Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

ESTIMATION OF CONDITIONAL VALUE-AT-RISK FOR INPUT UNCERTAINTY WITH BUDGET ALLOCATION

Helin Zhu
Enlu Zhou

H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

ABSTRACT

When simulating a complex stochastic system, the behavior of the output response depends on the input parameters estimated from finite real-world data, and the finiteness of data brings input uncertainty to the output response. The quantification of the impact of input uncertainty on output response has been extensively studied. However, most of the existing literature focuses on providing inferences on the mean output response with respect to input uncertainty, including point estimation and confidence interval construction of the mean response. To the best of our knowledge, risk assessment of input uncertainty has been rarely considered. In the present paper, we will introduce risk measures for input uncertainty, study a nested Monte Carlo estimator and construct an asymptotically valid confidence interval for a specific risk measure—Conditional Value-at-Risk of the mean response. We further study the associated budget allocation problem for more efficient nested simulation of the estimator.

1 INTRODUCTION AND MOTIVATION

For a complex real-world stochastic system, simulation is a powerful tool to analyze its behavior when real experiments on the system are expensive or difficult to conduct. For example, consider the system of a typical hospital emergency room (ER). When the administrators of the ER determine the number of on-call doctors, one of the main criteria is the expected number of waiting patients or the expected waiting time of an individual customer. An $M/M/n$ (Poisson patient arrival, Exponential treatment time, n doctors) queue is often used to simulate the system and provide inferences on the real system behavior. During the simulation experiments, the uncertainty on the simulation input parameters (e.g., the patient arrival rate and the treatment time for an individual patient) may need to be taken into account, since they are typically estimated from finite ER patient records. In general, there are two sources of uncertainty for a typical stochastic simulation experiment: the extrinsic uncertainty on input parameters (referred as to *input parameter uncertainty*) that reflects our belief on input parameters, and the intrinsic uncertainty on output response (referred to as *stochastic uncertainty*) that reflects the inherent stochastic variability of output. Many papers also consider a third source of uncertainty—*model uncertainty*, which refers to the uncertainty of the input probabilistic model.

The variability of simulation output response clearly depends on both stochastic uncertainty and input uncertainty. An important question to address is how to quantify the impact of input uncertainty on output response variability. A widely used procedure for such quantification is the point estimation of the mean response with respect to (w.r.t.) input uncertainty, combined with construction of the associated confidence interval (CI). Within this framework, the two main branches of approaches are the frequentist methods and the Bayesian methods. The frequentist methods include the Direct and Bootstrap Sampling methods by Barton and Schruben (1993), Barton and Schruben (2001), etc. The Bayesian methods include the Bayesian Model Averaging (BMA) methods by Chick (2001), Zouaoui and Wilson (2003), etc. In these

978-1-4673-9743-8/15/$31.00 ©2015 IEEE

methods, the Bayesian updating rule is applied on a chosen prior distribution of input parameters to obtain a posterior parameter distribution, which will be used as the sampling distribution of input parameters in the simulation experiment. In addition to the two aforementioned methods, Cheng and Holland (1997) also developed the δ-method, which is based on Taylor's Theorem to decompose the variance of simulation output response into two components that are caused by parameter uncertainty and stochastic uncertainty, respectively. Song and Nelson (2015) developed a method for quickly assessing the relative contribution of each input distribution to the overall variance due to input uncertainty. In recent years, with the rise of Stochastic Kriging meta-modeling method in stochastic simulation (e.g., Ankenman et al. (2010)), meta-model assisted methods have been developed for quantifying input uncertainty, see Barton et al. (2013), Xie et al. (2014), etc. Henderson (2003) provided an early review on the importance of input uncertainty and the common methods to deal with it. Barton (2012) provided a concise review on popular methods in output analysis with input uncertainty and highlighted some remaining challenges in this area.

Most of the aforementioned literature focuses on providing inferences on the mean response w.r.t. input uncertainty, i.e., the average behavior of the mean response, including point estimation and CI construction. The "quantification" of extreme behavior of the mean response w.r.t. input uncertainty is often overlooked. Such quantification could provide inferences on system sensitivity or stability, and thus is critical for risk assessment/control of the system. Consider the ER example previously mentioned, the risk assessment/control of the mean response (e.g., the expected number of waiting patients) w.r.t. input uncertainty is quite necessary, because the behavior of the system output under extreme input models indicates a large number of expected waiting patients, which might lead to delayed treatment of patients and possibly serious consequences in life-threatening situations.

Risk measures for input uncertainty are of great importance, because they provide rigorous quantifications of the behavior of the mean response under extreme input models. To the best of our knowledge, they have rarely been systematically studied in input uncertainty modeling. In some papers, the quantile of the mean response that is closely related to the risk measure *Value-at-Risk*(VaR), is used when percentile-type CIs are constructed. But it is rarely the main focus of analysis in the literature.

Common risk measures such as VaR and *Conditional Value-at-Risk* (CVaR) have been extensively studied in the financial industry, due to their great importance for financial organizations. Loosely speaking, VaR characterizes the extreme (e.g., 99%) quantile of the output distribution and CVaR characterizes the conditional mean of the very tail portion of the output distribution. An abundant literature has dedicated to studying the estimation and optimization of risk measures under various settings, see Hong et al. (2014) for an elegant review of Monte Carlo methods for VaR and CVaR.

Risk assessment of input certainty exhibits some overlaps with risk assessment in portfolio/credit risk management. In a broader sense, they both deal with simulating certain conditional expectations. Among the literature, Lee (1998) studied the point estimation of a quantile (VaR) of the distribution of a conditional expectation via a two-level simulation. Steckley (2006) considered estimating the density of a conditional expectation using kernel density estimation. In portfolio risk management, the risk assessment of a portfolio of securities at a given future date (risk horizon) is to re-evaluate the portfolio at the risk horizon for thousands of realizations of risk factors, and computing the portfolio value at each realization may further require the simulation of risk factors beyond the risk horizon by up to years. Common approaches for risk assessment in portfolio risk management include the delta-gamma method by Glasserman et al. (2000), etc; the nested simulation method by Gordy and Juneja (2010), Broadie et al. (2011), etc; the stochastic kriging method by Liu and Staum (2010), etc.

