Comput Econ (2018) 52:55–77
https://doi.org/10.1007/s10614-017-9661-0

Time Series Simulation with Randomized Quasi-Monte Carlo Methods: An Application to Value at Risk and Expected Shortfall

Yu-Ying Tzeng[1] · **Paul M. Beaumont**[2] · **Giray Ökten**[1]

Accepted: 20 January 2017 / Published online: 4 February 2017

Abstract Quasi-Monte Carlo methods are designed to produce efficient estimates of simulated values but the error statistics of these estimates are difficult to compute. Randomized quasi-Monte Carlo methods have been developed to address this shortcoming. In this paper we compare quasi-Monte Carlo and randomized quasi-Monte Carlo techniques for simulating time series. We use randomized quasi-Monte Carlo to compute value-at-risk and expected shortfall measures for a stock portfolio whose returns follow a highly nonlinear Markov switching stochastic volatility model which does not admit analytical solutions for the returns distribution. Quasi-Monte Carlo methods are more accurate but do not allow the computation of reliable confidence intervals about risk measures. We find that randomized quasi-Monte Carlo methods maintain many of the advantages of quasi-Monte Carlo while also providing the ability to produce reliable confidence intervals of the simulated risk measures. However, the advantages in speed of convergence of randomized quasi-Monte Carlo diminish as the forecast horizon increases.

Keywords Quasi-Monte Carlo · Randomized Quasi-Monte Carlo · Time series simulation · Value-at-risk · Expected shortfall

✉ Paul M. Beaumont
beaumont@fsu.edu

Yu-Ying Tzeng
ytzeng@math.fsu.edu

Giray Ökten
okten@math.fsu.edu

[1] Department of Mathematics, Florida State University, Tallahassee, FL 32306-4510, USA

[2] Department of Economics, Florida State University, Tallahassee, FL 32306-2180, USA

1 Introduction

Monte Carlo (MC) simulation is a popular and flexible numerical tool used in many applications, from physical sciences and engineering to social sciences, including finance and economics. A typical application involves simulating an assumed underlying probability model and then estimating the probabilities of events that are of interest as a result of the simulation. MC is flexible and easy to adapt to many problems, but its convergence is relatively slow.

Quasi-Monte Carlo (QMC), often called the deterministic version of Monte Carlo, is a close variant of Monte Carlo that emerged to deal with the slow convergence issues. The main difference between the two methods is in the way they simulate the underlying model. Classical MC uses pseudorandom numbers, whereas QMC uses the so-called low-discrepancy or uniformly distributed mod 1 sequences in the simulation. Low-discrepancy sequences are sequences designed to have high "uniformity", or "evenness", in their domain, whereas pseudorandom sequences are designed to imitate the true behavior of random numbers.

QMC has enjoyed increasing popularity among researchers in finance and economics, for example, in pricing financial derivatives (Paskov and Traub 1995; Papageorgiou and Traub 1996; Ninomiya and Tezuka 1996; Joy et al. 1996; Boyle et al. 1997; Caflisch et al. 1997; Owen and Tavella 1997; Fishman et al. 1997; Galanti and Jung 1997; Ökten et al. 2006; Zhang et al. 2013) and in econometric models (Göggelmann et al. 2000; Li and Winker 2003). In addition, various financial and mathematical software such as R, GAUSS, *Mathematica*, and Matlab, now include implementations of quasi-Monte Carlo methods.

One reason for the increasing interest in QMC is that its estimates are, in general, more accurate than the estimates produced with the same computational effort by Monte Carlo. The low discrepancy sequences used in QMC yield convergence rates of the estimates to the true solution with $O(N^{-1}(\log N)^s)$ (where s is the "dimension" of the problem and N is the number of samples) whereas the MC convergence rate using pseudorandom sequences is $O(N^{-1/2})$. However, in many applications, researchers have observed rates close to $O(N^{-1})$ for QMC. We will not discuss the reasons for this better than theoretical rate of convergence in some problems, which involve concepts like effective dimension and decreasing importance of variables (Caflisch et al. 1997; Sloan and Woźniakowski 1998). For an in-depth discussion of low discrepancy sequences and quasi-Monte Carlo methods, we refer the reader to Niederreiter (1992).

A disadvantage of QMC is that the deterministic nature in the way that the low discrepancy sequences are constructed makes it impossible to estimate reliable error bonds for estimated statistics. In the past decade randomized quasi-Monte Carlo (RQMC) methods have been developed to address this issue. Essentially, the low-discrepancy sequences from QMC are randomized so that many independent sets of QMC's may be run from which error statistics may then be computed. See Li and Winker (2003) for a survey of QMC methods with applications in time series simulation, Ökten and Eastman (2004) for a survey of RQMC methods with applications in pricing financial derivatives, and Sak and Başoğlu (2015) for a recent RQMC application in risk management. See Lemieux (2009) for a general treatment of QMC and RQMC with financial applications.

To illustrate the use of RQMC to compute error statistics for time series simulation estimates we use a realistic example from financial time series. We use the Markov switching stochastic volatility model of So et al. (1998) to model portfolio returns and to estimate the precision of simulated estimates of expected returns, value-at-risk and expected shortfall of the portfolio over various forecast horizons. The model has a highly nonlinear stochastic process and it provides a realistic test of the relative power of RQMC methods in time series simulations.

2 Quasi-Monte Carlo and Randomized Quasi-Monte Carlo Methods

It is convenient to describe the quasi-Monte Carlo (QMC) and randomized quasi-Monte Carlo (RQMC) methods in the setting of numerical integration. Consider the problem of estimating the integral

$$I = \int_{[0,1)^s} f(x)dx.$$

Monte Carlo (MC) and QMC methods both estimate this integral using sums of the form

$$\frac{1}{N}\sum_{n=1}^{N} f(q^{(n)})$$

where $q^{(n)}$ is an s-dimensional pseudorandom vector from the uniform distribution on $(0, 1)^s$ in the case of the MC method, and the nth term of an s-dimensional low-discrepancy sequence for the QMC method. In the former method, the convergence of the estimates to the correct answer I is probabilistic with a convergence rate of $O(N^{-1/2})$ and in the latter, the convergence is deterministic with a rate of $O(N^{-1}(\log N)^s)$.

The error of the QMC estimate is known to satisfy the Koksma-Hlawka (or, many of its variants) inequality

$$\left|\frac{1}{N}\sum_{n=1}^{N} f(q^{(n)}) - I\right| \leq V_{HK}(f)D_N^*(q^{(n)}).$$

V_{HK} is the variation of f, in the sense of Hardy and Krause, and D_N^* is the star-discrepancy of the vectors $q^{(1)}, \ldots, q^{(N)}$, defined as

$$D_N^*(q^{(n)}) = \sup_{\beta\in[0,1)^s}\left|\frac{A_N([0,\beta))}{N} - \beta_1 \times \cdots \times \beta_s\right|$$

where $\beta = (\beta_1, \ldots, \beta_s)$ is a vector on $[0, 1)^s$, and $A_N([0, \beta))$ is the number of vectors $q^{(1)}, \ldots, q^{(N)}$ that belong to the interval $[0, \beta) = [0, \beta_1) \times \cdots \times [0, \beta_s)$. Intuitively, the discrepancy of a sequence measures its deviation from the "ideal" distribution it is supposed to follow. The Koksma-Hlawka inequality suggests smaller estimation error can be obtained if the star-discrepancy of the underlying sequence is smaller.

The Koksma-Hlawka inequality, however, does not help in assessing the error because the right-hand side of the inequality, which in principle can be used as a measure for error, cannot, in general, be computed analytically. Numerical estimation is also not feasible because computing the star-discrepancy, for example, is an NP-hard problem which is likely more computationally burdensome than the original estimation problem (Gnewuch et al. 2009).

QMC is a simulation method with a fast rate of convergence but leaves us with no practical way to compute the precision of our estimates. The randomized quasi-Monte Carlo method (RQMC) was developed to deal precisely with this issue.

The RQMC method uses a randomization approach that enables constructing independent copies of a given QMC sequence $q^{(n)}$. Denote a random copy of the sequence $q^{(n)}$ by $q_u^{(n)}$ where u is a random parameter. The RQMC method approximates the integral $I = \int_{[0,1)^s} f(x)dx$, by

$$Q(q_u) = \frac{1}{N}\sum_{n=1}^{N} f(q_u^{(n)}) \tag{1}$$

where $Q(q_u)$ is a random variable, derived from the random parameter u.

