACM Transactions on Modeling and Computer Simulation*, 13(4):347–362, 2003.
- [10] W. Hörmann and J. Leydold. Random variate generation by numerical inversion when only the density is known, 2008. Department of Statistics and Mathematics Wirtschaftsuniversität Wien, Research Technical Report 78.
- [11] J. Hull and A. White. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 41:281–300, 1987.
- [12] J. Imai and K.S. Tan. A general dimension reduction technique for derivative pricing. *Journal of Computational Finance*, 10(2):129–155, 2006.
- [13] J. Imai and K.S. Tan. An accelerating quasi-Monte Carlo method for option pricing under the generalized hyperbolic Lévy process. *SIAM Journal on Scientific Computing*, forthcoming.
- [14] C. Joy, P.P. Boyle, and K.S. Tan. Quasi-Monte Carlo methods in numerical finance. *Management Science*, 42(6):926–938, 1996.
- [15] A. Kyprianou, W. Schoutens, and P. Wilmott, editors. *Exotic Option Pricing and Advanced Lévy Models*. Wiley, West Sussex, England, 2005.
- [16] G. Leobacher. Stratified sampling and quasi-Monte Carlo simulation of Lévy processes. *Monte Carlo Methods and Applications*, 12(3-4):231–238, 2006.
- [17] B. Moskowitz and R. Caflisch. Smoothness and dimension reduction in quasi-Monte Carlo methods. *Mathematical and Computer Modeling*, 23(8-9):37–54, 1996.
- [18] A.B. Owen. The dimension distribution and quadrature test functions. *Statistica Sinica*, 13:1–17, 2003.
- [19] C. Ribeiro and N. Webber. Valuing path-dependent options in the variance-gamma model by Monte Carlo with a gamma bridge. *Journal of Computational Finance*, 7(2):81–100, 2003.
- [20] P. Sabino. Efficient quasi-Monte simulations for pricing high-dimensional path-dependent options. *Decisions in Economics and Finance*, forthcoming.
- [21] P. Sabino. Implementing quasi-Monte Carlo simulations with linear transformations. *Computational Management Science*, forthcoming.
- [22] W. Schoutens. Meixner processes: theory and applications in finance. *EURANDOM Report 2002-004*, EURANDOM, Eindhoven, 2002.
- [23] W. Schoutens. *Lévy processes in Finance: Pricing Financial Derivatives*. Wiley, West Sussex, England, 2003.
- [24] I.M. Sobol. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55:271–280, 2001.
- [25] K.S. Tan and P.P. Boyle. Applications of randomized low discrepancy sequences to the valuation of complex securities. *Journal of Economic Dynamics and Control*, 24:1747–1782, 2000.
- [26] X. Wang and K-T Fang. The effective dimension and quasi-Monte Carlo. *Journal of Complexity*, 19:101–124, 2003.
- [27] X. Wang and I.H. Sloan. Why are high-dimensional finance problems often of low effective dimension? *SIAM Journal on Scientific Computing*, 27(1):159–183, 2005.
- [28] X. Wang and I.H. Sloan. Low discrepancy sequences in high dimensions: How well are their projections distributed? *Journal of Computational and Applied Mathematics*, 213(2):366–386, 2008.