Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

A SIMULATION-BASED PREDICTION FRAMEWORK FOR TWO-STAGE DYNAMIC DECISION MAKING

Wei Xie
Yuan Yi

Department of Industrial and System Engineering
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180-3590, USA

ABSTRACT

When we make operational decisions for high-tech manufacturing with products having short life times, there exist various challenges: (1) high uncertainty in the supply, production and demand; (2) limited amount of valid historical data; and (3) decision makers with a risk-averse attitude. We propose a simulation-based prediction framework to support real-time decision making. Specifically, we consider a generalized two-stage dynamic decision model accounting for both input uncertainty and system inherent stochastic uncertainty. Since the risk-adjusted cost objective involves nested risk measures, it could be computationally prohibitive to precisely estimate the system performance, especially for complex stochastic systems. Given a decision policy, in this paper, a metamodel-assisted approach is introduced to efficiently assess the system risk performance in the planning horizon, while delivering a credible interval quantifying the simulation estimation error. This information can guide us to select a good policy for real-time decision making.

1 INTRODUCTION

This paper is motivated by our on-going research on supply chain risk management collaborated with Regeneron, a bio-pharmaceutical manufacturing company. There exist various challenges for decision making, including *rapid change* in the technology and markets for end-products and *high uncertainty* in the supply, production and demand (Kaminsky and Wang 2015). Due to the significant capital investment required for the production process, globalization and specialization are often used to maximize the return on investment, which could lead to high supply uncertainty. Since some defects impacting the quality of products, e.g., contamination and cross-contamination in production lines (PLs), are hard to detect, there exist quality assurance tests following each production step. The testing cycle times could be lengthy and the results could have high variability. In addition, the demands of clinic products are much less predictable compared to those of commercial products. Therefore, how to dynamically make operational decisions, including decisions related to procurement, testing, inventory and production, under various sources of uncertainty becomes an important question.

Statistical models characterizing uncertainty in the supply, production and demand are called *input models*. Since the underlying true input models are often estimated by a finite amount of valid historical data, there exists input estimation uncertainty, called *input uncertainty*. Given an input model estimate, *stochastic uncertainty* represents the variability of system response, e.g., the total cost, induced by the input models. Thus, we should consider both input and stochastic uncertainty in the decision making.

The literature on stochastic optimization considers both sources of uncertainty *separately*. Stochastic uncertainty is a major concern in stochastic and robust optimization (Shapiro, Dentcheva, and Ruszczynski 2009). The distributionally robust optimization (Scarf, Arrow, and Karlin 1958; Delage and Ye 2010) considering the worse case situation for input uncertainty could lead to conservative decisions. In addition, the composite of two risk measures quantifying input and stochastic uncertainty was studied in Qian,

Wang, and Wen (2015) and Zhou and Xie (2015). Different combinations of these measures could lead to totally different system behaviors (Zhou and Xie 2015), and it is not clear how to select appropriate risk measures in the composite approach. Further, it could be hard to integrate this approach with dynamic optimization. Motivated by our previous study on input uncertainty (Xie, Nelson, and Barton 2014), *a compound approach* for stochastic optimization was proposed in Hu, Xie, and Xu (2016), which quantifies both input and stochastic uncertainty simultaneously. Since the analytic study over a newsvendor example demonstrates that the compound approach could lead to good and robust decisions (Hu, Xie, and Xu 2016), in this paper, we extend it to dynamic situations.

Optimal decisions depend on the choice of system performance measure. Even though most literature on supply chain management focuses on the mean performance, for high-tech products with high profit, decision makers tend to be *risk averse* (Chen et al. 2007). Thus, we consider risk measures to quantify the system performance. In this paper, we use the conditional value-at-risk (CVaR), the tail conditional expectation of the cost, for illustration; See Shapiro, Dentcheva, and Ruszczynski (2009) for more information on CVaR.