There are three important properties common to most RQMC methods:

1. $E[Q(q_u)] = I$; i.e., the estimate (1) indexed by the random parameter u is unbiased;
2. $Var(Q(q_u)) = O\left(N^{-2}(\log N)^{2s}\right)$, or better for certain integrands and RQMC methods;
3. $|Q(q_u) - I| \leq V_{HK}(f)D_N^*(q_u^{(n)})$; i.e., each RQMC estimate, obtained from a realization of the random parameter u, satisfies the Koksma-Hlawka inequality.

In practice, the random parameter u usually has the uniform distribution and $E[Q(q_u)] = I$ is estimated by the sample mean

$$\frac{Q(q_{u_1}) + \cdots + Q(q_{u_M})}{M} \tag{2}$$

where $u_1, \ldots, u_M$ are independent samples from u.

The RQMC methods currently used in the literature are:

Scrambled (t, m, s)-Nets and (t, s)-Sequences

(t, m, s)-nets and (t, s)-sequences are special constructions of QMC point sets and sequences; their theory is discussed in Niederreiter (1992). Owen (1994), introduced a way of randomizing these nets and sequences by scrambling and additional results were obtained in Hickernell (1998), Hickernell and Hong (1999), Loh (1996), and Owen (1997a, b). Alternative scrambling methods for (t, m, s)-nets and (t, s)-sequences were introduced by Matoušek (1998). These methods are computationally efficient, although they limit the extent to which the randomization is done in the original scrambling method.

Random Shifting

This method can be applied to any QMC sequence. An independent randomization

of the sequence $q^{(n)}$ is obtained by first generating a random vector u from the uniform distribution on $(0, 1)^s$ and then adding u to $q^{(n)}$ (for all n) component-wise and taking the fractional part of the sum. Properties and applications of random shifting are considered in Cranley and Patterson (1976), Joe (1990), Morohosi and Fushimi (2002), Morohosi et al. (2000) and Tuffin (1996).

Random-Start Halton Sequences

The van der Corput and Halton sequences are popular examples of low discrepancy sequences. One way to describe the van der Corput sequence is by using the von Neuman-Kakutani transformation, T_b, which is a well-known example of an ergodic mapping on $(0, 1)$. The van der Corput sequence in base b is simply the orbit of 0 under T_b. The Halton sequence in bases $b_1, \ldots, b_s$ is defined as the orbit of 0 under the s-dimensional von Neuman-Kakutani transformation $T_{\mathbf{b}}(\mathbf{x}) = \left(T_{b_1}(x_1), \ldots, T_{b_s}(x_s)\right)$. An independent realization of a random-start Halton sequence involves generating a random vector u from the uniform distribution on $(0, 1)^s$, and then constructing the orbit of $T_{\mathbf{b}}(u)$. For a given u, the orbit of $T_{\mathbf{b}}(u)$ is known to be a low-discrepancy sequence. Properties of these sequences and the random-start approach are studied in Lambert (1985), Struckmeier (1995), Wang and Hickernell (2000) and Ökten (2009). The random-start approach can also be applied to randomly permuted Halton sequences; details are discussed in Ökten (2009) and Ökten et al. (2012). These sequences are called RASRAP, and their implementation and applications to problems from computational finance are discussed in Xu and Ökten (2015).

3 AR(1) Model

QMC methods were employed in the forecasting of time series models in Li and Winker (2003). One of the time series models used in Li and Winker (2003) was a simple linear AR(1) model. In this section, we use the same model with the same parameters, to establish the following: (i) recreate the QMC results of Li and Winker (2003), and (ii) illustrate how error analysis of Li and Winker (2003) can be improved by switching to RQMC sequences.

The AR(1) model is

$$x_t = \alpha_0 + \alpha_1 x_{t-1} + \epsilon_t, \ \epsilon_t \sim \text{ iid } \mathcal{N}(0, \sigma_\epsilon^2), \ |\alpha_1| < 1 \tag{3}$$

and the question is to estimate the expected value of the 10-step forecast. Letting x_0 denote the last observed value and x_{10} the 10-step ahead forecast value, the analytical solutions for the mean and variance of the forecast are:

$$E[x_{10}] = \alpha_0 \left(\frac{1 - \alpha_1^{10}}{1 - \alpha_1} \right) + \alpha_1^{10} x_0 \tag{4}$$

$$Var[x_{10}] = \sigma_\epsilon^2 \left(\frac{1 - \alpha_1^{20}}{1 - \alpha_1^2} \right). \tag{5}$$

Following Li and Winker (2003) we set $x_0 = 1$, $\sigma_\epsilon^2 = 0.04$, $\alpha_0 = 0.1$, and $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$ and use MC and QMC methods to estimate $E[x_{10}]$

and then compare the estimate to the exact solution. Specifically, we generate N sample paths, $\{x_1^{(n)}, x_2^{(n)}, \ldots, x_{10}^{(n)}\}_{n=1}^N$, from (3), and compute $\widehat{x_{10}} = \frac{1}{N}\sum_{n=1}^N x_{10}^{(n)}$ as our estimate of $E[x_{10}]$.

The difference between MC and QMC is in the way the samples from the process ϵ_t are generated. In MC, a transformation method, such as the inverse transformation method or Box-Muller, is used to transform a pseudorandom sequence from the uniform distribution on (0, 1) to random samples from the standard normal distribution. In QMC, similar transformation methods are applied to the elements of a 10-dimensional low-discrepancy sequence. For example, if $q^{(n)} = (q_1^{(n)}, \ldots, q_{10}^{(n)})$ is the nth element of a 10-dimensional low-discrepancy sequence, and T denotes the transformation that maps a number from the uniform distribution on (0, 1) to a standard normal, then $x_{10}^{(n)}$ is obtained using the recursion (3) where $\epsilon_j^{(n)} = 0.2\,T(q_j^{(n)})$, for $j = 1, \ldots, 10$, where the "0.2" term is the standard deviation of the shock. For the MC simulations we use Mersenne twister (Matsumoto and Nishimura 1998) as the pseudorandom sequence and for the QMC simulations we use a realization of a random-start randomly permuted Halton sequence (RASRAP) (Xu and Ökten 2015).

Figure 1 plots the absolute error $|\widehat{x_{10}} - E[x_{10}]|$ against the sample size N for each of the four values $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$. To make comparisons easier the

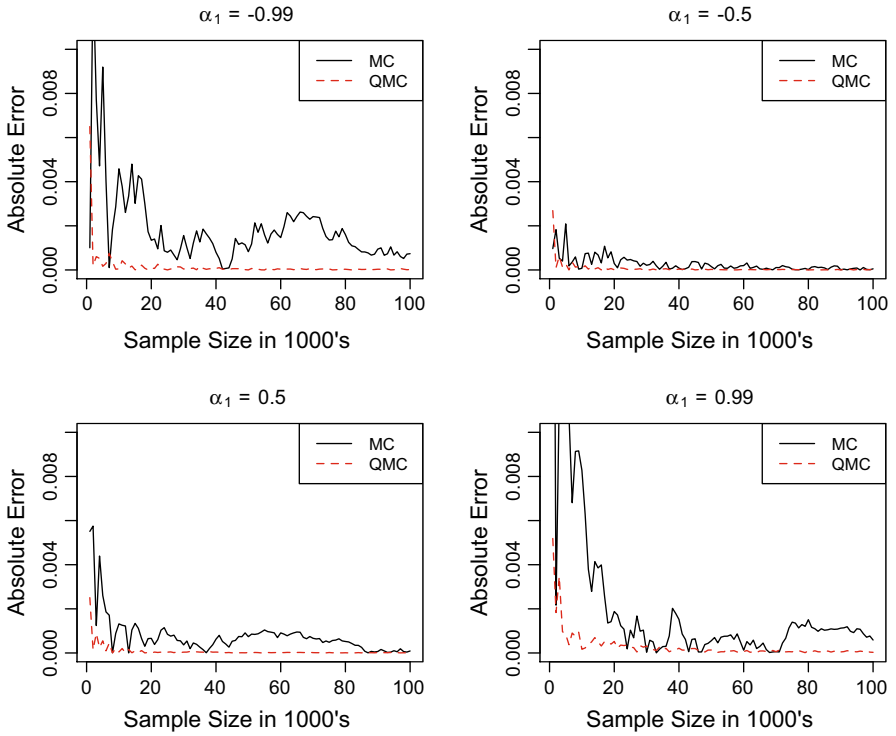

Fig. 1 Absolute error of $E[x_{10}]$ in the AR(1) model with $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$ simulated with MC and QMC with sample sizes N from 1000 to 100,000 by 1 K steps

scales of all four plots are the same but some upper values for the MC absolute errors are truncated to prevent extremely large values from distorting the plots. The plots show sample sizes only at 1 K steps to avoid too noisy an image. The QMC error converges quickly and is significantly smaller and more stable than the MC error as the simulation sample size grows.

