

A NESTED SIMULATION OPTIMIZATION APPROACH FOR PORTFOLIO SELECTION

Yifan Lin
Enlu Zhou

Industrial and Systems Engineering
Georgia Institute of Technology
755 Ferst Drive NW
Atlanta, GA 30332, USA

Aly Megahed

IBM Research - Almaden
650 Harry Rd
San Jose, CA 95120, USA

ABSTRACT

We consider the problem of portfolio selection with risk factors, where the goal is to select the portfolio position that minimizes the value at risk (VaR) of the expected portfolio loss. The problem is computationally challenging due to the nested structure caused by the risk measure VaR of the conditional expectation, along with the optimization over a discrete and finite solution space. We develop a nested simulation optimization approach to solve this problem. In the outer layer, we adapt the optimal computing budget allocation (OCBA) approach to sequentially allocate the simulation budget of the outer-layer to different portfolio positions. In the inner layer, we propose a new sequential procedure to efficiently estimate the VaR of the expected loss. We present a numerical example that shows that our approach achieves a higher probability of correct selection under the same computing budget compared to three other methods.

1 INTRODUCTION

Portfolio selection is an important problem in financial engineering that has come a long way since the work of Markowitz (1952). In this paper, we consider the problem of portfolio selection with risk factors (e.g., future interest rates). Each risk factor incorporates sufficient information so as to determine the assets prices at some future time (risk horizon). The objective is to select the optimal portfolio position that minimizes the extreme quantile (induced by the risk factors) of the expected portfolio loss. In order to measure the extreme quantile, an upper percentile of the loss distribution called value-at-risk (VaR) is used. VaR is a widely used risk measure in financial risk management, especially with practitioners, because its concept is easily understandable (see Jorion (1997) for a comprehensive introduction to VaR).

The objective function of our problem has a nested structure that consists of the outer risk measure VaR and the inner conditional expectation. We use VaR-Expectation to represent this objective function. To estimate the objective function, nested Monte Carlo simulation is a natural tool. Specifically, in the outer layer, samples of future scenarios of risk factors are generated; in the inner layer, sample paths of the asset price are generated, and the expected portfolio loss is then evaluated at the risk horizon, given the risk scenario generated in the outer layer. Such nested estimation problem has been studied extensively (e.g., Lee (1998), Lee and Glynn (2003), Sun et al. (2011), Gordy and Juneja (2010), Liu et al. (2010), Lan et al. (2010), Broadie et al. (2011), Zhu et al. (2020), and Dang et al. (2019)). In particular, both Gordy and Juneja (2010) and Zhu et al. (2020) consider the same estimation problem as ours, and they study asymptotic properties of the uniform nested estimator, which generates the same number of inner-layer sample paths for each outer scenario. Although uniform nested simulation is relatively easy to implement and understand, it is often computationally expensive, which hinders its usage in practice. Hence many non-uniform estimators have been proposed to improve efficiency by smartly allocating different number of inner sample paths to each outer scenario. In the literature of non-uniform estimators, Lee

and Glynn (2003) propose a non-uniform estimator of the distribution function of conditional expectation, where the inner sample size is outcome dependent; Liu et al. (2010) and Lan et al. (2010) both develop non-uniform estimation schemes for expected shortfall; Dang et al. (2019) propose an iterative procedure called importance-allocated nested simulation that determines the allocation of inner budget to different scenarios. The most relevant work to ours is Broadie et al. (2011). They propose a non-uniform estimator for probability of large loss and show it greatly reduces bias and variance of the uniform estimator under the same total budget. Inspired by their estimator, we propose a non-uniform estimator for VaR-Expectation based on the marginal change of the estimator when an additional inner sample budget is allocated. The main differences between their work and ours lie in the different risk measures (probability of large loss vs. VaR) and the goals for budget allocation. Specifically for the latter difference, they consider an estimation problem and focus on minimizing the mean squared error (MSE) of the estimator, while our goal is optimization and hence the budget allocation approach is different. Apart from nested simulation, another stream in the literature for estimating the unknown loss distribution function utilizes regression or meta-modeling techniques (e.g., Liu and Staum (2010), Broadie et al. (2015), Hong et al. (2017)).

Our portfolio selection problem goes beyond nested estimation of the objective function, because of the optimization over a solution space of possible portfolio positions. To make the problem more tractable, we assume the solution space is discrete and finite. This problem falls into the general framework of ranking and selection (R&S), a special category of simulation optimization. In R&S, the goal is to select the best alternative among a finite number of alternatives via simulation (see Goldsman and Nelson (2007) and Fu (2015) for comprehensive overviews of R&S). In nearly all of the existing R&S literature, the performance measure is the mean performance of the alternatives. However, in our problem, the best alternative is the one with the smallest VaR. Despite the wide use of quantile as a performance measure, quantile-based R&S has not received much attention in the literature with a few exceptions, such as the works of Bekki et al. (2007), Lee and Nelson (2014), Shin et al. (2016), and Peng et al. (2020). In our approach, we adapt the well-known optimal computing budget allocation (OCBA) scheme (Chen and Lee (2010)) to the performance measure based on VaR, and seamlessly incorporate it with our non-uniform nested estimator for VaR. Specifically, the extended OCBA scheme sequentially determines the outer-layer budget (i.e., number of scenarios generated for each solution), and the non-uniform nested estimator for VaR allocates the inner-layer budget (i.e., number of sample paths generated for each scenario). We note that the recent work by Peng et al. (2020) also extends OCBA to consider a quantile-based performance measure and derive specific allocation schemes for normal, exponential, and Pareto distributions of the solution performance. However, in the portfolio selection problem, the underlying distribution of a solution performance usually does not have a closed form because of the complex dynamics of asset prices, so we can only rely on sampling to approximate the solution performance.

To summarize, the contributions of this paper are two-fold. First, we propose a nested simulation optimization approach that adapts OCBA to the case where the performance measure is based on quantile (VaR) and the underlying distribution does not admit a closed form. Second, we develop a new non-uniform nested estimator of VaR-Expectation. The remainder of the paper is organized as follows: Section 2 introduces the problem of portfolio selection with risk factors and formulates the problem. Section 3 develops our proposed nested simulation optimization algorithm. Section 4 presents a numerical example of the portfolio selection problem, and Section 5 concludes the paper.

