STOCHASTIC ROOT FINDING FOR OPTIMIZED CERTAINTY EQUIVALENTS

Anna-Maria Hamm
Thomas Salfeld
Stefan Weber

Gottfried Wilhelm Leibniz Universität Hannover
Welfengarten 1
D-30167 Hannover, GERMANY

ABSTRACT

Global financial markets require suitable techniques for the quantification of the downside risk of financial positions. In the current paper, we concentrate on Monte Carlo methods for the estimation of an important and broad class of convex risk measures which can be constructed on the basis of optimized certainty equivalents (OCEs). This family of risk measures — originally introduced in Ben-Tal and Teboulle (2007) — includes, among others, the entropic risk measure and average value at risk. The calculation of OCEs involves a stochastic optimization problem that can be reduced to a stochastic root finding problem via a first order condition. We describe suitable algorithms and illustrate their properties in numerical case studies.

1 INTRODUCTION

Global financial markets require suitable techniques for the quantification of the downside risk of financial positions. The seminal papers of Artzner et al. (1999) and Föllmer and Schied (2002) triggered an intensive scientific discussion about sensible risk measures with economically meaningful properties. However, there are not many contributions in the literature on the Monte Carlo simulation and implementation of risk measures – an issue of crucial importance for the implementation in practice.

The current article introduces a new Monte Carlo simulation technique for an important and broad class of convex risk measures, the optimized certainty equivalents (OCE). This family of risk measures was suggested by Ben-Tal and Teboulle (2007) and includes, among others, the entropic risk measure and average value at risk. Average value at risk, also known as tail value at risk, conditional value at risk, or expected shortfall, plays already an important role in practice and provides e.g. the basis for the Swiss solvency test for insurance firms.

The computation of OCE-risk measures involves the solution of a stochastic optimization problem. The key contribution of the present paper is to realize that by a first order condition the Monte Carlo estimation of OCE-risk measures can be reduced to a two-step-procedure: the first step consists in solving a stochastic root finding problem, the second step amounts to a standard Monte Carlo simulation of a value function at the argument that was computed in the first step. Our observation allows us to use stochastic root finding techniques for the purpose of risk measurement that could previously only be applied to a very specific class of risk measures, i.e. utility-based shortfall risk, see Dunkel and Weber (2010). The current paper shows that stochastic root finding possesses much higher relevance for financial risk management than would previously have been expected in view of the existing literature.

The article is organized as follows. Section 2 reviews the (now standard) axioms of risk measures, see Artzner et al. (1999) and Föllmer and Schied (2002). Section 3 presents the OCE-risk measures. We review examples (see Ben-Tal and Teboulle (2007)) illustrating that well-known risk measures like the entropic

risk measure and the average value at risk are special cases of OCEs, corresponding to an exponential and piecewise linear utility function, respectively. The computation of OCEs requires the solution of a stochastic optimization problem that can be reduced to a stochastic root finding problem under suitable conditions. We present suitable algorithms in Section 4. The risk measures discussed in the paper and the suggested algorithms are illustrated in the context of numerical case studies in Section 5. We calculate both the entropic risk measure and the average value at risk as well as a third OCE that does not correspond to a classical risk measure and is defined in terms of quartic utility. Section 6 concludes with a short summary of the main findings.

2 RISK MEASURES

For the convenience of the reader, the current section reviews standard notions of the theory of static risk measures. For a detailed description we refer to the excellent book by Föllmer and Schied (2011). We assume that there is one time period characterized by the current time $t = 0$ and the future $t = 1$. Letting $(\Omega, \mathscr{F})$ be a measurable space of possible scenarios, a financial position is modeled by a measurable mapping $X : \Omega \to \mathbb{R}$. The value $X(\omega) \in \mathbb{R}$ for a specific $\omega \in \Omega$ represents the discounted net future value of X if the scenario ω occurs. We fix a family $\mathscr{X}$ of financial positions, assuming that $\mathscr{X}$ is a vector space containing the constants. Our aim is to summarize the risk of any financial position $X \in \mathscr{X}$ by a number $\rho(X) \in \mathbb{R}$.

Definition 1 A map $\rho : \mathscr{X} \to \mathbb{R}$ is called *monetary risk measure*, if it satisfies the following conditions for $X, Y \in \mathscr{X}$:

  (i)   *Monotonicity:* If $X \leq Y$, then $\rho(X) \geq \rho(Y)$.
  (ii)   *Cash-Invariance:* $\rho(X + m) = \rho(X) - m$, for all $m \in \mathbb{R}$.

This definition formalizes that (i) risk should increase, if positions become worse, and that (ii) risk is measured on a monetary scale.