Many challenges arise when trying to apply strictly analytical approaches to assess the risk performance of real complex stochastic systems. Since stochastic simulation is in every sense a statistical experiment, one of the most valuable features of simulation is its ability to characterize the risks inherent in complex stochastic systems (Nelson 2013). Kaminsky and Wang (2015) also pointed out that simulation-based approaches have great promise for supply chain risk management in the biopharmaceutical industry.

Therefore, in this paper, we propose a simulation-based prediction framework to support real-time decision making. We consider a generalized two-stage dynamic decision model. This class of problems was first proposed by Beale (1955) and Dentzig (1955), and gained wide usage subsequently. Shapiro and Homem-de Mello (1998) developed a simulation approach, and they mainly focused on two-stage linear programming problems with mean performance measure. In this paper, we consider more general situations with risk measures. Further, since the underlying true input models are unknown and estimated by real-world data, we also account for the input uncertainty.

Since the risk-adjusted cost objective in the two-stage model involves nested risk measures, it could be computationally prohibitive to precisely estimate the system risk performance in the planning horizon, especially for complex stochastic systems. Thus, in the simulation-based prediction framework, we introduce a metamodel-assisted approach that can make efficient use of the simulation resources to assess the risk-adjusted cost objective, while delivering a credible interval (CrI) quantifying the simulation estimation uncertainty. This information could be useful to guide the search for a good decision policy. Specifically, there are two decision-making points in the two-stage model. Given the historical data, the posterior predictive distribution is used to drive the simulation and predict the future system performance. As time evolves to the second decision point, new data arrive and our belief of input models is updated by following Bayes' rule. Then, the updated posterior predictive distribution can drive the simulation to estimate the system future performance after the second decision point. To assess the risk-adjusted cost objective, a Gaussian process (GP) metamodel is constructed to estimate the system future performance at each sample of updated input models and it also provides the posterior distribution characterizing the simulation estimation uncertainty.

In sum, the main contributions of this paper include: (1) To facilitate real-time decision making in supply chain risk management, we propose a generalized two-stage dynamic decision model that can account for both input and stochastic uncertainty simultaneously; (2) In our simulation-based prediction framework, a metamodel-assisted approach is introduced to make efficient use of the simulation resources so that we can quickly assess the system performance in the planning horizon for any given decision policy; and (3) Our approach also delivers a CrI quantifying the simulation estimation uncertainty.

The next section provides a formal problem description and describes the two-stage dynamic decision model for real-time decision making. In the simulation-based prediction framework proposed in Section 3, a metamodel-assisted approach is introduced to efficiently assess the system risk performance for a given decision policy, while delivering a CrI quantifying the simulation estimation error. A multi-stage inventory

example is used to study the finite-sample performance of our approach in Section 4. We conclude this paper in Section 5.

2 PROBLEM DESCRIPTION AND TWO-STAGE DYNAMIC DECISION MODEL

To make it easy to follow, in this section, we first consider *the static setting* with a fixed amount of real-world data, and then extend it to the dynamic setting with data streams. The input models, denoted by F, could include mutually independent models. For an inventory example, F could include input models for supply lead times and product demands. For notation simplification, suppose F is composed of a single input model.

The output of a stochastic system of interest is denoted by $Y = H(\mathbf{u}, \mathbf{X})$, where $\mathbf{u}$ represents decisions and $\mathbf{X}$ represents input samples. For example, in a newsvendor example, Y could be the cost, $\mathbf{u}$ and $\mathbf{X}$ could be the ordering decision and product demands. Most research in the existing supply chain literature focuses on the mean performance of Y. In recent years, some studies have started to use risk measures to characterize the system performance; see Chen, Xu, and Zhang (2009) and Gotoh and Takano (2007). Since decision makers for high-tech products in the bio-pharmaceutical industry tend to be risk averse, we will explore risk measures, denoted by $\rho(Y)$, to quantify the system performance. In this paper, we use CVaR for illustration and let $\rho = \text{CVaR}_\alpha$ with the right tail probability $\alpha \in (0,1)$.