2 PROBLEM FORMULATION

We consider the problem of portfolio selection with risk factors. Without loss of generality, we assume the portfolio consists of a risky asset and a risk-free asset such as T-bills. The investor makes the investment decision at time 0 and holds the portfolio until time $T > 0$. There is a risk factor θ , such as the interest rate, at a future time $\tau \in (0, T)$, which is called the risk horizon. The objective is to select the optimal investment level in the risky asset, in order to minimize VaR (with respect to the risk factor θ) of the expected portfolio loss at the risk horizon τ . For a random variable Z , $\text{VaR}^\alpha(Z)$ is defined

as the α -quantile of Z , i.e., $\text{VaR}^\alpha(Z) := \inf\{t : \mathbb{P}(Z \leq t) \geq \alpha\}$, where $\alpha \in (0, 1)$ is called the risk level. We use x to denote the investment level in the risky asset, and assume x can only take discrete values $\{x_1, x_2, \dots, x_M : 0 \leq x_1 < x_2 < \dots < x_M \leq 1\}$. Let X_t denote the value of the portfolio at time t . The asset price follows a certain stochastic process driven by an underlying random variable ξ that depends on the risk factor θ at time τ . Then, the expected portfolio loss is $H(x, \theta) = E_{\xi \sim \mathbb{P}_{\xi|\theta}}[l(X_T(x, \xi))]$, where l is a loss function, while $X_T(x, \xi)$ signifies the final portfolio value X_T and is a random variable that depends on the investment level x and the randomness ξ that underlies the asset price. Except for some simple models for asset prices that allow closed-form expressions, the asset price has to be evaluated by simulating the stochastic process starting from the initial price. Hence, the portfolio value X_T usually has to be evaluated through Monte Carlo simulation.

The above portfolio selection problem can be formulated as:

$$\min_{x \in \mathcal{X}} \text{VaR}_{\theta \sim \mathbb{P}_\theta}^\alpha [\mathbb{E}_{\xi \sim \mathbb{P}_{\xi|\theta}}[h(x, \xi)]] \tag{1}$$

where $x \in \mathcal{X}$ is the decision variable, θ is a random variable that follows the distribution \mathbb{P}_θ , ξ is a random variable that follows a θ -dependent distribution $\mathbb{P}_{\xi|\theta}$. We assume the distribution $\mathbb{P}_{\xi|\theta}$ has a known parametric form with θ as the parameter. The solution space \mathcal{X} is a finite set, i.e., $\mathcal{X} := \{x_1, x_2, \dots, x_M\}$, where M is the number of total alternatives. The function $h : \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$ is called the performance measure and has to be evaluated by simulation. For notational convenience, we define the function $H(x, \theta) := \mathbb{E}_{\xi \sim \mathbb{P}_{\xi|\theta}}[h(x, \xi)]$, which will be referred to as the mean response function, and define $v_\alpha(x) := \text{VaR}_{\theta \sim \mathbb{P}_\theta}^\alpha(H(x, \theta))$. Except in few special cases (when v_α , for a fixed x , can be computed analytically), it usually has to be estimated.

Problem (1) is essentially a R&S problem, where we aim to select the best solution from M alternatives x_1, \dots, x_M with a high probability of correct selection (PCS) under a given simulation budget. However, it differs from the traditional R&S in the sense that its objective function has a nested structure of VaR-Expectation (instead of expectation usually considered in R&S). Hence, the objective function naturally calls for nested simulation for risk measurement. In the next section, we will develop an approach that integrates a R&S procedure with nested simulation.

3 ALGORITHM DEVELOPMENT

3.1 Overview

In this subsection, we provide an overview of our proposed approach to solve (1). To reveal the structure and challenge of the problem, we first present a simple algorithm; the two-layer equal allocation (EA) algorithm.

- In the outer layer, for each fixed $x \in \mathcal{X}$, we generate N i.i.d. samples $\{\theta_1, \dots, \theta_N\}$ from \mathbb{P}_θ .
- In the inner layer, for each fixed θ_i ($i = 1, \dots, N$), generate K i.i.d. samples $\{\xi_{i,1}, \dots, \xi_{i,K}\}$ from $\mathbb{P}_{\xi|\theta_i}$, and evaluate $\{h(\xi_{i,j}) : j = 1, \dots, K\}$. Then, approximate $H(x, \theta_i)$ by the sample average $\bar{H}(x, \theta_i) = \frac{1}{K} \sum_{j=1}^K h(x, \xi_{i,j})$ for all θ_i 's, and sort them in the ascending order, denoted by $\bar{H}(x, \theta^{(1)}) \leq \bar{H}(x, \theta^{(2)}) \leq \dots \leq \bar{H}(x, \theta^{(N)})$. The estimator of $v_\alpha(x)$ is then $\hat{v}_\alpha(x) = \bar{H}(x, \theta^{(\alpha N)})$.
- Return the x that has the smallest $\hat{v}_\alpha(x)$ as the solution.

The EA algorithm above requires a total simulation budget $M * N * K$, where M is the total number of alternatives (i.e., cardinality of \mathcal{X}), N is the outer-layer sample size, and K is the inner-layer sample size. To achieve a high probability of correct selection, N and K need to be sufficiently large, and hence the total budget needs to be substantially large. In fact, the outer layer allocates budget to different alternatives in order to select the best alternative, and the inner layer aims to estimate the VaR-Expectation objective value of each alternative, which itself needs a two-layer nested simulation.

We will develop a sequential algorithm that efficiently allocates the budget in both layers. In a nutshell, we adapt the OCBA approach to generate samples of θ for each alternative x in the outer layer, and propose

a sequential non-uniform estimation procedure that generates samples of ξ one at a time in the inner layer. On a high level, the OCBA approach generates more samples of θ for the alternatives that are more likely to be the optimal ones or have more estimation uncertainty in their objective value; the proposed non-uniform estimation procedure sequentially generates the next sample of ξ that would have the largest impact on the improving the estimator for VaR-Expectation. Therefore, both layers are made more efficiently compared to the EA algorithm.

In the next two subsections, we will focus on the derivation of a nested estimator for the VaR of expectation, denoted as VaR-E for short in the remainder of the paper. In Subsection 3.4, we will adapt the OCBA procedure relying on the asymptotic normality of the estimators, and present our algorithm.

3.2 Uniform Estimator for VaR-E

In this and the next subsections, we focus on the inner layer estimation of VaR-E for a fixed x . Hence, we temporarily drop the notation x in all of the expressions.

