Insurance: Mathematics and Economics 43 (2008) 327–338

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

journal homepage: www.elsevier.com/locate/ime

Computation of optimal portfolios using simulation-based dimension reduction

Phelim Boyle [a], Junichi Imai [b], Ken Seng Tan [c,d,*]

[a] School of Business & Economics, Wilfrid Laurier University, Waterloo, N2L 3C5, Canada

[b] Keio University Faculty of Science and Technology, Yagami Campus 3-14-1 Hiyoshi Kohoku-ku, Yokohama-shi Kanagawa-ken 223-8522, Japan

[c] Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, N2L 3G1, Canada

[d] China Institute for Actuarial Science, Central University of Finance and Economics, Beijing, 10081, China

ARTICLE INFO

Article history:
Received May 2008
Accepted 4 May 2008

Keywords:
Optimal portfolio selection
Asset allocation
Dimension reduction
Quasi-Monte Carlo

ABSTRACT

This paper describes a simple and efficient method for determining the optimal portfolio for a risk averse investor. The portfolio selection problem is of long standing interest to finance scholars and it has obvious practical relevance. In a complete market the modern procedure for computing the optimal portfolio weights is known as the martingale approach. Recently, alternative implementations of the martingale approach based on Monte Carlo methods have been proposed. These methods use Monte Carlo simulation to compute stochastic integrals. This paper examines the efficient implementation of one of these methods due to [Cvitanic, J., Goukasian, L., Zapatero, F. 2003. Monte Carlo computation of optimal portfolios in complete markets. J. Econom. Dynam. Control 27, 971–986]. We explain why a naive application of the quasi-Monte Carlo method to this problem is often only marginally more efficient than the classical Monte Carlo method. Using the dimension reduction technique of [Imai, J., Tan, K.S., 2007. A general dimension reduction method for derivative pricing. J. Comput. Financ. 10 (2), 129–155] it is possible to significantly reduce the effective dimension of the problem. The paper shows why the proposed technique leads to a dramatic improvement in efficiency.

1. Introduction

The asset allocation decision is of great interest to the finance discipline. It deals with the optimal portfolio choice of the individual agent — a problem of theoretical importance in its own right. Asset allocation choice has always been of practical relevance to individuals. In recent years, it has assumed even greater importance because pension plans in many jurisdictions are now switching from defined benefit plans to defined contribution plans. This switch transfers the investment risk (and in many cases the asset allocation decision) from the company to the employees. Investment advisors[1] provide asset mix recommendations that indicate that the proportion of stocks to bonds should decrease as the investor becomes more risk averse.

Finance theorists have long been interested in providing a scientific framework for this decision. The earliest paper on portfolio selection is due to Markowitz (1952) who solved the asset allocation problem in a one period model for an investor with quadratic utility. One of the main lessons from this work is that investors should diversify, a lesson that is still sometimes ignored.[2] Also in 1952, Arrow (Arrow, 1952) wrote his foundational paper on the pricing of contingent claims in a complete market setting. In the late 1960's Robert Merton wrote a series[3] of important papers on the consumption investment problem using the powerful continuous-time framework. Merton assumed the investor adjusts the asset proportions and the consumption rate on a continuous basis to maximize expected utility. He formulated the problem in a dynamic programming framework and used techniques from optimal stochastic control to derive solutions. Merton was able to derive explicit solutions in a few cases by making strong assumptions about the investor's utility function and the asset price dynamics. For more general settings it is often difficult to construct solutions under the dynamic programming approach especially for high dimensional problems. Although

* Corresponding address: University of Waterloo, 200 University Avenue West, N2L 3G1 Waterloo, Ontario Canada. Tel.: +1 519 888 4567; fax: +1 519 746 1875.
E-mail addresses: pboyle@wlu.ca (P. Boyle), jimai@ae.keio.ac.jp (J. Imai), kstan@uwaterloo.ca (K.S. Tan).

[1] For a discussion of this point, see Brennan and Xia (2000) and Canner et al. (1997). Bajeux-Besnainou et al. (2001) show that intertemporal hedging and stochastic interest rates can resolve the Canner, Mankiv and Weil asset allocation puzzle. However Lioui (2007) shows that the puzzle may still persist.

[2] The dangers of the lack of diversification were dramatically illustrated in the Enron case. Enron was the seventh largest company in the United States prior to its collapse in December 2001. Many Enron employees had a significant proportion of their retirement plans invested in Enron's own stock.

[3] See Merton (1969, 1971).

0167-6687/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.insmatheco.2008.05.004

Brennan et al. (1997) (BLS) solve the HJB equation numerically using a finite difference approach, their solution is numerically unstable since the optimal weights bounce between 100% stock and 100% bonds within a small number of periods.

The so called martingale approach, derived independently by a number of authors,[4] has distinct advantages over the dynamic programming approach and it is easier to implement. The basic idea is to first obtain the investor's optimal wealth using the fact that in a complete market setting, the investor's budget constraint can be transformed into a static budget constraint. The investor's optimal final wealth can be viewed as a contingent claim which can be priced under the equivalent martingale measure. Under the complete market assumption, this payoff can also be hedged using the existing traded assets and this replicating portfolio furnishes the investor's optimal investment choice. The martingale approach requires the solution of a linear partial differential equation, in contrast to the non-linear partial differential equation associated with the dynamic programming approach. In a few specialized cases,[5] closed-form solutions can be obtained by making special assumptions about the asset price dynamics and/or the investor's utility specification. Closed-form solutions[6] are much prized because they are usually simpler to work with and serve as a guide to our intuition.

However for realistic asset price dynamics, closed-form solutions are generally unavailable and numerical methods must be used to compute the optimal portfolios. Recently three different approaches – all based on Monte Carlo (MC) simulation – have been proposed to compute the optimal portfolio weights. Detemple et al. (2003) (DGR) propose a simulation method that is based on a particular representation of the optimal portfolio weights as conditional expectations of random variables. This representation is derived from the Ocone and Karatzas (1991) formula for the weights under the martingale approach. In this case, the optimal portfolio weights are expressed in terms of conditional expectations that involve the optimal wealth and Malliavin derivatives of the state variables. These Malliavin derivatives capture the impact of innovations in the underlying drivers of uncertainty (the Brownian motions) on the state variables. The DGR method can handle realistic problems with complex dynamics for the state variables and several asset classes.

Cvitanic et al. (2003) (CGZ) noted that the covariation between the optimal wealth process and the uncertainty shocks provides an expression for the optimal portfolio weights. This covariation can be estimated numerically by simulating the explicit expressions for the optimal wealth process. Under the martingale approach, the investor's optimal wealth process can be written directly as an expectation of a stochastic integral. Hence, as in the DGR case the CGZ approach involves the numerical simulation of stochastic integrals. The third simulation based approach is due to Brandt et al. (2005) who propose an approximate procedure that uses a series expansion of the value function together with regressions of conditional expectations on powers of the state variables. This method is based on combining a number of different approximations and its implementation and convergence raises some interesting questions.[7]

Thus all three numerical methods are based on MC simulation. MC can accommodate a broad range of models and it can be used for high-dimensional problems, unlike the partial differential equation based approaches. In the first two methods, MC is used as a numerical technique to evaluate high dimensional integrals. Hence the efficiency of these two methods depends directly on an accurate evaluation of such integrals. However, MC is often criticized for its slow convergence. The conventional wisdom is that standard MC attains a convergence rate of $O(N^{-1/2})$ for sample size N and this is the case for many applications. However Detemple et al. (2005) point out that when the CGZ method is combined with an Euler discretization scheme the convergence is only $O(N^{-1/3})$. The $O(N^{-1/2})$ convergence rate will only be obtained if the exact transition densities are known and Euler discretization can be avoided.

The aim of the present paper is to explore methods of improving the efficiency of MC simulation in this context. We investigate the possibility of using quasi-Monte Carlo (QMC) methods to speed up the computations and we concentrate on the CGZ method.

Quasi-Monte Carlo (QMC) uses specially selected deterministic points rather than random points as in standard MC. These points have the property that they are more uniformly distributed than the random points. Sequences with this property are known as low discrepancy sequences. QMC promises a much higher asymptotic convergence rate than MC. Early applications of QMC to some finance problems have reported the advantages of QMC over MC. See, for example Paskov and Traub (1995) and Joy et al. (1996).

However for some high dimensional problems, the advantages of QMC are difficult to realize using the standard QMC approach. This is related to the fact that even though QMC sequences have good distributional properties, we do not get the benefit of this at higher dimensions unless we use a very large number of points. For high dimensional problems, standard QMC may not provide much improvement over classical MC. This point has been noted by several authors. See for example Akesson and Lehoczky (2000), Ninomiya and Tezuka (1996), and Tan and Boyle (2000).

In our application we will be concerned with estimating the value of stochastic path integrals and we generate the paths by simulation. If the path has d time steps, we can map the construction of a sample path into a point in the d-dimensional unit hypercube. For our applications, we can exploit the good distributional properties of QMC points in the lower dimensions if we can transform the problem of path generation so that the shape of the path is determined in large part by the first few dimensions. For some simple problems, the Brownian bridge construction provides an elegant method of doing this (see Caflisch et al. (1997)). The principal component construction proposed by Acworth et al. (1998) provides another way of enhancing QMC.

More recently Imai and Tan (2007) show that their proposed linear transformation (LT) leads to dramatic efficiency gains when it is combined with QMC. The LT construction makes explicit linkage between the efficiency of QMC and the effective dimension of the problem of interest. By optimally reducing the effective dimension, they demonstrate that the superior rate of QMC can be recovered. Their construction is powerful and general. In the current paper we use the LT construction to compute the optimal portfolio weights within the CGZ framework. We show that the Imai and Tan procedure leads to a significant improvement in computational efficiency over both MC and standard QMC.

The layout of the rest of the paper is as follows. Section 2 describes the basic problem and the solution procedure of CGZ. First we describe the martingale approach. Then we discuss the CGZ solution procedure. We illustrate the procedure in a special case where there is a closed-form solution. In Section 3, we explain how to adapt the QMC method to exploit the structure of a problem. In particular, we provide additional insight on the relationship between the effective dimension of a problem and the efficiency of QMC. Then we introduce the LT construction of Imai and Tan (2007). In Section 4, we discuss a basic example that

[4] Cox and Huang (1989), Karatzas et al. (1987), Pliska (1986).

[5] Brennan and Xia (2000), Kim and Omberg (1996), Liu (2007) and Wachter (2002) derive closed-form solutions.

[6] For a general discussion of the concept of closed-form solution in finance and some tentative attempts at defining it, see Boyle et al. (2002).

