

Dimension Reduction Approach To Simulating Exotic Options In A Meixner Levy Market

Junichi Imai* and Ken Seng Tan†

Abstract—In the past decade, quasi-Monte Carlo (QMC) method has become an important numerical tool in computational finance. This is driven, in part, by the sophistication of the models and, in part, by the complexity of the derivative securities. In this paper, we consider an enhanced QMC method recently proposed by Imai and Tan (2009). This method is known as the generalized linear transformation (GLT) and it increases the efficiency of QMC via dimension reduction. GLT can be used to simulate general stochastic processes and hence has a much wider range of applications. By assuming that the dynamics of the underlying asset price follows an exponential Meixner Lévy process and by resorting to some exotic options including average options and lookback options, we demonstrate the effectiveness and robustness of GLT and it substantially outperforms the standard applications of QMC and Monte Carlo methods.

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

1 Introduction

The Monte Carlo (MC) method is a powerful and flexible tool for providing numerical solutions to a large class of complex problems. In particular, since its introduction to computational finance by Boyle [4] in 1977, MC method has been gaining popularity and is becoming an indispensable tool in a variety of settings in computational finance. There are a few reasons for its usefulness in computational finance. Some of these reasons include the following:

1. In modern financial economics, security prices are modeled as stochastic processes to reflect future uncertainty. The current price of a security can be represented as the expected value of the future payouts

*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.

on the security. This follows from the assumption of no-arbitrage (see (2) in Section 2).

2. Increased complexity of the derivative securities
3. Increased sophistication on modeling the dynamics of the underlying stochastic processes.
4. Early finance applications of MC are mainly concerned with calculations related to the pricing of complex financial instruments and the computation of related hedging parameters. More recently, MC methods are used extensively in risk management (such as calculation of credit risk and market risk and value at risk computations), solvency analysis, and etc.

It follows from Point 1 that MC method lends itself naturally to this application since it involves estimating an expectation. A simple MC procedure for estimating the prices of derivative securities involves the following steps: First simulate the stochastic process that drives the underlying asset. Second, corresponding to the simulated asset path, record the discounted payoff. Third, repeat the simulation procedure independently N times to generate N independent estimates of the discounted payoffs. The crude MC estimate of the derivative security price is then given by the sample average of these discounted payoffs. The strong law of large number guarantees that the sample average converges to the true value as the sample size N tends to infinity. In addition, the central limit theorem assures us that it converges at a rate of $O(N^{-1/2})$, which is independent of dimension.

The complicated and exotic features of the derivative securities (Point 2) imply that only in rare cases where their prices can be expressed analytically, even under the simplest Black-Scholes [3] type framework. Examples of such exotic derivative securities include path-dependent options such as Asian option, lookback option, barrier option, and etc. There are other derivative securities which not only depend on a single asset but also on several underlying assets (such as basket option). For these exotic derivative securities, their prices can be formulated as multi-dimensional integrals. In many cases, the number of dimensions can be very large; for example, under mortgage-backed securities the number of dimensions can be as high as 360. Because of the inheritance

high-dimensional applications, other competitive numerical methods including the numerical solutions to partial difference equation, binomial lattice method and quadrature methods become computationally infeasible due to the curse of dimensionality. Consequently, MC becomes the only viable numerical tool.

Point 3 provides further impetus on the usefulness of MC methods. The Black-Scholes model is often criticized for assuming that the log-returns of the underlying asset price is normally distributed. Numerous empirical studies have provided ample of evidences that the dynamics of the underlying asset typically exhibits skewness and kurtosis. Motivated by the empirical evidences, a number of more elaborate models including GARCH models (e.g. see [9]) and models with stochastic volatility (e.g. see [14]) 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 [30] and [19]). These models, although they are better at capturing many of the stylized features of the assets, tend to be more sophisticated and more complicated. This implies that the chance of obtaining tractable pricing formulas are even more slim compared to the Black-Scholes case. Consequently this again advocates the use of MC method.

The key advantages of MC method lie on its flexibility and that its convergence rate is independent of dimension. This is particularly important since many finance applications typically have dimension of several hundreds. Despite its widespread use, MC is often criticized for its slow rate of convergence, particularly for large scale problems. Different methods for enhancing the underlying MC method have been proposed. These techniques are known as variance reduction techniques. In the past decade or so, the so-called quasi-Monte Carlo (QMC) method has been proposed as an alternate competitive numerical tool to the field of computational finance. This method relies on the specially constructed sequence which has the property that it is more evenly dispersed throughout the unit cube and is known as the (randomized) low discrepancy sequence. The monograph by [23] provides an excellent discussion of these sequences. Early applications of low discrepancy sequences to finance problems are advocated in [6], [18], [24] and [25]. See also the comprehensive and excellent survey paper of [5].

The surge of interest in QMC stems from its promised rate of convergence of $O(N^{-1} \log^d N)$, in dimension d . This rate is *asymptotically* more efficient than the corresponding MC rate of $O(N^{-1/2})$. However, the factor $\log^d N$ cannot be ignored for practical sample size N and moderate dimension d . Hence the superior convergence rate of QMC needs not be attained in practical applications due to its explicit dependent on dimension. Many numerical studies seem to suggest that the success of QMC is intricately related to the notion of *effective*

dimension. These studies show that when QMC is combined with dimension reduction techniques, the greater efficiency of QMC can be expected. Methods that enhance QMC by exploiting this feature include the Brownian bridge construction ([22] and [7]), the principal component construction [1], and the linear transformation (LT) method [15]. More recently, Imai and Tan (see [16] and [17]) consider an extension of the LT method known as the generalized linear transformation (GLT) method. The generalization is motivated by the recent interest in adapting QMC to other more exotic models, notably the Lévy models (see [12], [26], [2], [20]). The original LT method has the limitation that it only provides an efficient algorithm for simulating derivative prices when the underlying follows a Gaussian process. By inducing additional transformations, the GLT proposed in [16] and [17] provides a power alternate way of simulating arbitrary stochastic processes. Consequently, GLT has a much wider range of applications. When combined with QMC, its effectiveness on simulating Lévy processes is illuminated in the numerical examples conducted in [16] and [17]. The objective of this paper is to provide additional insights on the effectiveness of GLT. We consider more extensive test cases including plain-vanilla options and other exotic options such as average options and look-back options. The dimension associated with these applications ranging from 4 to 250. The robustness of the GLT is also addressed. Our numerical examples assume that the dynamics of the asset prices follows the exponential Meixner process.

