Insurance: Mathematics and Economics 76 (2017) 87–94

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

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

Efficient randomized quasi-Monte Carlo methods for portfolio market risk

Halis Sak [a,b,*], İsmail Başoğlu [c]

[a] International Business School Suzhou (IBSS), Xi'an Jiaotong-Liverpool University, Suzhou, China
[b] Department of Mathematical Sciences, Xi'an Jiaotong-Liverpool University, Suzhou, China
[c] School of Business Administration, Altınbaş University, Istanbul, Turkey

ARTICLE INFO

Article history:
Received October 2015
Received in revised form March 2017
Accepted 1 July 2017
Available online 19 July 2017

Keywords:
Risk management
Quasi-Monte Carlo
Importance sampling
Stratified sampling
t-copula

ABSTRACT

We consider the problem of simulating loss probabilities and conditional excesses for linear asset portfolios under the t-copula model. Although in the literature on market risk management there are papers proposing efficient variance reduction methods for Monte Carlo simulation of portfolio market risk, there is no paper discussing combining the randomized quasi-Monte Carlo method with variance reduction techniques. In this paper, we combine the randomized quasi-Monte Carlo method with importance sampling and stratified importance sampling. Numerical results for realistic portfolio examples suggest that replacing pseudorandom numbers (Monte Carlo) with quasi-random sequences (quasi-Monte Carlo) in the simulations increases the robustness of the estimates once we reduce the effective dimension and the impact of the non-smoothness of the integrands.

1. Introduction

Market risk management deals with the estimation of loss distribution of a portfolio of assets over a fixed time horizon. The widely used risk measures Value-at-Risk (VaR) and expected shortfall require accurate estimates of loss probability and conditional excess under a realistic model that captures dependence structure of the log-returns of multiple assets. As a flexible and accurate model for the logarithmic returns of stocks, we use the t-copula dependence structure and marginals following the generalized hyperbolic distribution (see Embrechts et al., 2002; Mashal et al., 2003; Prause, 1997; Glasserman et al., 2002). Furthermore, a more generalized version of the t-copula model is proposed by Demarta and McNeil (2005), and it has been also used in finance (see, e.g., Sun et al., 2008).

As there are no closed-form analytical results for loss probability and conditional excess under the t-copula model, we need to use a computational method like Monte Carlo simulation. In most cases, Monte Carlo simulation is a better alternative compared to other methods as it leads to error bounds on the estimated values. Due to the fact that Monte Carlo simulation has a slow convergence rate of $O(1/\sqrt{n})$, we need to increase the efficiency of the estimates using variance reduction techniques. There are papers proposing variance reduction methods in portfolio market risk estimation (see, e.g., Glasserman et al., 2002; Broadie et al., 2011; Başoğlu et al., 2013).

An alternative to Monte Carlo simulation is the quasi-Monte Carlo method (QMC), which uses low-discrepancy sequences instead of pseudorandom numbers. The rate of convergence of the quasi-Monte Carlo method is close to $O(1/n)$, which is faster than $O(1/\sqrt{n})$. However, an error bound under plain QMC cannot be estimated as low-discrepancy sequences do not have an i.i.d. property. Randomized quasi-Monte Carlo solves this problem by applying a randomization on low-discrepancy sequences.

Randomized quasi-Monte Carlo (RQMC) has been used in pricing extensively (see, e.g., Birge, 1995; Boyle et al., 1997). However, the application of RQMC to measure portfolio risk is rarely found in the literature (see Kreinin et al., 1998; Jin and Zhang, 2006). This can be explained by the fact that the integrand in risk management applications is a non-smooth function (e.g., indicator function) of high-dimensional random inputs. (As pointed out by Morokoff and Caflisch (1995), the performance of quasi-Monte Carlo method diminishes when the integrands are not smooth and high-dimensional.) To compute VaR using QMC, Kreinin et al. (1998) apply principal component analysis to reduce the dimensionality of the risk factor space. Jin and Zhang (2006) efficiently simulate VaR by smoothing the expectation of an indicator function via Fourier transformation and then applying RQMC.