Table 1 reports the average of the absolute of errors of MC, QMC and their ratios over sets of 20 K sample sizes. For instance, for $\alpha_1 = 0.5$, the average of the absolute errors over sample sizes 1 to 20,000 for the MC method is 1.87e−03 and over sample sizes 20,001 to 40,000 is 5.27e−04. The error ratio MC/QMC reported in columns 4 and 7 are the average of the 20 K ratios in each set and not the ratio of the averages for each method. For example, for $\alpha_1 = 0.5$, the average absolute error over the sample sizes 1 to 20 K is 1.87e−03 for the MC method and 8.30e−04 for the QMC method. To compute the error ratio MC/QMC of 38, we do not divide 1.87e−03 by 8.30e−04, the ratio of the averages, instead we average the ratios of all 20 K sample sizes.

Together, Fig. 1; Table 1 illustrate that QMC provides much faster convergence with less variation in simulation estimates than MC.

In the simple linear AR(1) model considered here we are able to compute exact errors because we can analytically compute the true value that we are trying to estimate by simulation. In more complicated time series models where there are no analytical formulae to compute the expectations, even though we may know the rate of convergence, it is more difficult to assess the accuracy of the simulation. In the case of MC simulation we can use the sample standard deviation, or other descriptive statistics, obtained from a number of independent estimates for the unknown quantity. However, in the QMC setting, assessing error in this way is not appropriate. Therefore, we now consider three RQMC sequences: random start Halton sequences (RAS); random-start randomly permuted Halton sequences (RASRAP); and scrambled SOBOL′ (using random linear scrambling of Matoušek (1998)) sequences. Using these RQMC sequences, we can compute independent estimates for $E[x_{10}]$. In Table 2, we present the mean absolute percent error (MAPE) using 100 independent estimates for $E[x_{10}]$, for two sample sizes $N = 1$ K and $N = 10$ K and for several values of the persistence parameter α_1.

The first column of Table 2 reports the value of α_1 and column two is the associated true value of $E[x_{10}]$ from (4). Columns 3–6 report the MAPE over the 100 replications and columns 7–9 report the ratios of the MC (the MC method uses the Mersenne twister generator) to the three RQMC methods used. The top panel reports the results for sample sizes of 1 K and the bottom panel reports the results for sample sizes of 10 K.

The RQMC methods are generally 5 to 10 times more precise than MC when $N = 1$ K and 20 to 40 times more precise for $N = 10$ K. The RQMC method based on the SOBOL′ sequences appears to be more accurate than the Halton-based sequence methods. These results are largely consistent with those presented in Table IV of Li and Winker (2003) using a single replication of QMC.

To get a better sense of the spread of the simulation estimations for the 100 replications, Fig. 2 shows box plots of the estimates of $E[x_{10}]$ from the four methods for the case $\alpha_1 = 0.9$ and Table 3 shows the ratios of the 25th through 75th percentile interquartile ranges (IQR) of the MC method over the three RQMC methods. For the

Table 1 Average absolute errors of MC, QMC, and the ratio of those errors over 20 K sets of sample sizes for $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$

Sample sizes (K)	Abs. error MC $\alpha_1 = -0.99$	Abs. error QMC	Error ratio MC/QMC	Abs. error MC $\alpha_1 = -0.5$	Abs. error QMC	Error ratio MC/QMC
1–20	4.74e−03	9.49e−04	1551	9.87e−04	5.36e−04	44
20–40	1.20e−03	8.38e−05	97	2.77e−04	4.03e−05	142
40–60	1.15e−03	4.67e−05	188	1.69e−04	2.31e−05	53
60–80	2.11e−03	3.43e−05	657	8.49e−05	1.65e−05	51
80–100	8.27e−04	2.85e−05	177	8.12e−05	1.32e−05	51
	$\alpha_1 = 0.5$			$\alpha_1 = 0.99$		
1–20	1.87e−03	8.30e−04	38	7.99e−03	1.78e−03	57
20–40	5.27e−04	5.23e−05	65	8.74e−04	1.59e−04	24
40–60	7.49e−04	3.09e−05	148	5.64e−04	9.57e−05	26
60–80	6.96e−04	2.22e−05	270	6.10e−04	6.42e−05	49
80–100	1.56e−04	1.64e−05	49	1.01e−03	4.96e−05	82

Table 2 MAPE statistics using 100 replications of MC and RQMC methods with sample sizes of $N = 1$ K (top panel) and $N = 10$ K (bottom panel) for various values of α_1 for the AR(1) model

α_1	True	MAPE				Ratio of MC over		
	$E[x_{10}]$	MC	RAS	RASRAP	SOBOL′	RAS	RASRAP	SOBOL′
$N = 1$ K								
−0.99	0.91	1.64	0.33	0.31	0.13	4.98	5.25	12.53
−0.90	0.38	2.95	0.59	0.67	0.28	5.01	4.39	10.49
−0.50	0.07	8.46	2.41	2.60	0.83	3.51	3.25	10.26
0.00	0.10	5.59	1.61	1.45	0.27	3.47	3.84	20.36
0.50	0.20	2.83	0.75	0.81	0.32	3.76	3.49	8.93
0.90	1.00	1.10	0.26	0.24	0.12	4.23	4.66	9.37
0.99	1.86	0.75	0.15	0.15	0.07	4.87	4.90	10.21
$N = 10$ K								
−0.99	0.91	0.52	0.03	0.03	0.02	19.07	20.65	29.20
−0.90	0.38	1.04	0.04	0.05	0.03	24.97	22.31	30.78
−0.50	0.07	2.59	0.15	0.14	0.10	17.00	18.79	27.17
0.00	0.10	1.54	0.09	0.08	0.04	18.03	18.17	36.03
0.50	0.20	0.85	0.05	0.05	0.03	18.32	17.99	25.60
0.90	1.00	0.35	0.02	0.02	0.01	18.45	18.31	32.51
0.99	1.86	0.26	0.01	0.01	0.01	19.04	22.46	29.67

For the value of α_1 in column 1, the true mean forecast is reported in column 2, the MAPEs for the 4 methods are in columns 3–6 and the ratios of the MC to RQMC methods are in columns 7–9. The RQMC methods are: RAS—random start Halton; RASRAP—random start–random permutation Halton; and SOBOL′—Sobol′ seqences

small sample sizes of $N = 1$ K, MC produces IQRs that are 5 to 10 times wider than the RQMC methods. When the sample size increases to $N = 10$ K the IQRs for MC are 15 to 30 times wider than those for RQMC.

To further investigate the convergence rate of RQMC relative to MC, we compute the root mean square error over 40 replications of MC using Mersenne twister pseudorandom numbers, and RQMC using RASRAP for simulation sample sizes from $N = 1$ to $N = 10$ K. In the AR(1) model (3), we use the following values for $\alpha = \{-0.99, -0.5, 0.5, 0.99\}$.

Figure 3 plots the $\log_{10}$ of the root mean square errors against the $\log_{10}$ of the sample sizes. The light gray lines show all 10 K points and the dashed lines show the linear trend through these points. The slope of the dashed line is reported for each method and approximates the rate of convergence of the method. As predicted, the convergence rate of MC is approximately of order $N^{1/2}$, and the convergence rate of RQMC is approximately of order N^{-1}; a rate better than the theoretical convergence rate of QMC. The convergence rates are robust across all values of α_1. Neither method is particularly accurate for sample sizes smaller than 10^2. The apparent volatility of the RMSE for larger sample sizes is illusory since this movement is primarily in the 4th decimal place.