In the present paper, we will introduce risk measures for input certainty and their estimation. In particular, we study nested Monte Carlo estimator for CVaR of the mean response w.r.t. input uncertainty. Furthermore, we show some asymptotic properties of the estimator and use them to construct an (asymptotically valid) CI for CVaR of the mean response. Due to space limit, we choose CVaR over VaR because CVaR is a coherent risk measure (see, e.g., Artzner et al. (1997)) and exhibits nice properties such as convexity for optimization (see, e.g., Rockafellar and Uryasev (2000)). Our work can be viewed as a starting point for research on more general risk measures or more sophisticated estimators of risk measures for input uncertainty. In particular, the contributions of this paper are three-folds: 1) for output analysis with input uncertainty, this paper is among the first to systematically study risk measures for input uncertainty, which is an important topic that has been largely overlooked in the literature; 2) we show some asymptotic properties of nested Monte Carlo estimator of CVaR, which are guarantees for constructing an (asymptotically valid) CI for CVaR of the mean response; 3) we develop a novel approach to solve the associated budget allocation problem for nested simulation of CVaR estimator, in order to improve the simulation efficiency.

The work most relevant to ours is probably Lan et al. (2010). They developed a procedure for constructing CIs of Expected Shortfall Risk (equivalent to CVaR) measurement via nested simulation in financial risk management, in which they use screening, a ranking and selection type technique, to improve the simulation efficiency. The differences between their work and ours are evident. First, their CI construction procedure is based on the empirical likelihood method, while our CI construction procedure is based on the asymptotic properties of the estimator; second, the budget allocation problem in their setting is more difficult to model and solve than ours. In fact, the budget allocation problem in our setting is straightforward to formulate and easy to solve with the proposed novel method.

The rest of the paper is organized as follows. In section 2, we will introduce risk measures for input uncertainty, describe the nested Monte Carlo estimation and CI construction of a specific risk measure—CVaR for input uncertainty. We formulate and solve the associated budget allocation problem in section 3. In section 4, we conduct numerical experiments to demonstrate the theoretical results in previous sections. Finally, conclusion and promising directions of future research are provided in section 5.

2 CVAR OF THE MEAN OUTPUT RESPONSE

2.1 Formulation

Let us first rigorously define the risk measures VaR and CVaR of the mean response w.r.t. input uncertainty. In a stochastic simulation experiment, consider an output response function in the form of $H(\theta;\xi)$, where θ denotes the input parameter(s) and ξ represents the noise (stochastic uncertainty) in the output response. Furthermore, suppose there is a probability distribution on θ that reflects our belief on input uncertainty, since θ needs to be inferred from historical data. Therefore, θ can be treated as a random variable, and we assume it is independent from ξ. From a Bayesian perspective, our belief on input uncertainty will be updated with new observations of parameters. Specifically, suppose there exists a prior distribution $p(\theta)$ on θ. The prior can be either non-informative or informative depending on subjective experiences, and the hyper-parameters of the prior can be estimated from historical data. With new observations $\mathbf{x}$ of θ, a posterior distribution $p(\cdot|\mathbf{x})$ on θ can be obtained via standard Bayesian updating. To facilitate the analysis, let us assume for fixed θ, $H(\theta;\xi)=\eta(\theta)+e(\theta;\xi)$, $\eta(\theta)=\mathbb{E}_{\xi}[H(\theta;\xi)]$ is the mean response, $e(\theta;\xi)$ is the output noise such that $\mathbb{E}[e(\theta;\xi)|\theta]=0$ and $Var[e(\theta;\xi)|\theta]=\tau_{\theta}^2$, where we assume τ_{θ}^2 is a finite deterministic function of θ and $\tau^2 \overset{\triangle}{=} \int \tau_{\theta}^2 p(\theta|\mathbf{x})d\theta$ is finite. Furthermore, to account for input uncertainty, we assume that $\eta(\theta)=\eta^{\mathbf{x}}+\delta(\theta)$, where the posterior mean $\eta^{\mathbf{x}}=\mathbb{E}_{p(\cdot|\mathbf{x})}[\eta(\theta)]=\int \eta(\theta)p(\theta|\mathbf{x})d\theta$, $\delta(\theta)$ is the input "noise" such that $\mathbb{E}_{p(\cdot|\mathbf{x})}[\delta(\theta)]=0$ and $Var_{p(\cdot|\mathbf{x})}[\delta(\theta)]=\sigma^2$, where σ^2 may depend on the observations $\mathbf{x}$.

Suppose $0<\alpha<1$ is a certain large probability level (usually, $\alpha=0.95$ or $\alpha=0.99$) of interest. $VaR_{\alpha}\left(\mathbb{E}_{\xi}[H(\theta;\xi)]\right)$ is defined by the α-quantile of the mean response $\eta(\theta)$, i.e.,

$$VaR_{\alpha}\left(\mathbb{E}_{\xi}[H(\theta;\xi)]\right) := \inf\{t : F(t)\geq\alpha\}, \tag{1}$$

where $F(\cdot)$ is the cumulative distribution function (c.d.f.) of $\eta(\theta)$. Further assume θ is a continuous random variable, the function $\eta(\cdot)$ is continuous and $\eta(\theta)$ is a continuous random variable. It follows that $F(\cdot)$ is continuous, and VaR defined in (1) can be simplified as $VaR_{\alpha}\left(\mathbb{E}_{\xi}[H(\theta;\xi)]\right)=F^{-1}(\alpha)$. Intuitively, VaR represents the cut-off level for the α-tail of the mean response. CVaR, on the other hand, is defined by the conditional mean of the α-tail distribution of the mean response. In particular, given that $\eta(\theta)$ is a continuous random variable,

$$CVaR_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right) := \frac{1}{1-\alpha}\mathbb{E}_{p(\cdot|\mathbf{x})}\left(\mathbb{E}_\xi[H(\theta;\xi)]I\{\mathbb{E}_\xi[H(\theta;\xi)] \geq VaR_\alpha\}\right), \tag{2}$$

where the indicator function $I\{A\}$ equals 1 when the statement A is true and 0 otherwise. With slight abuse of notations, we use VaR_α as an abbreviation for $VaR_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right)$ and $CVaR_\alpha$ as an abbreviation for $CVaR_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right)$ occasionally for convenience.