The estimator \hat{v}_α in the EA algorithm is in fact the uniform estimator for VaR-E, since it has the same number of ξ samples across all θ 's. Recall that in the outer stage, we first generate N i.i.d. samples $\{\theta_i, i = 1, 2, \dots, N\}$. In the inner stage, for each θ_i , we generate K i.i.d. random variables $\{\xi_j, j = 1, \dots, K\}$ from distribution $\mathbb{P}_{\xi|\theta_i}$ and evaluate $\{h(\xi_j) : j = 1, \dots, K\}$. Then, we approximate $H(\theta_i)$ by the sample average $\bar{H}(\theta_i) = \frac{1}{K} \sum_{j=1}^K h(\xi_j)$, and sort them in the ascending order, denoted by $\bar{H}(\theta^{(1)}) \leq \bar{H}(\theta^{(2)}) \leq \dots \leq \bar{H}(\theta^{(N)})$. The uniform estimator for VaR-E is then $\hat{v}_\alpha := \bar{H}(\theta^{(\alpha N)})$. As shown by Theorem 3.4 in Zhu et al. (2020), under the following Assumption 1, this uniform estimator is strongly consistent, i.e., $\lim_{N,K \rightarrow \infty} \hat{v}_\alpha = v_\alpha$ w.p.1.

Assumption 1 (Assumption 3.2 in Zhu et al. (2020)) Let $\bar{H}(\theta) = \frac{1}{K} \sum_{j=1}^K (H(\theta) + \varepsilon(\xi_j)) = H(\theta) + \frac{1}{K} \sum_{j=1}^K \varepsilon(\xi_j)$, where $\xi_j \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_{\xi|\theta}, j = 1, \dots, K$. Define $\bar{\varepsilon}_K \triangleq \sqrt{K} \cdot \frac{1}{K} \sum_{j=1}^K \varepsilon(\xi_j)$. Then,

- The performance measure function $h(\xi)$ has finite conditional second moment, i.e., $\tau_\theta^2 = \mathbb{E}[h^2(\xi)|\theta] < \infty$ with probability 1 (w.p.1) and $\tau^2 = \int \tau_\theta^2 p(\theta) d\theta < \infty$.
- The joint density $p(h, e)$ of $H(\theta)$ and $\bar{\varepsilon}_K$, and its partial derivatives $\frac{\partial}{\partial h} p(h, e)$ and $\frac{\partial^2}{\partial h^2} p(h, e)$ exist for all the pairs of (h, e) .
- There exist non-negative functions $g_0(\cdot), g_1(\cdot), g_2(\cdot)$ such that $p(h, e) \leq g_0(e), \left| \frac{\partial}{\partial h} p(h, e) \right| \leq g_1(e), \left| \frac{\partial^2}{\partial h^2} p(h, e) \right| \leq g_2(e)$ for all (h, e) . Further, $\int |e|^r g_i(e) de < \infty$ for $i = 0, 1, 2$, and $0 \leq r \leq 4$.

The first item in Assumption 1 ensures that $\bar{\varepsilon}_K$ has a limiting distribution as $K \rightarrow \infty$. Letting $\tilde{f}_K(\cdot)$ denote the probability density function (PDF) of $\bar{H}(\theta)$, then the second and third assumptions ensure the difference between $\tilde{f}_K(\cdot)$ and $f(\cdot)$ is of the order $O(\frac{1}{K})$. The next theorem shows that the estimator \hat{v}_α is asymptotically normally distributed when the sample sizes satisfy $N = O(K^2)$.

Theorem 1 (Theorem 3.6 in Zhu et al. (2020)) **Asymptotic Normality of Estimator.** Define $\Lambda(t) = \frac{1}{2} f(t) \mathbb{E}[\tau_\theta^2 | H(\theta) = t]$, $\sigma := \frac{\sqrt{\alpha(1-\alpha)}}{f(v_\alpha)}$, $\mu := \frac{-\Lambda'(v_\alpha)}{f(v_\alpha)}$, and $\mathcal{N}(0, 1)$ stands for standard normal random variable. Then, under Assumption 1, the existence of limit $R^2 = \lim_{N,K \rightarrow \infty} N/K^2$ is a sufficient and necessary condition for

$$\lim_{N,K \rightarrow \infty} \sqrt{N}(\hat{v}_\alpha - v_\alpha) \stackrel{\mathcal{D}}{=} \sigma \mathcal{N}(0, 1) + |R|\mu.$$

If we denote by θ^α the θ -value at the quantile v_α , it is easy to see that the uniform estimator \hat{v}_α wastes a large portion of simulation budget on the θ samples that are not in the close neighborhood of θ^α and thus are unlikely to affect the estimator. This motivates us to develop a non-uniform estimator, where we

consider sequentially allocating the inner simulation budget (i.e., generating an ξ sample) one at a time so that it will have the largest impact on the estimator.

3.3 Non-Uniform Sequential Estimation for VaR-E

Our proposed non-uniform estimator for VaR-E is inspired by the non-uniform estimator for probability of large loss from Broadie et al. (2011). To estimate $\mathbb{P}_\theta(\mathbb{E}_{\xi|\theta}[h(\xi)] > c)$, where c is a constant, they sequentially select a θ that maximizes

$$\mathbb{P}\{(\bar{H}^l(\theta) - c)(\bar{H}^{l+1}(\theta) - c) \leq 0\},$$

where $\bar{H}^l(\theta)$ is the estimate of loss at scenario θ at iteration l , and $\bar{H}^{l+1}(\theta)$ is the estimate of loss at scenario θ at the next iteration $l + 1$ if there is one more ξ sample generated in that scenario θ . The above inequality ensures that the additional generated sample ξ would change the relative position of $\bar{H}^{l+1}(\theta)$ with respect to c . This change will impact the estimate for the probability of large loss. This idea is very intuitive but not directly applicable in our problem, where the outer measure is VaR (instead of the probability of large loss) and we need to account for the uncertainty associated with the estimation of VaR-E (unlike the deterministic number c).

Consider a fixed N number of θ samples, which will also be referred to as scenarios in what follows. Suppose $\bar{H}^l(\theta_i)$ is the performance estimate of the scenario θ_i at iteration l , and in the next iteration $l + 1$, one more ξ sample is generated in θ_i and updates the performance estimate to $\bar{H}^{l+1}(\theta_i)$. Intuitively, we want to give more simulation budget to the scenarios θ that are around the true quantile value. Following the idea in Broadie et al. (2011), we want to choose a θ_i that maximizes the following probability:

$$\mathbb{P}\{(\bar{H}^l(\theta_i) - v_\alpha)(\bar{H}^{l+1}(\theta_i) - v_\alpha) \leq 0\}.$$

The inequality above means that the performance estimate of θ_i jumps over v_α given one additional ξ sample and hence changes the relative order of the θ samples around the true quantile. Maximizing the probability of such a change, therefore, assigns the additional ξ sample to the most important θ_i scenario. However, the exact value v_α is not available, so it has to be replaced by its estimator \tilde{v}_α and the uncertainty associated with the estimator should also be accounted for. The change of the performance estimate of θ_i is only statistically significant when it jumps over a confidence interval of \tilde{v}_α . We define this change by dividing it into several cases below.