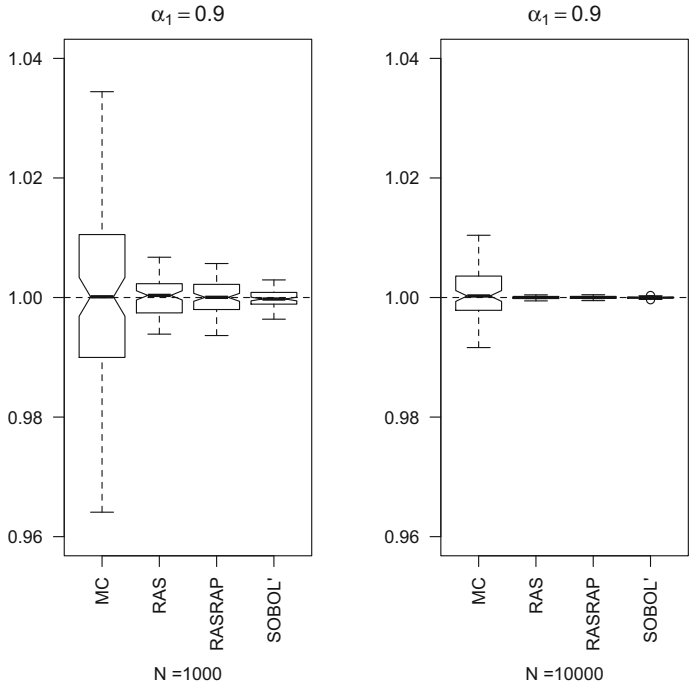

Fig. 2 *Box plot* of the estimated values of $E[x_{10}]$ for 100 replications of sample sizes of 1 and 10 K. Given $\alpha_0 = 0.1$, $x_0 = 1$, and $\alpha = 0.9$, the true expected value for the 10-step ahead forecast is $E[x_{10}] = 1$. The *solid lines* are the medians, the notches in each *box* represent the approximate 95% confidence interval for the differences in the medians in each panel, the *boxes* range from the 25th to the 75th percentile and the whiskers extend 1.5 times the box length with outliers reported as *circles* (only visible for SOBOL′ with $N = 10$ K)

Table 3 Ratios of interquartile ranges (IQR) from 100 replications of MC and RQMC methods for estimates of $E[x_{10}]$ from the AR(1) model with $\alpha_1 = 0.9$

Sample size (K)	IQR ratio		
	MC/RAS (K)	MC/RASRAP	MC/SOBOL′
1	4.2	4.9	10.5
10	16.2	16.7	33.2

The IQR is the length of the boxes from the 25th to the 75th percentile in the box plots in Fig. 2. The ratios are for the MC IQR divided by the RQMC IQR

Table 4 reports the averages of RMSE values for each method and their ratios over 2 K ranges of the simulation sample sizes. The ratios are computed as the average of the 2 K points in each range and not as the ratio of the averages for each method. As an example, for $\alpha_1 = -0.99$ and simulation sizes $N = 8001$ through $N = 10{,}000$, the average RMSE for the MC method is 6.29e−03 and for the RQMC method is 4.47e−04 with the average of the 2000 ratios being approximately 15.

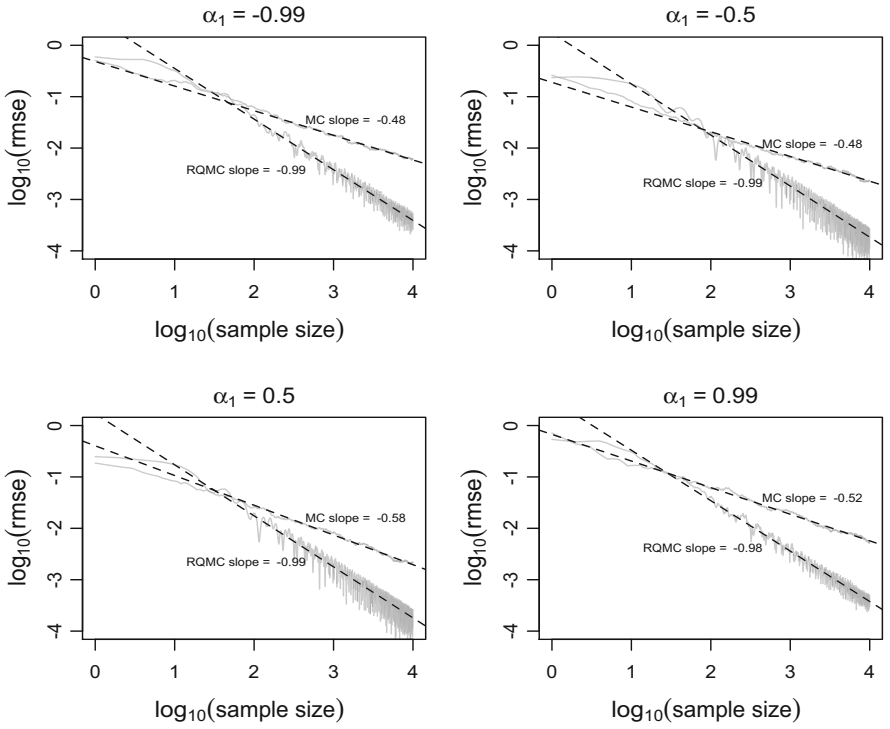

Fig. 3 Root mean square error of 40 MC and RQMC estimates for $E[x_{10}]$ in the AR(1) model for sample sizes $N = 1, \ldots, 10$ K and for $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$. The *plots* show the $\log_{10}(\text{rmse})$ versus $\log_{10}(N)$. The *light gray lines* are for each N and the *dashed lines* are the linear fit through the data with the slopes approximating the convergence rates of the methods. Approximately, MC converges at the rate $N^{-1/2}$ and RQMC at the rate N^{-1}

In general, the precision of the RQMC estimates is an order of magnitude better than those for MC and, as can be seen from columns 4 and 7 of Table 4 and the slopes of the lines in Fig. 3, the relative precision increases with the sample size N. The RQMC method allows us to compute error statistics for our simulated estimates and maintains the fast convergence rate of QMC.

4 An Application of RQMC: Estimating Value-at-Risk and Expected Shortfall

The AR(1) model example in the previous sections demonstrates the relatively fast convergence rate of QMC over MC and the use of RQMC for computing the precision of simulation estimates. Because the simple linear model that we used admits an analytical solution, we were able to compute the RMSE and MAPE statistics using the true expected forecast values. Of course, if we know the true values there is little need for simulation methods.

Table 4 Average root mean square errors of MC, RQMC and the ratio of those errors over 2 K ranges of simulation sizes for $\alpha_1 = \{-0.99, -0.5, 0.5, 0.99\}$

Sample range (K)	RMSE MC $\alpha_1 = -0.99$	RMSE RQMC	Ratio MC/RQMC	RMSE MC $\alpha_1 = -0.5$	RMSE RQMC	Ratio MC/RQMC
1–2	2.49e−02	1.22e−02	4	9.30e−03	6.08e−03	4
2–4	1.02e−02	1.39e−03	8	4.27e−03	6.64e−04	7
4–6	8.50e−03	7.74e−04	11	3.40e−03	3.85e−04	10
6–8	7.01e−03	5.69e−04	13	2.63e−03	2.76e−04	11
8–10	6.29e−03	4.47e−04	15	2.27e−03	2.14e−04	12
	$\alpha_1 = 0.5$			$\alpha_1 = 0.99$		
1–2	1.09e−02	5.93e−03	4	2.74e−02	1.16e−02	5
2–4	4.25e−03	6.43e−04	7	9.96e−03	1.28e−03	8
4–6	2.80e−03	3.78e−04	8	8.20e−03	7.66e−04	11
6–8	2.29e−03	2.69e−04	9	7.35e−03	5.40e−04	14
8–10	2.18e−03	2.10e−04	11	5.94e−03	4.32e−04	14

Li and Winker (2003) consider several nonlinear time series models for which true forecast values are difficult to compute and they demonstrate that QMC maintains a relatively fast convergence rate. Less clear, since they are using QMC rather than RQMC, is how Li and Winker (2003) compute the statistical accuracy of their QMC simulations. In this section we demonstrate the speed and precision of RQMC for realistic financial applications of time series models. We use a nonlinear regime switching, stochastic volatility process to model portfolio returns and estimate the expected portfolio return and two measures of tail risk.