2.2 A Nested Monte Carlo Estimator

CVaR, as one of the mostly used risk measures in financial applications, has been extensively studied. Calculating CVaR for input uncertainty can be straightforward when the system is simple. For example, when the mean response w.r.t. stochastic uncertainty can be evaluated exactly and the mean response (as a random variable due to input uncertainty) admits an explicit density function, CVaR of the mean response w.r.t. input uncertainty can be easily calculated via numerical integration. For complex real-world stochastic systems, numerical calculation of CVaR might not be applicable, especially when the mean output cannot be evaluated exactly or the mean output distribution is complex. In this case, Monte Carlo simulation is a powerful alternative approach to obtain good estimates of CVaR. To gain more intuition, let us first consider the estimation of $CVaR_\alpha$ in (2) without stochastic uncertainty. That is, $\eta(\theta)$ can be evaluated exactly given any θ. Now that the probabilistic model is one-layer, a natural approach to an estimator of $CVaR_\alpha$ is by naive Monte Carlo sampling described as follows. First, draw N i.i.d. scenarios $\theta_1,...,\theta_N$ from the posterior distribution $p(\theta|\mathbf{x})$; second, evaluate the response $\eta(\theta_i)$ for $i = 1,...,N$ and sort the resulting response scenarios $\eta(\theta_1),....,\eta(\theta_N)$ in ascending order, denoted by $\eta(\theta_{(1)}) \leq \eta(\theta_{(2)}) \leq \cdots \leq \eta(\theta_{(N)})$; finally, a naive Monte Carlo estimator of $CVaR_\alpha$ is given by

$$\widehat{CVaR}_\alpha(\eta(\theta)) = \frac{1}{(1-\alpha)N+1}\sum_{i=\alpha N}^{N}\eta(\theta_{(i)}), \tag{3}$$

where we assume αN is an integer. Intuitively, $\widehat{CVaR}_\alpha$ in (3) is the average of "effective" response scenarios that are greater than or equal to $\widehat{VaR}_\alpha$, where $\widehat{VaR}_\alpha$ is defined by the α-quantile of the empirical distribution of the response scenarios that functions as an approximation of VaR_α. The properties of $\widehat{CVaR}_\alpha$ have been well studied. For example, although $\widehat{CVaR}_\alpha$ is biased, it is proven to be consistent under mild regularity conditions by Sun and Hong (2010). We will elaborate more on this point later.

Now back to the original model with input uncertainty that we are interested in, where stochastic uncertainty needs to be taken into account as well. Since the mean response $\eta(\theta)$ now cannot be evaluated exactly, it is estimated from samples of the response. To obtain an estimator of $CVaR_\alpha$, we can extend the sampling procedure by replacing $\eta(\theta_i)$'s with the corresponding sample average estimates $\widehat{\eta}(\theta_i)$'s. Specifically, for $i = 1,...,N$, draw M i.i.d. samples $\xi_{i1},...,\xi_{iM}$ from the distribution of ξ and evaluate the responses $H(\theta_i;\xi_{ij}), j = 1,...,M$, approximate the mean response $\eta(\theta_i)$ by $\widehat{\eta}(\theta_i) = \frac{1}{M}\sum_{j=1}^{M}H(\theta_i;\xi_{ij})$ and sort them in ascending order, denoted by $\widehat{\eta}(\theta^{(1)}) \leq \widehat{\eta}(\theta^{(2)}) \leq \cdots \leq \widehat{\eta}(\theta^{(N)})$. Here note that $(\theta^{(1)},...,\theta^{(N)})$ and $(\theta_{(1)},...,\theta_{(N)})$ are different order sequences due to sampling error. In fact, for fixed scenarios $\theta_1,...,\theta_N$, $(\theta_{(1)},...,\theta_{(N)})$ is a constant vector, while $(\theta^{(1)},...,\theta^{(N)})$ is a random vector that depends on the sample realizations of ξ. Finally, a nested Monte Carlo estimator of $CVaR_\alpha$ is given by

$$\widetilde{CVaR}_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right) = \frac{1}{(1-\alpha)N+1}\sum_{i=\alpha N}^{N}\widehat{\eta}(\theta^{(i)}). \tag{4}$$

With the complication of stochastic uncertainty, the properties of $\widetilde{CVaR}_\alpha$ in (4) become more difficult to analyze. Nevertheless, we will show that it remains to be consistent under mild regularity conditions, and thus using it as an inference for $CVaR_\alpha$ is still reasonable.

2.3 Consistency of CVaR Estimator

In this subsection, we will analyze the asymptotic behavior of $\widetilde{CVaR}_\alpha$ in (4). In particular, we will prove the consistency of the estimator. To facilitate the analysis, let us ignore the stochastic uncertainty first, and the problem reduces to standard CVaR estimation. Some well-established results on the analysis of $\widehat{CVaR}_\alpha$ in (3) will be useful. In particular, Sun and Hong (2010) have the following proposition on the asymptotic representation of $\widehat{CVaR}_\alpha$.

Proposition 1 Assume $\eta(\theta)$ admits a positive and continuously differentiable density around VaR_α. Then

$$\widehat{CVaR}_\alpha(\eta(\theta)) - CVaR_\alpha = \underbrace{\left(\frac{1}{N}\sum_{i=1}^{N}\left[VaR_\alpha + \frac{1}{1-\alpha}(\eta(\theta_i) - VaR_\alpha)^+\right] - CVaR_\alpha\right)}_{\text{error of an unbiased estimator of } CVaR_\alpha} + A_N, \tag{5}$$

where $A_N = O_{a.s.}(N^{-1}\log N)$, $(x)^+ \overset{\triangle}{=} \max\{x,0\}$. Here note that the statement $f(N) = O_{a.s.}(g(N))$ means that $f(N) \le C \cdot g(N)$ for some constant C almost surely.

Reformulate (2) as

$$CVaR_\alpha = VaR_\alpha + \frac{1}{1-\alpha}\mathbb{E}_{p(\cdot|\mathbf{x})}\left[(\eta(\theta) - VaR_\alpha)^+\right]. \tag{6}$$

Therefore, $\frac{1}{N}\sum_{i=1}^{N}\left[VaR_\alpha + \frac{1}{1-\alpha}(\eta(\theta_i) - VaR_\alpha)^+\right]$ is an unbiased estimator of $CVaR_\alpha$. Hence, Proposition 1 implies that the bias of $\widehat{CVaR}_\alpha$, $\mathbb{E}[A_N]$, is asymptotically insignificant compared with the error of the unbiased estimator. By Strong Law of Large Numbers (SLLN) and Central Limit Theorem (CLT), Sun and Hong (2010) have the following corollary on the strong consistency and asymptotic normality of $\widehat{CVaR}_\alpha$.

Corollary 2 Under the assumptions in Proposition 1, and $Var\left[(\eta(\theta) - VaR_\alpha)^+\right]$ is finite, the estimator $\widehat{CVaR}_\alpha(\eta(\theta))$ is strongly consistent and asymptotic normally distributed. In particular, $\widehat{CVaR}_\alpha \overset{N\to\infty}{\to} CVaR_\alpha$, $w.p.1$, and

$$\sqrt{N}\left(\widehat{CVaR}_\alpha - CVaR_\alpha\right) \Rightarrow \frac{\sqrt{Var\left[(\eta(\theta) - VaR_\alpha)^+\right]}}{(1-\alpha)}\mathscr{N}(0,1), \text{ as } N \to \infty, \tag{7}$$

where $\mathscr{N}(0,1)$ represents the standard normal distribution.