Definition 2 A monetary risk measure $\rho : \mathscr{X} \to \mathbb{R}$ is called *convex risk measure*, if it satisfies the following condition for $X, Y \in \mathscr{X}$ and $\lambda \in [0, 1]$:
$$\text{Convexity: } \rho(\lambda X + (1 - \lambda)Y) \leq \lambda \rho(X) + (1 - \lambda)\rho(Y).$$

Convexity formalizes that diversification of financial positions does not increase risk measurements.

Definition 3 A convex risk measure $\rho : \mathscr{X} \to \mathbb{R}$ is called *coherent risk measure*, if it satisfies the following condition for $X \in \mathscr{X}$:
$$\text{Positive Homogeneity: } \text{If } \lambda \geq 0, \text{ then } \rho(\lambda X) = \lambda \rho(X).$$

Example 4 We recall three well-known risk measures.

  1. **Value at risk.** For $\lambda \in (0, 1)$ we define

$$\text{VaR}_\lambda(X) := \inf \{m \in \mathbb{R} \mid P[X + m < 0] \leq \lambda\}.$$

  The value at risk at level λ yields the smallest monetary amount m that needs to be added to the financial position X in order to avoid that the probability of a loss exceeds λ. VaR is generally not convex, but only positively homogeneous.
  2. **Average value at risk.** For $\lambda \in (0, 1)$ we define the coherent risk measure

$$\text{AVaR}_\lambda(X) := \frac{1}{\lambda} \int_0^\lambda \text{VaR}_\nu(X) \, d\nu,$$

  which corresponds, under weak technical conditions, to the conditional expectation of a loss beyond the $\text{VaR}_\lambda(X)$. AVaR is also known as CVaR, i.e. conditional value at risk.

3. **Utility-based shortfall risk.** Letting $l : \mathbb{R} \to \mathbb{R}$ be an increasing, non-constant, convex loss function and $\lambda \in [0,1]$ be a threshold level, we define the convex risk measure

$$\text{UBSR}_{l,\lambda}(X) := \inf \{m \in \mathbb{R} \mid \mathbb{E}[l(-(X+m))] \leq \lambda\}.$$

Utility-based shortfall risk of a financial position X equals the smallest monetary amount m that needs to be added to X in order to avoid that the expected utility $-\mathbb{E}[l(-(X+m))]$ is less than the threshold level $-\lambda$. The special case $l(-X) = e^{-\beta X}$ admits the representation

$$\text{UBSR}_{l,\lambda}(X) = \frac{1}{\beta} \left(\log \left(\mathbb{E}\left[e^{-\beta X} \right] \right) - \log(\lambda) \right).$$

3 OPTIMIZED CERTAINTY EQUIVALENTS

The current section summarizes important results on OCEs that provide the basis for the algorithms suggested below. For further details we refer to the original reference Ben-Tal and Teboulle (2007).

In the following we let $u : \mathbb{R} \to [-\infty, \infty)$ be a concave and non-decreasing utility function satisfying $u(0) = 0$, $u(x) \geq 0 \ \forall x \geq 0$, and $u(x) < x \ \forall x$.

Definition 5 Letting u be a utility function as above and $X \in L^\infty$ be a random variable, the *optimized certainty equivalent (OCE)* is defined by the map $\text{OCE}_u : L^\infty \to \mathbb{R}$ with

$$\text{OCE}_u(X) = \sup_{\eta \in \mathbb{R}} \ \{\eta + \mathbb{E}[u(X - \eta)]\}. \tag{1}$$

As shown by Ben-Tal and Teboulle (2007), the negative of a OCE defines a convex risk measure. This is stated by the following corollary.

Corollary 6 Under the assumptions of Definition 5,

$$\rho(X) := -\text{OCE}_u(X), \quad X \in L^\infty,$$

defines a convex risk measure.

OCE thus yields a simple method for the construction of convex risk measures associated to a utility function u. As demonstrated in Ben-Tal and Teboulle (2007), the optimal η^* that maximizes (1) satisfies a first order condition – provided that u is continuously differentiable and strictly concave. The first order condition is given by the equation

$$\mathbb{E}[u'(X - \eta^*)] = 1. \tag{2}$$

This equation immediately leads to a stochastic root finding problem. An alternative approach would consist in a gradient-based simulation optimization procedure, as described in Kim (2006), Hong and Nelson (2009), Fu (2006), Fu (2008), Andradóttir (1998), and Henderson and Nelson (2006). In the current setting, the derivative of the function u that defines the risk measure can be computed directly.

Note that the derivation of Equation (2) involves an exchange of the order of differentiation and integration. If $X \in L^\infty$ and u is continuously differentiable and strictly concave, $u'(X - \eta)$ is almost surely bounded from above and below for any $\eta \in \mathbb{R}$ which justifies the first order condition. In the general case, the assumptions of the differentiation lemma need to be verified, see e.g. Lemma 16.2. in Bauer (2001).