The rest of the paper is organized as follows. Section 2 provides an overview of a special family of Lévy process known as the Meixner process. Section 3 describes the dimension reduction GLT method. Extensive test results on the relative effectiveness of GLT are presented in Section 4. Section 5 concludes the paper.

2 The Meixner Process

In this section, we provide a brief overview of the Meixner process as well as its relation to modeling risky assets and derivative pricing. Detailed description of the Meixner process can be found in [11] and [29] where we have largely drawn the material from.

The density of the Meixner distribution, denoted by $\text{Meixner}(a, b, d, m)$, is 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 $a > 0$, $-\pi < b < \pi$, $d > 0$, $m \in \mathbb{R}$, and $\Gamma(\cdot)$ is the Euler gamma function. Several important characteristics associated with this distribution are:

1. Moments of all order of this distribution exist. In particular, its mean, variance, and kurtosis are given

by $m + ad \tan(b/2)$, $\frac{a^2 d}{2} \cos^{-2}(b/2)$, and $3 + [3 - 2 \cos^2(b/2)]/d$, respectively. Clearly, the kurtosis of the Meixner distribution is always larger than the corresponding kurtosis from the normal distribution.

- The characteristic function of the Meixner(a, b, d, m) distribution, denoted by $\phi(u)$, has the form

$$\phi(u) = \left(\frac{\cos(b/2)}{\cosh \frac{au-ib}{2}} \right)^{2d} \exp(imu).$$

It is easy to see that the Meixner(a, b, d, m) is infinitely divisible.

- The Meixner(a, b, d, m) distribution has semi-heavy tails. This means that the tails of the density function behave as:

$$f^{\text{Meixner}}(x; a, b, d, m) \sim$$

$$\begin{cases} C_- |x|^{\rho_-} \exp(-\sigma_- |x|) & \text{as } x \rightarrow -\infty \\ C_+ |x|^{\rho_+} \exp(-\sigma_+ |x|) & \text{as } x \rightarrow +\infty, \end{cases}$$

where

$$\rho_- = \rho_+ = 2d - 1, \quad \sigma_- = \frac{\pi - b}{a}, \quad \text{and} \quad \sigma_+ = \frac{\pi + b}{a}.$$

The infinite divisibility property of the Meixner(a, b, d, m) distribution implies that we can associate with it a Lévy process which we denote as the Meixner process. More formally, a Meixner 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 where the distribution of X_t is given by the Meixner(a, b, dt, mt) distribution. Unlike the general Lévy process, the Meixner process $\{X_t, t \geq 0\}$ has no Brownian part and a pure jump part governed by the Lévy measure

$$v(dx) = d \frac{\exp(bx/a)}{x \sinh(\pi x/a)} dx.$$

Recall that in the celebrated Black-Scholes model, the price process of the underlying is given by the geometric Brownian motion:

$$S_t = S_0 \exp \left(\left(\mu - \frac{\sigma^2}{2} \right) t + \sigma B_t \right),$$

where S_t is the price of the asset at time t , μ and σ are the mean and volatility of the underlying asset, and $\{B_t, t \geq 0\}$ is the standard Brownian motion; i.e. B_t is normally distributed with mean 0 and variance t . As pointed out in the introduction that there is ample of empirical evidence criticizing the deficiency of the normality assumption of the log-returns. A natural remedy is to replace the Brownian motion in the Black-Scholes model by a more sophisticated Lévy process. For example, [11]

and [29] propose the following dynamics of the asset price based on the Meixner process:

$$S_t = S_0 \exp(X_t) \tag{1}$$

where $X_t \sim \text{Meixner}(a, b, d, m)$. [29] successfully fits the above model to the daily log-returns of the Nikkei-225 Index.

Let $h((S_{t_1}, \dots, S_{t_d}))$ denote the payoff at maturity $T = t_d$ of a European-style derivative security. Then the fundamental theorem of asset pricing asserts that its time-0 no-arbitrage price is given by (see [8])

$$E_Q[e^{-rT} h(S_{t_1}, \dots, S_{t_d})], \tag{2}$$

where r is the risk-free rate of return and the expectation is taken with respect to an equivalent martingale measure Q . Assuming more complex dynamics of the asset prices induces market incompleteness. This implies there exists multiple equivalent martingale measures which lead to nonuniqueness of the no-arbitrage price of the derivative. In our simulation studies, we identify the martingale measure via the Esscher transform proposed in [10]. Using this method, the equivalent martingale measure Q follows a Meixner($a, a\theta + b, d, m$) distribution where θ is given by

$$\theta = -\frac{1}{a} \left(b + 2 \arctan \left(\frac{-\cos(\frac{a}{2}) + e^{(m-r)/(2d)}}{\sin(a/2)} \right) \right).$$

Additional details on option pricing with respect to this model can be found in [11].

3 Generalized Linear Transformation (GLT) Method

Some researches (see for example [1], [7], [16] and [17], [22]) have suggested that the success of QMC depends critically on the underlying *effective dimension*, as opposed to the *nominal dimension*, of the problem of interest. Using “analysis of variance” (ANOVA) decomposition of a function, [7] defines two notions of effective dimension of an integrand as follows: Consider an integrand with d nominal dimensions (i.e. depends on d variables). The effective dimension of f , in the superposition sense, is the smallest integer d_S such that $\sum_{|u| \leq d_S} \sigma^2(f_u) \geq p\sigma^2(f)$, where f_u is a function depends on the components in the set u , $\sigma^2(\cdot)$ denotes the variance of the given function, and p is an arbitrary confidence level such as 99%. The effective dimension of f , in the truncation sense, is the smallest integer d_T such that $\sum_{u \in \{1, 2, \dots, d_T\}} \sigma^2(f_u) \geq p\sigma^2(f)$. Essentially, the truncation dimension indicates the number of important variables which predominantly captures the given function f . The superposition dimension, on the other hand, measures to what extent the low-order ANOVA terms dominate the function.