Our specification for portfolio returns is the Markov Switching Stochastic Volatility (MSSV) model introduced by So et al. (1998). As in stochastic volatility models, the mean and variance processes have independent shocks but in addition the variance process is generalized to allow for regime-switching between high-, medium- and low-volatility states where the transitions between states are modeled using a Markov transition matrix. The MSSV model is better able to simulate the extreme tail events necessary to accurately measure tail risk of the portfolio. The nonlinearity of the MSSV model makes analytical solutions for risk measures intractable so that simulation methods are necessary.

We compute two measures of tail risk for the portfolio: Value at Risk (VaR) and Expected Shortfall (ES). VaR(p) measures the minimum expected return that the portfolio will earn p% of the time. If VaR(5)= -6%, the portfolio would lose at least six percent of its value 5% of the time. ES(p) measures the expected return of the portfolio in the worst p% of the possible outcomes. Because ES(p) measures the area in the returns left tail, it is less sensitive to the tail shape of the returns density than is VaR(p) which only measures the pth-quantile point. If VaR(5)= -6%, the ES(5) might be more like -8%. ES(p) is intended to provide a better measure of how large expected losses may be if markets do crash.

4.1 Markov Switching Stochastic Volatility Model of Portfolio Returns

The So et al. (1998) model for returns r_t of an asset portfolio is

$$r_t = \mu_t + \sqrt{h_t} u_t. \tag{6}$$

The time-varying conditional variance is given by

$$\log h_{t+1} = \alpha_{s_{t+1}} + \phi \log h_t + \eta_t, \tag{7}$$

where $u_t \sim$ iid $\mathcal{N}(0, 1)$, $\eta_t \sim$ iid $\mathcal{N}(0, \sigma_\eta^2)$ and s_t denotes the state, or the regime, of the economy denoted by integers $\{1, 2, \ldots, K\}$. Following So et al. (1998), our model will use $K = 3$ volatility states. The economy makes transitions between regimes, according to a Markov process with transition probability matrix

$$P = \begin{pmatrix} p_{11} & p_{12} & p_{13} \\ p_{21} & p_{22} & p_{23} \\ p_{31} & p_{32} & p_{33} \end{pmatrix}, \tag{8}$$

i.e., $p_{ij} = \Pr(s_t = j | s_{t-1} = i)$ with $\sum_{j=1}^{3} p_{ij} = 1$ for $i = 1, 2, 3$.

Given the state variable s_t, the parameter α_{s_t} is given by

$$\alpha_{s_t} = \gamma_1 + \sum_{j=2}^{3} \gamma_j I_{jt} \tag{9}$$

where

$$I_{jt} = \begin{cases} 1, & \text{if } s_t \geq j \\ 0, & \text{otherwise} \end{cases} \tag{10}$$

is an indicator function that is equal to 1 when s_t is greater than or equal to j. We constrain parameters $\gamma_j \leq 0, \ j = 2, 3$ so that state 1 has the interpretation of the high-volatility state and states 2 and 3 represent the medium and low volatility states.

Estimation of Markov switching stochastic volatility models is computationally complex because both the volatilities and the regime states are unobserved. The method used by So et al. (1998) is a forward-filter–backward sampling algorithm with a Kalman filter smoother. A sequential algorithm using the particle filter method was shown by Carvalho and Lopes (2007) and Rios and Lopes (2013) to produce more stable estimates. To our knowledge, no one has implemented a QMC algorithm to estimate this type of model and, since the main focus of the present paper is on time series simulation issues, we will defer this complex estimation problem to a subsequent paper. For our simulations we will use the estimation results of the MSSV(3) model by So et al. (1998) using the S&P 500 Index weekly returns from the first week of 1961 to the last week of 1987. Complete estimation results are reported in Tables 2 and 4 of So et al. (1998) but, for convenience, we summarize them here:[1] the estimated mean in (6) is $\hat{\mu} = 0.00103$; the estimated persistence in the conditional variance (7) is $\hat{\phi} = 0.472$; the estimated regime parameters are $\hat{\gamma}_1 = -3.536$, $\hat{\gamma}_2 = -0.764$, and $\hat{\gamma}_3 = -0.541$; the estimated variance of the volatility shock is $\hat{\sigma}_\eta^2 = 0.074$; and the estimated Markov transition probability matrix is

$$\hat{P} = \begin{pmatrix} .859 & .130 & .011 \\ .029 & .967 & .004 \\ .009 & .017 & .974 \end{pmatrix}. \tag{11}$$

A full interpretation of the model and estimation results are provided by So et al. (1998). We note that the mean time before switching out of the high-volatility state is about 7 weeks, and 30 and 38 weeks for the medium- and low volatility states. Regime switching is a major source of variance in volatility as can be seen by the estimated persistence of a volatility shock, ϕ, which is much lower than would typically be estimated in a model without regime switching.

[1] So et al. (1998) filter their returns data to eliminate a small negative persistence term in the mean equation (6) induced by the market crash at the end of the estimation period. For the purposes of our simulations, we treat the conditional mean returns as constant.

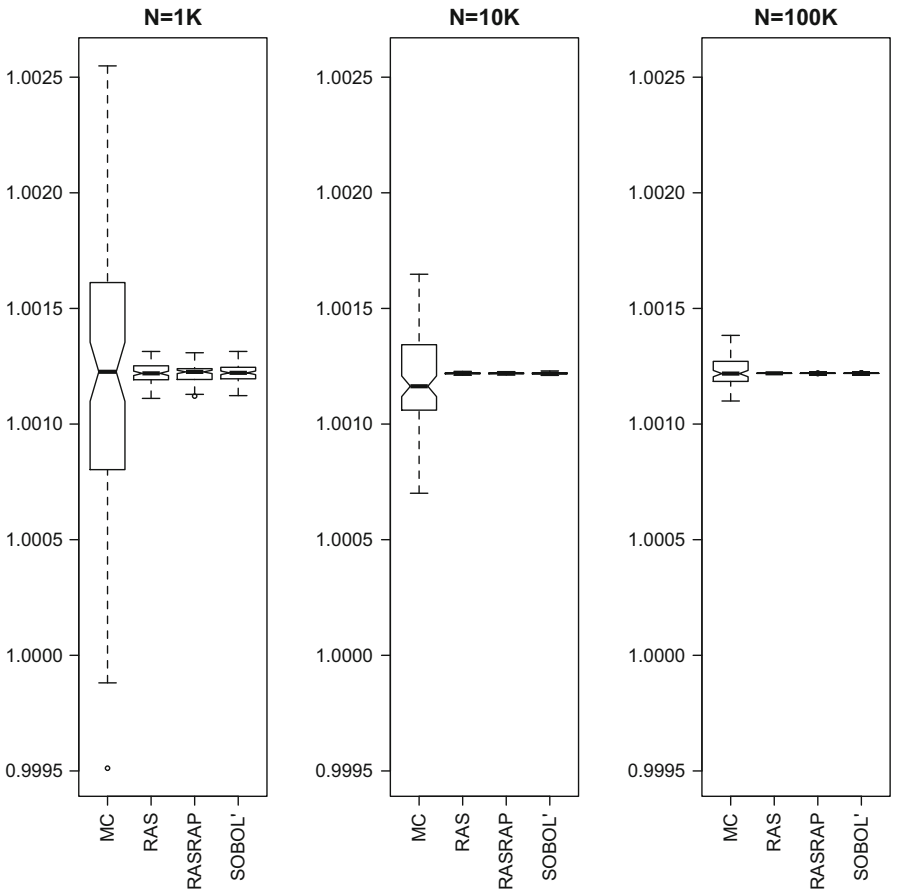

Fig. 4 *Box plot* of the estimated values of $E[1 + r_1]$ for 100 replications of sample sizes of 1, 10, and 100 K

4.2 Simulating Expected Returns

The MSSV model of portfolio returns is highly nonlinear and it provides a good test of the relative performance of MC and RQMC methods. To examine the effect of increasing dimension on the simulation accuracy of the (R)QMC method we consider two cases: (i) 1-step ahead (1 week simulation horizon) portfolio returns from the MSSV model, and (ii) 4-step ahead (a 4 week or about 1 month simulation horizon) portfolio returns.