Let $\beta > 0$ be the parameter of the confidence level about the estimator \tilde{v}_α , which is typically set between 0 and 3. It controls the concentration of the budget around the quantile. Let $\hat{\sigma}_\alpha$ denote the sample standard deviation of $\bar{H}^l(\theta^{(\alpha N)})$. We define the following events:

$$A_{11} := \left\{ \bar{H}^{l+1}(\theta_i) \geq \tilde{v}_\alpha + \beta \hat{\sigma}_\alpha \mid \tilde{v}_\alpha - \beta \hat{\sigma}_\alpha \leq \bar{H}^l(\theta_i) \leq \tilde{v}_\alpha \right\}$$

$$A_{12} := \left\{ \bar{H}^{l+1}(\theta_i) \leq \tilde{v}_\alpha - \beta \hat{\sigma}_\alpha \mid \tilde{v}_\alpha < \bar{H}^l(\theta_i) \leq \tilde{v}_\alpha + \beta \hat{\sigma}_\alpha \right\}$$

$$A_{21} := \left\{ \bar{H}^{l+1}(\theta_i) < \tilde{v}_\alpha \mid \bar{H}^l(\theta_i) > \tilde{v}_\alpha + \beta \hat{\sigma}_\alpha \right\}$$

$$A_{22} := \left\{ \bar{H}^{l+1}(\theta_i) > \tilde{v}_\alpha \mid \bar{H}^l(\theta_i) < \tilde{v}_\alpha - \beta \hat{\sigma}_\alpha \right\}.$$

Notice that these events are disjoint, and their union $A = A_{11} \cup A_{12} \cup A_{21} \cup A_{22}$ captures all the events at which the performance estimate of a scenario θ has a statistically significant change from one side of the quantile to the other side and hence changes the relative order of the θ samples around the quantile. So, we would like to allocate the next ξ sample to maximize the probability of event A .

If at iteration $l + 1$, we are given an additional inner simulation budget to the scenario θ_i , this would result in

$$\bar{H}^{l+1}(\theta_i) = \frac{1}{K_i + 1} \sum_{j=1}^{K_i+1} h(\xi_j) = \frac{1}{K_i + 1} h(\xi_{K_i+1}) + \frac{K_i}{K_i + 1} \bar{H}^l(\theta_i),$$

where K_i is the inner simulation budget allocated to the scenario θ_i , and ξ is sampled from distribution $\mathbb{P}_{\xi|\theta_i}$. Using the approximation $E[h(\xi_{K_i+1})] \approx \bar{H}^l(\theta_i) \approx \tilde{v}_\alpha \pm \beta \hat{\sigma}_\alpha$, which follows from the assumption that $K_i \gg 1$, and applying Cantelli's inequality (see Cantelli (1928)), we have

$$\begin{aligned} \mathbb{P}(A_{11}) &\approx \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] > K_i(\tilde{v}_\alpha + \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i))\right) \\ &\leq \left(1 + \frac{K_i^2}{\hat{\sigma}_i^2} \left(\tilde{v}_\alpha + \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i)\right)^2\right)^{-1}, \end{aligned}$$

$$\begin{aligned} \mathbb{P}(A_{12}) &\approx \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] < K_i(\tilde{v}_\alpha - \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i))\right) \\ &\leq \left(1 + \frac{K_i^2}{\hat{\sigma}_i^2} \left(\tilde{v}_\alpha - \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i)\right)^2\right)^{-1}. \end{aligned}$$

Similarly, we have

$$\begin{aligned} \mathbb{P}(A_{21}) &= \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] < -K_i(\bar{H}^l(\theta_i) - \tilde{v}_\alpha) - (\mathbb{E}[h(\xi_{K_i+1})] - \tilde{v}_\alpha)\right) \\ &\leq \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] < K_i(\tilde{v}_\alpha - \bar{H}^l(\theta_i))\right) \\ &\leq \left(1 + \frac{K_i^2}{\hat{\sigma}_i^2} \left(\tilde{v}_\alpha - \bar{H}^l(\theta_i)\right)^2\right)^{-1}, \end{aligned}$$

where the first inequality follows from $\mathbb{E}[h(\xi_{K_i+1})] \approx \bar{H}^l(\theta_i) > \tilde{v}_\alpha$ in the event A_{21} , and the second inequality applies Cantelli's inequality. Additionally,

$$\begin{aligned} \mathbb{P}(A_{22}) &= \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] > -K_i(\bar{H}^l(\theta_i) - \tilde{v}_\alpha) - (\mathbb{E}[h(\xi_{K_i+1})] - \tilde{v}_\alpha)\right) \\ &\leq \mathbb{P}\left(h(\xi_{K_i+1}) - \mathbb{E}[h(\xi_{K_i+1})] > K_i(\tilde{v}_\alpha - \bar{H}^l(\theta_i))\right) \\ &\leq \left(1 + \frac{K_i^2}{\hat{\sigma}_i^2} \left(\tilde{v}_\alpha - \bar{H}^l(\theta_i)\right)^2\right)^{-1}, \end{aligned}$$

where the first inequality follows from $\mathbb{E}[h(\xi_{K_i+1})] \approx \bar{H}^l(\theta_i) < \tilde{v}_\alpha$ in the event A_{22} , and the second inequality applies Cantelli's inequality.

Noticing $\mathbb{P}(A_{21})$ and $\mathbb{P}(A_{22})$ have the same upper bound, we consider three subsets of scenarios I_1, I_2, I_3 , which are disjoint and whose union is the entire set of scenarios:

$$\begin{aligned} I_1 &= \left\{i = 1, 2, \dots, N \mid \tilde{v}_\alpha - \beta \hat{\sigma}_\alpha \leq \bar{H}^l(\theta_i) \leq \tilde{v}_\alpha\right\} \\ I_2 &= \left\{i = 1, 2, \dots, N \mid \tilde{v}_\alpha < \bar{H}^l(\theta_i) \leq \tilde{v}_\alpha + \beta \hat{\sigma}_\alpha\right\} \\ I_3 &= \left\{i = 1, 2, \dots, N \mid |\bar{H}^l(\theta_i) - \tilde{v}_\alpha| > \beta \hat{\sigma}_\alpha\right\}. \end{aligned}$$

To maximize the probability of event A , we resort to maximizing the upper bounds that we obtained above. The following proposition summarizes our proposed inner simulation budget allocation.