In addition to introducing the concept of effective dimension, [7] argues that the efficiency of QMC is intricately tied to the effective dimension. One of the potential explanations for the success of QMC on problems of low effective dimension can be attributed to the well-known phenomenon that the uniformity of some low discrepancy sequences deteriorates with dimensions. QMC can be expected to be more effective when the greater uniformity portion of the low discrepancy sequences is applied to the dominant dimensions of the function. This is exactly the strategy recommended by [7] for enhancing QMC, as can be seen from the following excerpt (page 45) of that article:

- *First analyze the problem, mathematically or numerically, to determine the most important input dimensions.*
- *Where possible, reformulate the problem to concentrate the variation in fewer dimensions.*
- *When a small number of dominant dimensions can be identified or induced, apply quasi-random or randomized quasi-random sequences to those dimensions.*

Methods such as the linear transformation (LT, [15]) and the generalized linear transformation (GLT, [16] and [?]) increase the efficiency of QMC by adopting these strategies. There are other methods, including the Brownian bridge construction ([22]) and the principal component construction ([1]), which also enhance QMC via dimension reduction. In the remaining of the section, we focus on the LT method, in particular the GLT method. The latter method will form the basis of our numerical studies in the next section.

Suppose we are interested in solving

$$E[g(\mathbf{Z})], \tag{3}$$

where $g(\mathbf{Z})$ is a function which depends on d -dimensional standardized normal random vector \mathbf{Z} . The above expectation is of particular interest to us as many of the problems associated with derivative pricing can be formulated as (3). In these cases, the expectation is taken with respect to an equivalent martingale measure and $g(\mathbf{Z})$ corresponds to the discounted payoff of the derivative security (see (2)). We are further assuming that a simple analytical expression for (3) does not exist and hence we need to resort to numerical methods such as MC or QMC due to the high dimension d .

Before describing LT and GLT methods, let us first recall how QMC is used to estimate (3). The standard QMC estimator of (3) is given by

$$\frac{1}{N} \sum_{i=1}^N g(\boldsymbol{\varepsilon}_i), \tag{4}$$

where the normal vector $\boldsymbol{\varepsilon}_i \equiv (\varepsilon_{i1}, \varepsilon_{i2}, \dots, \varepsilon_{id})$ is generated via inverse transforms from a d -dimensional (randomized) low discrepancy sequence. More precisely, we have $\varepsilon_{ij} = \Phi^{-1}(x_{ij})$ where x_{ij} denotes the i -th point and j -th dimension of a (randomized) low discrepancy sequence and $\Phi(\cdot)$ is the normal cumulative density function. To contrast with the MC method, the normal vectors $\{\boldsymbol{\varepsilon}_i, i = 1, \dots, N\}$ would have been generated independently and randomly from a pseudo-random sequence.

Instead of using (4) to estimate (3), [15] proposes the following QMC-based estimator:

$$\frac{1}{N} \sum_{i=1}^N g(\mathbf{A}\boldsymbol{\varepsilon}_i), \tag{5}$$

where \mathbf{A} is an orthogonal matrix satisfying $\mathbf{A}'\mathbf{A} = \mathbf{I}$ and \mathbf{I} is the identity matrix. They denote this method as LT, the linear transformation method. They also recommend the following algorithm for determining the optimal columns of \mathbf{A} iteratively; i.e. $\mathbf{A}_{\cdot k}^*$, for $k = 1, \dots, N$:

$$\max_{\mathbf{A}_{\cdot k} \in \mathbb{R}^d} \left(\left. \frac{\partial g(\mathbf{A}\boldsymbol{\varepsilon})}{\partial \varepsilon_k} \right|_{\boldsymbol{\varepsilon} = \hat{\boldsymbol{\varepsilon}}_k} \right)^2$$

subject to $\|\mathbf{A}_{\cdot k}\| = 1$ and $\langle \mathbf{A}_{\cdot j}^*, \mathbf{A}_{\cdot k} \rangle = 0, \quad j = 1, \dots, k-1.$

(6)

In the above algorithm, $\mathbf{A}_{\cdot k}$ denotes the k -th column of \mathbf{A} , $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the inner product between vectors \mathbf{a} and \mathbf{b} , and $\hat{\boldsymbol{\varepsilon}}_k = (v_1, \dots, v_{k-1}, 0, \dots, 0)^\top$ denotes the d -dimensional vector with arbitrary chosen random variables v_1, \dots, v_{k-1} .

We now summarize some properties pertaining to the method of LT:

- Since \mathbf{A} is an orthogonal matrix, this implies that $E[g(\mathbf{Z})] = E[g(\mathbf{A}\mathbf{Z})]$ and hence (5) is a consistent estimator of (3). Furthermore, under the special case $\mathbf{A} = \mathbf{I}$, the LT estimator (5) reduces to the standard QMC estimator (4).
- Algorithm (6) aims at reducing the dimension reduction of a given function of interest and is motivated by its intricate connection to variance decomposition. By maximizing the variance contribution of each column of \mathbf{A} iteratively, this ensures that the dominant effect of the function is concentrated on the earlier dimensions. This in turn has the impact of dimension reduction.
- We emphasize that algorithm (6) is carried out iteratively for $k = 1, 2, \dots, d$. This implies that in the k -th iteration, $\mathbf{A}_{\cdot j}^*, j = 1, \dots, k-1$ are already optimally determined in the earlier iteration steps.

- Methods such as the Brownian bridge construction [22] and the principal component construction [1] reduce the effective dimension by only focusing on the discrete Brownian paths. LT, on the other hand, reduces the effective dimension by directly exploiting the given function.
- A carefully chosen \mathbf{A} enhances QMC. The numerical examples in [15], [28] and [27] even suggest that the LT-based QMC, with \mathbf{A} obtained from (6), can be more effective than the Brownian bridge construction and the principal component construction.