[7] See footnote 21 of Detemple et al. (2003).

will serve as a prototype and we also demonstrate how to adapt the LT construction to solve the optimal portfolio problems. Section 5 provides numerical results. We compare different simulation methods and show the superiority of the LT method over standard QMC. We also quantify explicitly the extent to which the LT method reduces the dimensionality of the problem since this is the source of its computational superiority. Section 6 concludes the paper.

2. The Monte Carlo covariation method

This section describes the Monte Carlo Covariation Method for optimal portfolio selection. We first describe the basic optimization problem and the details of the financial market and investor preferences. Then we discuss the martingale method and describe the approach proposed by Cvitanic et al. (2003) (CGZ). We illustrate their method using a simple example where the optimal solution is known and there is no need for a numerical approach.

2.1. Portfolio selection using the Martingale method

In this subsection we give a brief summary of the martingale approach to portfolio selection. For more details of this approach see Cox and Huang (1989). We assume a complete market where there is no arbitrage. There are m risky assets. The price of asset i satisfies the following stochastic differential equation

$$\frac{\mathrm{d}S_t^i}{S_t^i} = \mu_t^i \mathrm{d}t + (\boldsymbol{\sigma}_t^i)' \mathrm{d}W_t, \tag{2.1}$$

where W is a vector of m standard Brownian motions and μ^i and $\boldsymbol{\sigma}^i$ denote, respectively, the drift and the volatility of the asset price process. The bank account process B_t follows

$$\mathrm{d}B_t = r_t B_t \mathrm{d}t, \tag{2.2}$$

where r_t is the locally riskless interest rate.

The market price of risk, $\boldsymbol{\theta}_t$, is an $m \times 1$ vector, defined by

$$\boldsymbol{\theta}_t = (\boldsymbol{\sigma}_t)^{-1}(\boldsymbol{\mu}_t - r_t \mathbf{1}), \tag{2.3}$$

where $\mathbf{1}$ is the unit vector, $\boldsymbol{\mu}_t$ is the $m \times 1$ vector whose i-th component is μ_t^i, and $\boldsymbol{\sigma}_t$ is the $m \times m$ matrix whose i-th column is $\boldsymbol{\sigma}_t^i$. We assume that $\boldsymbol{\theta}_t$ is continuously differentiable and satisfies the Novikov condition.

The state price density represents the continuous time generalization of Arrow–Debreu prices and it can be defined in terms of $\boldsymbol{\theta}$. By defining the following process

$$\xi_t = \xi_0 \exp\left\{-\frac{1}{2}\int_0^t \boldsymbol{\theta}_s'\boldsymbol{\theta}_s \mathrm{d}s - \int_0^t \boldsymbol{\theta}_s' \mathrm{d}W_s\right\}, \tag{2.4}$$

then the discounted value of ξ, as given by

$$\mathrm{e}^{-\int_0^t r_s \mathrm{d}s}\xi_t,$$

is the state price density.

We consider an investor, with initial wealth x_0, who wishes to maximize expected utility by selecting a dynamic portfolio with positions in the risky assets and the riskless asset. For example, if the investor were to maximize the expected utility of terminal wealth, X_T, over time horizon T, then the optimal dynamic trading strategy, $\boldsymbol{\pi}_t^*$, is the solution to the following optimization problem:

$$\max_{\boldsymbol{\pi}} \mathrm{E}\left[u(X_T)|\mathcal{F}_0\right] \tag{2.5}$$

$$\text{subject to } X_0 = x_0 \tag{2.6}$$

$$\mathrm{d}X_t = [\boldsymbol{\pi}_t'\boldsymbol{\mu}_t + (X_t - \boldsymbol{\pi}_t'\mathbf{1})r_t]\mathrm{d}t + \boldsymbol{\pi}_t'\boldsymbol{\sigma}_t \mathrm{d}W_t \tag{2.7}$$

$$X_t \geq 0, \quad t \in [0, T]. \tag{2.8}$$

In the above formulation, the $m \times 1$ vector $\boldsymbol{\pi}_t$ denotes the amounts invested in the n risky assets at time t and $u(\cdot)$ is the utility function which is assumed to be strictly increasing, strictly concave and twice differentiable. The budget constraint (2.6) ensures that we begin with the initial exogenous level of wealth x_0 while the constraint (2.7) corresponds to the wealth process of the investor that must be satisfied over time $t \in [0, T]$.

Alternatively, if the investor were to maximize the expected utility from consumption and c_t represents the investor's consumption at time t, then the problem faced by the investor is to solve the following optimization problem for the optimal consumption plan c_t^* and the optimal investment strategy $\boldsymbol{\pi}_t^*$:

$$\max_{\pi, c} \mathrm{E}\left[\int_0^T u(c_s)\mathrm{d}s \middle| \mathcal{F}_0\right] \tag{2.9}$$

$$\text{subject to } X_0 = x_0$$

$$\mathrm{d}X_t = [\boldsymbol{\pi}_t'\boldsymbol{\mu}_t + (X_t - \boldsymbol{\pi}_t'\mathbf{1})r_t - c_t]\mathrm{d}t + \boldsymbol{\pi}_t'\boldsymbol{\sigma}_t \mathrm{d}W_t$$

$$X_t \geq 0, \quad c_t \geq 0, \quad t \in [0, T].$$

We now summarize the martingale approach for determining the optimal portfolio. We only focus on the investment optimization problem (2.5) as a similar technique can be applied to the consumption-investment problem (2.9).

Together with (2.3) and (2.7), the dynamics of X_t can be rewritten as

$$\begin{aligned}\mathrm{d}X_t &= (\boldsymbol{\pi}_t'\boldsymbol{\sigma}_t\boldsymbol{\theta}_t + r_t X_t)\mathrm{d}t + \boldsymbol{\pi}_t'\boldsymbol{\sigma}_t \mathrm{d}W_t \\ &= (r_t X_t)\mathrm{d}t + \boldsymbol{\pi}_t'\boldsymbol{\sigma}_t(\mathrm{d}W_t + \boldsymbol{\theta}_t \mathrm{d}t) \\ &= (r_t X_t)\mathrm{d}t + \boldsymbol{\pi}_t'\boldsymbol{\sigma}_t \mathrm{d}W_t^Q,\end{aligned}$$

where W^Q is a vector of standard Brownian motions with respect to the equivalent martingale measure Q. The Brownian motions W and W^Q are related through the Girsanov theorem (see Karatzas and Shreve (1991)):

$$W_t^Q = W_t + \int_0^t \boldsymbol{\theta}_s \mathrm{d}s.$$

The investor's optimal terminal wealth is the solution to the optimization problem (2.5) subject to the constraints (2.6)–(2.8). Because we are assuming a complete market, the optimal terminal wealth can be regarded as a contingent claim with payoff at time T equal to X_T. This implies that the investor's initial wealth and her final wealth at time T must satisfy

$$X_0 = x_0 = \mathrm{E}^Q\left[\mathrm{e}^{-\int_0^T r_s \mathrm{d}s} X_T \middle| \mathcal{F}_0\right] = \mathrm{E}\left[\xi_T \mathrm{e}^{-\int_0^T r_s \mathrm{d}s} X_T \middle| \mathcal{F}_0\right]. \tag{2.10}$$

More generally, the wealth process for any time t can be expressed as

$$X_t = \mathrm{E}^Q\left[\mathrm{e}^{-\int_t^T r_s \mathrm{d}s} X_T \middle| \mathcal{F}_t\right] = \mathrm{E}[H_{t,T} X_T | \mathcal{F}_t] \tag{2.11}$$

where

$$H_{t,u} = \mathrm{e}^{-\int_t^u r_s \mathrm{d}s}\frac{\xi_u}{\xi_t}, \quad u > t. \tag{2.12}$$

The above results imply that the dynamic constraints considered in our optimization problems can be transformed into equivalent constraints that are static and hence easier to deal with. In fact they can be incorporated directly into the maximization problem and this leads to a solution for the investor's optimal wealth in terms of the state price density and the investor's utility function. We now provide a brief summary of the derivation of this result.

Suppose $g(\cdot)$ denotes the inverse of the agent's marginal utility function in (2.5); i.e.

$$u'(g(y)) = y.$$

Our assumptions on the utility function imply that g is strictly decreasing and continuous. For a more general definition of this inverse function, see Karatzas and Shreve (1998, p. 95). Using standard optimization techniques, the investor's optimal terminal wealth, X_T^*, corresponding to problem (2.5) is given by

$$X_T^* = g\left(\lambda \mathrm{e}^{-\int_0^T r_s \mathrm{d}s}\xi_T\right), \tag{2.13}$$

where λ is a Lagrange multiplier. Since the optimal terminal wealth X_T^* must satisfy (2.10), this provides a way of obtaining the Lagrange multiplier λ; namely solving the following equation:

$$x_0 = \mathrm{E}\left[\xi_T \mathrm{e}^{-\int_0^T r_s \mathrm{d}s} g\left(\lambda \mathrm{e}^{-\int_0^T r_s \mathrm{d}s}\xi_T\right)\middle|\, \mathcal{F}_0\right]. \tag{2.14}$$

Recall that our objective is to determine the optimal portfolios π_t^*. We now present the method proposed by Cvitanic et al. (2003), which is known as the Monte Carlo covariation method. First note that from the Martingale Representation Theorem,[8] the wealth process (2.11) must satisfy the following diffusion process

$$\mathrm{d}X_t = \alpha_t \mathrm{d}t + \boldsymbol{\nu}_t' \mathrm{d}W_t, \tag{2.15}$$

where α_t is the drift and $\boldsymbol{\nu}_t$ is the m-dimensional diffusion coefficient. It follows from (2.7) that the dynamics for the investor's optimal wealth also satisfies

$$\mathrm{d}X_t^* = [\boldsymbol{\pi}_t^{*'}\boldsymbol{\mu}_t + (X_t^* - \boldsymbol{\pi}_t^{*'}\mathbf{1})r_t]\mathrm{d}t + \boldsymbol{\pi}_t^{*'}\boldsymbol{\sigma}_t \mathrm{d}W_t.$$

Because the last two stochastic differential equations represent the same process, their diffusion terms must be equal. Hence

$$\boldsymbol{\nu}_t' = \boldsymbol{\pi}_t^{*'}\boldsymbol{\sigma}_t,$$

which immediately leads to the following formula for $\boldsymbol{\pi}_t^*$:

$$\boldsymbol{\pi}_t^* = (\boldsymbol{\sigma}_t)^{-1}\boldsymbol{\nu}_t. \tag{2.16}$$

This means that if we were able to calculate the vector $\boldsymbol{\nu}_t$, then we could obtain the optimal portfolio weights $\boldsymbol{\pi}_t^*$. Apart from a few special cases,[9] one cannot obtain an explicit expression for $\boldsymbol{\nu}_t$. Hence we need to rely on some numerical procedures. CGZ point out that there are two simulation-based approaches to estimate $\boldsymbol{\nu}_t$.