Proposition 1 Define $\Gamma: I_1 \cup I_2 \cup I_3 \rightarrow \mathbb{R}$ such that $\Gamma(i) = \frac{K_i^2}{\hat{\sigma}_i^2} (\tilde{v}_\alpha + \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i))^2$ if $i \in I_1$, $\Gamma(i) = \frac{K_i^2}{\hat{\sigma}_i^2} (\tilde{v}_\alpha - \beta \hat{\sigma}_\alpha - \bar{H}^l(\theta_i))^2$ if $i \in I_2$, and $\Gamma(i) = \frac{K_i^2}{\hat{\sigma}_i^2} (\tilde{v}_\alpha - \bar{H}^l(\theta_i))^2$ if $i \in I_3$. Set $i^* = \min_{i=1, \dots, N} \Gamma(i)$. Then, we should generate one additional sample of ξ for the scenario θ_{i^*}

Proposition 1 provides an algorithm to conduct the sequential simulation. Since θ_{i^*} is the scenario that minimizes $\Gamma(i)$, it maximizes the upper bound on the probability of event A . Thus, scenario θ_{i^*} ensures that the additional ξ sample will have the maximal impact on the estimator for VaR-E. Note that the expression $\Gamma(\cdot)$ reveals that the new inner simulation budget will be allocated to a scenario that either has a large variance, is close to the current empirical quantile, or has very few inner simulation budget allocated to it before. We provide the asymptotic normality result of this non-uniform estimator for VaR-E in the theorem below.

Theorem 2 Asymptotic Normality of Our Non-Uniform Estimator. Let $K(\theta^{(\alpha N)})$ denote the inner simulation budgets allocated to the scenario $\theta^{(\alpha N)}$. Define $\sigma := \frac{\sqrt{\alpha(1-\alpha)}}{f(v_\alpha)}$, and $\mu := \frac{-\Lambda'(v_\alpha)}{f(v_\alpha)}$. Under Assumption 1, the existence of limit $R^2 = \lim_{N, K(\theta^{(\alpha N)}) \rightarrow \infty} N/K(\theta^{(\alpha N)})^2$ is a sufficient and necessary condition for

$$\lim_{N, K(\theta^{(\alpha N)}) \rightarrow \infty} \sqrt{N} (\tilde{v}_\alpha - v_\alpha) \stackrel{\mathcal{D}}{=} \sigma \mathcal{N}(0, 1) + |R|\mu.$$

Notice that there are two main differences between the uniform and non-uniform estimators. First, for the uniform estimator, the inner budget K allocated to each scenario is the same. Theorem 1 shows that the bias of the estimator is inversely proportional to K . For the non-uniform estimator, the inner budget allocated to each scenario is different according to our sequential algorithm, and we use $K(\theta)$ to denote the inner budget allocated to the scenario θ . Theorem 2 shows that the bias of the estimator is inversely proportional to $K(\theta^{(\alpha N)})$, i.e., the inner budget allocated to the scenario $\theta^{(\alpha N)}$. Our sequential algorithm achieves the same rate as the uniform estimator by only requiring $K(\theta^{(\alpha N)})$, the budget allocated to the empirical quantile scenario, to be equal to the budget K in the uniform estimator. Intuitively, the budget allocated to scenarios far from $\theta^{(\alpha N)}$ is much smaller than $K(\theta^{(\alpha N)})$, and therefore, the total budget K_{total} of the non-uniform estimator is less than that of the uniform estimator. Rigorous analysis of total budget of the non-uniform estimator will be included in our future work. Another difference lies in the additional computational overhead in the non-uniform estimator, i.e., we need to compute $\Gamma(\cdot)$ for each scenario in order to decide the allocation of one additional inner budget. However, in practice, evaluating the (black-box) function $h(\cdot)$ is time consuming. The additional computational overhead in the non-uniform estimator can be neglected compared with the saved inner budgets.

3.4 OCBA with VaR-E Estimation

In this subsection, we first adapt the OCBA approach to allocate the outer simulation budget to the alternatives x_1, \dots, x_M . Most of the OCBA literature assumes the objective function takes an expectation form, which is often referred to as the mean response. The derivation of OCBA generally requires two main assumptions: (1) independence across all alternatives; (2) that the mean response follows a normal distribution. The normality assumption on the mean response is usually justified by the Central Limit Theorem (CLT), as the mean response is often obtained from a sample average estimate or a batch mean. In our problem, the objective function takes the form of VaR-Expectation, which is more challenging than the mean response. Luckily, both the uniform and non-uniform estimators of VaR-E have asymptotic normality (as shown in Theorems 1 and 2), and therefore, we can extend OCBA to our problem.

Recall that the objective value is $v_\alpha(x_i) = VaR(H(x_i, \theta))$, $i = 1, \dots, M$, its uniform estimator is $\hat{v}_\alpha(x_i)$ and its non-uniform estimator is $\tilde{v}_\alpha(x_i)$. As shown in Theorem 1, the uniform estimator has an asymptotic bias $\frac{\mu_i}{K}$ and variance $\frac{\sigma_i^2}{N_i}$, where N_i is the number of scenarios (i.e., θ samples) generated for the i th alternative,

K is the number of ξ samples generated for each scenario, and μ_i and σ_i are as defined in Theorem 1. As shown in Theorem 2, the non-uniform estimator has an asymptotic bias $\frac{\mu_i}{K(\theta_i^{(\alpha N_i)})}$ and variance $\frac{\sigma_i^2}{N_i}$, where $K(\theta_i^{(\alpha N_i)})$ is the number of ξ samples generated in the scenario $\theta_i^{(\alpha N_i)}$. Using these results, we present the posterior distribution of $v_\alpha(x_i)$ given a normal prior distribution. The proposition below shows the result for the non-uniform estimator, where the analogous result for the uniform estimator holds by replacing the bias and variance terms in the posterior distribution.

Proposition 2 Assume $v_\alpha(x_i)$ has a conjugate normal prior distribution $\mathcal{N}(\mu_i^0, (\sigma_i^0)^2)$. Denote by $\bar{v}_\alpha(x_i)$ the posterior distribution of $v_\alpha(x_i)$, then given $(\bar{H}(x_i, \theta_{i,j}), j = 1, 2, \dots, N_i)$, and assume $\sigma_i^0 \rightarrow \infty$, we have

$$\bar{v}_\alpha(x_i) \sim \mathcal{N} \left(\bar{v}_\alpha(x_i) - \frac{\mu_i}{K(\theta_i^{(\alpha N_i)})}, \frac{\sigma_i^2}{N_i} \right).$$

Proposition 2 shows that the posterior distribution of $v_\alpha(x_i)$ is still normal. Furthermore, the variance of the posterior distribution is reversely proportional to the number of θ samples of the estimator. We assume the inner simulation budget will be allocated independently from the outer simulation budget, such that the focus of our adapted OCBA approach will be on the outer layer budget allocation.