The motivation of this paper is to investigate whether RQMC and variance reduction techniques can be combined efficiently for simulating loss probability and conditional excess under the t-copula model. In order to solve the problem of high-dimensionality

* Corresponding author at: International Business School Suzhou (IBSS), Xi'an Jiaotong-Liverpool University, Suzhou, China.
E-mail addresses: halis.sak@xjtlu.edu.cn (H. Sak), ismail.basoglu@kemerburgaz.edu.tr (İ. Başoğlu).

http://dx.doi.org/10.1016/j.insmatheco.2017.07.001
0167-6687/

of the integrands, we apply a linear transformation on the random input to reduce the effective dimension. Furthermore, we reduce the area of sampling regions where the simulation integrand is equal to zero using importance sampling. This also decreases the impact of the non-smoothness of the simulation integrand. We finally apply stratification to further improve the accuracy of the estimates. Numerical experiments illustrate the effectiveness of RQMC implementations of variance reduction methods over their Monte Carlo implementations. Although, the methodology of the paper is explained on market risk management under t-copula model, it is much more generally applicable to other fields like credit risk, insurance, and operational risk where t-copula models are widely used. We also implemented the methods in Google's TensorFlow *Python* package (Abadi et al., 2015) so that we are also able to give the execution time performance of the methods using GPU computing.

The rest of the paper is organized as follows. Section 2 describes the t-copula model for portfolio market risk. Section 3 presents background on efficient Monte Carlo simulation methods for estimating loss probability and conditional excess. Section 4 combines the RQMC method with importance sampling and stratified importance sampling for estimating loss probability and conditional excess. We present numerical results in Section 5.

2. Portfolio market risk in the t-copula model

The essence of any model of portfolio market risk is its ability to capture dependence among assets. In this section, we describe the widely used t-copula model (see, e.g., Glasserman et al., 2002; Sak et al., 2010).

We are interested in the distribution of losses caused by depreciation of stocks over a fixed time period. The following notation is used in order to represent this distribution.

- $D =$ the number of stocks in portfolio
- $w_d =$ the weight of the dth stock
- $X_d =$ the log-return of the dth stock
- $L = 1 - \sum_{d=1}^{D} w_d e^{X_d} =$ portfolio loss (initial value of portfolio is assumed to be equal to one).

We assume that we are given a portfolio of stocks with known weights $(w_1, \ldots, w_D)'$ and unknown future log-returns $(X_1, \ldots, X_D)'$. The main objective is to estimate loss probability $P(L > \tau)$, and conditional excess $E[L|L > \tau]$, especially at large values of τ.

To model dependence among stocks, we need to introduce dependence among the log-returns. The log-return vector $(X_1, \ldots, X_D)'$ of the stocks is assumed to follow a t-copula with ν degrees of freedom. The dependence is introduced through a multivariate t-vector $\mathbf{T} = (T_1, \ldots, T_D)'$ with ν degrees of freedom. Each log-return is represented as

$$X_d = c_d G_d^{-1} \left(F_\nu \left(T_d \right) \right), \quad d = 1, \ldots, D, \tag{1}$$

in which

- F_ν denotes the cumulative distribution function (CDF) of a t-distribution with ν degrees of freedom;
- G_d denotes the CDF of the marginal distribution of the dth log-return;
- c_d is the scaling factor for the dth log-return.