While LT enjoys many of the above advantages, one severe limitation is that it is restricted to a class of function which depends on a vector of normal random variables. This implies that the LT cannot be used directly to other non-Gaussian process including Lévy process. It is therefore of significant interest to providing an effective QMC-based algorithm for the more general problem of estimating $E[g(\mathbf{X})]$, where $\mathbf{X} = (X_1, \dots, X_d)^\top$ is a vector of d iid random variables with arbitrary probability density function (pdf) $f(x)$ and cumulative distribution function (cdf) $F(x)$. Motivated by this, [16] 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 $Z = \Phi^{-1}(Y)$. Then $E[g(\mathbf{X})]$ can be expressed as follows:

$$\begin{aligned} & \int \cdots \int_{-\infty}^{\infty} g(F^{-1}(\Phi(z_1)), \dots, F^{-1}(\Phi(z_d))) \times \\ & \quad \phi(z_1) \cdots \phi(z_d) dz_1 \cdots dz_d \\ & = E[g(F^{-1}(\Phi(\mathbf{Z}_1)), \dots, F^{-1}(\Phi(\mathbf{Z}_d)))] \end{aligned} \quad (7)$$

where ϕ is the pdf of the standard normal, and $\mathbf{Z} = (Z_1, \dots, Z_d)^\top$ is a vector of independent standard normal random variable. The significance of (7) 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} . They refer this approach as the generalized LT method or GLT. Figure 1 describes the implementation of GLT.

We now make the following remarks with regards to the GLT-based QMC:

Initialization step: Initialize the orthogonal matrix \mathbf{A} .

For $i = 1, 2, \dots, N$

step 1: Draw $(x_1, \dots, x_d) \in [0, 1]^d$ from a d -dimensional low discrepancy sequence.

step 2: Generate $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_d)^\top$, where $\varepsilon_j = \Phi^{-1}(x_j), j = 1, \dots, d$

step 3: Set $y_j = \Phi(\mathbf{A}_j \boldsymbol{\varepsilon}), j = 1, \dots, d$

step 4: Set $g_i = g(F^{-1}(y_1), \dots, F^{-1}(y_d))$.

$$\text{GLT estimate of } E[g(\mathbf{X})] = \frac{1}{N} \sum_{i=1}^N g_i.$$

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

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 [13] in our numerical illustrations in Section 4.

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.

4 Numerical Illustrations

In this section, we provide additional numerical evidences on the effectiveness of GLT relative to standard MC and standard QMC. The first half of Subsection 4.1 describes the setup of the numerical experiments. The second half of the subsection presents some numerical results on the relative efficiencies of various simulation methods. Subsection 4.2 addresses the robustness of GLT. Subsection 4.3 examines the effectiveness of GLT on dimension

reduction.

4.1 The effectiveness of dimension reduction on numerical accuracy

In this subsection, we provide numerical evidences on the relative efficiency of MC, QMC, and the dimension reduction GLT method. We are primarily concerned with the performance of these methods on non-Gaussian applications. Hence, we assume that the dynamic of the log-returns of the underlying asset follows the Meixner process as described in Section 2. We use the following set of parameter values: $a = 0.02982825$, $b = 0.12716244$, $d = 0.57295483$, $m = -0.00112426$. for simulating the Meixner process in our simulation studies. These values are obtained by fitting Meixner process to the daily log-returns of the Nikkei-225 Index for the period January 1,1997 to December 31, 1999. More details can be found in [29].

In terms of our test cases, we resort to the following representative European-style options:

- Plain-vanilla European call options with payoff at maturity $T = t_d$ given by

$$h(S_{t_1}, \dots, S_{t_d}) = (S_{t_d} - K)^+ \quad (8)$$

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 $t_d = T$.

- European average call options with payoff at maturity T given by

$$h(S_{t_1}, \dots, S_{t_d}) = \left(\sum_{i=1}^d w_i S_{t_i} - K \right)^+, \quad (9)$$

where w_i denotes the weight assigned to asset price at time t_i . In our test cases, we consider three specifications of w_i ; namely

- Equally weighted average case with $w_i = \frac{1}{d}$. This special case is commonly known as the Asian option.
- Decreasing weighted average case with $w_i = c(d - i + 1)^2$ where c is a normalizing constant such that $\sum_{i=1}^d w_i = 1$.
- Increasing weighted average case with $w_i = ci^2$.

Note that the average option is an example of a path-dependent option as its payoff depends explicitly on the past asset prices, in addition to the asset price at maturity.

In our option specifications, we set the initial asset price $S_0 = 100$, strike price $K = 100$, interest rate $r = 4\%$, maturity $T = 1$ year. Furthermore, the 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 corresponds to the nominal dimension of the option since each trajectory of the asset price requires d variables. This provides us a convenient way of assessing the effect of nominal dimension on various simulation approaches by simply increasing the frequency of the monitoring time points.

For each option contract, we estimate its price using three simulation techniques: the standard MC, the standard QMC, and the GLT. Recall that the method of GLT requires pre-computation of the optimal orthogonal matrix \mathbf{A} , which in turn depends on the target function $g(\mathbf{X})$. In our option examples, the target function g corresponds to the option payoff so that for the plain-vanilla option, we use the function

$$g_T(X_1, \dots, X_d) = S_{t_d} = S_0 e^{\sum_{i=1}^d X_i}, \quad (10)$$

while for the average options, we have

$$g_A(X_1, \dots, X_d) = \sum_{i=1}^d w_i S_{t_i} = S_0 \sum_{i=1}^d w_i e^{\sum_{j=1}^i X_j}. \quad (11)$$

Here $X_j, j = 1, \dots, d$ are independent variates from the Meixner distribution. Optimal orthogonal matrix \mathbf{A} corresponding to each of the above functions is first determined using the algorithm described in Section 3. The resulting optimal \mathbf{A} is then used to simulate trajectories of the underlying process for estimating option values. We denote the application of QMC-based GLT method that is optimal for g_T and g_A as, respectively, GLT- g_T and GLT- g_A . Note that g_T represents the terminal asset price at time T , while g_A captures the average price of the underlying asset over d monitoring time points. We emphasize that the optimality of GLT- g_A depends on the prescribed weights $\{w_j, j = 1, \dots, d\}$.

As formally established in [15], LT is most effective when the function of interest is linear. Regardless of the nominal dimension, in this case the function becomes one-dimension after the application of LT. This property is exploited in our search of optimal orthogonal matrix \mathbf{A} for GLT- g_T . Because $\log g_T$ is also linear in $X_j, j = 1, \dots, d$, this implies a plausible optimal orthogonal matrix \mathbf{A} can be obtained by only optimizing its first column. On the other hand, we optimize no more than ten columns of \mathbf{A} for GLT- g_A , regardless of the nominal dimensions of the problems. The use of sub-optimal orthogonal matrix is driven, in part, by the saving in computational time (see Remark 2 of the last section) and, in part, by the remarkable dimension reduction with the GLT (see Subsection 4.3).