As explained by CGZ the i-th component of the vector $\boldsymbol{\nu}_t$ is given by the following limit:

$$\begin{aligned}\nu_t^i &= \lim_{\Delta t\to 0} \mathrm{E}\left[\frac{\left(X_{t+\Delta t}^* - X_t^*\right)\left(W_{t+\Delta t}^i - W_t^i\right)}{\Delta t}\middle|\, \mathcal{F}_t\right] \\ &= \lim_{\Delta t\to 0} \mathrm{E}\left[\frac{\left(X_{t+\Delta t}^* - X_t^*\right) Z_t^i}{\sqrt{\Delta t}}\middle|\, \mathcal{F}_t\right],\end{aligned} \tag{2.17}$$

where Z_t^i is a standard normal random variable. Furthermore, combining both (2.11) and (2.13), we have

$$\begin{aligned}X_{t+\Delta t}^* &= \mathrm{E}\left[H_{t+\Delta t,T} X_T^*\middle|\, \mathcal{F}_{t+\Delta t}\right] \\ &= \mathrm{E}\left[H_{t+\Delta t,T} g\left(\lambda \mathrm{e}^{-\int_0^T r_s \mathrm{d}s}\xi_T\right)\middle|\, \mathcal{F}_{t+\Delta t}\right].\end{aligned} \tag{2.18}$$

The expectations in both (2.17) and (2.18) can be evaluated using Monte Carlo simulation. However, the former expectation is conditional on information up to time t while the latter is conditional on information up to time $t + \Delta t$. This implies that we need to use a two-tier Monte Carlo simulation approach to estimate both (2.17) and (2.18), as explained by CGZ.

The second Monte Carlo based approach is more direct and we call it the *direct method*. From the law of iterated conditional expectations, (2.17) can be expressed as

$$\nu_t^i = \lim_{\Delta t\to 0} \mathrm{E}\left[\frac{\left(H_{t+\Delta t,T} X_T^* - X_t^*\right) Z_t^i}{\sqrt{\Delta t}}\middle|\, \mathcal{F}_t\right]. \tag{2.19}$$

The direct method allows us to estimate $\boldsymbol{\nu}_t$ directly and is computationally more efficient than the two-tier[10] simulation procedure. The standard errors of the estimate are easier to compute if we use the direct method and as we shall explain later there is an additional computational advantage in using this approach for the simulation technique considered here. Hence, in this paper, we use the direct approach to estimate $\boldsymbol{\nu}_t$.

It is interesting to note the connection between Eq. (2.19) and Malliavin calculus applications in finance (see Fournié et al. (1999, 2001)). Eq. (2.19) is an approximation to the Malliavin weight estimator based on the Euler discretization.

2.2. An example with an analytical solution

It is instructive to discuss a simple and familiar example where there is an analytical solution for $\boldsymbol{\nu}_t$ and $\boldsymbol{\pi}_t^*$. This provides insight into the procedure in a simple setting. In addition this closed-form expression is useful for examining the convergence of the limit in Eq. (2.17) as Δt becomes small. We assume that the investor has a power utility function and for simplicity there is just one risky asset whose return dynamics are given by geometric Brownian motion. The risk-free rate is constant and equal to r_0. We assume the investor wishes to solve problem (2.5). In other words, the investor selects the optimal portfolio that will maximize her expected utility of wealth at time T, given initial wealth of x_0.

The risky asset dynamics are given by

$$\frac{\mathrm{d}S_t}{S_t} = \mu \mathrm{d}t + \sigma \mathrm{d}W_t,$$

where μ and σ are constants and W_t is a standard Brownian motion. The investor's utility function is

$$u(x) = \frac{x^\gamma}{\gamma}, \quad \gamma \le 1,\ \gamma \neq 1.$$

It is well known that the optimal amount in the risky asset in this case is a constant proportion of the optimal wealth; i.e.,

$$\pi_t^* = \frac{\mu - r_0}{\sigma^2(1-\gamma)} X_t^*, \tag{2.20}$$

where X_t^* is the investor's optimal wealth at time t. The above result was first derived by Merton (1971) and the expression

$$\frac{\mu - r_0}{\sigma^2(1-\gamma)}$$

is often known as the Merton ratio.

In this particular case, there is an explicit expression for X_t^* because the stochastic integral in (2.11) has a simple closed-form solution. Thus the complete solution to the problem can be derived directly using the martingale method. Although this solution is well known, we present it again here since it serves to illustrate the martingale approach to the portfolio selection problem. In addition it shows how the CGZ approach works in a simple setting.

[8] See Oksendal (1998).

[9] One of these cases is when the investor has power utility and the risky assets have a multivariate lognormal distribution. We discuss this case below.

[10] As noted the two-tier simulation method is the one recommended by CGZ. The reasons are described in their paper and derive from the inequality $\mathrm{Var}(E[Y \mid X]) < \mathrm{Var}(Y)$.

First note that the market price of risk is constant since

$$\theta_t = \frac{\mu - r_0}{\sigma} = \theta_0.$$

This implies that the process ξ_t at time t simplifies to

$$\xi_t = \mathrm{e}^{-\frac{1}{2}\int_0^t \theta_0^2 \mathrm{d}s - \int_0^t \theta_0 \mathrm{d}W_s} = \mathrm{e}^{-\frac{1}{2}\theta_0^2 t - \theta_0(W_t - W_0)}.$$

Using the fact that the investor has power utility, the optimal terminal wealth (2.13) can be expressed as

$$\begin{aligned} X_T^* &= (\lambda \mathrm{e}^{-\int_0^T r_0 \mathrm{d}s} \xi_T)^{\frac{1}{\gamma-1}} \\ &= \lambda^{\frac{1}{\gamma-1}} \mathrm{e}^{-\frac{1}{\gamma-1}[\frac{1}{2}\theta_0^2 T + r_0 T + \theta_0(W_T - W_0)]}. \end{aligned}$$

Also, it follows from (2.14) that we can obtain an explicit expression for the Lagrange multiplier λ as

$$\lambda^{\frac{1}{\gamma-1}} = X_0 \mathrm{e}^{r_0 \rho T - \frac{1}{2}\theta_0^2 \rho T\left(\frac{1}{\gamma-1}\right)}, \tag{2.21}$$

where $\rho = \gamma/(\gamma - 1)$.

Furthermore using (2.11), it is easy to verify that the optimal wealth at time t can be expressed explicitly as

$$X_t^* = \lambda^{\frac{1}{\gamma-1}} G(t) \mathrm{e}^{\frac{\theta_0}{1-\gamma}(W_t - W_0)}, \tag{2.22}$$

where $G(t)$ is a deterministic function of time given by

$$G(t) = \mathrm{e}^{\frac{1}{2}\theta_0^2(\rho^2(T-t) - \rho T + t) + r_0(-\rho T + t)}. \tag{2.23}$$

Armed with the explicit expressions, we revisit the CGZ approach of deriving the optimal portfolio weight. By defining

$$\nu_t(h) = \mathrm{E}\left[\left.\frac{\left(X_{t+h}^* - X_t^*\right)(W_{t+h} - W_t)}{h}\right| \mathcal{F}_t\right]$$

for $h \geq 0$, then it follows from (2.17) that

$$\nu_t = \lim_{h \to 0} \nu_t(h).$$

For our example, straightforward calculations show that

$$\nu_t(h) = X_t^* \frac{\theta_0}{1-\gamma}\left[\frac{G(t+h)}{G(t)}\, \mathrm{e}^{\frac{\theta_0^2 h}{2(1-\gamma)^2}}\right] \tag{2.24}$$

and taking the limit $h \to 0$, we obtain

$$\nu_t = X_t^* \frac{\theta_0}{1-\gamma}.$$

Hence the optimal investment in the risky asset is

$$\pi_t^* = \sigma^{-1}\nu_t = \sigma^{-1}\frac{\theta_0}{1-\gamma} X_t^* = \frac{\mu - r_0}{\sigma^2(1-\gamma)} X_t^*.$$

This solution is of course is the Merton ratio we referred to earlier.

We were able to obtain an analytical expression in the present example because of the specialized assumptions. In more general cases, we will not be able to obtain an analytical expression for ν and we will have to rely on numerical methods.

3. Adapting the quasi-Monte Carlo method

In this section we explain the Imai and Tan adaption of the quasi-Monte Carlo method. As noted earlier their method has computational advantages over standard QMC in evaluating stochastic integrals. First, we introduce the concept of effective dimension and discuss the connection between the effective dimension of a problem and the efficiency of a QMC implementation. Then we describe the LT method of Imai and Tan and explain the key ideas in this method.

It is well known that the classical Monte Carlo (MC) method relies on random sequences and attains a convergence rate of $O(N^{-1/2})$ for sample size N for many[11] applications. The quasi-Monte Carlo (QMC) method, on the other hand, achieves a convergence rate of $O(N^{-1}\log^d N)$ in dimension d. While QMC converges asymptotically at a much faster rate than MC, it is important to appreciate their relative performance for the sample sizes that are feasible in practice, especially in high dimensional applications. It turns out that for large d and for practical values of N, the factor $\log^d N$ in the rate of convergence of QMC is not negligible. The deterioration of QMC with increasing dimensions is well documented. In the following subsection, we provide additional insight on the role of dimensions on the performance of QMC. We distinguish between the effective and the nominal dimensions, and their impact on QMC. Motivated by these discussions, Section 3.2 describes a general technique of recovering the superior rate of convergence of QMC. This is the linear transformation proposed by Imai and Tan (2007).

3.1. Effective dimension and nominal dimension

We first provide some background on the ANOVA (analysis of variance) of a function. For a detailed description of the ANOVA decomposition of a function, see Efron and Stein (1981). The set $\mathcal{A} = \{1, 2, \ldots, d\}$ denotes the coordinate axes of $[0, 1)^d$. Then for any subset $u \subseteq \mathcal{A}$, we define $|u|$ as its cardinality and $\mathcal{A} - u$ as its complementary set. A generic point of $[0, 1)^d$ is written as $\boldsymbol{x} = (x_1, \ldots, x_d)'$ and $\boldsymbol{x_u}$ denotes the $|u|$-vector of components x_j for $j \in u$.