For each simulation we start with a return and variance equal to the unconditional values, $r_0 = \hat{\mu} = 0.00103$ and $h_0 = 0.0217^2$, with the market currently in the medium-volatility state regime, state $s_0 = 2$. In case (i) we need to draw the mean equation shock u_1, the variance equation shock η_1, and the state of the system, s_1, which requires a 3-dimensional vector of pseudorandom, or low-discrepancy sequence, for each simulation. Since the simulation horizon in case (ii) is 4 periods, we must draw

Table 5 Interquartile range ratios (IQR) from 100 replications of MC and RQMC methods for estimates of gross returns from the MSSV model

Sample size (K)	IQR ratio		
	MC/RAS	MC/RASRAP	MC/SOBOL′
Case (i): 1-step ahead simulations			
1	13.4	17.3	16.3
10	62.0	62.6	54.6
100	115.9	134.2	185.8
Case (ii): 4-step ahead simulations			
1	6.6	5.9	7.3
10	12.8	7.8	17.7
100	9.9	12.0	14.7

The IQR is the length of the boxes from the 25th to the 75th percentile in the box plots. The ratios are for the MC IQR divided by the RQMC IQR for sample sizes $N = 1$ K, 10 K and 100 K. The upper panel are the IQRs for the 1-step ahead estimates of $E[1 + r_1]$, and and the lower panel are the IQRs for the 4-step ahead estimates of $E[1 + r_4]$

a 12-dimensional vector of pseudorandom, or low-discrepancy sequence, for each simulation.

In Fig. 4 we show box plots of 100 independent simulations of the 1-step ahead estimates of the expected portfolio gross return, $E[1 + r_1]$, using the MC and RQMC methods with the same sequences considered in Section 2 with the AR(1) model. The rapid convergence of the RQMC methods relative to MC is evident.

We compare the relative precision of the MC to the RQMC methods using the ratios of the interquartile ranges (IQR) from the box plots of the 100 simulations used to estimate the expected portfolio returns. The upper panel of Table 5 reports the IQRs for case (i) from the box plots in Fig. 4. The width of the RQMC boxes are approximately 15 times smaller than the MC boxes for sample sizes of $N = 1$ K, 50 to 60 times smaller for sample sizes of $N = 10$ K, and over 2 orders of magnitude smaller for sample sizes of $N = 100$ K.

The lower panel of Table 5 reports the IQRs for case (ii), whose box plots we do not show. Although the precision of the estimates for the RQMC methods is still approximately an order of magnitude better than for MC, the convergence rate for RQMC is much slower for the 4-step ahead simulations that use 12-dimensional sequences than for the 1-step ahead simulations of case (i) that use 3-dimensional sequences in the simulations. The precision of the RQMC estimates is still quite good, with estimated standard deviations across the 100 replications for $N = 100$ K on the order of 10^{-6}, but the two order of magnitude increase precision of RQMC over MC is lost in the higher dimensional case. This reduced advantage of RQMC over MC with higher dimensional sequences is consistent with the findings of Ökten et al. (2006).

4.3 Simulating the Measures of Tail Risk

To compute VaR(p) and ES(p) we simulate the h-step ahead forecast of portfolio returns. It is conventional to report returns in terms of h-step ahead losses. Let $F_h(\ell)$

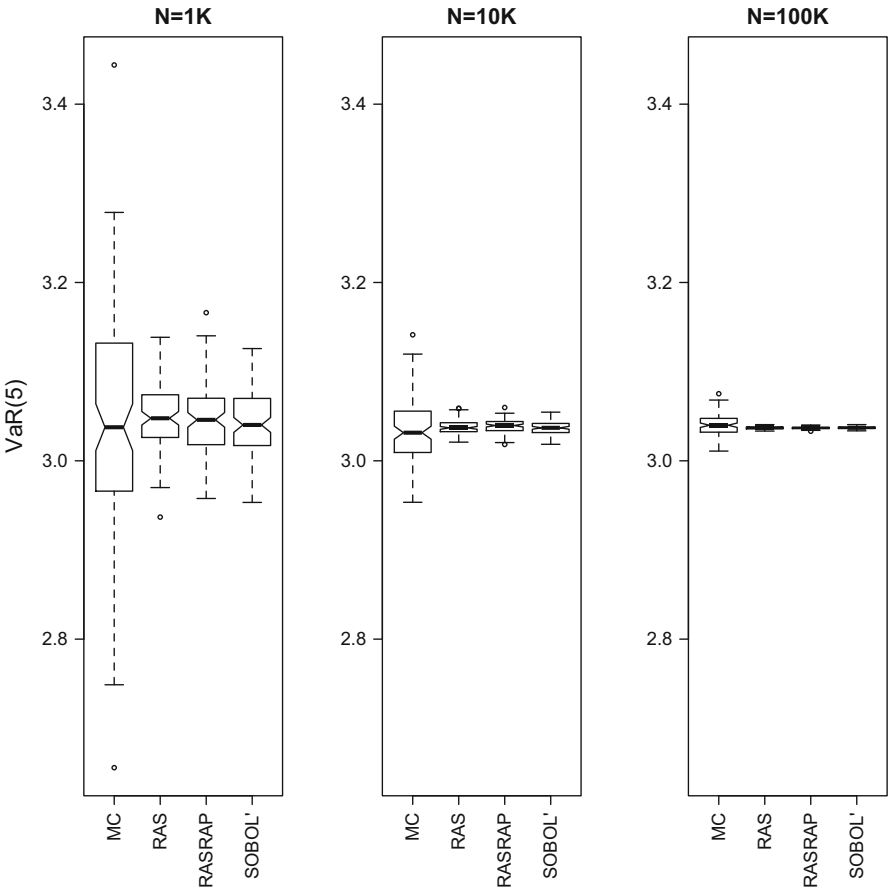

Fig. 5 Portfolio value for the 1-step ahead estimates of VaR(5) from 100 replications of MC and three RQMC methods for simulation sample sizes $N = 1$ K, 10 K, 100 K

denote the distribution function of h-step ahead losses ℓ. Then the VaR(p) of a portfolio is defined as

$$p = \Pr[L(h) \geq \text{VaR}(p)] = 1 - \Pr[L(h) < \text{VaR}(p)] \tag{12}$$

or

$$\text{VaR}(p) = \inf \{\ell \mid F_h(\ell) \geq p\} \tag{13}$$

so the probability of a loss of VaR(p) or greater over the h-period horizon is p. The expected shortfall is defined as

$$\text{ES}(p) = E\left[\ell \mid \ell \geq \text{VaR}(p)\right] = \frac{1}{p} \int_{\text{VaR}(p)}^{\infty} \ell \, dF_h(\ell). \tag{14}$$

Table 6 Interquartile range ratios (IQR) from 100 replications of MC and RQMC methods for estimates of VaR(5) from the MSSV model

Sample size (K)	IQR ratio		
	MC/RAS	MC/RASRAP	MC/SOBOL′
Case (i): 1-step ahead simulations			
1	3.5	3.2	3.1
10	4.6	4.5	4.5
100	7.3	9.4	8.4
Case (ii): 4-step ahead simulations			
1	1.3	1.2	1.4
10	1.4	1.3	1.4
100	1.1	1.1	1.4

The IQR is the length of the boxes from the 25th to the 75th percentile in the box plots. The ratios are for the MC IQR divided by the RQMC IQR for sample sizes $N = 1$ K, 10 K, 100 K. The upper panel are the IQRs for the 1-step ahead estimates of VaR(5), and and the lower panel are the IQRs for the 4-step ahead estimates of VaR(5)

Beginning with an initial portfolio value of $V_0 = \$1$, the portfolio value at step $t = 1, \ldots, T$ is given by

$$V_t = V_{t-1} \times e^{r_t}, \; t = 1, \ldots, T \tag{15}$$

where r_t is the simulated return process from the MSSV model, and $T = 1$ for case (i), and $T = 4$ for case (ii). In our simulations we use $p = 5\%$. To measure the precision of the VaR(5) and ES(5) estimates we use the same $M = 100$ simulations of sample sizes $N = 1$ K, 10 K, 100 K using the MC method and the three RQMC methods (RAS, RASRAP and SOBOL′) described earlier.

Figure 5 displays box plots of the 100 independent replications of the 1-step ahead simulation estimates of VaR(5) for the MC and each of the RQMC methods, for simulation sample sizes of $N = 1$ K, 10 K, 100 K. The estimates predict there is a 5% chance that this portfolio will lose at least 3.04% of its value during the next 1 week horizon. The standard deviation of these estimates in the $N = 10$ K case is about 4 basis points for the MC method and less than 1 basis point for the RQMC methods.