Example 7 (a) Choosing the exponential utility function $u(t) = 1 - e^{-t}$, the negative of the OCE yields the entropic risk measure with parameter 1, i.e.

$$\rho(X) = \log \left(\mathbb{E}\left[e^{-X} \right] \right).$$

If we choose a utility function of the form $u(t) = 1 - e^{-\beta t}$, $\beta > 1$, the optimal η^* – computed according to (2) – coincides with the negative of utility-based shortfall risk with loss function $l(-X) = e^{-\beta X}$ and threshold level $\lambda = \frac{1}{\beta}$. The risk measure defined in Corollary 6 is thus given by

$$\rho(X) = \text{UBSR}_{l,\lambda}(X) - \frac{\beta - 1}{\beta}.$$

(b)   Setting

$$u(t) = \begin{cases} \alpha_2 t, & t \leq 0, \\ \alpha_1 t, & t > 0, \end{cases}$$

with $0 \leq \alpha_1 < 1 < \alpha_2$, we obtain

$$u(t) = \alpha_1 [t]_+ - \alpha_2 [-t]_+,$$

using the abbreviation $[z]_+ := \max\{0, z\}$, $z \in \mathbb{R}$. Calculating the optimal η^* according to (2), we obtain

$$\eta^* = F^{-1}\left(\frac{1 - \alpha_1}{\alpha_2 - \alpha_1}\right).$$

The special case $\alpha_1 = 0$ corresponds to

$$\eta^* = F^{-1}\left(\frac{1}{\alpha_2}\right) = \text{VaR}_{\frac{1}{\alpha_2}}(X).$$

In this case the OCE-risk measure defined in Corollary 6 coincides with the average value at risk at level $1/\alpha_2$, i.e.

$$\rho(X) = -\text{OCE}_u(X) = \text{AVaR}_{\frac{1}{\alpha_2}}(X).$$

Monte Carlo Methods   The main contribution of the current paper is a new Monte Carlo method for the computation of OCE-risk measures that are defined according to Corollary 6. The first order condition (2) allows for the application of stochastic root finding methods. These can be used instead of stochastic optimization procedures that are directly applied to the defining equation (1). We will explain two alternative approaches to compute the OCE-risk measure on the basis of stochastic root finding.

(i)   **Method 1:** According to equation (2) the optimal η^* in (1) is the unique root of the function

$$g(\eta) = \mathbb{E}[u'(X - \eta)] - 1. \tag{3}$$

The root η^* can thus be computed by a stochastic root finding scheme, as explained in the next section. In a second step, the associated OCE and the associated OCE-risk measure can be estimated using a standard Monte Carlo procedure, since

$$\text{OCE}_u(X) = \eta^* + \mathbb{E}[u(X - \eta^*)].$$

(ii)   **Method 2:** Alternatively, both the optimal η^* in (1) and the corresponding OCE could jointly be estimated by solving a 2-dimensional stochastic root finding problem. Denoting $\text{OCE}_u(X) = \xi^*$, the pair (ξ^*, η^*) satisfies

$$\begin{aligned} 0 &= \xi^* - \eta^* - \mathbb{E}[u(X - \eta^*)] \\ 0 &= \mathbb{E}[u'(X - \eta^*)] - 1 \end{aligned}$$

However, a 2-dimensional root finding algorithm turns out to be less efficient than Method 1. This is not surprising, since the equation characterizing η^* does not depend on ξ^*. It is thus preferable to estimate η^* first, and to compute ξ^* only for an approximately correct value of η^*. Intermediate computations of ξ^* for less precise estimates of η^* do not accelerate the computation of η^* – as it would be the case in a fully coupled system. For this reason, we will concentrate on Method 1 in the following sections.

4   STOCHASTIC APPROXIMATION

In this section we provide a brief survey of stochastic root finding algorithms and explain how these can be applied to the computation of OCEs and the corresponding risk measures. For further information on stochastic root finding, we refer to Dunkel and Weber (2010) and the references therein. A very detailed description of the topic is provided by Kushner and Yin (2003).

Robbins-Monro Algorithm. Our purpose is to construct an algorithm yielding a sequence of estimators (η_n), $n \in \mathbb{N}$, which converges quickly to the sought root η^* of the function $g(\eta)$ defined in (3). The OCE can then be obtained in a second step by a Monte Carlo procedure.

In many situations, stochastic root finding will not be based on naive Monte Carlo procedures, but on sophisticated variance reduction techniques. It is thus advisable to describe the algorithms in a way that allows for variance reduction. This approach is, for example, explained in Dunkel and Weber (2010). Moreover, in practice random numbers are typically drawn uniformly on the interval $(0,1)$. We will thus use uniform random variables as the starting point for the simulation scheme.