Through this representation, the dependence among the log-returns, X_d, can be determined by the correlations among T_d. Suppose, we are given the correlation matrix $\boldsymbol{\Sigma}$ of vector $\mathbf{T}$ and let $\boldsymbol{\Lambda} \in \mathbb{R}^{D \times D}$ be the lower triangular Cholesky factor of $\boldsymbol{\Sigma}$ satisfying $\boldsymbol{\Lambda}\boldsymbol{\Lambda}' = \boldsymbol{\Sigma}$. Then, $\mathbf{T}$ can be generated using

$$\mathbf{T} = \frac{\boldsymbol{\Lambda}\mathbf{Z}}{\sqrt{Y/\nu}}, \tag{2}$$

where $\mathbf{Z} = (Z_1, \ldots, Z_D)'$ is a standard multi-normal random vector and Y is an independent chi-squared random variable with ν degrees of freedom.

3. Efficient Monte Carlo simulation methods

In this section, we provide a brief summary of efficient Monte Carlo simulation algorithms designed for the estimation of portfolio market risk. Before that, we start with the implementation of the naive Monte Carlo algorithm. The naive identity of the loss probability is $P(L > \tau) = E\left[\mathbf{1}_{\{L>\tau\}}\right]$, where $\mathbf{1}\{.\}$ denotes the indicator of the event in braces. We are also interested in the conditional excess that can be represented as the ratio of two expectations

$$E[L|L > \tau] = \frac{E\left[L\mathbf{1}_{\{L>\tau\}}\right]}{P(L > \tau)} = \frac{E\left[L\mathbf{1}_{\{L>\tau\}}\right]}{E\left[\mathbf{1}_{\{L>\tau\}}\right]}, \tag{3}$$

which can be estimated in a single simulation run.

Each replication of the naive Monte Carlo algorithm follows the steps given below:

1. Generate D independent standard normal random variables, $\mathbf{Z} = (Z_1, \ldots, Z_D)'$, and a chi-squared random variable Y with ν degrees of freedom, independent of $\mathbf{Z}$.
2. Calculate $\mathbf{T}$ in (2).
3. Calculate the log-returns X_d, $d = 1, \ldots, D$ in (1).
4. Compute the portfolio loss $L = 1 - \sum_{d=1}^{D} w_d \exp(X_d)$ and return the estimators $\mathbf{1}_{\{L>\tau\}}$ and $L\mathbf{1}_{\{L>\tau\}}$.

3.1. Importance sampling

At a large threshold value τ, most of the replications of the naive simulation algorithm return the value zero for the estimator $\mathbf{1}_{\{L>\tau\}}$. To increase the number of replications that fall in the region $L > \tau$, importance sampling modifies the joint density of the random input.

Suppose $f(.)$ is the joint probability density function (PDF) of input variables $\mathbf{Z}$ and Y, and $\tilde{f}(.)$ is the modified density. Importance sampling uses the following identity to estimate the loss probability

$$E\left[\mathbf{1}_{\{L>\tau\}}\right] = \tilde{E}\left[\mathbf{1}_{\{L>\tau\}} \frac{f(\mathbf{Z}, Y)}{\tilde{f}(\mathbf{Z}, Y)}\right],$$

where $\tilde{E}$ is the expectation taken using the modified density $\tilde{f}(.)$.

Finding an importance sampling density that minimizes the variance of Monte Carlo estimators is a subtle problem. But it is possible to use the zero-variance IS function in search of an effective IS density (see, e.g., Glasserman et al., 1999 and Arouna, 2004). Glasserman et al. (1999) add the mode of the zero-variance IS function as a mean shift to the original density for pricing path-dependent options. Sak et al. (2010) utilize the same idea to find a close-to-optimal optimal parameters for simulating loss probabilities in the t-copula model of portfolio market risk.

Sak et al. (2010) add a mean shift vector with negative entries to the normal vector $\mathbf{Z}$ and use a scale parameter less than two for the chi-square (i.e., Gamma) random variable Y to construct the IS density. The shift vector and the scale value are selected so that the mode of the resulting IS density is equal to the mode of the zero-variance IS function. For more details on the determination of the IS parameters and implementation of the simulation algorithm, see Section 4 of Sak et al. (2010).