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

For standard MC, we use the 64-bit Mersenne Twister pseudo-random generator (see [21]). For the methods involving QMC, we use the same set of randomized Sobol' low discrepancy sequence. To simulate the trajectories of the asset prices, we need a fast and efficient Meixner variate generator. Here we use the numerical inversion method of [13] for inverting both MC and QMC points. In addition to producing results that are based on strictly random sequences and strictly (randomized) low discrepancy sequences, we also generate results using hybrid sequences. The hybrid sequence is a sequence that is a concatenation of a (randomized) low discrepancy sequence and a random sequence. More precisely, suppose we are interested in a d -dimensional sequence for large d . We can avoid using such a high-dimensional (randomized) low discrepancy sequence by using, say, k -dimensional (randomized) low discrepancy sequence where $k \ll d$. The remaining $d - k$ dimensions are then padded with random sequences to produce the required d -dimensional sequence. We refer the concatenation approach of generating high-dimensional sequences as padding and we use the notation P- k to emphasize that it is a hybrid sequence whereby the first k dimensions are (randomized) low discrepancy sequences. In addition to providing an easy way of avoiding high-dimensional (randomized) low discrepancy sequence, the hybrid sequence can also be useful when the underlying problem of interest has low effective dimension.

The simulated option results based on MC and standard application of QMC are summarized in Table 1 for $d = 4, 12, 50$, and 250. The reported values are based on 30 independent batches, with each batch consists of 4096 trajectories. The MC estimates, with the standard errors in parentheses, are reported in second column of the table. To gauge the efficiency relative to MC, we tabulate the efficiency ratio which is the ratio of the standard error of the MC estimate to the standard error of the QMC. Hence a ratio greater than 1 indicates an efficiency gain relative to MC. Columns 3 to 6, with labels P-1, P-2, P-5, and P-10, respectively, give the corresponding results from the hybrid sequences. The last column, labeled as P- d , is the result from using d -dimensional randomized Sobol' and hence strictly is QMC. We also control our experiments by ensuring that the same set of sequences, whenever possible, are used consistently. For example, the hybrid sequence from P-1 corresponds to the same set of random sequence used in MC except by replacing the first dimension with the randomized Sobol' sequence. Hence any difference in the results is attributed to the enhanced uniformity of the randomized Sobol' sequence.

The simulated results in Table 1 indicate that there is an advantage of using standard QMC relative to standard MC. Its competitive advantage, however, diminishes quickly with dimensions. Consider, for instance, the Asian option example. The standard QMC attains

Table 1: Efficiency of standard applications of QMC (using padded and randomized Sobol' sequences) relative to MC. Second column gives the MC estimate (with its standard error in parentheses) of the respective derivative security. Last five columns display the efficiency ratios of QMC (using padded and randomized Sobol' sequences) relative to MC.

d	MC	P-1	P-2	P-5	P-10	P- d
Plain-vanilla call options						
4	11.918(0.060)	1.5	1.6	-	-	20.2
12	11.898(0.048)	1.0	1.0	1.0	1.8	6.4
50	11.975(0.051)	1.1	1.1	1.1	1.3	2.4
250	12.027(0.050)	1.0	1.0	1.0	1.0	1.6
Asian call options						
4	7.985(0.045)	1.9	3.2	-	-	35.9
12	7.112(0.030)	1.2	1.2	1.6	7.7	10.5
50	6.802(0.031)	1.1	1.1	1.3	1.5	4.1
250	6.743(0.025)	1.0	1.0	1.0	1.0	2.1
Average call options (decreasing weights)						
4	6.287(0.037)	2.9	8.2	-	-	87.1
12	4.900(0.022)	1.4	1.7	4.6	18.0	16.2
50	4.369(0.019)	1.3	1.3	1.4	2.1	6.4
250	4.228(0.016)	1.0	1.1	1.0	1.1	2.4
Average call options (increasing weights)						
4	10.100(0.053)	1.6	2.1	-	-	23.9
12	9.633(0.041)	1.1	1.1	1.2	4.3	7.4
50	9.501(0.043)	1.1	1.1	1.2	1.3	2.7
250	9.505(0.035)	1.0	1.0	1.0	1.0	1.7

Table 2: Efficiency of GLT (using padded and randomized Sobol' sequences) relative to MC. Second column gives the MC estimate (with its standard error in parentheses) of the respective derivative security. Last five columns display the efficiency ratios of GLT (using padded and randomized Sobol' sequences) relative to MC.

d	MC	P-1	P-2	P-5	P-10	P- d
Plain-vanilla call options						
4	11.918(0.060)	64.8	95.9	-	-	87.1
12	11.898(0.048)	40.4	53.7	46.7	42.1	67.4
50	11.975(0.051)	24.1	24.6	21.0	26.5	51.5
250	12.027(0.050)	4.6	6.6	7.0	5.6	17.2
Asian call options (equally weighted)						
4	7.985(0.045)	12.4	65.7	-	-	117.3
12	7.112(0.030)	21.7	28.8	67.9	46.9	66.7
50	6.802(0.031)	22.9	20.6	22.3	24.7	25.4
250	6.743(0.025)	6.0	4.8	4.3	6.2	5.2
Average call options (decreasing weights)						
4	6.287(0.037)	9.8	107.6	-	-	136.1
12	4.900(0.022)	19.5	28.8	82.7	70.0	77.5
50	4.369(0.019)	24.4	20.6	24.5	22.2	42.9
250	4.228(0.016)	6.5	6.8	6.9	5.8	7.3
Average call options (increasing weights)						
4	10.100(0.053)	15.5	31.9	-	-	104.2
12	9.633(0.041)	30.5	24.9	63.0	46.1	51.2
50	9.501(0.043)	20.1	18.8	20.9	22.6	24.0
250	9.505(0.035)	5.6	5.4	7.6	4.9	6.9

an impressive improvement of almost 36 times relative to MC for $d = 4$. As we increase the dimension from 4 to 250, the underlying QMC loses its attractiveness and is only marginally better than the corresponding MC (with 2.1 efficiency ratio). The effect of using the padded hybrid sequence seems negligible.