For practical and mathematical reasons we design algorithms that are restricted to a bounded domain. Letting $-\infty < a < \eta^* < b < \infty$, we define for $t \in \mathbb{R}$ the projection

$$\pi(t) = \begin{cases} a, & t \le a, \\ t, & a < t < b, \\ b, & b \le t. \end{cases}$$

Assume that for any $\eta \in [a,b]$ there exists a known function $Y_\eta : (0,1) \to \mathbb{R}$ such that for any on $(0,1)$ uniformly distributed random variable U we have

$$g(\eta) = \mathbb{E}[Y_\eta(U)].$$

The Robbins-Monro algorithm for the estimation of η^* is constructed as follows:

1. Choose
    – a constant $\gamma \in \left(\frac{1}{2}, 1\right]$,
    – a constant $c > 0$, and
    – a starting point $\eta_1 \in [a,b]$.
2. Calculate the sequence of estimators recursively:

$$\eta_{n+1} = \pi\left(\eta_n + \frac{c}{n^\gamma} Y_{\eta_n}(U_n)\right), \quad n \in \mathbb{N},$$

for a sequence (U_n) of independent, $\mathscr{U}(0,1)$-distributed random variables.

red In the case considered in the paper, a simple choice for Y_η could consist in $Y_\eta(U) = 1 - u'(q(U) - \eta)$, where q denotes the quantile function of X.

Theorem 8 Suppose that the following assumptions hold:

(i)   The function $g(\eta)$ is well-defined and finite for all $\eta \in [a,b]$.

(ii) $\sup_{\eta \in [a,b]} \text{Var}[Y_\eta(U_n)] < \infty.$

Then the sequence $(\eta_n)_n$ converges *P*-almost surely to the root η^* of function (3).

Besides the consistency of the algorithm, its convergence rate can be further analyzed on the basis of limit theorems for the rescaled quantities $\sqrt{n^\gamma}(\eta_n - \eta^*)$, $n \in \mathbb{N}$. For precise conditions we refer to Dunkel and Weber (2010).

Theorem 9 Under suitable conditions, the rescaled quantities exhibit the following asymptotic behavior: If $\gamma = 1$, then

$$\sqrt{n}(\eta_n - \eta^*) \to \mathcal{N}\left(0, \frac{-c^2\text{Var}[Y_{\eta^*}(U)]}{2cg'(\eta^*)+1}\right).$$

If $\gamma \in \left(\frac{1}{2}, 1\right)$, then

$$\sqrt{n^\gamma}(\eta_n - \eta^*) \to \mathcal{N}\left(0, \frac{-c\text{Var}[Y_{\eta^*}(U)]}{2g'(\eta^*)}\right).$$

5   NUMERICAL CASE STUDIES

The aim of the current section is to illustrate the properties of Method 1 in numerical case studies. We focus on three examples of OCEs: the entropic risk measure, average value at risk, and an OCE defined in terms of quartic utility. In the first two cases, we analyze the example in which X is standard normally distributed. In the third case in which the OCE corresponds to quartic utility, we investigate the impact of the properties of the distribution. We provide results for X being distributed either according to a normal distribution or a Student t-distribution. Both distributions are assumed to have the same mean and variance. Observe that t-distributions are heavy tailed.

The estimation of the root η^* of (3) is compared for different Robbins-Monro algorithms (RM). We generated results from the recursive algorithms for values $n = 100, 300, 1\,000, 3\,000, 10\,000, 30\,000$. For each recursion depth n we repeated the simulation $N = 10^4$ times in order to obtain an empirical distribution of the RM-estimator of the root. The starting point η_1 of each run was sampled uniformly from the interval $[a,b]$, where a and b are the lower and upper projection-bounds. In the cases of the entropic risk measure and average value at risk we set $[a,b] := [\eta^* - 5, \eta^* + 5]$; in the case of quartic utility we choose $[a,b] := [\eta^* - 10, \eta^* + 10]$. In real-world applications η^* is, of course, unknown. Choosing the bounds a and b will in this case require pre-estimates of η^* before running the algorithms. The fact that we assumed that the projection interval is symmetric about η^* is not crucial for our results. The Robbins-Monro algorithm was simulated for the parameters $\gamma = 0.7$ and $\gamma = 1$ setting $c = 1$.

5.1 Entropic Risk Measure

Let u be the exponential utility function $u(t) = 1 - e^{-2t}$. Assuming that X is standard normally distributed we can compute the optimal η^* exactly, i.e. $\eta^* = -\frac{\log(2)}{2} - 1 \approx -1.34657$. This yields $\text{OCE}_u(X) = -\frac{\log(2)}{2} - \frac{1}{2} \approx -0.846574$. Consequently, this example provides a good basis for testing our methods.