From Proposition 2, a direct application of Theorem 3.2 in Chen and Lee (2010) gives the following OCBA scheme:

$$N_b = \sigma_b \sqrt{\sum_{i=1, i \neq b}^M \left(\frac{N_i}{\sigma_i} \right)^2}, \text{ where } b = \arg \min_{i \in [M]} \left(\bar{v}_\alpha(x_i) - \frac{\mu_i}{K(\theta_i^{(\alpha N_i)})} \right) \quad (2)$$

$$\frac{N_i}{N_j} = \frac{\sigma_i^2 \left(\left(\bar{v}_\alpha(x_b) - \frac{\mu_b}{K(\theta_b^{(\alpha N_b)})} \right) - \left(\bar{v}_\alpha(x_j) - \frac{\mu_j}{K(\theta_j^{(\alpha N_j)})} \right) \right)^2}{\sigma_j^2 \left(\left(\bar{v}_\alpha(x_b) - \frac{\mu_b}{K(\theta_b^{(\alpha N_b)})} \right) - \left(\bar{v}_\alpha(x_i) - \frac{\mu_i}{K(\theta_i^{(\alpha N_i)})} \right) \right)^2}, \text{ for all } i \neq j \neq b. \quad (3)$$

Although OCBA is derived to achieve asymptotic optimality as the total budget goes to infinity, in practice, it is often implemented in a sequential way through increasing the budget by a fixed amount each time. We take this sequential approach, i.e., at each iteration, the algorithm generates Δ_i number of additional θ samples for each alternative x_i according to the budget allocation scheme above.

Integrating the outer-layer OCBA with the inner-layer nonuniform estimation of VaR-E, we propose the following algorithm. Here, Δ denotes additional number of scenarios generated in each iteration, N_0 denotes initial number of scenarios for each alternative, K_0 denotes the initial number of ξ -samples in each scenario, K denotes number of ξ -samples in each scenario in the uniform estimator, T denotes the total budget, and γ is a parameter that determines the shift of the inner budget to the outer layer.

Algorithm 1: OCBA with Non-Uniform VaR-E Estimation

- **Input.** $\Delta, N_0, K_0, K, T, \gamma$.
- **Step 1: Initialization.** Set $l = 0$. For each x , generate N_0 initial scenarios $\theta_1, \dots, \theta_{N_0}$. For each scenario θ , generate K_0 initial ξ samples. For each x , compute $\bar{H}^l(x, \theta)$ for all scenarios, and compute \bar{v}_α^l .
- **Step 2: VaR-E Estimation.** Allocate inner simulation budget to each scenario sequentially according to Proposition 1, until the average inner simulation budget reaches $(K - K_0)\gamma + K_0$. For x_i , compute $\bar{v}_\alpha^l(x_i)$, sample estimate of bias $\bar{\mu}_i$ and sample estimate of standard deviation $\bar{\sigma}_i$ of the estimator.

- Step 3: OCBA.** Find $b = \arg \min_i \left(\tilde{v}_\alpha^l(x_i) - \frac{\tilde{\mu}_i}{K \left(\theta_i^{(\alpha N_i^l)} \right)} \right)$. Increase the outer layer simulation budget by Δ and calculate the new outer layer budget allocation $N_1^{l+1}, \dots, N_M^{l+1}$, according to (2) and (3), where we replace the bias term and standard deviation term by their sample estimates.
- Step 4: Stopping or Looping.** Calculate used total budget T_{used} . If $T_{\text{used}} < T$, generate additional $\max(N_i^{l+1} - N_i^l, 0)$ outer scenarios for alternative $i = 1, \dots, M$, set $l = l + 1$, and go to **Step 2**. Otherwise, **stop**, and return alternative b as the best alternative.

In Algorithm 1, the non-uniform estimation in the inner layer concentrates budgets on the samples around the quantile, and thus saves the inner budget so that more budget can be shifted to the outer layer, which then uses OCBA to smartly allocate the outer budget. Notice that the algorithm can be terminated either by limiting the total budget to T or by setting a threshold on the approximate probability of correct selection (APCS). APCS is a lower bound of PCS, that is, $\text{APCS} \equiv 1 - \sum_{i=1, i \neq b}^M \mathbb{P} \{ \bar{v}_\alpha(x_b) > \bar{v}_\alpha(x_i) \}$. In addition, we could replace the non-uniform estimator in Step 2 by the uniform estimator (i.e., the number of ξ -samples for all scenarios would be replaced by K), and will refer to this modified algorithm as OCBA with Uniform VaR-E Estimation, which will be used in the comparisons within our numerical experiments in the next section.

4 NUMERICAL RESULTS

We present the numerical results for the portfolio selection problem outlined in Section 2. Specifically, the portfolio consists of a stock and a risk-free asset. The risk-free asset has an interest rate r . The stock price is S_0 at time 0, and evolves according to a Geometric Brownian motion with drift coefficient μ and volatility σ . Given the risk factor θ at risk horizon τ , the stock price at τ is $S_\tau(\theta) = S_0 e^{(\mu - \sigma^2/2)\tau + \sigma\sqrt{\tau}\theta}$. Hence, at some finite time horizon $T \geq \tau$, the stock price is $S_T(\theta) = S_\tau(\theta) e^{(r - \sigma^2/2)(T - \tau) + \sigma\sqrt{T - \tau}W}$, where W follows the standard normal distribution. We assume θ also follows a standard normal distribution. The investor previously had x_0 portion of the capital X_0 in the stock and would like to decide the position x at time 0, in order to minimize the VaR (with respect to θ) of the expected portfolio loss. The portfolio loss is $h(x, S_T) = xX_0 \left[1 - \frac{e^{-r(T-\tau)}}{S_0} S_T \right] + c(x - x_0)^2$, where the first term comes from our assumption that the price of the risky asset is path dependent and can be discounted from finite time T to risk horizon τ , and the second term is the transaction cost for the change of the portfolio position. Note that S_T is in fact the random variable ξ in our formulation (1), and it follows a log-normal distribution whose location parameter is determined by θ . In this model, we can analytically compute the expected loss $H(x, \theta) = xX_0 \left[1 - \frac{S_\tau(\theta)}{S_0} \right] + c(x - x_0)^2$ and the objective value $\text{VaR}^\alpha(H(x, \theta)) = xX_0 \left[1 - \exp^{-[(\mu - \frac{\sigma^2}{2})\tau + z_\alpha \sigma \sqrt{\tau}]} \right] + c(x - x_0)^2$, where z_α satisfies $\Phi(z_\alpha) = \alpha$ and Φ is the c.d.f. of the standard normal distribution.