When stochastic uncertainty also needs to be taken into account, the error of estimator $\widetilde{CVaR}_\alpha$ in (4) becomes more complicated. Nevertheless, we will show that the estimator remains to be strongly consistent. In particular, we have the following theorem on the asymptotic consistency of $\widetilde{CVaR}_\alpha$.

Theorem 3 Under the same assumptions in Proposition 1, we have

$$\lim_{N\to\infty}\lim_{M\to\infty}\widetilde{CVaR}_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right) = CVaR_\alpha, \quad w.p.1. \tag{8}$$

In particular, $\widetilde{CVaR}_\alpha\left(\mathbb{E}_\xi[H(\theta;\xi)]\right)$ is a strongly consistent estimator of $CVaR_\alpha$.

We refer to Zhu and Zhou (2015) for the proof, and same for the rest of theorems in this paper. We also point out that right now the order of N and M in (8) cannot be interchanged or relaxed such that N and M go to infinity simultaneously. Relaxing such restriction on the limit is the direction of future research.

2.4 Confidence Interval

Now let us describe the procedure to construct a CI for $CVaR_\alpha$ and prove its asymptotic validity. A natural idea is to establish the asymptotic normality of $\widetilde{CVaR}_\alpha$ by providing an asymptotic representation of $\widetilde{CVaR}_\alpha$ similar to the one in (5). In view of the following error decomposition

$$\widetilde{CVaR}_\alpha - CVaR_\alpha = \left(\widetilde{CVaR}_\alpha - \widehat{CVaR}_\alpha\right) + \left(\widehat{CVaR}_\alpha - CVaR_\alpha\right) \stackrel{\triangle}{=} Err_1 + Err_2, \tag{9}$$

it is quite challenging to do so because Err_1, the error component mainly accounts for stochastic uncertainty, also depends on input uncertainty and thus not independent from Err_2. Alternatively, we can establish the asymptotic normality result for Err_1. Combining with a similar result for Err_2 in Corollary 2, we could construct CIs for Err_1 and Err_2 respectively, and integrate the two CIs into a wider CI for $CVaR_\alpha$. It remains to show that Err_1 is asymptotic normally distributed. In particular, we have the following theorem.

Theorem 4 Assume $\eta(\theta)$ is a continuous random variable. Then for any finite N and conditional on $\theta_1, ..., \theta_N$,

$$\sqrt{[(1-\alpha)N+1]M}\left(\widetilde{CVaR}_\alpha - \widehat{CVaR}_\alpha\right) \stackrel{M\to\infty}{\Rightarrow} \tau_I \mathcal{N}(0,1), \tag{10}$$

where $\tau_I \stackrel{\triangle}{=} \sqrt{\sum_{i=\alpha N}^{N} \tau_{(i)}^2 / [(1-\alpha)N+1]}$ and $\tau_{(i)}^2$ is used to denote $\tau_{\theta_{(i)}}^2$ with slight abuse of notations.

We refer to Zhu and Zhou (2015) for the proof. Basically, Theorem 4 implies that the disruption on the order statistics $(\eta(\theta_{(1)}), \eta(\theta_{(2)}), \cdots, \eta(\theta_{(N)}))$ caused by stochastic uncertainty vanishes asymptotically. With Corollary 2 and Theorem 4, let us construct a two-sided CI $[\widetilde{CL}, \widetilde{CU}]$ for $CVaR_\alpha$ with confidence level $1-\beta$. Following the error decomposition (9), the error level β is decomposed into β_O and β_I (hence $\beta = \beta_O + \beta_I$) as well, representing the errors due to input uncertainty (outer-layer simulation) and stochastic uncertainty (inner-layer simulation), respectively.

By Corollary 2, the two-sided (unknown variance) CI for $\widehat{CVaR}_\alpha - CVaR_\alpha$ with confidence level $1-\beta_O$ is

$$\widehat{CVaR}_\alpha - CVaR_\alpha \in \left[\frac{t_{\beta_O/2,N-1}\widehat{\sigma}_{cvar}}{\sqrt{N}}, \frac{t_{1-\beta_O/2,N-1}\widehat{\sigma}_{cvar}}{\sqrt{N}}\right], \tag{11}$$

where $\widehat{\sigma}_{cvar}$ is the sample estimate of $\sigma_{cvar} \stackrel{\triangle}{=} \sqrt{Var\left[(\eta(\theta) - VaR_\alpha)^+\right]}/(1-\alpha)$, $t_{\beta_O/2,N-1}$ and $t_{1-\beta_O/2,N-1}$ represent the $\beta_O/2$-level and $(1-\beta_O/2)$-level quantiles of the t-distribution with degree of freedom $N-1$, respectively. Note that here using the estimate $\widehat{\sigma}_{cvar}$ instead of σ_{cvar} is necessary because σ_{cvar} is usually unknown and can be obtained using the same samples generated in the simulation experiment.

Similarly, by Theorem 4, the two-sided (unknown variance) CI for $\widetilde{CVaR}_\alpha - \widehat{CVaR}_\alpha$ with confidence level $1-\beta_I$ is

$$\widetilde{CVaR}_\alpha - \widehat{CVaR}_\alpha \in \left[\frac{t_{\beta_I/2,[(1-\alpha)N+1]M-1}\widehat{\tau}_I}{\sqrt{[(1-\alpha)N+1]M}}, \frac{t_{1-\beta_I/2,[(1-\alpha)N+1]M-1}\widehat{\tau}_I}{\sqrt{[(1-\alpha)N+1]M}}\right], \tag{12}$$

where $\widehat{\tau}_I$ is the sample estimate of τ_I, $t_{\beta_I/2,[(1-\alpha)N+1]M-1}$ and $t_{1-\beta_I/2,[(1-\alpha)N+1]M-1}$ represent the $\beta_I/2$-level and $1-\beta_I/2$-level quantiles of the t-distribution with degree of freedom $[(1-\alpha)N+1]M-1$, respectively.