Figure 1 illustrates the distributions of 10^4 runs of the stochastic root finding algorithm for each recursion depth. The drop shapes in the charts are generated based on a smooth kernel density estimation. The estimated PDFs are mirrored along the η-axis (vertical axis) and scaled in the width. For increasing step size the distributions are getting more Gaussian like with decreasing variance.

For $n = 100$ the right panel ($\gamma = 1$) of Figure 1 shows a bimodal distribution. The same effect can be observed in the other examples presented below. This is due to the fact that the initial root estimates hit the bounds of the simulation interval quite frequently. If γ is large, the correction terms of the iterative scheme decrease faster than for smaller γ. The initial accumulation of mass at the boundaries of the projection

interval converges to the sought root as *n* grows. For $\gamma = 1$ a bimodal distribution remains visible in the histograms.

Table 1 shows the mean and the variance of the simulation depending on the number of iterations and the constant γ. The bias decreases in both cases in a qualitatively similar way, but the variance decreases faster for $\gamma = 1$.

5.2 Average Value at Risk

Let *u* be defined according to Example 7(b), i.e. $u(t) = \alpha_1[t]_+ - \alpha_2[-t]_+$. As explained above, we will calculate η^* as the value at risk at level $\frac{1}{\alpha_2}$ according to the first order condition. Afterwards we can simulate the OCE which yields the average value at risk at level $\frac{1}{\alpha_2}$ in the case $\alpha_1 = 0$.

Figure 2 illustrates the distributions of 10^4 runs of the stochastic root finding algorithm for a given iteration size between 100 and 30 000. The figure shows that the variance decreases rapidly and that for larger number of iterations the performance of the algorithm is better for $\gamma = 1$. Mean and variance of the empirical distributions are provided in Table 2.

Figure 3 illustrates a path of the calculation of the OCE which is the average value at risk for the piecewise linear utility function. The path was generated by simple Monte Carlo simulation.

Table 1: Mean and variance of the empirical distribution of η_n depending on the number of iteration steps *n* and the constant γ.

iteration steps *n*	$\gamma = 0.7$		$\gamma = 1$	
	mean	variance	mean	variance
100	-1.8373	0.2214	-1.9000	0.5051
300	-1.7296	0.0946	-1.5774	0.2096
1000	-1.4639	0.2552	-1.4245	0.0798
3000	-1.4032	0.1083	-1.3725	0.0251
10000	-1.3722	0.0445	-1.3543	0.0047
30000	-1.3589	0.0195	-1.3493	0.0017
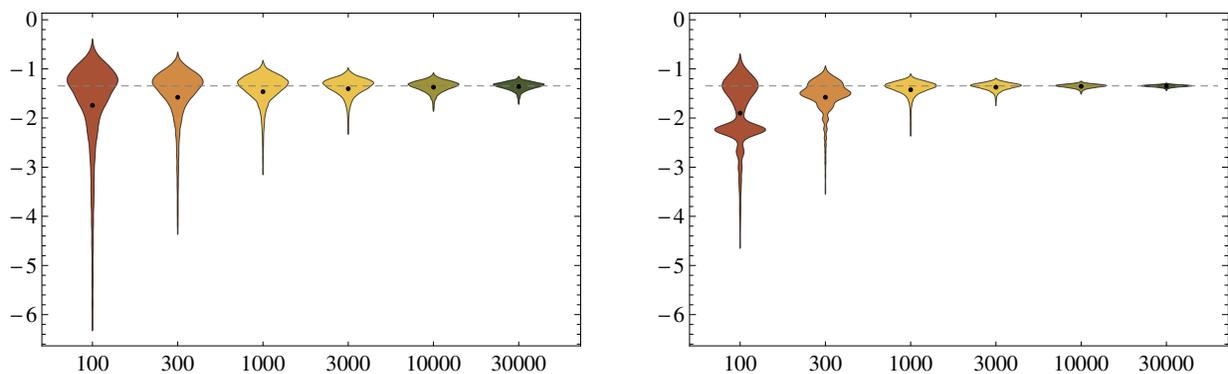
Figure 1: Distribution charts for the RM algorithm with recursion depth $n \in \{100, ..., 30\,000\}$ and (left: $\gamma = 0.7$, right: $\gamma = 1.0$) for the exponential utility function. The black dots represent the expectations of the distributions and the gray dashed line marks the exact value of the root.

The precision of the Monte Carlo estimation of the OCE or corresponding risk measure does, of course, depend on the precision of the estimate of η^*. For a fixed simulation depth *n*, both bias and variance

contribute to the error in the estimation of η^*. Figure 4 illustrates how bias and variance in the estimation of η^* influence the average estimate of average value at risk in step 2 of the simulation procedure.