In our numerical experiments, the parameters are set as follows: Risk level of VaR $\alpha = 0.9$, initial stock price $S_0 = 0.1$, drift in risk horizon $\mu = 8\%$, volatility $\sigma = 10\%$, risk-free interest rate $r = 10\%$, initial position in stock $x_0 = 50\%$, initial wealth $X_0 = 3$, risk horizon $\tau = 1/52$ years (one week), finite time horizon $T = 0.25$ years (three months), coefficient of transaction cost $c = 1$. The portfolio position in stock can be chosen from a discrete set $\mathcal{X} = \{10\%, 20\%, \dots, 90\%, 1\}$. Figure 1 shows the structure of this portfolio selection problem, where the minimum is reached at $x = 0.5$, i.e., the optimal level of investment in stock is 50%. We now present two experiments.

Experiment 1: We first compare three estimators for VaR-E. In addition to the uniform estimator (presented in Section 3.2) and non-uniform estimator (proposed in Section 3.3), we consider a third baseline estimation procedure based on OCBA-m. The OCBA-m approach, proposed by Chen et al. (2008), is an efficient simulation budget allocation scheme for selecting a subset of m best alternatives. To use the OCBA-m approach for estimating VaR-E, we set $m = (1 - \alpha)N$, apply OCBA-m to select the top m scenarios that have largest sample mean response, and then set the scenario with the smallest sample mean

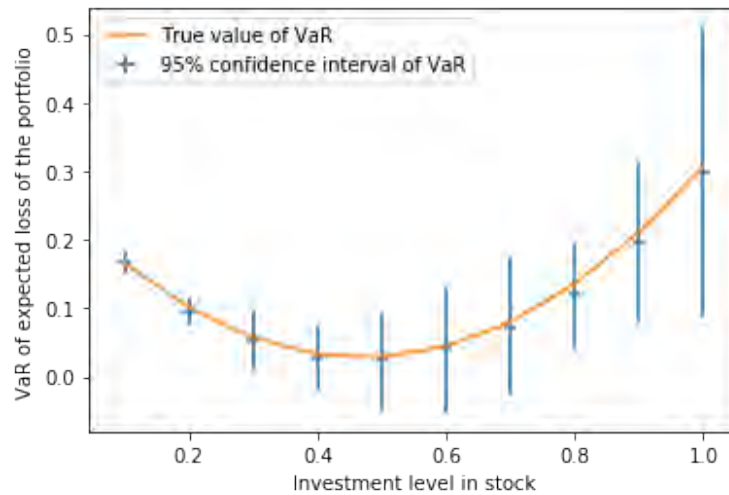
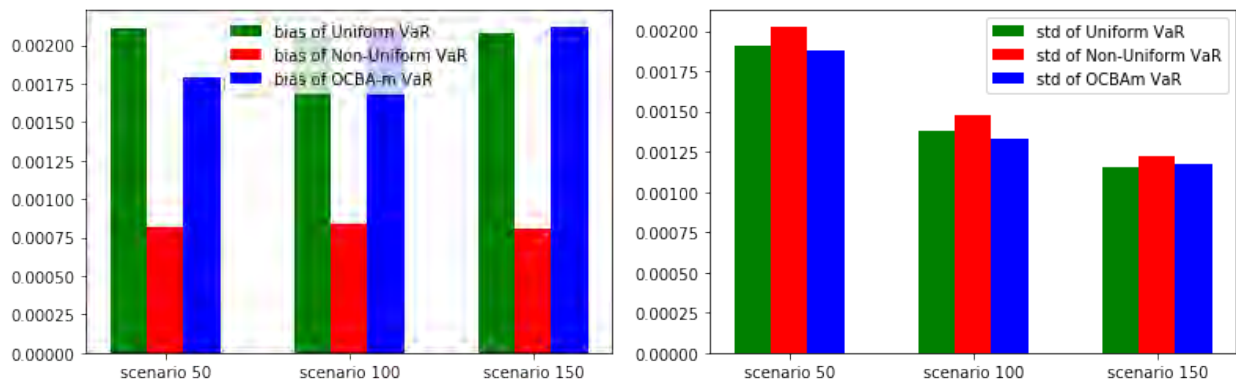


Figure 1: Structure of the portfolio selection problem

response in the selected subset as the VaR-E estimator. By the design of OCBA-m, this estimation approach concentrates the budget on the θ -samples around the quantile. The parameters of the algorithms are set as $K = 20, K_0 = 10, \beta = 1$, and the decision variable is fixed to $x = 50\%$. Figure 2 shows the performance of the uniform estimator, the non-uniform estimator, and the OCBA-m-based estimator, under a fixed number of outer scenarios and a fixed average number of inner simulation budget. In particular, one can see that the non-uniform estimator has much smaller bias than the other two estimators, while it has slightly higher (but still comparable) standard deviation compared to the other estimators. When the number of outer scenarios increases, the standard deviations of all three estimators decrease and the biases remain almost the same, as expected.



(a) Bias of three estimators of VaR

(b) Standard deviation of three estimators of VaR

Figure 2: Performance of uniform, non-uniform, and OCBA-m estimators of VaR-E

Experiment 2: We then solve the portfolio problem by four algorithms: Equal allocation (presented in Section 3.1), OCBA with non-uniform VaR-E estimation (Algorithm 1), OCBA with uniform VaR-E estimation (Algorithm 1, after replacing the non-uniform VaR-E estimator with the uniform one), OCBA with OCBA-m VaR-E estimation algorithm (Algorithm 1, after replacing the non-uniform VaR-E estimator with OCBA-m estimator). The parameters of the algorithms are set as $K = 40, K_0 = 10, N_0 = 50, \delta = 10, \beta = 1, \gamma = 0.4$, and all four algorithms are run for 500 macro-replications. Figure 3 presents the empirical PCS

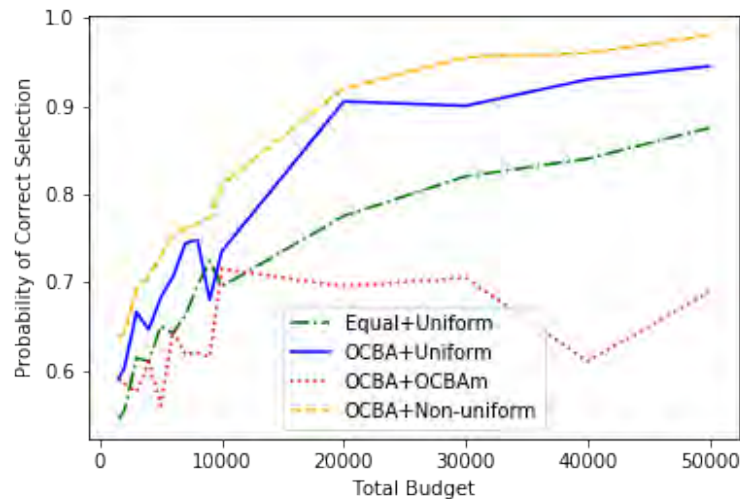


Figure 3: PCS versus total budget under four budget allocation algorithms

(over the 500 macro-replications) as the total budget increases. It is clear that the OCBA with non-uniform VaR-E estimation algorithm achieves the highest empirical PCS, followed by the algorithm with uniform VaR-E estimation. To our surprise, the OCBA with OCBA-m VaR-E estimation algorithm does not perform well. With further analysis/investigation, we found that the reason is that, with fixed number of outer scenarios, the budget is concentrated on the θ -sample at the quantile; however, as additional outer scenarios generated in later iterations, the θ -sample at the quantile may change, but there is not enough computing budget allocated to the new θ -sample at the quantile. In contrast, the OCBA with uniform VaR-E estimation algorithm allocates the budget equally among all outer scenarios, and the OCBA with non-uniform VaR-E estimation algorithm can be viewed as a balanced budget allocation scheme, such that it neither concentrates too much on a fixed scenario nor spread too much over all scenarios.