Table 2 is similar to Table 1 except that the trajectories of asset prices are simulated using GLT. More precisely, GLT- g_T is used to simulate the plain-vanilla options while GLT- g_A is used to simulate the average options. The simulated results clearly highlight the effectiveness of the GLT dimension reduction method. Using the same Asian option example, the efficiency gain from GLT is slightly more than 117 times in the quarterly sampling case. As we increase the dimension by increasing the sampling frequency, the GLT method also shows sign of deterioration as in the standard QMC. The GLT method, however, still able to maintain an improvement of 5.2 times in the high-dimensional example of $d = 250$.

The effectiveness of dimension reduction is further illuminated by comparing MC to P-1. Recall that the only difference between the sequences used in these two methods lie on the first dimension. The first dimension of MC is generated from the usual random sequence while the first dimension of P-1 is from the randomized Sobol' sequence. By simply changing the technique from MC to GLT and by using a more uniformly distributed sequence in the first dimension, we observe a significant improvement in P-1 relative to MC. We can further isolate the difference induced by the sequences by comparing the results from hybrid sequences in Table 1 to the corresponding results in Table 2. Any difference between these results is therefore solely attributed to the difference on how the trajectories of the asset price are generated. The much higher efficiency ratios of Table 2 again support the outstanding performance of GLT.

4.2 The robustness of GLT

An optimal use of GLT entails us to exploit explicitly the given function of interest by optimally determining its orthogonal matrix \mathbf{A} . This feature can also be perceived as a drawback of GLT in that it becomes problem dependent. For example, the orthogonal matrix derived from GLT- g_T that is optimal for simulating the plain-vanilla option may not be optimal for simulating other options. A logical question to ask is that to what extent is the loss of precision, if any, induced by the sub-optimal application of GLT. We attempt to address this issue by conducting the following numerical experiments. The method associated with GLT- g_T is not only used to estimate the plain-vanilla options, it is also used to simulate the three variants of average options described in the last subsection. These results are then compared to the optimal GLT that is based on GLT- g_A to assess the impact of sub-optimal applications. Similarly, GLT- g_A that

is optimal for simulating Asian options is in turn used to simulate the plain-vanilla options. The simulated results are depicted in Table 3. Some remarks with respect to these results are:

- In most cases, there is incentive of using GLT optimally. For instance, the efficiency gains for the plain-vanilla option is higher with GLT- g_T while for the Asian option, GLT- g_A is more effective than the corresponding GLT- g_T . An exception is the average option with increasing weights. For this particular type of option, GLT- g_T turns out to be more effective, albeit marginally, than the GLT- g_A .
- Sub-optimal application of GLT can be less efficient than QMC. This is not surprising since QMC is a special case of GLT with the orthogonal matrix coincides with the identity matrix. In other words, other sub-optimal choice of \mathbf{A} needs not always outperform the corresponding GLT with identity matrix. This is demonstrated in the average options with decreasing weights. Intuitively, this anomaly seems reasonable since GLT- g_T is putting all the emphasize on the asset price at maturity. Yet this is the least importance due to the geometrically decreasing weights. It is reassuring to note that if GLT were used optimally, a much higher precision, relative to both MC and QMC, can be recovered.
- Both GLT- g_T and GLT- g_A are also used to simulate another option contract known as the floating strike lookback call option. Its payoff at maturity T is given by

$$(S_{t_d} - S_{\min})^+, \quad (12)$$

where S_{\min} is the lowest asset price observed over d monitoring time points t_1, \dots, t_d , i.e., $S_{\min} = \min\{S(t_1), \dots, S(t_d)\}$. Even though GLT- g_T and GLT- g_A are not optimal to simulate the lookback options, both methods substantially outperform MC and QMC.

- Although the relative efficiency of GLT depends on its optimal application, it appears that it is reasonably robust to rely on GLT- g_A for other kinds of path-dependent options.

4.3 The efficiency in terms of dimension reduction

Recall that truncation dimension of a function is one of the measures of effective dimension. It constructively identifies the subset of the dimensions that predominantly explains the function. To elaborate, let us consider a function which depends on 100 variables. Nominally, this function is said to have 100 dimensions. Suppose further that at 99% confidence level (i.e., setting $p = 0.99$ in the definition of truncation dimension), the

Table 3: Simulated prices of various European derivative securities based on MC, QMC, GLT- g_T and GLT- g_A . Efficiency ratios of QMC, GLT- g_T and GLT- g_A , relative to MC, are reported in parentheses.

d	MC	QMC	GLT- g_T	GLT- g_A
Plain-vanilla call options				
4	11.918(0.060)	11.947(20.2)	11.952(87.1)	11.951(85.8)
12	11.898(0.048)	11.967(6.4)	11.952(67.4)	11.950(33.7)
50	11.975(0.051)	11.917(2.4)	11.954(51.5)	11.951(12.6)
250	12.027(0.050)	11.961(1.6)	11.949(17.2)	11.956(6.0)
Asian call options				
4	7.985(0.045)	8.016(35.9)	8.015(59.6)	8.015(117.3)
12	7.112(0.030)	7.136(10.5)	7.132(23.7)	7.133(66.7)
50	6.802(0.031)	6.787(4.1)	6.795(10.1)	6.796(25.4)
250	6.743(0.025)	6.737(2.1)	6.713(5.1)	6.710(5.2)
Average call options (decreasing weights)				
4	6.287(0.037)	6.307(87.1)	6.305(31.0)	6.306(136.1)
12	4.900(0.022)	4.908(16.2)	4.910(13.0)	4.909(77.5)
50	4.369(0.019)	4.363(6.4)	4.360(3.8)	4.365(42.9)
250	4.228(0.016)	4.239(2.4)	4.225(2.1)	4.225(7.3)
Average call options (increasing weights)				
4	10.100(0.053)	10.133(23.9)	10.136(120.8)	10.135(104.2)
12	9.633(0.041)	9.681(7.4)	9.672(51.4)	9.672(51.2)
50	9.501(0.043)	9.472(2.7)	9.493(24.5)	9.492(24.0)
250	9.505(0.035)	9.480(1.7)	9.446(10.1)	9.442(6.9)
lookback options				
4	15.286(0.058)	15.303(36.0)	15.307(61.8)	15.306(59.4)
12	17.169(0.052)	17.216(8.5)	17.205(16.4)	17.200(20.9)
50	18.715(0.047)	18.660(2.4)	18.686(8.3)	18.697(10.1)
250	19.608(0.049)	19.516(1.7)	19.507(5.3)	19.519(6.1)

truncation dimension of the function is found to be 3. This suggests that even though we begin with a function that is high-dimension, the first three dimensions capture at least 99% of the total variation (as measured by the variance). Hence effectively the function is said to have low dimension.