The top panel of Table 6 shows the interquartile ranges for the box plots in Fig. 5. The MC method IQRs are 3 to 9 times wider than those for the RQMC methods depending upon the sample size. The standard deviation of the estimates by sample size $N = 1$ K, 10 K, 100 K are 12, 4 and 1 basis points for the MC method, and 4, 0.8 and 0.2 basis points for the RQMC methods.

The lower panel of Table 6 reports the IQRs for the 4-step ahead simulated estimates of VaR(5). The estimates predict there is a 5% chance that this portfolio will lose at least 5.6% of its value during the next 1 month horizon. For this higher dimensional problem, the RQMC methods have almost completely lost their advantage over the MC method. It is worth noting that the VaR(5) estimate is the 95% percentile of the distribution of losses and for each simulation that point estimate is computed from

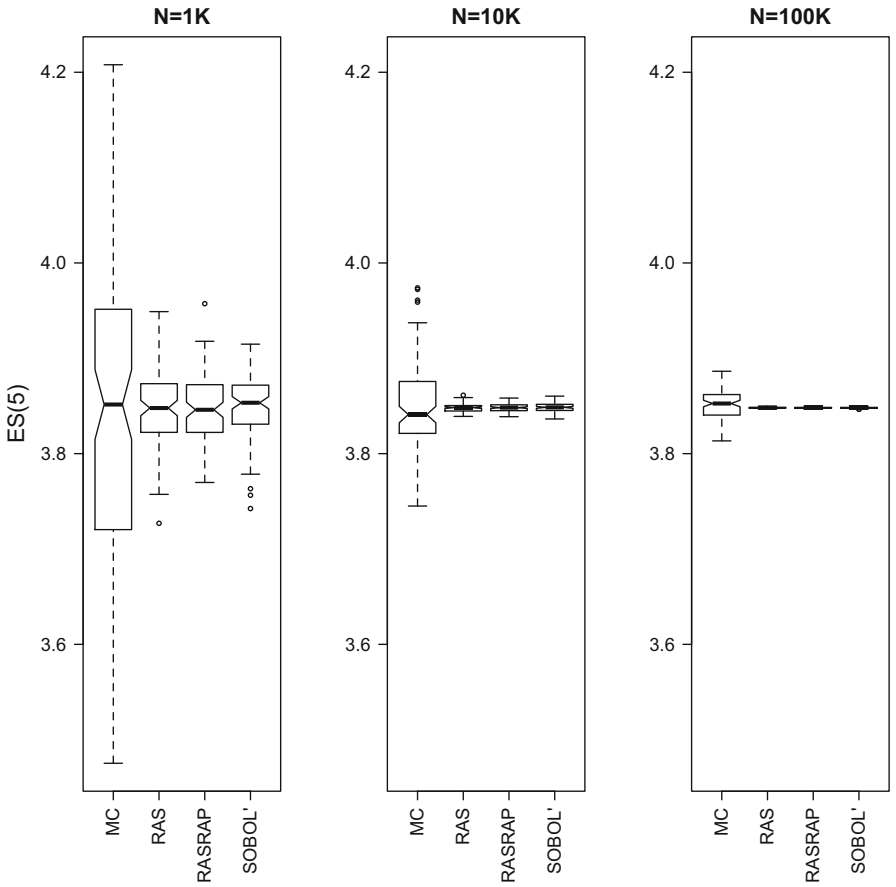

Fig. 6 Portfolio value for the 1-step ahead estimates of ES(5) from 100 replications of MC and three RQMC methods for simulation sample sizes $N = 1$ K, 10K, 100 K

the order statistics of the returns. Even when the simulation sizes are large there are relatively few points in the upper tail of losses so the number of regime switches in the simulations can play an important role in the computation of the VaR. This, of course, is precisely why the VaR is difficult to estimate with precision and this fact remains true for RQMC methods.

Box plots of the the 1-step ahead simulation estimates of the 5% expected shortfall, ES(5), are shown in Fig. 6, and the IQRs for the 1-step and 4-step ahead simulations are shown in Table 7. The estimates of ES(5) are around 3.85% which, as expected, are larger than the estimates of VaR. RQMC methods are somewhat more precise than MC at estimating ES(5) than VaR(5), particularly for the 1-step ahead simulations where the IQRs are 20 to 30 times smaller for RQMC than MC in the largest sample size simulations. For the 1-step ahead simulations, the standard deviations of the ES(5) estimates are generally about an order of magnitude smaller for RQMC than for MC. The standard deviation of the estimates based on RQMC methods for the case (i)

Table 7 Interquartile range ratios (IQR) from 100 replications of MC and RQMC methods for estimates of ES(5) from the MSSV model

Sample size (K)	IQR ratio		
	MC/RAS	MC/RASRAP	MC/SOBOL′
Case (i): 1-step ahead simulations			
1	4.5	4.6	5.6
10	9.6	8.9	8.4
100	22.4	24.4	29.1
Case (ii): 4-step ahead simulations			
1	1.1	1.1	1.2
10	1.7	2.0	1.3
100	1.2	2.0	1.6

The IQR is the length of the boxes from the 25th to the 75th percentile in the box plots. The ratios are for the MC IQR divided by the RQMC IQR for sample sizes $N = 1$ K, 10 K, 100 K. The upper panel are the IQRs for the 1-step ahead estimates of ES(5), and the lower panel are the IQRs for the 4-step ahead estimates of ES(5)

simulations with $N = 100$ K are about one-tenth of one basis point and for the case (ii) simulations are about 6 basis points. Our estimates of expected shortfall ES(5) are more precise than those for the associated VaR(5), confirming the recommendation of the Basel Committee that risk measures should switch from VaR to ES.

We note that the estimation of VaR is a notoriously difficult problem and has been approached in many ways including IGARCH (Longerstaey and More 1995), extreme value theory (Longin 1999; Tsay 1999), cluster analysis (Chang et al. 2007), and neural networks (Chen and Hsieh 2010). Estimates of VaR can be improved using variance reduction methods, or smoothing methods such as the Fourier transform approach introduced by Jin and Zhang (2006) that smooths the expectation estimation of the indicator function in the VaR estimation. Numerical results of Jin and Zhang (2006) suggest that QMC, when used in conjunction with the Fourier transform approach, gives more accurate results than plain MC and QMC implementations. The RQMC results we have reported for the estimation of VaR can be improved using these techniques, but since the main purpose of this paper is not to estimate VaR, we do not pursue these techniques here.

5 Conclusion

In this article we have demonstrated the use of randomized quasi-Monte Carlo methods (RQMC) in time series simulations. Quasi-Monte Carlo (QMC) methods use low-discrepancy sequences to improve the spatial coverage of simulated processes over classical Monte Carlo methods. QMC methods have theoretical convergence rate of $O(N^{-1}(\log N)^s)$, although in many applications, researchers have observed rates close to $O(N^{-1})$, compared to the MC convergence rate of $O(N^{-1/2})$. One drawback of the QMC method is that its estimates are not amenable to practical error analy-

sis. RQMC methods overcome this problem by introducing a randomization step into the selection of the low-discrepancy sequences. Descriptive statistics can be used to analyze the error of estimates obtained by RQMC.

We consider a simpler time series application estimating the expected value of an h-step forecast of a linear time series model, and a realistic time series application of estimating the values and precision of expected returns, value-at-risk, and expected shortfall of a portfolio, whose returns follow a highly nonlinear Markov switching stochastic volatility model. We demonstrate how RQMC can be used to analyze estimation error statistically, and compare errors obtained by several RQMC methods with Monte Carlo. Our numerical results show that RQMC can offer substantial error reduction, up to two orders of magnitude, in estimating the expected returns of the Markov switching model. We also observe that the factor of improvements offered by RQMC diminish as the the dimension of the underlying sequences increase, especially in the case of VaR estimation, where 4-step ahead simulations via RQMC offer little improvement over MC. Techniques such as variance reduction or smoothing can be used to improve RQMC, however, the detrimental effect of increasing dimension cannot be completely avoided; see Ökten et al. (2006) for a discussion of high dimensional simulation.

Acknowledgements We thank Dr. Ahmet Göncü for his helpful comments.

References

Boyle, P. P., Broadie, M., & Glasserman, P. (1997). Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, *21*(8–9), 1267–1321.

Caflisch, R. E., Morokoff, W. J., & Owen, A. B. (1997). Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension. *Journal of Computational Finance*, *1*(1), 27–46.