5 CONCLUSION

We consider the problem of portfolio selection with risk factors, where the goal is to find the optimal portfolio position that minimizes the VaR (with respect to the risk factor) of the expected portfolio loss. To solve the problem, we propose a nested simulation optimization approach by integrating OCBA with a new non-uniform estimator of VaR-Expectation. The non-uniform estimator sequentially allocates the inner budget such that the budget is concentrated on the samples around the quantile, which prevents the waste of budget that would have been allocated to the samples far away from the quantile. We provide the asymptotic normality of the non-uniform estimator. Our numerical results demonstrate the effectiveness of the non-uniform estimator of VaR-Expectation and the proposed nested simulation optimization approach for solving the portfolio selection problem.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support by the National Science Foundation under Grant CAREER CMMI-1453934 and the Air Force Office of Scientific Research under Grant FA9550-19-1-0283.

REFERENCES

- Bekki, J. M., J. W. Fowler, G. T. Mackulak, and B. L. Nelson. 2007. "Using quantiles in ranking and selection procedures". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1722–1728. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

- Broadie, M., Y. Du, and C. C. Moallemi. 2011. "Efficient Risk Estimation via Nested Sequential Simulation". *Management Science* 57(6):1172–1194.
- Broadie, M., Y. Du, and C. C. Moallemi. 2015. "Risk Estimation via Regression". *Operations Research* 63(5):1077–1097.
- Cantelli, F. P. 1928. "Sui confini della probabilita". In *Atti del Congresso Internazionale dei Matematici*, Volume 6, 47–60. Bologna, Italy: Zanichelli.
- Chen, C.-H., D. He, M. Fu, and L. H. Lee. 2008. "Efficient Simulation Budget Allocation for Selecting an Optimal Subset". *INFORMS Journal on Computing* 20(4):579–595.
- Chen, C.-H., and L. H. Lee. 2010. *Stochastic Simulation Optimization*. Singapore: World Scientific.
- Dang, J. O., B. M. Feng, and M. R. Hardy. 2019. "Efficient Nested Simulation of Tail Risk Measures". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 938–949. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Fu, Michael C, E. 2015. *Handbook of Simulation Optimization*, Volume 216. New York: Springer.
- Goldsman, D., and B. L. Nelson. 2007. "Comparing Systems via Simulation". In *Handbook of Simulation*, edited by J. Banks, 273–306. Hoboken, New Jersey: John Wiley & Sons, Ltd.
- Gordy, M. B., and S. Juneja. 2010. "Nested Simulation in Portfolio Risk Measurement". *Management Science* 56(10):1833–1848.
- Hong, L., S. Juneja, and G. Liu. 2017. "Kernel Smoothing for Nested Estimation with Application to Portfolio Risk Measurement". *Operations Research* 65(3):657–673.
- Jorion, P. 1997. *Value at Risk: The New Benchmark for Controlling Market Risk*. New York: McGraw-Hill.
- Lan, H., B. L. Nelson, and J. Staum. 2010. "A Confidence Interval Procedure for Expected Shortfall Risk Measurement via Two-Level Simulation". *Operations Research* 58(5):1481–1490.
- Lee, S., and B. L. Nelson. 2014. "Bootstrap Ranking Selection Revisited". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3857–3868. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Lee, S.-H. 1998. *Monte Carlo Computation of Conditional Expectation Quantiles*. Ph.D. thesis, Stanford University, Stanford, CA.
- Lee, S.-H., and P. W. Glynn. 2003. "Computing the Distribution Function of a Conditional Expectation via Monte Carlo: Discrete Conditioning Spaces". *ACM Transactions on Modeling and Computer Simulation* 13(3):238–258.
- Liu, M., B. L. Nelson, and J. Staum. 2010. "An Efficient Simulation Procedure for Point Estimation of Expected Shortfall". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, 2821–2831. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Liu, M., and J. Staum. 2010. "Stochastic Kriging for Efficient Nested Simulation of Expected Shortfall". *Journal of Risk* 12:3–27.
- Markowitz, H. 1952. "Portfolio Selection". *The Journal of Finance* 7(1):77–91.
- Peng, Y., C.-H. Chen, M. Fu, J. Hu, and I. Ryzhov. 2020. "Efficient Sampling Allocation Procedures for Optimal Quantile Selection". *INFORMS Journal on Computing*. Advance online publication. doi:10.1287/ijoc.2019.0946.
- Shin, D., M. Broadie, and A. Zeevi. 2016. "Tractable Sampling Strategies for Quantile-based Ordinal Optimization". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, 847–858. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Sun, Y., D. W. Apley, and J. Staum. 2011. "Efficient Nested Simulation for Estimating the Variance of a Conditional Expectation". *Operations Research* 59(4):998–1007.
- Zhu, H., T. Liu, and E. Zhou. 2020. "Risk Quantification in Stochastic Simulation under Input Uncertainty". *ACM Transactions on Modeling and Computer Simulation* 30(1):1–24.

AUTHOR BIOGRAPHIES

YIFAN LIN is a Ph.D. student in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. His email address is ylin429@gatech.edu.

ENLU ZHOU is an Associate Professor in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. She received the B.S. degree from Zhejiang University, China, in 2004, and received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, in 2009. Her research interests include stochastic control and simulation optimization. Her email address is enlu.zhou@isye.gatech.edu and her web page is <http://enluzhou.gatech.edu/>.

ALY MEGAHED is currently an AI Research Manager and Master Inventor at IBM's Almaden Research Center. He received his Ph.D. in industrial engineering from Georgia Institute of Technology, in 2014. He currently leads the development of analytical tools for complex service engagements. His email address is aly.megahed@us.ibm.com and his web page is <https://researcher.watson.ibm.com/researcher/view.php?person=us-aly.megahed>.