Now consider an integrand $f(\boldsymbol{x})$ where $\boldsymbol{x} \in [0, 1)^d$. Under the (mild) condition that f is a square integrable function, the ANOVA decomposition expresses the integrand f as a sum of 2^d additive functions as follows:

$$f(\boldsymbol{x}) = \sum_{u \subseteq \{1,2,\ldots,d\}} f_u(\boldsymbol{x}), \tag{3.1}$$

where the function f_u, which depends only on the components of $\boldsymbol{x}$ in the set u, is defined recursively by

$$f_u(\boldsymbol{x}) = \int_{[0,1)^{\mathcal{A}-u}} f(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}_{\mathcal{A}-u} - \sum_{v \subsetneq u} f_v(\boldsymbol{x}), \tag{3.2}$$

with the usual convention that $f_\emptyset(\boldsymbol{x}) = \int_{[0,1)^d} f(\boldsymbol{x})\mathrm{d}\boldsymbol{x} = I(f)$. The ANOVA decomposition is orthogonal in that $\int f_u(\boldsymbol{x}) f_v(\boldsymbol{x})\mathrm{d}\boldsymbol{x} = 0$, for $u \neq v$. Let $\sigma^2(f)$ and $\sigma_u^2(f)$ denote the variance of f and f_u, respectively. Formally, these two quantities are defined as $\sigma^2(f) = \int_{[0,1)^d}(f(\boldsymbol{x}) - I)^2 \mathrm{d}\boldsymbol{x}$, and $\sigma_u^2(f) = \int_{[0,1)^u}[f_u(\boldsymbol{x})]^2 \mathrm{d}\boldsymbol{x}$, for $|u| > 0$, respectively. Also, $\sigma_\emptyset^2 = 0$ and an alternate way of computing $\sigma^2(f)$ is via $\sigma^2(f) = \sum_{|u|>0} \sigma_u^2(f)$. Furthermore, based on the above variance decomposition we define D_u as the total variance corresponding to the subset u; i.e.,

$$D_u = \sum_{v \subseteq u} \sigma_v^2(f). \tag{3.3}$$

In the context of QMC, it is important to distinguish between the notion of *nominal* and *effective* dimensions of a function. When a function $f(\boldsymbol{x})$ depends on d variables, it is typically said to have a nominal dimension d whereas its effective dimension can be quite small. Motivated by the ANOVA decompositions, Caflisch

[11] As mentioned in the introduction to this paper, Detemple et al. (2005) show that a lower convergence rate will obtain in the CGZ applications if the Euler approximation is involved. However it is convenient to frame the convergence discussion in terms of the classical convergence rate.

et al. (1997) introduced two definitions of effective dimension: The effective dimension of f, in the superposition sense, is the smallest integer d_S such that $\sum_{|u|\le d_S}\sigma_u^2(f)\ge p\sigma^2(f)$. The effective dimension of f, in the truncation sense, is the smallest integer d_T that satisfies $D_{\{1,2,\ldots,d_T\}}(f)\ge p\sigma^2(f)$. The critical level p is usually close to one. The truncation dimension indicates the number of important variables which essentially capture the given function f. The superposition dimension measures to what extent the low-order ANOVA terms dominate the function.

Consider a QMC point set $P=\{\boldsymbol{x}_i\}_{i=1}^N, \boldsymbol{x}_i\in[0,1)^d$ is used to approximate the value of $I(f)$. Then the error associated with such a set of point P satisfies the following bound:

$$\left|\frac{1}{N}\sum_{i=1}^{N}f(\boldsymbol{x}_i)-I(f)\right|\le\sum_{|u|>0}D_{N,u}(P_u)\|f_u\| \tag{3.4}$$

where P_u is the projection of the point set P on $[0,1)^{|u|}$, $D_{N,u}(P_u)$ is the discrepancy corresponding to P_u of N points, and $\|f_u\|$ is the variation of f_u. See Hickernell (1998) for various suitable choices of discrepancy and variation.

The bound in (3.4) explicitly associates the QMC error with the uniformity of all the projections P_u as well as all the low-dimensional structures f_u. The significance of this error bound is that QMC relies on the low discrepancy sequences which are constructed to have greater uniformity than random sequences. However, for finite number of points, such "greater" uniformity is not preserved for all dimensions and for all projections. It is well known that (see Morokoff and Caflisch (1994)) as dimension increases, the uniformity of low discrepancy sequences decreases. Nevertheless, as argued in Wang and Fang (2003) QMC can still be more effective than MC, particularly on problems with low truncation dimension. This can be justified by decomposing the bound (3.4) as

$$\left|\frac{1}{N}\sum_{i=1}^{N}f(\boldsymbol{x}_i)-I(f)\right|\le\sum_{u\subseteq\{1,\ldots,d_T\}}D_{N,u}(P_u)\|f_u\|+\sum_{u\cap(\mathcal{A}-\{1,\ldots,d_T\})\neq\emptyset}D_{N,u}(P_u)\|f_u\|, \tag{3.5}$$

assuming the truncation dimension of f is d_T. Note the role of d_T in the above representation. When d_T is small, the discrepancies of all the low-dimensional projections of low discrepancy point sets P_u are much smaller relative to those of the random point sets. This implies that the first summation in (3.5) is much smaller for QMC than for MC. As we further increase the dimension, the uniformity of the low discrepancy point sets deteriorates, which implies that for higher values of $\|u\|$, $D_{N,u}(P_u)$ of QMC can be larger than MC. Yet the second summation in (3.5) can be insignificant since the quantities $\|f_u\|$ are often small. The overall effect is that for QMC, if terms on the right-hand side of (3.5) are small the error bound is low. In this case QMC can be more effective than MC when f has low truncation dimension.

3.2. The linear transformation (LT) construction

We have just argued there is a strong connection between the effective dimension and the efficiency of QMC. In particular, as long as the truncation dimension of a problem is low, we can expect QMC to outperform MC, even though its nominal dimension can be very large. This also suggests that if a problem can be transformed into an *equivalent* problem with a lower truncation dimension, then QMC will be more efficient in evaluating the transformed problem. This is the key idea of the linear transformation (LT) construction proposed by Imai and Tan (2007). They show that their method significantly recovers the superior rate of convergence of QMC, even for very high nominal dimensional applications. We now provide a brief discussion of the standard application of QMC and use this to describe the LT construction.

Let $N_d(\boldsymbol{\mu},\boldsymbol{\Sigma})$ denote a d-variate normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Suppose we are interested in computing $\mathrm{E}[f(\mathbf{Z})]$, where f is a differentiable function and $\mathbf{Z}\sim N_d(\mathbf{0},\boldsymbol{I})$ with d-dimensional zero vector $\mathbf{0}$ and d-dimensional identify matrix $\boldsymbol{I}$. The procedure for estimating $\mathrm{E}[f(\mathbf{Z})]$ using QMC can be summarized as follows:

Step 1: Generate a point from a d-dimensional low discrepancy sequence. Let $\boldsymbol{x}=(x_1,\ldots,x_d)'\in[0,1)^d$ be the point generated.
Step 2: Generate normal vector $\boldsymbol{\varepsilon}=(\varepsilon_1,\ldots,\varepsilon_d)'$ from $\boldsymbol{x}$ via the inverse transformation; i.e., $\epsilon_i=\Phi^{-1}(x_i)$, $i=1,\ldots,d$, where $\Phi(\cdot)$ is the cumulative standard normal distribution.
Step 3: Compute $f(\boldsymbol{\varepsilon})$. This is the value of f corresponding to the simulated normal vector $\boldsymbol{\varepsilon}$.
Step 4: Repeat Steps 1–3 to obtain a large number of realizations of f and the QMC estimate of the expectation is the average of these realizations.

The essence of the LT construction is to recognize that for any d-dimensional orthogonal matrix $\boldsymbol{A}$, if $\boldsymbol{\varepsilon}\sim N_d(\mathbf{0},\boldsymbol{I})$, then $\boldsymbol{z}=\boldsymbol{A}\boldsymbol{\varepsilon}\sim N_d(\mathbf{0},\boldsymbol{I})$ since $\boldsymbol{A}'\boldsymbol{A}=\boldsymbol{I}$. This implies that we can produce another consistent QMC estimate of $\mathrm{E}[f(\mathbf{Z})]$ by simulating the realization in Step 3 using $f(\boldsymbol{A}\boldsymbol{\varepsilon})$, instead of $f(\boldsymbol{\varepsilon})$. The orthogonal matrix $\boldsymbol{A}$ is non-unique and hence to enhance QMC, $\boldsymbol{A}$ can be carefully selected so that the truncation dimension of the transformed $f(\boldsymbol{A}\boldsymbol{\varepsilon})$ is minimized.

Imai and Tan (2007) provide a systematic procedure for obtaining the matrix $\boldsymbol{A}$. This entails solving the following optimization problem:

$$\max_{\boldsymbol{A}_k\in\mathcal{R}e^d}\left(\left.\frac{\partial f(\boldsymbol{A}\boldsymbol{\varepsilon})}{\partial\varepsilon_k}\right|_{\boldsymbol{\varepsilon}=\hat{\boldsymbol{\varepsilon}}_k}\right)^2 \tag{3.6}$$

subject to $\|\boldsymbol{A}_k\|=1$ and $\langle\boldsymbol{A}_j^*,\boldsymbol{A}_k\rangle=0$, $j=1,\ldots,k-1$.

In the above algorithm, $\boldsymbol{A}_j$ denotes the j-th column of $\mathbf{A}$, $\langle\mathbf{a},\boldsymbol{b}\rangle$ denotes the inner product between vectors $\mathbf{a}$ and $\boldsymbol{b}$, and $\hat{\boldsymbol{\varepsilon}}_k=(1,\ldots,1,0\ldots,0)'$ denotes a d-dimensional vector with $k-1$ leading ones.

The above algorithm is carried out iteratively for $k=1,2,\ldots,d$. In the k-th optimization step, the objective is to determine the optimal $\boldsymbol{A}_k$ given the optimal columns $\boldsymbol{A}_j^*, j=1,\ldots,k-1$ that have already been determined in the previous iterations. The objective function in the algorithm can be interpreted as the variance contribution due to the k-th dimension and hence iteratively maximizing this quantity ensures that the truncation dimension of f under the prescribed transformation is minimized. Imai and Tan (2007) refer to this approach as the linear transformation (LT) construction. The power of the above procedure lies in its flexibility. It can be applied to an arbitrary function $f(\mathbf{Z})$. More importantly when the LT construction is combined with QMC, a significant increase in efficiency of QMC can be achieved. In the following section, we will demonstrate how to apply the LT construction for computing the optimal portfolio weights.

4. Implementation of the LT construction

In the last section, we described the LT construction which is a method of increasing the efficiency of QMC. In this section, we demonstrate how to adapt this algorithm to solve the optimal portfolio problem. We will use the example considered in Cvitanic et al. (2003) as the benchmark case. The following subsection describes the benchmark example and introduces some additional notation. The remaining two subsections describe how to implement the LT construction in this case.