By integrating the CIs in (11) and (12), the two-sided (unknown variance) CI for $CVaR_\alpha$ with confidence level $1-\beta$ is

$$\begin{aligned} CVaR_\alpha \in & \left[\frac{1}{(1-\alpha)N+1}\sum_{i=\alpha N}^{N}\widehat{\eta}(\theta^{(i)}) + \frac{t_{\beta_O/2,N-1}\widehat{\sigma}_{cvar}}{\sqrt{N}} + \frac{t_{\beta_I/2,[(1-\alpha)N+1]M-1}\widehat{\tau}_I}{\sqrt{[(1-\alpha)N+1]M}},\right. \\ & \left.\frac{1}{(1-\alpha)N+1}\sum_{i=\alpha N}^{N}\widehat{\eta}(\theta^{(i)}) + \frac{t_{1-\beta_O/2,N-1}\widehat{\sigma}_{cvar}}{\sqrt{N}} + \frac{t_{1-\beta_I/2,[(1-\alpha)N+1]M-1}\widehat{\tau}_I}{\sqrt{[(1-\alpha)N+1]M}}\right]. \end{aligned} \tag{13}$$

The following theorem shows the asymptotic validity of the CI constructed above. The proof can be found in Zhu and Zhou (2015).

Theorem 5 Under the assumptions in Corollary 2 and Theorem 3, the CI defined in (13) is asymptotically valid, i.e.,

$$\lim_{N\to\infty}\lim_{M\to\infty} P\{\widetilde{CL} \le CVaR_\alpha \le \widetilde{CU}\} \ge 1-\beta, \tag{14}$$

where $\widetilde{CL}$, $\widetilde{CU}$ are the corresponding low-up CI boundaries for $CVaR_\alpha$ in (13), respectively.

Right now the restriction of iterated limits cannot be relaxed so that N and M in (14) is interchangeable, or N and M go to infinity simultaneously. This is a direct result of the restriction of iterated limits in (8). Furthermore, it seems to be difficult to obtain a narrower CI that remains to be asymptotically valid. The reason is that the error components Err_1 and Err_2 are not independent from each other. Therefore, we choose to split the overall error level into two error levels corresponding to the two error terms and constructing CI for each error term, which results in a wider CI than a typical result in CI construction.

3 BUDGET ALLOCATION

In a practical simulation experiment, there is usually a budget limit on the total computation consumption that is mainly influenced by the experiment parameters N and M. Intuitively, the number of outer-layer simulation N determines the error of CVaR estimator from sampling input uncertainty, while the number of inner-layer simulation M determines the error of CVaR estimator from sampling stochastic uncertainty. Therefore, choosing appropriate values for N and M is critical to balance the trade-off between capturing input uncertainty and capturing stochastic uncertainty, and improve the overall experiment performance. A natural criterion for evaluating the experiment performance is the (half) width of the CI, and a smaller CI width indicates better performance. Ideally, we want to choose values for N and M such that the resulted CI has the smallest width. Notice that the CI width is a random variable that depends on the realizations of scenarios and samples generated in the simulation experiment. Therefore, optimizing the expected CI width w.r.t. the parameters N and M before the experiment will provide us the guideline to determine a good budget allocation scheme for the simulation experiment. The key question is how to formulate and solve the CI width minimization problem.

Let us use $W_{cvar}(M,N)$ to denote the (approximate) expected half width of the CI in (13), i.e.,

$$W_{cvar}(N,M) := \frac{t_{1-\beta_O/2,N-1}\sigma_{cvar}}{\sqrt{N}} + \frac{t_{1-\beta_I/2,[(1-\alpha)N+1]M-1}\tau_{cvar}}{\sqrt{[(1-\alpha)N+1]M}}, \tag{15}$$

where $\tau^2_{cvar} \triangleq \mathbb{E}[\tau^2_\theta|\eta(\theta) \ge VaR_\alpha]$ is approximately equal to $\mathbb{E}[\tau^2_I]$. $W_{cvar}(M,N)$ is the objective function in the budget allocation problem. Notice that there are four experiment parameters β_O,β_I,N,M to be determined. To reduce the number of decision variables and ease the optimization, we pre-select β_O and β_I (a typical selection is $\beta_O = \beta_I = \beta/2$). Next, let us describe the constraints in the minimization problem. We use $\Delta(N,M)$ to denote the total computation time $S(N)+c\cdot NM$, where $S(N)$ is the time for sorting a vector of length N in ascending order, and $c\cdot NM$ is the time for generating and evaluating NM output response samples. Of course, there could be other criteria such as computation complexity and these can be minimized in a similar manner. CB is used to denote the total computation budget. Furthermore, there are lower bound constraints on N and M such that the statistical properties of the estimators are valid. To this end, let us consider the following minimization problem for CVaR estimation

$$\begin{aligned} \min_{N,M} \quad & W_{cvar}(N,M) \\ s.t. \quad & \Delta(N,M) \le CB \\ & N \ge \Gamma_0,\ M \ge \Gamma_0,\ [(1-\alpha)N+1]M \ge \Gamma_0 \\ & N,M \in \mathbb{Z}^+ \end{aligned} \tag{16}$$

The constraints $N \geq \Gamma_0$, $M \geq \Gamma_0$ and $[(1-\alpha)N+1]M \geq \Gamma_0$ are imposed to ensure the validity of the t-statistics, and a typical choice for Γ_0 is 30.

The practical challenge in solving the minimization problem (16) is the lack of information for the "variance parameters" σ_{cvar} and τ_{cvar} in the representation of $W_{cvar}(M,N)$. This challenge is not unique to our problem. In fact, for general computation budget allocation problems in simulation, the lack of knowledge for some important parameters is quite common. A natural fix is to run a pilot experiment with a small fraction of the budget, obtain crude approximations of those unknown parameters, and replace the unknown parameters with the corresponding approximations in the optimization problem. Finally, solving the approximate version of the budget allocation problem instead of the original one.

We will adopt this procedure in our problem. That is, we will perform a pilot run with a few samples to obtain crude approximations of σ_{cvar} and τ_{cvar}, denoted by $\widetilde{\sigma}_{cvar}$ and $\widetilde{\tau}_{cvar}$, respectively, and then solve the approximate version of budget allocation problem (16). Theoretically, these approximations of variance parameters can be the corresponding Monte Carlo estimates. However, the estimates might be extremely inaccurate or unstable because the Monte Carlo estimation in this case is a rare-event type simulation with few samples. For example, recall that $\sigma_{cvar} \triangleq \sqrt{Var\left[(\eta(\theta)-VaR_\alpha)^+\right]}/(1-\alpha)$, and its square can be reformulated as

$$
\begin{aligned}
\sigma^2_{cvar} &= \frac{Var\left[(\eta(\theta)-VaR_\alpha)^+\right]}{(1-\alpha)^2} \\
&= \frac{1}{(1-\alpha)^2}\left\{\mathbb{E}\left[(\eta(\theta)-VaR_\alpha)^2 I\{\eta(\theta)\geq VaR_\alpha\}\right]-\left(\mathbb{E}\left[(\eta(\theta)-VaR_\alpha)^+\right]\right)^2\right\}. \qquad (17)
\end{aligned}
$$