In our numerical analysis, we use the cumulative explanatory ratio (CER) to gauge the effectiveness of dimension reduction. Formally, CER is defined as

$$CER(k) = \frac{\sum_{u \in \{1, 2, \dots, k\}} \sigma^2(f_u)}{\sigma^2(f)}. \quad (13)$$

This ratio gives the proportion of the variance captured by the first k dimensions relative to the total variance. Consequently, $CER(k)$ is between 0 and 1 and that $CER(d_T) \geq p$ for truncation dimension d_T at confidence level p .

Table 4 produces $\{CER(k); k = 1, \dots, 5\}$ for the option examples considered previously. The ratios are estimated numerically (based on MC with 100,000 sample size) using the procedure described in [31]. In addition to reporting CER of the standard QMC application, we also tabulate CER based on GLT- g_T and GLT- g_A . We now make the following observations based upon the reported CER:

- The effectiveness of GLT at dimension reduction, relative to standard QMC, is exemplified by the results depicted in Table 4. Let us consider, for example, the plain-vanilla option with $d = 250$. Using GLT- g_T , the first dimension impressively captures at least 97% of the total variance, in sharp contrast to the standard QMC which, even if we were to include the first five dimensions, it only manages to capture about 1% of the total variance. Similarly for the Asian option with $d = 250$, the first dimension of GLT- g_A accounts for at least 97% of the total variance while the first dimension of the standard QMC explains less than 1%.
- For the plain-vanilla options, the reported CERs based on GLT- g_T are consistently higher than the corresponding values from GLT- g_A . On the other hand, the CERs of the three types of average options are consistently higher for GLT- g_A , relative to GLT- g_T . This is to be expected since by design, GLT- g_T is optimal for the plain vanilla options while GLT- g_A is optimal for the average options.
- It is instructive to note the role of weights of the average options on GLT- g_T . In terms of CER, GLT- g_T is most effective for the average options with increasing weights while the least effective for the corresponding options with decreasing weights. This is again quite intuitive since GLT- g_T is optimally designed for simulating terminal asset prices. The payoff of the average option depends on the monitoring asset prices at time $t_i, i = 1, \dots, d$. However, due to the distribution of the weights, the terminal asset price becomes the most dominant for the increasing weighted average option and the least for the decreasing weighted average option. Consequently, GLT- g_T is best suited for the average option with increasing weights and the least with decreasing weights. This in part explains why GLT- g_T is less efficient than the corresponding QMC for the Asian option with decreasing weights (see Table 3).
- For the lookback options, the CERs of GLT are calculated by assuming the terminal asset price g_T and the equally weighted price path g_A . Therefore, both implementations of GLT are not optimal for simulating lookback options. Nevertheless, the CERs of both GLT- g_T and GLT- g_A are much larger than the standard QMC. Interestingly, GLT- g_T yields a much higher CER than the corresponding method based on g_A .
- There is a slight discrepancy in the reported CER in the sense in some cases $CER(k)$ decreases with increasing k . This anomaly is induced by the MC error in estimating CER.

Table 4: A comparison of CER (up to five dimensions) for standard QMC, GLT- g_T and GLT- g_A .

Dimension u		1	2	3	4	5
$d = 4$						
Plain-vanilla option	standard	19.3	42.5	69.4	100.0	-
	GLT- g_T	100.0	100.0	100.0	100.0	-
	GLT- g_A	85.0	94.0	98.8	100.0	-
Asian option	standard	44.2	78.0	95.3	100.0	-
	GLT- g_T	80.2	91.9	98.2	100.0	-
	GLT- g_A	99.1	99.9	100.0	100.0	-
Decreasing weighted ave. option	standard	73.7	96.8	99.9	100.0	-
	GLT- g_T	49.2	65.8	86.5	100.0	-
	GLT- g_A	98.4	100.0	100.0	100.0	-
Increasing weighted ave. option	standard	27.0	59.1	87.0	100.0	-
	GLT- g_T	94.4	99.4	100.0	100.0	-
	GLT- g_A	99.4	99.8	99.9	100.0	-
Lookback option	standard	11.1	31.8	61.7	100.0	-
	GLT- g_T	93.0	96.3	98.7	100.0	-
	GLT- g_A	71.4	91.1	97.6	100.0	-
$d = 12$						
Plain-vanilla option	standard	5.5	12.2	19.0	26.3	33.5
	GLT- g_T	100.0	100.0	100.0	100.0	100.0
	GLT- g_A	72.8	86.0	91.7	94.3	95.2
Asian option	standard	16.4	32.8	47.4	60.4	71.4
	GLT- g_T	73.7	79.8	84.4	89.4	98.5
	GLT- g_A	99.7	99.8	99.9	100.0	100.0
Decreasing weighted ave. option	standard	35.3	61.8	78.7	89.2	95.1
	GLT- g_T	40.6	45.5	50.1	57.4	63.8
	GLT- g_A	99.7	99.9	100.0	100.0	100.0
Increasing weighted ave. option	standard	8.8	19.1	29.6	40.7	51.7
	GLT- g_T	88.4	93.6	96.5	98.5	99.3
	GLT- g_A	99.8	99.9	99.9	99.9	100.0
Lookback option	standard	1.4	5.0	9.4	15.3	21.8
	GLT- g_T	91.4	93.1	94.2	95.1	96.2
	GLT- g_A	55.6	79.7	87.0	91.2	92.7
$d = 50$						
Plain-vanilla option	standard	1.2	3.0	5.1	6.5	7.6
	GLT- g_T	99.9	99.7	99.8	99.7	99.9
	GLT- g_A	69.9	82.6	87.9	90.7	90.5
Asian option	standard	3.9	8.1	12.3	16.8	21.1
	GLT- g_T	71.8	72.3	74.0	75.9	99.8
	GLT- g_A	99.8	99.7	99.8	99.8	99.8
Decreasing weighted ave. option	standard	9.6	18.4	27.3	36.0	43.7
	GLT- g_T	39.0	38.7	39.5	41.2	42.1
	GLT- g_A	99.7	99.7	99.8	99.8	99.8
Increasing weighted ave. option	standard	2.1	4.4	7.0	9.5	12.0
	GLT- g_T	85.6	86.8	88.6	90.2	91.7
	GLT- g_A	99.7	99.7	99.8	99.8	99.9
Lookback option	standard	0.0	0.8	2.2	2.3	2.4
	GLT- g_T	91.1	91.1	91.7	92.0	92.8
	GLT- g_A	51.5	74.6	81.7	86.2	86.0
$d = 250$						
Plain-vanilla option	standard	0.2	0.4	0.5	0.7	1.1
	GLT- g_T	97.8	97.9	97.9	97.7	97.8
	GLT- g_A	67.1	80.4	85.1	86.4	88.6
Asian option	standard	0.9	1.4	2.5	3.0	3.4
	GLT- g_T	68.7	69.6	70.1	70.1	70.7
	GLT- g_A	97.5	97.5	97.7	97.9	97.9
Decreasing weighted ave. option	standard	2.1	3.5	6.1	7.7	9.0
	GLT- g_T	36.9	37.4	37.6	37.2	38.0
	GLT- g_A	97.7	97.4	97.7	97.8	97.8
Increasing weighted ave. option	standard	0.5	0.6	1.2	1.4	1.7
	GLT- g_T	82.4	83.1	83.8	83.6	84.4
	GLT- g_A	97.8	97.5	97.6	97.7	97.6
Lookback option	standard	0.0	0.0	0.0	0.0	0.0
	GLT- g_T	89.3	89.3	89.8	89.3	89.3
	GLT- g_A	48.7	72.6	81.2	81.2	84.5