4.1. Description of the Benchmark example

In this subsection, we describe the example studied by Cvitanic et al. (2003). We assume the economy has only one risky asset and one riskfree asset. The risky asset follows geometric Brownian motion with a constant drift μ_S and a constant volatility σ_S:

$$\frac{\mathrm{d}S_t}{S_t} = \mu_S \mathrm{d}t + \sigma_S \mathrm{d}W_t. \tag{4.1}$$

The money market (bank account) account process B_t follows

$$\mathrm{d}B_t = r_t B_t \mathrm{d}t, \tag{4.2}$$

where r_t is the stochastic interest rate and has the following dynamics

$$\mathrm{d}r_t = \mu_r(r_t)\,\mathrm{d}t + \sigma_r(r_t)\,\mathrm{d}W_t. \tag{4.3}$$

In this example, we set $\mu_r(r_t) = \kappa_r(\bar{r} - r_t)$ and $\sigma_r(r_t) = \sigma_r\sqrt{r_t}$ so that (4.3) becomes the Cox, Ingersoll and Ross (CIR) (Cox et al., 1985) interest rate model. Similarly, we assume that the market price of risk θ follows the diffusion process

$$\mathrm{d}\theta_t = \mu_\theta(\theta_t)\,\mathrm{d}t + \sigma_\theta(\theta_t)\,\mathrm{d}W_t, \tag{4.4}$$

where $\mu_\theta(\theta_t) = \kappa_\theta(\bar{\theta} - \theta_t)$ and $\sigma_\theta(\theta_t) = \sigma_\theta$. Hence the market price of risk follows an Ornstein Uhlenbeck process. Note that S_t, r_t and θ_t are all driven by the same Brownian motion. This is a rather simple example but the method we describe is applicable to problems involving several Brownian motions.

We assume the investor has the power utility with parameter γ and wishes to maximize the expected utility of terminal wealth. This implies that the investor is faced with the optimization problem as formulated in (2.5).

It is convenient to introduce some new notation. By defining Θ_t as

$$\Theta_t = \frac{1}{2}\int_0^t \theta_s^2 \mathrm{d}s + \int_0^t \theta_s \mathrm{d}W_t$$

and R_t as

$$R_t = \int_0^t r_s \mathrm{d}s,$$

we express the state price density as

$$\xi_t = \xi_0 \exp\left(-\frac{1}{2}\int_0^t \theta_s^2 \mathrm{d}s - \int_0^t \theta_s \mathrm{d}W_s\right) = \mathrm{e}^{-\Theta_t} \tag{4.5}$$

since $\xi_0 = 1$. Under the additional assumption of power utility, it is easy to verify that the optimal terminal wealth (2.13) becomes

$$X_T^* = \lambda^{\frac{1}{\gamma-1}} \mathrm{e}^{-\frac{(R_T+\Theta_T)}{\gamma-1}}, \tag{4.6}$$

while the optimal wealth at time t is given by

$$X_t^* = \lambda^{\frac{1}{\gamma-1}} \frac{1}{\mathrm{e}^{-(\Theta_t+R_t)}}\,\mathrm{E}\left[\mathrm{e}^{-\rho(\Theta_T+R_T)}|\mathcal{F}_t\right]. \tag{4.7}$$

Setting $t = 0$ in the above equation provides us with an expression for the Lagrange multiplier λ:

$$\lambda = \left\{\frac{x_0}{\mathrm{E}\left[\mathrm{e}^{-\rho(R_T+\Theta_T)}|\mathcal{F}_0\right]}\right\}^{\gamma-1}. \tag{4.8}$$

Substituting the above result into (4.7) leads to

$$X_t^* = \frac{x_0}{\mathrm{e}^{-(\Theta_t+R_t)}} \frac{\mathrm{E}\left[\mathrm{e}^{-\rho(\Theta_T+R_T)}|\mathcal{F}_t\right]}{\mathrm{E}\left[\mathrm{e}^{-\rho(\Theta_T+R_T)}|\mathcal{F}_0\right]}. \tag{4.9}$$

To conclude this subsection, let us summarize the simulation procedure for estimating the optimal portfolio π_t^* at time t. This involves the following three steps:

Step 1: Estimate the Lagrange multiplier from (2.14). For our base case example, we use (4.8).

Step 2: Estimate ν_t using either the two-tier simulation approach of (2.17) or the direct method of (2.19). For our numerical examples, we use (2.19), together with (4.6) and (4.9) and the estimated Lagrange multiplier from the first stage.

Step 3: Compute π_t^* using (2.16).

4.2. Computation of λ using the LT method

Recall that before deriving the optimal portfolio weight, we first need to estimate the Lagrange multiplier. This involves estimating the following expectation (see (4.8))

$$\mathrm{E}\left[\mathrm{e}^{-\rho(R_T+\Theta_T)}\,\middle|\,\mathcal{F}_0\right].$$

Both R_T and Θ_T are stochastic integrals and to numerically evaluate these integrals, one approach is to generate a large number of sample paths corresponding to the interest rate and the market price of risk. We achieve this by discretizing the respective diffusion processes. For example given the processes (4.3) and (4.4) with initial values r_0 and θ_0, and using the Euler discretization, each sample path of the interest rate and the market price of risk can be simulated recursively via

$$r_n = r_{n-1} + \mu_r(r_{n-1})\,\Delta t + \sigma_r(r_{n-1})\sqrt{\Delta t} z_n, \tag{4.10}$$

$$\theta_n = \theta_{n-1} + \mu_\theta(\theta_{n-1})\,\Delta t + \sigma_\theta(\theta_{n-1})\sqrt{\Delta t} z_n, \tag{4.11}$$

for $n = 1, \ldots, d$. Here d denotes the number of time steps, $\Delta t = \frac{T}{d}$ is the size of a time step, and $\boldsymbol{z} = (z_1, \ldots, z_d)' \sim N(\mathbf{0}, \boldsymbol{I})$. Note again that both the interest rate and the market price of risk are driven by the same random factor. Furthermore the trajectories of $\{r_n, n = 1, \ldots, d\}$ and $\{\theta_n, n = 1, \ldots, d\}$ require d standardized independently distributed normal samples. For each simulated sample path, the corresponding realizations of R_T and Θ_T are approximated by

$$R_T = \int_0^T r_s \mathrm{d}s \approx \sum_{n=0}^{d-1} r_n \Delta t \tag{4.12}$$

$$\begin{aligned}\Theta_T &= \frac{1}{2}\int_0^T \theta_s^2 \mathrm{d}s + \int_0^T \theta_s \mathrm{d}W_s \\ &\approx \sum_{n=0}^{d-1}\left(\frac{1}{2}\theta_n^2 \Delta t + \theta_n \sqrt{\Delta t} z_{n+1}\right).\end{aligned} \tag{4.13}$$

We now let $Y(z_1, \ldots, z_d) \equiv \mathrm{e}^{-\rho(R_T+\Theta_T)}$ to emphasize the fact that the function Y depends explicitly on d normally distributed samples $(z_1, \ldots, z_d)$. As described in the last section, an efficient estimator of $\mathrm{E}[Y]$ can be obtained by using the LT method which generates $(z_1, \ldots, z_d)$ from $(\varepsilon_1, \ldots, \varepsilon_d)$ via an optimal orthogonal matrix $\boldsymbol{A}$. In the remaining subsection, we explain how to implement the LT construction by considering the optimization algorithm (3.6) with $k = 1$.

First we need to derive the partial derivative $\frac{\partial Y}{\partial \varepsilon_l}$, $l = 1, \ldots, d$. This is given by

$$\frac{\partial Y}{\partial \varepsilon_l} = -\rho \mathrm{e}^{-\rho(R_T+\Theta_T)}\left(\frac{\partial R_T}{\partial \varepsilon_l} + \frac{\partial \Theta_T}{\partial \varepsilon_l}\right). \tag{4.14}$$

It follows from (4.12) that

$$\frac{\partial R_T}{\partial \varepsilon_l} \approx \sum_{n=0}^{d-1} \frac{\partial r_n}{\partial \varepsilon_l} \Delta t$$

so that for $n = 1, \ldots, d-1$, $\frac{\partial r_n}{\partial \varepsilon_l}$ can be defined recursively as

$$\begin{aligned}\frac{\partial r_n}{\partial \varepsilon_l} &= \frac{\partial r_{n-1}}{\partial \varepsilon_l}\left(1 + \frac{\partial \mu_r(r_{n-1})}{\partial r_{n-1}}\Delta t + \frac{\partial \sigma_r(r_{n-1})}{\partial r_{n-1}}\sqrt{\Delta t} z_n\right) \\ &\quad + \sigma_r(r_{n-1})\sqrt{\Delta t} a_{n,l} \\ &= \frac{\partial r_{n-1}}{\partial \varepsilon_l}\beta_r(r_{n-1}, z_t) + \sigma_r(r_{n-1})\sqrt{\Delta t} a_{n,l}, \end{aligned} \tag{4.15}$$

where $a_{i,j}$ denotes the (i, j) entry of the matrix $\boldsymbol{A}$ and

$$\beta_r(r_{n-1}, z_n) = \left(1 + \frac{\partial \mu_r(r_{n-1})}{\partial r_{n-1}}\Delta t + \frac{\partial \sigma_r(r_{n-1})}{\partial r_{n-1}}\sqrt{\Delta t} z_n\right).$$

Note that $\frac{\partial r_0}{\partial \varepsilon_l} = 0$ and $\frac{\partial z_n}{\partial \varepsilon_l} = a_{n,l}$ for $n, l = 1, \ldots, d$. In our example with CIR interest rate dynamics, we have $\frac{\partial \mu_r(r_n)}{\partial r_n} = -\kappa_r$ and $\frac{\partial \sigma_r(r_n)}{\partial r_n} = -\frac{1}{2}\sigma_r r_n^{-\frac{1}{2}}$.

In the same way, it follows from (4.13) that

$$\begin{aligned}\frac{\partial \Theta_T}{\partial \varepsilon_l} &\approx \sum_{n=0}^{d-1}\left\{\frac{\partial \theta_n}{\partial \varepsilon_l}\left(\theta_n \Delta t + \sqrt{\Delta t} z_{n+1}\right) + \theta_n \Delta t a_{n+1,l}\right\} \\ &= \sum_{n=0}^{d-1}\left\{k_1(\theta_n, z_{n+1})\frac{\partial \theta_n}{\partial \varepsilon_l} + k_2(\theta_n) a_{n+1,l}\right\} \end{aligned} \tag{4.16}$$

where $k_1(\theta_n, z_{n+1}) = \theta_n \Delta t + \sqrt{\Delta t} z_{n+1}$ and $k_2(\theta_n) = \theta_n \Delta t$.