By definition we have $CVaR_\alpha = VaR_\alpha + \frac{1}{1-\alpha}\mathbb{E}\left[(\eta(\theta)-VaR_\alpha)^+\right]$. Hence, estimation of σ^2_{cvar} is at least as difficult as estimation of $CVaR_\alpha$—the initial goal in our model. If we use the naive Monte Carlo estimation to approximate the two expectation terms in (17), most of samples will be ineffective. In fact, theoretically only $100(1-\alpha)$ percent of the total samples will be effective. Therefore, the approximation might be extremely inaccurate because the total number of samples in the pilot run is very limited. To put it into perspective, suppose $\alpha = 0.99$ and 100 scenarios of mean response ($N = 100$) are generated in the pilot run. Then theoretically only one scenario will be used to estimate σ_o since the rest 99 scenarios result in a simple value of 0. Obviously, the one-scenario approximation is very likely to be far from the true value.

Intuitively, the naive Monte Carlo estimation is problematic because the information about the distribution carried by the ineffective samples is not used. When the number of samples is small, the problem become even more severe because the inference on the very tail of the distribution is negligible with little/no information used. Therefore, in theory, a good approximation will try to make use of the information carried by all the samples. For example, using (adaptive) importance sampling will turn most of ineffective samples to effective samples, and therefore improve the accuracy of the estimation. This approach is not readily applicable in our model because we lack the knowledge for the density function of the mean response distribution and the number of total available samples is too limited to be used to learn the distribution.

Next, we will describe a new approach to approximations of the variance parameters, denoted by $\widetilde{\sigma}_{cvar}$ and $\widetilde{\tau}_{cvar}$, that exploits the information carried by all the generated samples. In view of the definitions of the variance parameters, i.e.,

$$
\sigma_{cvar} \triangleq \sqrt{Var\left[(\eta(\theta)-VaR_\alpha)^+\right]}/(1-\alpha),\ \tau_{cvar} \triangleq \sqrt{\mathbb{E}[\tau^2_\theta|\eta(\theta)\geq VaR_\alpha]},
$$

the challenges in computing these variance parameters are two folds: (i) The lack of explicit formula for the density of $\eta(\theta)$ in computing σ_{cvar}. (ii) The lack of functional representation of τ^2 in $\eta(\theta)$, i.e., $\tau^2(y) \triangleq \mathbb{E}[\tau^2_\theta|\eta(\theta)=y]$, in computing τ_{cvar}.

To address the first challenge, we apply the "density projection" technique to project the empirical distribution of the mean response $\eta(\theta)$ onto a parameterized family of distributions. To be more specific, a *projection mapping* from a space of probability distributions $\mathscr{P}$ to another space consisting of a parameterized family of densities $\mathscr{F}$, denoted as $Proj_{\mathscr{F}} : \mathscr{P} \to \mathscr{F}$, is defined by

$$Proj_{\mathscr{F}}(g) \stackrel{\triangle}{=} \arg\min_{f \in \mathscr{F}} D_{KL}(g \parallel f), \quad \forall g \in \mathscr{P}, \tag{18}$$

where $D_{KL}(g \parallel f)$ denotes the *Kullback-Leibler (KL) divergence* (or *relative entropy*) between g and f, which is

$$D_{KL}(g \parallel f) \stackrel{\triangle}{=} \int g(x) \log \frac{g(x)}{f(x)} dx. \tag{19}$$

Here note that the densities g and f are assumed to have the same support space. Hence, the projection of g on $\mathscr{F}$ has the minimum KL divergence from g among all the densities in $\mathscr{F}$. Loosely speaking, the projection of g on $\mathscr{F}$ is the "best" approximation of g one can find in $\mathscr{F}$. When $\mathscr{F}$ is an exponential family of densities, the minimization problem (18) has an analytical solution and can be carried out easily. The exponential families include many common families of densities, such as Gaussian, Binomial, Poisson, Gamma, etc. Therefore, by choosing an exponential family of densities as the target space in projecting the empirical distribution of $\eta(\theta)$, the resulting density admits an explicit formula and it can be regarded as an approximation of the probability density function of $\eta(\theta)$. This significantly facilitates the computation of σ_{cvar} by simple numerical calculation instead of Monte Carlo sampling. More importantly, this technique makes use of the information carried by all the samples.

As for the second challenge, i.e., $\tau^2(y) \stackrel{\triangle}{=} \mathbb{E}[\tau_\theta^2 | \eta(\theta) = y]$ is not available in computing τ_{cvar}. We can apply regression for $\tau^2(y)$ onto $\eta(\theta)$ to construct the response surface of $\tau^2(y)$ with sample responses of $\tau^2(y)$ (which can be easily computed from the scenarios generated in the pilot run). For example, a polynomial regression with basis functions consisting of polynomial functions of $\eta(\theta)$ is sufficient for a good approximation of $\tau^2(y)$. Finally, $\widetilde{\tau}_{cvar}$ is computed via numerical integration since $\tau^2(y)$ admits an (approximate) explicit formula and $\eta(\theta)$ admits an (approximate) explicit probability density function.

After plugging in the approximated variance parameters $\widetilde{\sigma}_{cvar}$ and $\widetilde{\tau}_{cvar}$ into problem (16), the remaining challenge is how to solve the optimization problem efficiently. Obviously, solving it analytically to optimality is very unlikely because the objective function does not exhibit structure properties (e.g., convexity). Alternatively, one can enumerate a reasonable amount of candidate solutions (for examples, a two-dimensional grid of feasible solutions) easily based on the constraints, and evaluating the objective function at these candidate solutions can be carried out straightforwardly. Finally, among the candidate solutions, we can choose the best parameter design that yields the smallest expected CI width.

Note that, due to the special structure of the CI width minimization problem (16), if the total budget limit is instead defined by $CB = N \times M$, then the objective function $W_{cvar}(N,M)$ is independent of M and decreases in N. It follows that the optimal solution is to take N as large as possible. Mathematically, it means the optimal allocation scheme is $(N,M) = (\lfloor CB/\Gamma_0 \rfloor, \Gamma_0)$ such that the t-statistics are valid as well. Our numerical experiments (not included in this paper due to space limit) demonstrate that this indeed is the case. Nevertheless, here we consider a more general problem formulation and a more sophisticated solution technique because they could be readily generalized for risk measures other than CVaR.