Carvalho, C. M., & Lopes, H. F. (2007). Simulation-based sequential analysis of markov switching stochoastic volatility models. *Computational Statistics & Data Analysis*, *51*, 4526–4542.

Chang, H. -C., Chen, C. -T., & Hsieh, C. -S. (2007). *Forecasting of value at risk by using percentile of cluster method* (pp. 480–486). World Scientific, Singapore, Chapter 68.

Chen, C. -T., & Hsieh, C. -S. (2010). Measuring of value at risk (VAR) on emerging stock markets by neural networks method. In *2010 International Symposium on Computer Communication Control and Automation (3CA)* (vol. 2, pp. 137–140).

Cranley, R., & Patterson, T. (1976). Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, *13*(6), 904–914.

Fishman, V., Fitton, P., & Galperin, Y. (1997). Hybrid low-discrepancy sequences: Effective path reduction for yield curve scenario generation. *Journal of Fixed Income*, *7*, 75–84.

Galanti, S., & Jung, A. (1997). Low-discrepancy sequences: Monte Carlo simulation of option prices. *The Journal of Derivatives*, *5*(1), 63–83.

Gnewuch, M., Srivastav, A., & Winzen, C. (2009). Finding optimal volume subintervals with k points and calculating the star discrepancy are NP-hard problems. *Journal of Complexity*, *25*(2), 115–127.

Göggelmann, K., Winker, P., Schellhorn, M., & Franz, W. (2000). Quasi-Monte Carlo methods in stochastic simulations: An application to policy simulations using a disequilibrium model of the West German economy 1960–1994. *Empirical Economics*, *25*(2), 247–259.

Hickernell, F. J. (1998). What affects the accuracy of quasi-Monte Carlo quadrature? In Niederreiter, H. & Spanier, J. (eds) *Monte Carlo and Quasi-Monte Carlo Methods* (pp. 16–55), New York. Proceedings of a Conference held at the Claremont Graduate University, Claremont, CA, USA, June 22–26, Springer-Verlag.

Hickernell, F. J., & Hong, H. S. (1999). The asymptotic efficiency of randomized nets for quadrature. *Mathematics of Computation*, *68*(226), 767–791.

Jin, X., & Zhang, A. X. (2006). Reclaiming Quasi-Monte Carlo efficiency in portfolio value-at-risk simulation through Fourier transform. *Management Science*, *52*(6), 925–938.

Joe, S. (1990). Randomization of lattice rules for numerical multiple integration. *Journal of Computational and Applied Mathematics*, *31*(2), 299–304.

Joy, C., Boyle, P. P., & Tan, K. S. (1996). Quasi-Monte Carlo methods in numerical finance. *Management Science*, *42*(6), 926–938.

Lambert, J. (1985). Quasi-Monte Carlo, low discrepancy sequences, and ergodic transformations. *Journal of Computational and Applied Mathematics*, *12*, 419–423.

Lemieux, C. (2009). *Monte Carlo and Quasi-Monte Carlo sampling*., Springer Series in Statistics Berlin: Springer.

Li, J. X., & Winker, P. (2003). Time series simulation with Quasi-Monte Carlo methods. *Computational Economics*, *21*(1–2), 23–43.

Loh, W.-L. (1996). A combinatorial central limit theorem for randomized orthogonal array sampling designs. *The Annals of Statistics*, *24*(3), 1209–1224.

Longerstaey, J., & More, L. (1995). *Introduction to risk metrics*. New York: Morgan Guaranty and Trust Company.

Longin, F. M. (1999). Optimal margin level in futures markets: Extreme price movements. *Journal of Futures Markets*, *19*, 127–152.

Matoušek, J. (1998). On the L2-discrepancy for anchored boxes. *Journal of Complexity*, *14*(4), 527–556.

Matsumoto, M., & Nishimura, T. (1998). Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, *8*(1), 3–30.

Morohosi, H., & Fushimi, M. (2002). Experimental studies on the error estimation of the numerical integration by generalized Niederreiter sequences. In *Proceedings of the Fourth International Symposium, Operations Research and Its Applications (ISORA) '02* (pp. 46–53). New York: World Publishing.

Morohosi, H., Kobayashi, K., & Fushimi, M. (2000). *Estimating the error of quasi-Monte Carlo integrations–experimental studies in financial problems* (pp. 1–11). Kokyuroku: Research Institute for Mathematical Sciences: Kyoto University .

Niederreiter, H. (1992). *Random number generation and quasi-Monte Carlo methods*. Philadelphia: Society for Industrial and Applied Mathematics.

Ninomiya, S., & Tezuka, S. (1996). Toward real-time pricing of complex financial derivatives. *Applied Mathematical Finance*, *3*(1), 1–20.

Ökten, G. (2009). Generalized von Neumann–Kakutani transformation and random-start scrambled Halton sequences. *Journal of Complexity*, *25*(4), 318–331.

Ökten, G., & Eastman, W. (2004). Randomized quasi–Monte Carlo methods in pricing securities. *Journal of Economic Dynamics and Control*, *28*(12), 2399–2426.

Ökten, G., Shah, M., & Goncharov, Y. (2012). Random and deterministic digit permutations of the Halton sequence. In *Monte Carlo and Quasi-Monte Carlo methods 2010* (pp. 609–622). Berlin: Springer.

Ökten, G., Tuffin, B., & Burago, V. (2006). A central limit theorem and improved error bounds for a hybrid-Monte Carlo sequence with applications in computational finance. *Journal of Complexity*, *22*(4), 435–458.

Owen, A. B. (1994). Randomly permuted (t, m, s)-nets and (t, s)-sequences. In Niederreiter, H. & Shiue, P. J. -S. (eds) *Monte Carlo and Quasi-Monte Carlo methods in scientific computing* Lecture Notes in Statistics, ((vol. 106, pp. 299–317). New York. Proceedings of a Conference at the University of Nevada, Las Vegas, Nevada, USA, June 23–25, 1994, Springer-Verlag.

Owen, A. B. (1997a). Monte Carlo variance of scrambled net quadrature. *SIAM Journal on Numerical Analysis*, *34*(5), 1884–1910.

Owen, A. B. (1997b). Scrambled net variance for integrals of smooth functions. *The Annals of Statistics*, *25*(4), 1541–1562.

Owen, A. B., & Tavella, D. (1997). Scrambled nets for value at risk calculations. In S. Grayling (Ed.), *VaR: Understanding and applying value-at-risk*. London: Risk.

Papageorgiou, A., & Traub, J. F. (1996). Beating Monte Carlo. *Risk*, *9*(6), 63–65.

Paskov, S. H., & Traub, J. F. (1995). Faster valuation of financial derivatives. *The Journal of Portfolio Management*, *22*(1), 113–123.

Rios, M. P., & Lopes, H. F. (2013). The extended liu and west filter: Parameter learning in markov switching stochastic volatility models. In Y. Zeng & S. Wu (Eds.), *State-space models: Applications in economics and finance* (pp. 23–61). New York, NY: Springer New York.

Sak, H. & Başoğlu, İ. (2015). Efficient randomized quasi-Monte Carlo methods for portfolio market risk. arXiv preprint: arXiv:1510.01593.

Sloan, I. H., & Woźniakowski, H. (1998). When are quasi–monte carlo algorithms efficient for high dimensional integrals? *Journal of Complexity*, *14*(1), 1–33.

So, M. E. P., Lam, K., & Li, W. K. (1998). A stochastic volatility model with Markov switching. *Journal of Business and Economic Statistics*, *16*(2), 244–253.

Struckmeier, J. (1995). Fast generation of low-discrepancy sequences. *Journal of Computational and Applied Mathematics*, *61*(1), 29–41.

Tsay, R. S. (1999). Extreme value analysis of financial data. Technical report, Graduate School of Business, University of Chicago.

Tuffin, B. (1996). On the use of low discrepancy sequences in Monte Carlo methods. *Monte Carlo Methods and Applications*, *2*, 295–320.

Wang, X., & Hickernell, F. J. (2000). Randomized Halton sequences. *Mathematical and Computer Modelling*, *32*(7), 887–899.

Xu, L., & Ökten, G. (2015). High-performance financial simulation using randomized quasi–Monte Carlo methods. *Quantitative Finance*, *15*(8), 1425–1436.

Zhang, L., Zhang, W., Xu, W., & Shi, X. (2013). A modified least-squares simulation approach to value American barrier options. *Computational Economics*, *44*(4), 489–506.