By setting $\boldsymbol{\varepsilon} = (0, \ldots, 0)'$, it is easy to verify that both $\frac{\partial R_T}{\partial \varepsilon_l}$ and $\frac{\partial \Theta_T}{\partial \varepsilon_l}$ are linear functions of $a_{i,l}$; $i = 1, \ldots, d$. This implies that we can find a coefficient vector $\boldsymbol{b} = (b_1, \ldots, b_d)'$ that satisfies the following equation:

$$\begin{aligned}\left.\frac{\partial Y}{\partial \varepsilon_l}\right|_{\boldsymbol{\varepsilon}=(0,\ldots,0)'} &= -\rho \mathrm{e}^{-\rho(R_T(\boldsymbol{\varepsilon})+\Theta_T(\boldsymbol{\varepsilon}))}\left(\frac{\partial R_T}{\partial \varepsilon_l} + \frac{\partial \Theta_T}{\partial \varepsilon_l}\right)\Bigg|_{\boldsymbol{\varepsilon}=(0,\ldots,0)'} \\ &= \langle \boldsymbol{b}, \boldsymbol{A}_l \rangle. \end{aligned}$$

Consequently under the LT construction, the optimal $\boldsymbol{A}_1^*$ is the solution to the following optimization problem:

$$\max_{\boldsymbol{A}_1 \in \Re^d} \langle \boldsymbol{b}, \boldsymbol{A}_l \rangle^2 \quad \text{subject to } \|\boldsymbol{A}_1\| = 1.$$

This corresponds to (3.6) when we set $k = 1$. It follows from Theorem 1 of Imai and Tan (2007) that the optimal solution to the above optimization problem is

$$\boldsymbol{A}_1^* = \pm \frac{\boldsymbol{b}}{\|\boldsymbol{b}\|}.$$

This completes the procedure for optimizing the first column of $\boldsymbol{A}$. Subsequent columns of $\boldsymbol{A}$ can similarly be optimized by repeating algorithm (3.6) with other points of expansion.

4.3. Computation of v_0 using the LT method

In this subsection, we proceed to the second step of the simulation procedure which involves estimating v_t using the direct formulation (2.19). We are interested in determining v_0. It is convenient to denote, assuming $t = 0$, the function in the expectation of (2.19) as U; i.e.

$$\begin{aligned}U(z_1, \ldots, z_d) &= \frac{1}{\sqrt{\Delta t}}(H_{\Delta t,T} X_T^* - x_0) z_1 \\ &= \frac{1}{\sqrt{\Delta t}}(CV_T - x_0) z_1, \end{aligned} \tag{4.17}$$

where $C = \mathrm{e}^{r_0 \Delta t + (\rho-1)\delta T}\lambda^{\rho-1}$ is a constant and $V_T = \mathrm{e}^{-\rho(R_T+\Theta_T)+\Theta_{\Delta t}}$ is a function of $(z_1, \ldots, z_d)$.

As noted before, the key to the LT construction is the determination of the optimal orthogonal matrix $\boldsymbol{A}$. The LT construction can similarly be applied to the above function. To see this, let us revisit algorithm (3.6) with $k = 1$.

First note that for $l = 1, \ldots, d$,

$$\frac{\partial V_T}{\partial \varepsilon_l} = V_T\left\{-\rho\left(\frac{\partial R_T}{\partial \varepsilon_l} + \frac{\partial \Theta_T}{\partial \varepsilon_l}\right) + \theta_0\sqrt{\Delta t} a_{1,l}\right\}.$$

Then

$$\begin{aligned}\frac{\partial U}{\partial \varepsilon_l} &= \frac{1}{\sqrt{\Delta t}}\left[Cz_1\frac{\partial V_T}{\partial \varepsilon_l} + (CV_T - X_0) a_{1,l}\right] \\ &= \frac{1}{\sqrt{\Delta t}}\left[Cz_1 V_T\left\{-\rho\left(\frac{\partial R_T}{\partial \varepsilon_l} + \frac{\partial \Theta_T}{\partial \varepsilon_l}\right) + \theta_0\sqrt{\Delta t} a_{1,l}\right\}\right. \\ &\quad \left. + (CV_T - X_0) a_{1,l}\right] \\ &= \frac{1}{\sqrt{\Delta t}}\left[-\rho C V_T z_1\left(\frac{\partial R_T}{\partial \varepsilon_l} + \frac{\partial \Theta_T}{\partial \varepsilon_l}\right)\right. \\ &\quad \left. + \left\{CV_T\left(\theta_0\sqrt{\Delta t} z_1 + 1\right) - X_0\right\} a_{1,l}\right]. \end{aligned}$$

By setting $\boldsymbol{\varepsilon} = (0, \ldots, 0)'$, we have

$$\left.\frac{\partial U}{\partial \varepsilon_l}\right|_{\boldsymbol{\varepsilon}=(0,\ldots,0)'} = \frac{1}{\sqrt{\Delta t}}(CV_T - X_0) a_{1,l}.$$

Therefore it is easy to see that the optimal solution is $\boldsymbol{A}_1^* = (1, 0, \ldots, 0)'$ and this completes the procedure for optimizing the first column of $\boldsymbol{A}$. Other columns of $\boldsymbol{A}$ can similarly be optimized by iteratively applying algorithm (3.6).

As pointed out earlier the LT construction is general and flexible. We can also apply this technique to estimate v_t from (2.17) which is based on the two-tier simulation. This is the procedure recommended by Cvitanic et al. (2003). Because of the two-tier simulation, we would need to re-optimize the orthogonal matrix $\boldsymbol{A}$ at each second tier simulation. Consequently this significantly increases the computational effort and hence when using the LT construction, it is more efficient to use the direct approach.

5. Numerical results

In this section, we assess the effectiveness of the various simulation techniques. We compare the efficiency of classical Monte Carlo, standard quasi Monte Carlo and our LT enhanced quasi-Monte Carlo. Section 5.1 provides the simulated results on estimating the optimal portfolio weights. We show that the LT method is superior to both standard QMC and classical MC in this example. Section 5.2 calculates the dimension reduction for the first ten dimensions under the LT approach and hence provides an attribution of the superior convergence of the LT approach.

5.1. Estimating optimal portfolio weights

In this subsection, we give the results of using different simulation methods for finding the optimal portfolio weights. We use the same example as Cvitanic et al. (2003) which we have described in Section 4.1. Table 1 provides the parameter values for our base case examples. Recall that the power utility function is well defined for $\gamma \leq 1, \gamma \neq 1$ and the relative risk aversion is given by $1-\gamma$. In our simulation studies, we have excluded any results for $\gamma \in (0, 1)$ as we experienced numerical instabilities for γ in this range. This finding is consistent with the results of Korn and Kraft (2004) who showed that the solution to the portfolio optimization problem can blow up for $\gamma \in (0, 1)$ for a problem very similar to ours.

Table 1
Base example parameter values

Parameter	Value
$\bar{r}$	0.0600
σ_r	0.0364
κ_r	0.0824
$\bar{\theta}$	0.0871
σ_θ	0.2100
κ_θ	0.6950
r_0	0.0600
θ_0	0.1000
σ_S	0.2000
γ	-1, -2, -5, and -10
T	1, 5, and 10 years

For each set of parameter values, we produced four simulation based results. The first two results, which we labeled as MC1 and MC2, were based on the classical Monte Carlo method. The optimal portfolio weights π_0^* were estimated directly from (2.19) with discretization time step $\Delta t = 1/100$ years. The key difference between the reported values MC1 and MC2 lies on how the standard errors of the estimates were estimated. The MC1 results are based on a single large number of simulations ($N = 2^{20} =$ 1,048,576) and the standard errors are estimated based on this single sequence of estimates. Because of the large number of the sample paths that have been simulated, this value provides a benchmark for the standard Monte Carlo method. They also show how much work is needed to obtain accurate results with this method.

The results correspond to MC2 are obtained by conducting the simulations in 30 independent batches. For each batch, 16,384 sample paths are generated and (2.19) is again used to provide an estimate for π_0^*. The overall estimate of π_0^* is obtained by averaging over the 30 independent batch averages. The standard error of the overall estimate is calculated based upon the 30 independent batch averages. Note that the total number of simulation runs is $30 \times 16{,}384$, which is slightly less than half of that based on MC1. The rationale for reporting MC2 is to provide us with a level playing field when we compare MC and QMC. This is because we need to carry out the simulation in batches in order to estimate the standard errors of the QMC estimates.

The third set of results, labeled QMC, correspond to the standard application of the randomized QMC together with scrambled Sobol' low discrepancy sequence and Latin Supercube Sampling. Since we have set $\Delta t = 1/100$, this implies that for the examples with $T = 1$ year we need a 100-dimensional low discrepancy sequence. Similarly, for $T = 5$ and 10 years, the respective nominal dimensions of the problems are 500 and 1000. The Latin Supercube Sampling proposed by Owen (1998) is a convenient way of creating a high-dimensional low discrepancy sequence by appropriately concatenating the lower dimensional low discrepancy sequences. Similar to MC2, we produce the QMC results based on 30 independent batches each with 16,384 sample paths.

The fourth set of results, which we denote by QMC-LT, are based on the LT construction with QMC. To compare the results between the standard QMC and LT-based QMC on a consistent basis, the same low discrepancy sequences are used for both of these methods. As explained in the earlier sections, the LT method requires a preliminary calculation of the optimal orthogonal matrix $\boldsymbol{A}$ in order to fully exploit the uniformity of the low discrepancy sequence. This introduces some additional computational work, especially for high-dimensional portfolio selection problems. As pointed out in Imai and Tan (2007), one way of alleviating the computational effort is by only optimizing the early dimensions of the orthogonal matrix. Consider a d-dimensional problem. Instead of optimizing the entire d columns of the orthogonal matrix, we can optimize only up to d^* columns by iteratively using (3.6). The remaining columns are then generated randomly as long as the resulting matrix is orthogonal. This translates into a significant reduction in computational effort when $d^* \ll d$. In our examples with $T = 1$, we only optimize the first 10 columns of the orthogonal matrix with the remaining 90 dimensions randomly generated. Similarly for $T = 5$ and $T = 10$, we optimize the orthogonal matrix up to 45 and 75 dimensions, respectively. The loss of efficiency from using such suboptimal matrices is likely to be negligible since under the LT-based QMC, the first few dimensions already capture most of the variation. We confirm this observation in the following subsection.