5 Conclusion

In this paper, a comprehensive numerical experiment is conducted to validate the GLT dimension reduction reduction technique recently proposed in [16] and [17]. We assumed that the dynamics of the asset prices was governed by the Meixner Lévy process. We used a number of representative exotic options as our test cases. Furthermore, we evaluated the effectiveness of GLT using criteria such as the efficiency ratios, robustness and CER. These studies demonstrated the competitive advantage of the GLT relative to the the standard MC and QMC. Testing GLT on the hybrid sequences further illuminated the effectiveness GLT on dimension reduction. In conclusion, GLT offers a powerful and yet general approach of simulating a wide range of stochastic processes. It will be of interest to compare the relative efficiency of GLT to other dimension reduction techniques. 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] P.P. Boyle. Options: A Monte Carlo approach. *Journal of Financial Economics*, 4(4):323–338, 1977.
- [5] P.P. Boyle, M. Broadie, and P. Glasserman. Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21(8-9):1276–1321, 1997.
- [6] R.E. Caflisch, W. Morokoff, and A. Owen. Valuation of mortgaged-backed securities using brownian bridges to reduce effective dimension. *Journal of Computational Finance*, 1(1):27–46, 1997.
- [7] 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.
- [8] F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300:463–520, 1994.

- [9] J-C Duan. The GARCH option pricing model. *Mathematical Finance*, 5:13–32, 1995.
- [10] H.U. Gerber and E.S.W. Shiu. Option pricing by esscher-transforms. *Transactions of the Society of Actuaries*, 46:99–191, 1994.
- [11] B. Grigelionis. Process of Meixner type. *Lithuanian Mathematics Journal*, 39(1):33–41, 1999.
- [12] 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.
- [13] 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.
- [14] J. Hull and A. White. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 41:281–300, 1987.
- [15] J. Imai and K.S. Tan. A general dimension reduction technique for derivative pricing. *Journal of Computational Finance*, 10(2):129–155, 2006.
- [16] 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*, 31(3):2282–2302, 2009.
- [17] J. Imai and K.S. Tan. A generalized linear transformation method for simulating Meixner Lévy processes. In S.I. Ao, D.W.L. Hukins, A. Hunter, and A.M. Korsunsky, editors, *Proceedings of the World Congress on Engineering 2009, WCE 2009, July 1 - 3, 2009, London, U.K.*, volume II, pages 1406–1411. Newswood Limited, International Association of Engineers, 2009.
- [18] C. Joy, P.P. Boyle, and K.S. Tan. Quasi-Monte Carlo methods in numerical finance. *Management Science*, 42(6):926–938, 1996.
- [19] A. Kyprianou, W. Schoutens, and P. Wilmott, editors. *Exotic Option Pricing and Advanced Lévy Models*. Wiley, West Sussex, England, 2005.
- [20] G. Leobacher. Stratified sampling and quasi-Monte Carlo simulation of Lévy processes. *Monte Carlo Methods and Applications*, 12(3-4):231–238, 2006.
- [21] M. Matsumoto and T. Nishimura. Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM Transactions on Modeling and Computer Simulation*, 8(1):3–30, 1998.
- [22] B. Moskowitz and R. Caffisch. Smoothness and dimension reduction in quasi-Monte Carlo methods. *Mathematical and Computer Modeling*, 23(8-9):37–54, 1996.
- [23] H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*. S.I.A.M., Philadelphia, 1992.
- [24] S. Ninomiya and S. Tezuka. Toward real-time pricing of complex financial derivatives. *Applied Mathematical Finance*, 3:1–20, 1996.
- [25] S. H. Paskov and J. F. Traub. Faster valuation of financial derivatives. *Journal of Portfolio Management*, 22(1):113–120, 1995.
- [26] 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.
- [27] P. Sabino. Implementing quasi-Monte Carlo simulations with linear transformations. *Computational Management Science*, 2009.
- [28] P. Sabino. Efficient quasi-Monte simulations for pricing high-dimensional path-dependent options. *Decisions in Economics and Finance*, 32(1):49–65, forthcoming.
- [29] W. Schoutens. Meixner processes: theory and applications in finance. *EURANDOM Report 2002-004, EURANDOM, Eindhoven*, 2002.
- [30] W. Schoutens. *Lévy processes in Finance: Pricing Financial Derivatives*. Wiley, West Sussex, England, 2003.
- [31] X. Wang and K-T Fang. The effective dimension and quasi-Monte Carlo. *Journal of Complexity*, 19:101–124, 2003.