The variance of the estimation of η^* can be reduced to zero, if we replace its random estimator by the mean of the empirical distributions shown in Figure 2. The error of the resulting estimate of η^* for various recursion depths is then essentially only due to the bias of the stochastic root finding algorithms. The purple dots in Figure 4 show the corresponding Monte Carlo estimates for average value at risk, if the mean of the empirical distribution is used as input parameter η^* in step 2 of the computation of the OCE.

Table 2: Mean and variance of the empirical distribution of η_n depending on the number of iteration steps n and the constant γ.

iteration steps n	$\gamma = 0.7$		$\gamma = 1$	
	mean	variance	mean	variance
100	-1.8432	0.2219	-2.2636	0.2590
300	-1.7211	0.0927	-1.8101	0.0413
1000	-1.6782	0.0388	-1.6696	0.0066
3000	-1.6601	0.0176	-1.6504	0.0020
10000	-1.6509	0.0073	-1.6462	0.0006
30000	-1.6484	0.0035	-1.6453	0.0002
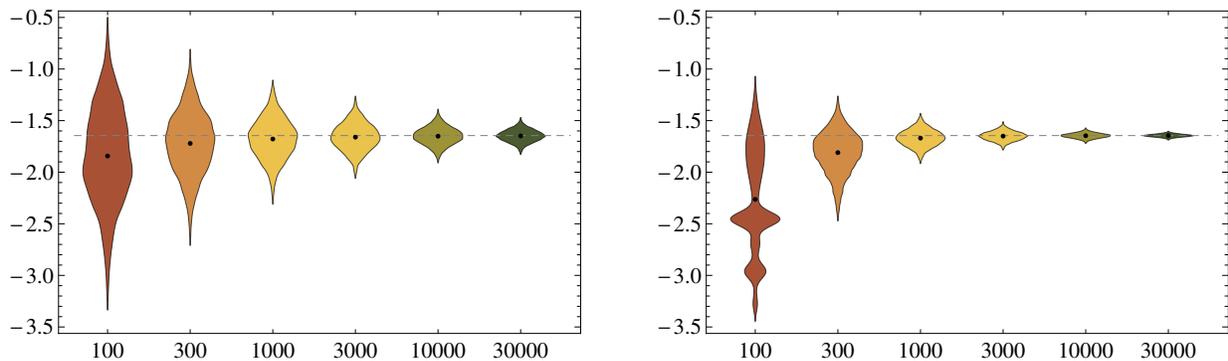
Figure 2: Distribution chart for the RM algorithm with recursion depth $n \in \{100, ..., 30\,000\}$ and (left: $\gamma = 0.7$, right: $\gamma = 1.0$) for the piecewise linear utility function. The black dots represent the mean of the distributions and the gray dashed line marks the exact value of the root.
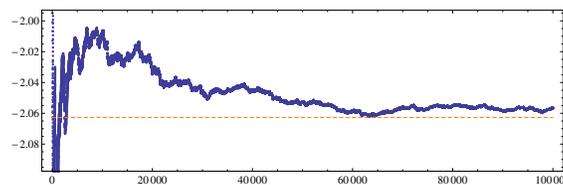
Figure 3: A path of a Monte Carlo estimation of the OCE for the piecewise linear utility function. The dashed orange line represents the exact value.

In practice, the estimation of η^* will typically not be repeated several times. A single estimate will be used as the input parameter for the Monte Carlo simulation of the OCE. This corresponds to sampling η^* from the empirical distributions shown in Figure 2. The blue dots in Figure 4 show the sample average

of the corresponding Monte Carlo estimates for average value at risk. As expected, the comparison in Figure 4 shows that the bias of the estimator of AVaR is significantly larger in the second case.

5.3 Quartic Utility

Assume that

$$u(t) := \begin{cases} 1 - (t-1)^4, & \text{if } t \leq 1, \\ 1, & \text{otherwise.} \end{cases}$$