Table 2
Estimates of π_0^* based on different simulation methods

γ	CGZ	MC1	MC2	QMC	QMC-LT
$T = 1$ year (100 nominal dimensions)					
-1	0.252	0.2554	0.2457	0.2484	0.2541
		(0.0010)	(0.0050)	(0.0042)	(0.0007)
-2	0.175	0.1810	0.1689	0.1721	0.1793
		(0.0013)	(0.0065)	(0.0054)	(0.0009)
-5	0.110	0.1099	0.0954	0.0991	0.1077
		(0.0016)	(0.0079)	(0.0066)	(0.0008)
-10	0.059	0.0785	0.0631	0.0669	0.0762
		(0.0017)	(0.0085)	(0.0071)	(0.0008)
$T = 5$ years (500 nominal dimensions)					
-1	0.295	0.3059	0.3153	0.2866	0.3153
		(0.0037)	(0.0120)	(0.0151)	(0.0013)
-2	0.230	0.2513	0.2642	0.2265	0.2519
		(0.0048)	(0.0154)	(0.0189)	(0.0018)
-5	0.170	0.2006	0.2166	0.1704	0.1990
		(0.0059)	(0.0186)	(0.0223)	(0.0026)
-10	0.139	0.1786	0.1959	0.1462	0.1769
		(0.0063)	(0.0200)	(0.0237)	(0.0029)
$T = 10$ years (1000 nominal dimensions)					
-1	0.328	0.3572	0.3625	0.3520	0.3571
		(0.0052)	(0.0227)	(0.0198)	(0.0021)
-2	0.270	0.3170	0.3249	0.3110	0.3167
		(0.0064)	(0.0293)	(0.0252)	(0.0030)
-5	0.190	0.2796	0.2902	0.2728	0.2753
		(0.0074)	(0.0354)	(0.0303)	(0.0041)
-10	0.167	0.2635	0.2753	0.2563	0.2582
		(0.0079)	(0.0381)	(0.0326)	(0.0046)

The values in parentheses denote the standard errors of the estimates.

Table 2 reports our simulated results for $T \in \{1, 5, 10\}$ and $\gamma \in \{-1, -2, -5, -10\}$. Note that the time horizon T controls the nominal dimensions of the problem since we have fixed the discretization time step. Along with our results (MC1, MC2, QMC and QMC-LT), we also tabulate the corresponding estimates reported in Cvitanic et al. (2003).[12] These authors did not report individual standard errors. Recall that their values were estimated (2.17) using two-tier simulation based on the standard Monte Carlo method.

We summarize our results as follows:

- The QMC-LT estimates lie within the confidence limits of the MC1 estimates confirming their accuracy.
- Using the standard errors as a measure of efficiency, the standard QMC is just marginally better than those of the standard Monte Carlo (i.e. MC2) for the 100 dimensional examples with $T = 1$. However, as we increase the dimension to 500, there are cases where QMC performs worse than MC2. This clearly illustrates the potential problem with the standard

[12] Our results in Table 2 indicates that the numbers reported by CGZ become somewhat less accurate for higher values of risk aversion $(1 - \gamma)$ and longer times to maturity.

Table 3
Computational time (in seconds) for MC2, QMC and QMC-LT methods

T	Method	Setup time	Marginal executive time					
		LT	$N = 1024$		$N = 4096$		$N = 16{,}384$	
			LSS	Batch	LSS	Batch	LSS	Batch
	MC2	–	–	0.047	–	0.187	–	0.750
1	QMC	–	0.297	0.453	0.562	1.828	2.172	7.390
	QMC–LT	0.047	0.297	0.547	0.562	3.031	2.188	8.937
	MC2	–	–	0.234	–	0.937	–	3.766
5	QMC	–	0.328	2.281	0.719	9.031	4.531	36.860
	QMC–LT	4.375	0.312	4.422	0.766	18.046	4.517	70.781
	MC2	–	–	0.468	–	1.875	–	7.515
10	QMC	–	0.344	4.594	0.907	18.406	7.422	73.438
	QMC-LT	39.906	0.313	12.906	0.922	51.625	7.375	207.328

application of QMC. In high dimensions, the standard QMC does not always outperform classical Monte Carlo; see for example Akesson and Lehoczky (2000) who describe this phenomenon.

- It is of interest to note the loss of efficiency of the standard QMC with increasing dimensions disappears when the underlying method is combined with the LT construction. The efficiency gain, as measured by the ratio of the standard error of the MC2 estimate to the standard error of the QMC-LT estimate, is around 10 in most cases. For example, consider the 1000-dimensional problem with $\gamma = -2$, the standard error based on MC2 is 0.0293. The corresponding standard error based on QMC-LT is 9.8 times smaller while the standard error using QMC is only 1.2 times smaller. The superiority of QMC-LT can be attributed to its effectiveness in dimension reduction. Under the LT construction, the effective dimension of the problem becomes significantly lower, regardless of its nominal dimension (which are 100, 500 and 1000 in our examples).

The results in Table 2 clearly demonstrate the greater rate of convergence of the LT-based QMC method relative to MC and QMC. The greater efficiency of the proposed method, however, is achieved at the expense of the higher computational cost. To have a fair assessment of the method, it is important to take into consideration the additional computational time. Table 3 reports the CPU time (in seconds) for the various methods based on the computer platform Intel Xeon™ CPU3.60 GHz, 2.00 GB RAM. Note that we have divided the CPU time into "set up time" and "marginal execution time". The latter refers to the time it takes to generate one batch estimate for sample sizes $N = 1024, 4096, 16{,}384$. The column labeled "LSS" denotes the time required to Latin Supercube Sampling the scrambled Sobol low discrepancy sequence. This table indicates that the additional time it requires for the LT-based QMC, particularly in higher dimensions and with larger sample runs, can be quite significant. The LT-based QMC has the further disadvantage of requiring initial setup cost for determining the optimal orthogonal matrix $\boldsymbol{A}$, as reported in Table 3.

Table 3 illustrates the deficiency of the proposed QMC-LT method if we were to focus exclusively on the computational effort. A more appropriate comparison is to take into account of both computational effort and the reduction of variance. Suppose an algorithm A produces a variance σ_A^2 of the estimator in computational time t_A. Similarly for algorithm B which generates another pair (σ_B^2, t_B). One possible measure of gauging the efficiency of algorithm A relative to algorithm B is to compute the following efficiency ratio

$$\frac{\sigma_A^2 t_A}{\sigma_B^2 t_B}.$$

An efficiency ratio greater than one implies the algorithm B is more efficient than the algorithm A and vice versa if the ratio is less than one. Table 4 depicts the efficiency ratio of MC2 relative to QMC and QMC-LT, respectively. The ratios are calculated assuming 30 replications with each batch consists of 16,384 sample runs. After incorporating the computational budget, the efficiency ratio for the QMC is much smaller than one, indicating its inferiority relative to the MC estimates. The proposed LT-based QMC, on the other hand, is uniformly more efficient even though it introduces additional computational burden.

Table 4
Efficiency ratio of MC2 relative to QMC and QMC-LT, respectively, and assuming 30 replications with 16,384 sample run in each batch

γ	$T = 1$ year		$T = 5$ years		$T = 10$ years	
	QMC	QMC-LT	QMC	QMC-LT	QMC	QMC-LT
−1	0.11	3.44	0.06	4.25	0.12	4.07
−2	0.11	3.52	0.06	3.65	0.13	3.32
−5	0.11	6.57	0.06	2.55	0.13	2.59
−10	0.11	7.61	0.06	2.37	0.13	2.39

5.2. The effectiveness of dimension reduction

We now delve into the origins of the good performance of the QMC-LT method by analyzing the dimension reduction associated with each of the first ten dimensions for the different problems. Recall that for a given function f, its truncation dimension is the smallest integer d_T such that $D_{\{1,2,\ldots,d_T\}}(f) \geq p\sigma^2(f)$. For example if $p = 99\%$ and $d_T = 2$, then this implies that the first two dimensions capture more than 99% of the total variations, even though the nominal dimension of f can be very large. It was also argued that the efficiency of the QMC is intricately related to the effective dimension. The simulation results from the last subsection signifies the clear advantage of the LT-based QMC. In this subsection, we analyze the source of the good performance of the LT construction.

We analyze the efficiency by considering the following ratio

$$\frac{D_{\{1,2,\ldots,u\}}}{\sigma^2}. \tag{5.1}$$

We refer to this quantity as the cumulative explanatory ratio since it gives the proportion of the variance captured by the first u dimensions (i.e. $D_{\{1,2,\ldots,u\}}$) relative to the total variance (i.e. σ^2). In general it is not possible to provide an explicit expression for D_u (see (3.3)). However, as shown in Sobol' (2001) this quantity can be calculated (using notation from Section 3) using

$$D_u = \int_{[0,1)^{2d-|u|}} f(\boldsymbol{x}) f(\boldsymbol{x}_u, \boldsymbol{y}_{\mathcal{A}-u}) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{y}_{\mathcal{A}-u} - [I(f)]^2, \tag{5.2}$$

where $\boldsymbol{x} = (\boldsymbol{x}_u, \boldsymbol{x}_{\mathcal{A}-u})$ and $\boldsymbol{y} = (\boldsymbol{y}_u, \boldsymbol{y}_{\mathcal{A}-u})$. The above representation allows us to estimate D_u using simulation methods.

Table 5 depicts the estimated proportions (in percentage) for $T \in \{1, 5, 10\}$, $d \in \{2, 4, 6, 8, 10\}$, and for both standard and LT-based simulation approaches. These values are estimated using the Monte Carlo method with 100,000 simulation trials. Recall that

Table 5
Estimates of the ratio $D_{\{1,2,\ldots,d\}}/\sigma^2$ (in percentage) for $T \in \{1, 5, 10\}$, $\gamma \in \{-1, -2, -5, -10\}$ and based on "Standard" and "LT" methods

d	$\gamma = -1$		$\gamma = -2$		$\gamma = -5$		$\gamma = -10$	
	Standard	LT	Standard	LT	Standard	LT	Standard	LT
$T = 1$ year (100 nominal dimensions)								
2	1.2	95.3	0.5	95.1	0.3	72.4	0.3	72.3
4	2.0	98.0	1.4	97.9	1.2	97.7	1.2	97.7
6	3.5	98.0	2.9	98.0	2.7	98.1	2.7	98.1
8	4.9	98.1	4.3	98.0	4.2	98.3	4.1	98.3
10	6.1	98.2	5.5	98.1	5.3	98.3	5.3	98.3
$T = 5$ years (500 nominal dimensions)								
2	0.6	72.3	0.5	63.9	0.5	57.5	0.5	56.8
4	0.1	97.7	0.0	86.0	0.0	89.0	−0.1	88.7
6	−0.5	98.1	−0.5	96.7	−0.5	94.7	−0.4	94.6
8	0.2	98.3	0.1	98.5	0.1	97.9	0.1	97.7
10	−0.4	98.3	−0.5	98.7	−0.6	98.6	−0.6	98.6
$T = 10$ years (1000 nominal dimensions)								
2	0.7	41.1	0.6	39.3	0.6	38.1	0.6	37.4
4	1.2	61.7	1.2	60.0	1.2	67.1	1.2	66.3
6	0.6	83.6	0.5	82.3	0.5	80.9	0.5	80.3
8	0.0	94.5	0.0	93.9	−0.1	92.1	−0.1	91.8
10	1.0	96.9	1.0	96.5	1.0	93.6	1.0	93.3

the LT construction involves optimizing the orthogonal matrix $\boldsymbol{A}$ while the standard simulation approach can be considered as a special case of LT construction by setting $\boldsymbol{A} = \boldsymbol{I}$. Note that some of the estimated values are negative. This can be attributed to the statistical sampling errors as the true values for these cases are close to zero.