4 NUMERICAL RESULTS

We will use the $M/M/1$ queueing system considered in Zouaoui and Wilson (2003) to illustrate the theoretical results in previous sections. In particular, we will consider the estimation of CVaR of the mean sojourn time w.r.t. input uncertainty. In the $M/M/1$ queueing system, assume the "true" Poisson customer arrival rate is λ_o, which means the inter-arrival times between customers are independently sampled from an exponential distribution with rate λ_o. Further assume the "true" Exponential service rate is μ_o, which means the service time for each customer is sampled from an exponential distribution with rate μ_o. Here "true" means that the values of λ_o and μ_o are known to us (the judges) but unknown to the experimenter. Specifically, we choose $\mu_o = 500$ and $\lambda_o = 50, 250, 450$—a range of values corresponding to increasing levels of "true" arrival intensity. Assuming non-informative priors for both the Poisson arrival rate λ and the Exponential service rate μ, i.e., $p(\lambda) \propto 1/\lambda$ and $p(\mu) \propto 1/\mu$. The experimenter takes a Bayesian approach towards the construction of the posterior distributions of λ and μ. That is, based on $n = 10, 100, 10000$ historical observations of λ and μ (drawn from the corresponding distributions with "true" parameters), standard Bayesian updating is applied to obtain the posterior distributions of λ and μ, which are the experimenter's belief on input uncertainty. In particular, assume the historical observations of λ are $\mathbf{x} = (x_1, ..., x_n)$. Then the Bayesian updating of the posterior distribution of λ can be carried out analytically and results in $p(\lambda|\mathbf{x}) = \lambda^{n-1} \exp(-\lambda \sum_{i=1}^{n} x_i)$, which is a Gamma distribution with shape parameter n and scale parameter $1/(\sum_{i=1}^{n} x_i)$. Similarly, let $\mathbf{y} = (y_1, ..., y_n)$ denote the historical observations for μ. Then the posterior distribution of μ is $p(\mu|\mathbf{y}) = \mu^{n-1} \exp(-\mu \sum_{i=1}^{n} y_i)$, which is also a Gamma distribution with shape parameter n and scale parameter $1/(\sum_{i=1}^{n} y_i)$.

The objective is to use nested Monte Carlo simulation to estimate $CVaR_\alpha$ ($\alpha = 0.90, 0.95, 0.99$) of the mean sojourn time w.r.t. the posterior distributions $p(\lambda|\mathbf{x})$ and $p(\mu|\mathbf{y})$, and construct the associated $100(1-\beta)\%$ CIs ($\beta = 0.05$). Specifically, we draw $N = 5000$ scenarios from the posterior distributions of the parameters such that every parameter scenario results in a stable queue ($\lambda < \mu$), and for each scenario of the parameters, we further draw $M = 200$ samples by simulating the first 200 sojourn cycles of that queue to estimate the mean sojourn time. Finally, the $CVaR_\alpha$ estimator and the associated CIs are computed via (4) and (13), in which $\widehat{\sigma}_{cvar}$ and $\widehat{\tau}_l$ are computed using the same samples generated. Note that we choose a large value for N since CVaR estimation is a rare-event type simulation.

Table 1: CVaR (with 95% CI) for the mean sojourn time in an M/M/1 queue.

λ_0	n	Mean $\pm$ Half CI Width	$CVaR_{\alpha_1} \pm$ Half CI Width	$CVaR_{\alpha_2} \pm$ Half CI Width	$CVaR_{\alpha_3} \pm$ Half CI Width
50	10	$2.4 \times 10^{-3} \mp 3.4 \times 10^{-5}$	$5.0 \times 10^{-3} \mp 2.8 \times 10^{-4}$	$6.0 \times 10^{-3} \mp 4.7 \times 10^{-4}$	$9.0 \times 10^{-3} \mp 1.6 \times 10^{-3}$
50	100	$2.2 \times 10^{-3} \mp 9.7 \times 10^{-6}$	$2.8 \times 10^{-3} \mp 4.9 \times 10^{-5}$	$2.9 \times 10^{-3} \mp 6.9 \times 10^{-5}$	$3.2 \times 10^{-3} \mp 1.6 \times 10^{-4}$
50	10000	$2.2 \times 10^{-3} \mp 6.9 \times 10^{-6}$	$2.6 \times 10^{-3} \mp 3.3 \times 10^{-5}$	$2.6 \times 10^{-3} \mp 4.7 \times 10^{-5}$	$2.8 \times 10^{-3} \mp 9.8 \times 10^{-5}$
250	10	$5.2 \times 10^{-3} \mp 2.1 \times 10^{-4}$	$2.1 \times 10^{-2} \mp 2.4 \times 10^{-3}$	$3.1 \times 10^{-2} \mp 4.2 \times 10^{-3}$	$5.3 \times 10^{-2} \mp 9.6 \times 10^{-3}$
250	100	$4.2 \times 10^{-3} \mp 4.1 \times 10^{-5}$	$7.0 \times 10^{-3} \mp 3.3 \times 10^{-4}$	$7.8 \times 10^{-3} \mp 5.6 \times 10^{-4}$	$1.0 \times 10^{-2} \mp 2.0 \times 10^{-3}$
250	10000	$3.9 \times 10^{-3} \mp 1.8 \times 10^{-5}$	$4.8 \times 10^{-3} \mp 9.7 \times 10^{-5}$	$4.9 \times 10^{-3} \mp 1.4 \times 10^{-4}$	$5.3 \times 10^{-3} \mp 3.2 \times 10^{-4}$
450	10	$9.9 \times 10^{-3} \mp 3.3 \times 10^{-4}$	$3.8 \times 10^{-2} \mp 3.0 \times 10^{-3}$	$4.7 \times 10^{-2} \mp 4.6 \times 10^{-3}$	$6.8 \times 10^{-2} \mp 1.1 \times 10^{-2}$
450	100	$1.8 \times 10^{-2} \mp 3.6 \times 10^{-4}$	$4.3 \times 10^{-2} \mp 2.4 \times 10^{-3}$	$4.9 \times 10^{-2} \mp 3.3 \times 10^{-3}$	$5.8 \times 10^{-2} \mp 7.8 \times 10^{-3}$
450	10000	$2.1 \times 10^{-2} \mp 2.6 \times 10^{-4}$	$3.5 \times 10^{-2} \mp 1.7 \times 10^{-3}$	$3.8 \times 10^{-2} \mp 2.4 \times 10^{-3}$	$4.4 \times 10^{-2} \mp 5.7 \times 10^{-3}$

The numerical results are summarized in Table 1, and based on which we have the following observations.