We compare two different distributions for the financial position X. In the first case we assume that X is normally distributed with variance $\sigma^2 = \frac{5}{3}$, in the second case we suppose that X has a Student t-distribution with $v = 5$ degrees of freedom. Both random variables have the same mean and variance, but different tail behavior. In both cases, the expression $\eta + \mathbb{E}[u(X - \eta)]$ can be written as a closed-form expression of η. This allows us to obtain the exact value of the OCE. We compute that $\eta^* = -2.16359$ and $\text{OCE}_u(X) = -1.6511$ in the first case, i.e. if X is normally distributed. In the second case, i.e. if X is Student t-distributed, we get $\eta^* = -3.73624$ and $\text{OCE}_u(X) = -5.93075$. As expected, the heavy tailed Student t-distribution does indeed lead to a higher OCE-risk measure.
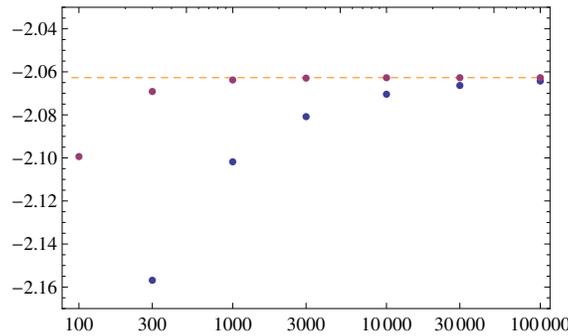
Figure 4: Convergence of the Monte Carlo estimation of the OCE for η having the empirical distribution (blue) or being Dirac distributed at the mean of the empirical distributions (purple).

Figures 5 and 6 illustrate the empirical distributions of the estimates of η^* which are sampled according to RM algorithms. The projection interval was set to $\eta^* \pm 10$. The figures show that for quartic utility and finite sample sizes the estimation errors can be smaller for $\gamma = 0.7$ than for $\gamma = 1$. Finite sample properties might, thus, deviate from the asymptotic behavior suggested by the limit theorems above.

Figure 6 illustrates that the algorithm might actually perform quite badly for heavy tailed distributions like the t-distribution. Variance reduction techniques are in this case a crucial tool that needs to be employed in order to secure reasonable simulation results. For a discussion of this insight in a different context we refer to Dunkel and Weber (2010).

6 CONCLUSION

A large family of OCE-risk measures has been introduced by Ben-Tal and Teboulle (2007) which includes well-known examples like average value at risk and the entropic risk measure. The OCE-risk measures can be computed by solving a stochastic optimization problem. Under weak conditions, by a first order condition, the calculation can be reduced to a two step procedure: a stochastic root finding problem in the first step, and the computation of an expected value in the second step. In the current paper we suggest algorithms for the estimation of OCE-risk measures and test them numerically.
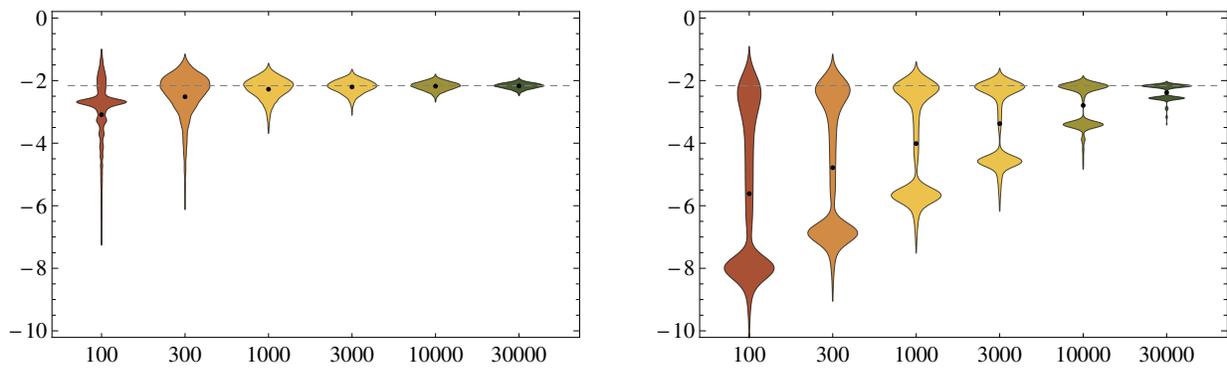
Figure 5: Distribution chart for the RM algorithm with recursion depth $n \in \{100, ..., 30\,000\}$ and (left: $\gamma = 0.7$, right: $\gamma = 1.0$) for the piecewise quartic utility function and a normally distributed random variable.
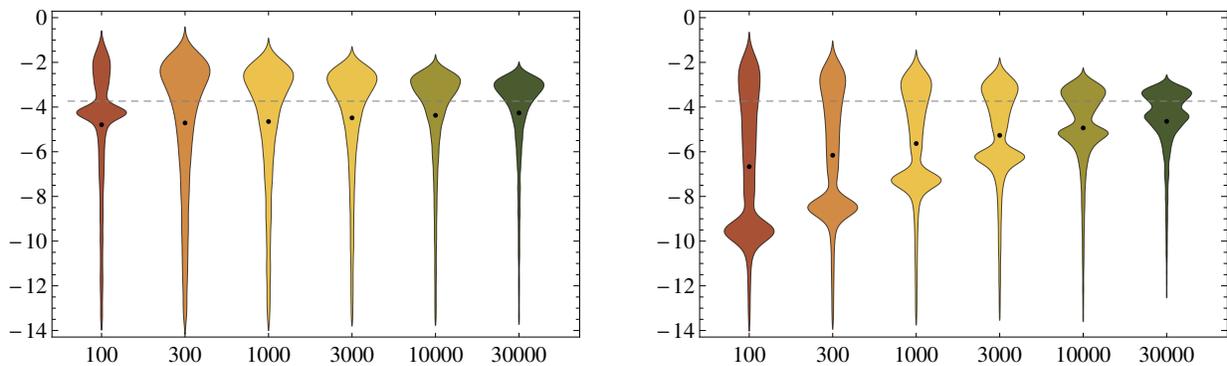
Figure 6: Distribution chart for the RM algorithm with recursion depth $n \in \{100, ..., 30\,000\}$ and (left: $\gamma = 0.7$, right: $\gamma = 1.0$) for the piecewise quartic utility function and a Students t-distribution with $v = 5$ degrees of freedom.