The immediate conclusion is that the numerical results present overwhelming evidence favoring the transformed problems based on the LT construction over the original problems. The LT construction is remarkably effective in reducing the truncation dimension of the original problem. For example, consider the case with nominal dimensions of 100 (i.e. $T = 1$). The upper panel of Table 5 indicates that, under the LT construction, the first four dimensions already contribute more than 97% of the total variation. If we were to set the confidence level p at 97%, the truncation dimension of these examples is no more than four, for the ranges of γ considered here. This is in sharp contrast to the standard implementation of the optimal portfolio problems where the cumulative explanatory of the first four dimensions is no more than 2%. Even if we were to expand the set to ten dimensions, the standard simulation method still captures no more than 6.1% of the total variance. Consequently in the truncation sense, the effective dimension of the original problem is significantly higher than the effective dimension of the transformed problem. This largely explains the greater efficiency of the LT-based QMC that we observed in the last subsection. This also justifies the use of the suboptimal orthogonal matrix $\boldsymbol{A}$.

The cumulative explanatory ratio also provides a useful way of quantifying the difficulty of a problem. In particular, it could be used to predict the performance of classical QMC. Notice that as we increase the time horizon T, the proportion of the variance captured by the same number of dimensions declines. Similarly as we decrease γ from -1 to -10, there is also a slight deterioration of the ratios. This suggests that the problem becomes increasingly more challenging with increasing T and decreasing γ. This is consistent with our simulation results from the last subsection; the estimates of the standard errors are progressively larger. However, it should be emphasized for the example with $T = 10$ and $\gamma = -10$, even though its nominal dimension is 1000, the cumulative ratio up to ten dimensions is well over 90% for the LT construction. This compares favorably to the standard approach which is only capable of capturing about 1% of the total variability.

6. Summary

This paper discussed an enhanced numerical procedure to solve the portfolio selection problem. We showed how to improve the efficiency of simulation based methods using the approach proposed by Cvitanic et al. (2003). We demonstrated that the Linear Transformation method proposed by Imai and Tan (2007) significantly improves the efficiency of quasi-Monte Carlo in this connection. Specifically the Linear Transformation method achieves standard errors that are about one tenth of those obtained by standard quasi-Monte Carlo methods. It would be interesting to apply this method to the procedure advocated by Detemple et al. (2003). Based on standard Monte Carlo methods, Detemple Garcia and Rindisbacher conclude that their method is more efficient than the CGZ procedure. It seems probable the approach developed in this paper will lead to significant efficiency improvements over the standard Monte Carlo method in this case as well. This is left for future research.

Acknowledgements

The authors acknowledge the constructive comments of Raman Uppal and Marcel Rindisbacher and they thank Martina Wilhelm for research assistance. They are grateful to two anonymous referees for their constructive comments. Phelim Boyle acknowledges support from the Natural Sciences and Engineering Research Council of Canada. Junichi Imai is grateful for research supports from the Institute for Quantitative Finance and Insurance at the University of Waterloo and from Grant-in-Aid for Scientific Research in Japan Society for the Promotion of Science. Ken Seng Tan acknowledges funding from the Canada Research Chairs Program and the Natural Sciences and Engineering Research Council of Canada.

References

Acworth, P., Broadie, M., Glasserman, P., 1998. A comparison of some Monte Carlo and quasi-Monte Carlo methods for option pricing. In: Niederreiter, H., et al. (Eds.), Monte Carlo and quasi-Monte-Carlo Methods 1996: Proceedings of a conference at the University of Salzburg. In: Lecture notes in statistics, vol. 127. Springer Verlag, New York.

Akesson, F., Lehoczky, J.P., 2000. Path generation for quasi-Monte Carlo simulation of mortgage-backed securities. Management Science 46 (9), 1171–1187.

Arrow, K., 1952. Le rôle des valeurs boursières pour la repartition meilleure des risques. Econometrie. Colloq. Internat. du CNRS, Paris 40, 41–47. (with discussion, 47–48). English translation in Review of Economic Studies 31 (1964) 91–96.
Bajeux-Besnainou, I., Jordan, J., Portait, R., 2001. An asset allocation puzzle: Comment. American Economic Review 91, 1170–1179.
Boyle, P.P., Tian, W., Guan, F., 2002. The Riccati equation in mathematical finance. Journal of Symbolic Computation 33 (3), 343–356.
Brennan, M., Schwartz, E., Lagnado, R., 1997. Strategic asset allocation. Journal of Economic Dynamics and Control 21, 1377–1403.
Brennan, M., Xia, Y., 2000. Stochastic interest rates and the bond stock mix. The European Finance Review 4, 197–201.
Brandt, M.W., Goyal, A., Santa Clara, P., Stroud, J.R., 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Review of Financial Studies 18, 831–873.
Caflisch, R.E., Morokoff, W., Owen, A., 1997. Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension. Journal of Computational Finance 1 (1), 27–46.
Canner, N., Mankiw, N.G., Weil, D.N., 1997. An asset allocation puzzle. American Economic Review 87, 181–191.
Cox, J.C., Huang, Chi-fu, 1989. Optimum consumption and portfolio policies when asset prices follow a diffusion process. Journal of Economic Theory 49, 33–83.
Cox, J.C., Ingersoll, J., Ross, S., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–408.
Cvitanic, J., Goukasian, L., Zapatero, F., 2003. Monte Carlo computation of optimal portfolios in complete markets. Journal of Economic Dynamics and Control 27, 971–986.
Detemple, J.B., Garcia, R., Rindisbacher, M., 2003. A Monte Carlo method for optimal portfolios. Journal of Finance 58 (1), 401–446.
Detemple, J.B., Garcia, R., Rindisbacher, M., 2005. Asymtptotic properties of Monte Carlo estimation of derivatives. Management Science 51 (11), 1657–1765.
Efron, B., Stein, C., 1981. The jackknife estimate of variance. Annals of Statistics 9, 586–596.
Fournié, E., Lasry, J.-M., Lebouchoux, J., Lions, P.-L., Touzi, M., 1999. Applications of Malliavin calculus to Monte Carlo methods in finance. Finance and Stochastics 3 (4), 391–412.
Fournié, E., Lasry, J.-M., Lebouchoux, J., Lions, P.-L., Touzi, M., 2001. Applications of Malliavin calculus to Monte Carlo methods in finance II. Finance and Stochastics 5 (2), 201–236.
Hickernell, F.J., 1998. A generalized discrepancy and quadrature error bound. Mathematics of Computation 67 (221), 299–322.
Imai, J., Tan, K.S., 2007. A general dimension reduction method for derivative pricing. Journal of Computational Finance 10 (2), 129–155.
Joy, C., Boyle, P.P., Tan, K.S., 1996. Quasi-Monte Carlo methods in numerical finance. Management Science 42 (6), 926–938.
Karatzas, I., Lehoczky, J.P., Shreve, S., 1987. Optimum portfolio and consumption decisions for a small investor on a finite horizon. SIAM Journal of Control and Optimization 25, 1557–1586.
Karatzas, I., Shreve, S., 1998. Methods of mathematical finance. Springer Verlag, New York.
Karatzas, I., Shreve, S., 1991. Brownian motion and stochastic calculus. Springer Verlag, New York.
Kim, T.S., Omberg, E., 1996. Dynamic nonmyopic portfolio behavior. Review of Financial Studies 9, 141–161.
Korn, R., Kraft, H., 2004. Counter examples and stability in continuous-time portfolio optimization. Mathematical Finance 14 (3), 403–414.
Lioui, A., 2007. The asset allocation puzzle is still a puzzle. Journal of Economic Dynamics and Control 31, 1185–1216.
Liu, J., 2007. Portfolio selection in stochastic environments. Review of Financial Studies 20 (1), 1–39.
Markowitz, H., 1952. Portfolio selection. Journal of Finance 7, 77–91.
Merton, R.C., 1969. Lifetime portfolio selection under uncertainty: The continuous time case. Review of Economics and Statistics 51, 247–257.
Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous time model. Journal of Economic Theory 3, 273–413.
Morokoff, W., Caflisch, R.E., 1994. Quasi random sequences and their discrepancy. SIAM Journal of Scientific Computing 15, 1251–1270.
Ninomiya, S., Tezuka, S., 1996. Toward real-time pricing of complex financial derivatives. Applied Mathematical Finance 3, 1–20.
Ocone, D., Karatzas, I., 1991. A generalized Clark representation formula, with application to optimal portfolios. Stochastics and Stochastics Reports 34, 187–220.
Oksendal, B., 1998. Stochastic differential equations: An introduction with applications, fifth ed. Springer Verlag, New York.
Owen, A., 1998. Latin supercube sampling for very high-dimensional simulations. ACM Transactions on Modeling and Computer Simulation 8 (1), 71–102.
Paskov, S., Traub, J.F., 1995. Faster valuation of financial derivatives. Journal of Portfolio Management 22 (1), 113–120.
Pliska, S., 1986. A stochastic calculus model of continuous trading. Mathematics of Operations Research 11, 371–383.
Sobol', I.M., 2001. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation 55, 271–280.
Tan, K.S., Boyle, P.P., 2000. Applications of randomized low discrepancy sequences to the valuation of complex securities. Journal of Economic Dynamics and Control 24, 1747–1782.
Wachter, J., 2002. Portfolio and consumption decisions under mean reverting returns: An exact solution for complete markets. Journal of Financial and Quantitative Analysis 37, 63–91.
Wang, X., Fang, K.-T., 2003. The effective dimension and quasi-Monte Carlo integration. Journal of Complexity 2, 101–124.