- In general, there is a significant gap between mean of the mean sojourn time (values in column 3) w.r.t. input uncertainty and CVaR of the mean sojourn time (values in columns 4 to 6) w.r.t. input uncertainty. That implies using the risk estimators instead of mean estimators is quite necessary for accurate risk assessment/control.

- As the number of historical observations increases, CVaR of the mean sojourn time is decreasing, which indicates that the effect of input uncertainty on the extreme behavior of the mean sojourn time is reducing. Intuitively, the posterior distribution of the input parameter will put more weights on values close to the “true” parameter and less weights on the values far from the “true” parameter as the number of historical observations increases. Therefore, loosely speaking, the posterior distribution of the mean response will also put more weights on the values close to “true” mean response and essentially reduce the risk of input uncertainty.
- As the arrival traffic intensifies (λ_o increases) and approaches the service rate μ_o, the “true” mean sojourn time (equals $1/(\mu_o - \lambda_o)$) is increasing, and the effect of input uncertainty on the extremely behavior of the mean sojourn time is more significant. Intuitively, it indicates that, as the arrival rate λ_o approaches the service rate μ_o, the mean system output response becomes more and more sensitive to input uncertainty.

In conclusion, the primary numerical results from simulating the simple $M/M/1$ queueing system provide empirical evidences for the importance/neccessity of investigating CVaR (or other risk measures) of the mean response w.r.t. input uncertainty in a stochastic system. We will further conduct numerical experiments with budget allocation scheme on more complex systems in future study.

5 CONCLUSION

In the present paper, we introduce risk measures for input certainty, which rigorously quantify the extreme behavior of the mean output response under all possible input models. In particular, we propose nested Monte Carlo estimator for CVaR of the mean response w.r.t. input uncertainty. We show some asymptotic properties of the estimator and use them to construct an (asymptotically valid) CI. The work in this paper can be viewed as a starting point of research on more general risk measures for input uncertainty. On the other hand, the naive estimator considered could be very restrictive for risk assessment of input uncertainty for large-scale systems, due to the fact that most of the scenarios generated for the input parameters will fall into the non-tail portion of the mean response distribution and essentially become ineffective. The budget allocation problem we considered and solved partially addresses this issue in the sense that it leads to good budget allocation schemes. Developing more sophisticated estimators of various risk measures and budget allocation schemes will be a promising direction of future research.

REFERENCES

Ankenman, B., B. L. Nelson, and J. Staum. 2010. “Stochastic Kriging for Simulation Metamodeling”. *Operations Research* 58 (2): 371–382.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. 1997. “Coherent Measures of Risk”. *Mathematical Finance*:203–228.

Barton, R. R. 2012. “Tutorial: Input Uncertainty in Outout Analysis”. In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R., B. L. Nelson, and W. Xie. 2013. “Quantifying Input Uncertainty via Simulation Confidence Intervals”. *INFORMS Journal on Computing* 26 (1): 74–87.

Barton, R. R., and L. W. Schruben. 1993. “Uniform and Bootstrap Resampling of Empirical Distributions”. In *Proceedings of the 1993 Winter Simulation Conference*, edited by G. W. Evans, M. Mollaghasemi, W. E. Biles, and E. C. Russell, 503–508. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R., and L. W. Schruben. 2001. "Resampling Methods for Input Modeling". In *Proceedings of the 2001 Winter Simulation Conference*, edited by B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 372–378. 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.

Cheng, R. C., and W. Holland. 1997. "Sensitivity of Computer Simulation Experiments to Errors in Input Data". *Journal of Statistical Computation and Simulation* 57 (1-4): 219–241.

Chick, S. E. 2001. "Input Distribution Selection for Simulation Experiments: Accounting for Input Uncertainty". *Operations Research* 49 (5): 744–758.

Glasserman, P., P. Heidelberger, and P. Shahabuddin. 2000. "Variance Reduction Techniques for Estimating Value-at-Risk". *Management Science* 46 (10): 1349–1364.

Gordy, M. B., and S. Juneja. 2010. "Nested Simulation in Portfolio Risk Measurement". *Management Science* 56 (10): 1833–1848.

Henderson, S. G. 2003. "Input Model Uncertainty: Why Do We Care and What Should We Do about It?". In *Proceedings of the 2003 Winter Simulation Conference*, edited by D. F. S. Chick, P. J. Sdnchez and D. J. Morrice, 90–100. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hong, L. J., Z. Hu, and G. Liu. 2014. "Monte Carlo Methods for Value-at-Risk and Conditional Value-at-Risk: A Review". *ACM Transactions on Modeling and Computer Simulation* 24 (4): 22.

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.-H. 1998. *Monte Carlo Computation of Conditional Expectation Quantiles*. Ph. D. thesis, Stanford University, Stanford, California.

Liu, M., and J. Staum. 2010. "Stochastic Kriging for Efficient Nested Simulation of Expected Shortfall". *Journal of Risk* 12 (3): 3–27.

Rockafellar, R. T., and S. Uryasev. 2000. "Optimization of Conditional Value-at-Risk". *Journal of Risk* 2:21–42.

Song, E., and B. L. Nelson. 2015. "Quickly Assessing Contributions to Input Uncertainty". *IIE Transactions* 47:893–909.

Steckley, G. 2006. *Estimating the Density of a Conditional Expectation*. Ph. D. thesis, Cornell University, Ithaca, New York.

Sun, L., and L. J. Hong. 2010. "Asymptotic Representations for Importance-Sampling Estimators of Value-at-Risk and Conditional Value-at-Risk". *Operations Research Letters* 38 (4): 246–251.

Xie, W., B. L. Nelson, and R. R. Barton. 2014. "A Bayesian Framework for Quantifying Uncertainty in Stochastic Simulation". *Operations Research* 62 (6): 1439–1452.

Zhu, H., and E. Zhou. 2015. "Risk Assessment of Input Uncertainty". *Working Paper*.

Zouaoui, F., and J. R. Wilson. 2003. "Accounting for Parameter Uncertainty in Simulation Input Modeling". *IIE Transactions* 35 (9): 781–792.

AUTHOR BIOGRAPHIES

HELIN ZHU is a Ph.D student in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. His research interests include stochastic control and simulation optimization. His email address is hzhu67@gatech.edu.

ENLU ZHOU is an Assistant Professor in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. She received the B.S. degree with highest honors in electrical engineering 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, with applications towards financial engineering. Her e-mail address is enlu.zhou@isye.gatech.edu and her web page is http://enluzhou.gatech.edu/.