The performance of the algorithms depends largely on the properties of the distribution of the financial position whose risk is estimated. The performance of the algorithms is reasonable for distributions with bounded variance and light tails. However, the suggested simulation procedures are usually not optimal, if distributions are heavy tailed. In this case, it seems to be crucial to employ variance reduction techniques in order to improve performance.

The suggested techniques provide promising algorithms for the estimation of a broad class of convex risk measures based on utility theory. Future research needs to investigate how their performance can be improved in the case of distributions with high variance or heavy tails and how they can efficiently be applied to real-world examples.

REFERENCES

Andradóttir, S. 1998. "A Review of Simulation Optimization Techniques". In *Proceedings of the 1998 Winter Simulation Conference*, edited by D. Medeiros, E. Watson, J. Carson, and M. Manivaninnan, 151–158. Pistacaway, New Jersey: IEEE Press.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. 1999. "Coherent Measures of Risk". *Mathematical Finance* 9 (3): 203–228.

Bauer, H. 2001. *Measure and Integration Theory*, Volume 26. Berlin: Walter de Gruyter.

Ben-Tal, A., and M. Teboulle. 2007. "An old-new Concept of Convex Risk Measures: The Optimized Certainty Equivalent". *Mathematical Finance* 17 (3): 449–476.

Dunkel, J., and S. Weber. 2010. "Stochastic Root Finding and Efficient Estimation of Convex Risk Measures". *Operations Research* 58 (5): 1505–1521.

Föllmer, H., and A. Schied. 2002. "Convex Measures of Risk and Trading Constraints". *Finance and Stochastics* 6 (4): 429–447.

Föllmer, H., and A. Schied. 2011. *Stochastic Finance. An Introduction in Discrete Time*. Berlin: Walter de Gruyter & Co.

Fu, M. C. 2006. "Gradient Estimation". *Handbooks in Operations Research and Management Science* 13:575–616.

Fu, M. C. 2008. "What You Should Know About Simulation and Derivatives". *Naval Research Logistics (NRL)* 55 (8): 723–736.

Henderson, S. G., and B. L. Nelson. 2006. *Simulation*, Volume 13 of *Handbooks in Operations Research and Management Science*. Amsterdam: North-Holland.

Hong, L. J., and B. L. Nelson. 2009. "A Brief Introduction to Optimization via Simulation". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 75–85. Pistacaway, New Jersey: IEEE Press.

Kim, S. 2006. "Gradient-Based Simulation Optimization". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 159–167. Pistacaway, New Jersey: IEEE Press.

Kushner, H.-J., and G.-G. Yin. 2003. *Stochastic Approximation and Recursive Algorithms and Applications*. 2nd ed. New York: Springer.

AUTHOR BIOGRAPHIES

ANNA-MARIA HAMM studied mathematics and economics at the University of Hannover, Germany, from October 2006 to January 2012. Since February 2012 she is a Ph.D. student at the Institute of Probability and Statistics in Hannover. Furthermore she holds a position as research fellow at the Competence Center for Risk and Insurance, Hannover. Her email address is hamm@stochastik.uni-hannover.de.

THOMAS SALFELD studied physics and mathematics at the University of Hannover. Since April 2011 he is a Ph.D. student at the Institute of Probability and Statistics in Hannover. His email address is salfeld@stochastik.uni-hannover.de.

STEFAN WEBER is a Professor of Insurance and Financial Mathematics at the Institute of Probability and Statistics in Hannover, Germany, and a Managing Director at the Center for Risk and Insurance, Hannover, Germany. Prior to this, he was an Assistant Professor in Financial Engineering and Applied Probability in the School of Operations Research and Information Engineering at Cornell University. He holds a Ph.D. in Mathematics from Humboldt-Universität zu Berlin. His email address is sweber@stochastik.uni-hannover.de.