If the underlying "correct" input model, denoted by F^c specified by the input parameters $\boldsymbol{\theta}^c$, is known, the system risk performance is

$$\rho_{\mathbf{X}}[Y(\boldsymbol{\theta}^c)] = \min_{\xi \in \Re} \int_{\mathbf{x}} \left[\xi + \frac{1}{\alpha} (H(\mathbf{u}, \mathbf{x}) - \xi)^+ \right] f(\mathbf{x}|\boldsymbol{\theta}^c) d\mathbf{x} \tag{1}$$

where $f(\mathbf{x}|\boldsymbol{\theta})$ denotes the input density function. Since the variability in the response Y is induced by $\mathbf{X}$ through the logic of system of interest and $H(\cdot)$ is a (unknown) deterministic function, the subscript $\mathbf{X}$ in $\rho_{\mathbf{X}}$ emphasizes that the risk measure ρ depends on the distribution of $\mathbf{X}$. Therefore, given the true input model F^c, Equation (1) quantifies the impact of *stochastic uncertainty* on the system risk performance.

In reality, the underlying input model F^c is unknown. In this paper, we assume that the parametric family is known and unknown true parameters $\boldsymbol{\theta}^c$ are estimated using valid historical data, denoted by $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0}\}$ with $\mathbf{x}_i \overset{i.i.d.}{\sim} F^c$ for $i = 1, 2, \ldots, t_0$. Following Bayes' rule, the input uncertainty or our belief about $\boldsymbol{\theta}^c$ can be characterized by a posterior distribution, denoted by $g(\boldsymbol{\theta}) \equiv p(\boldsymbol{\theta}|\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0})$. Then, we use the input model estimate to predict the new data $\mathbf{X}$. Thus, the compound random vector $\mathbf{X}(\boldsymbol{\Theta})$ accounts for both input uncertainty, $\boldsymbol{\Theta} \sim g(\boldsymbol{\theta})$, and stochastic uncertainty, $\mathbf{X}|\boldsymbol{\Theta} \sim F(\cdot|\boldsymbol{\Theta})$. We can quantify their impact on the system risk performance by

$$\rho_{\mathbf{X}}[Y(\boldsymbol{\Theta})] = \min_{\xi \in \Re} \int_{\mathbf{x}} \left[\xi + \frac{1}{\alpha} (H(\mathbf{u}, \mathbf{x}) - \xi)^+ \right] f(\mathbf{x}|\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0}) d\mathbf{x} \tag{2}$$

where $f(\mathbf{x}|\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0}) \equiv \int f(\mathbf{x}|\boldsymbol{\theta}) g(\boldsymbol{\theta}) d\boldsymbol{\theta}$ is *the posterior predictive density*. Therefore, the risk measure in Equation (2) accounts for both input and stochastic uncertainty. Notice that the posterior predictive distribution is often used in the Bayesian approaches to check if the selected input model is appropriate or it can capture the important properties in the historical data (Gelman et al. 2004).

Equations (1) and (2) have similar forms. However, the measure $f(\mathbf{x}|\boldsymbol{\theta}^c)$ only accounts for stochastic uncertainty, and the measure $f(\mathbf{x}|\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0})$ accounts for both input and stochastic uncertainty. The input model only provides a tool to extract the important information from the historical data and predict the new observation $\mathbf{x}$. When we assess the system risk performance, it is not appropriate to separate input and stochastic uncertainty because the estimation over F^c is mainly for predicting the new observation $\mathbf{x}$ based on the information collected from the historical data $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0}\}$.

Then, we consider *the dynamic setting* with data streams, denoted by $\{\mathbf{X}_t\}_{t=1}^{\infty}$, from the underlying true input model F^c. In each time period t, we first observe the new data $\mathbf{X}_t$, update the physical state $\mathbf{S}_t$ and

the knowledge state or our belief on the input model, and then implement the decision $\mathbf{u}_t$. For example, for an inventory system with a single product, in the time period t, we first observe the realization of demand X_t. Then, we update the inventory level either by satisfying the demand with available stock or recording it as the backlog if there is out of stock. Based on the new information from X_t, we also update our belief on the demand distribution. After that, we make the ordering decision u_t.

At any time t_0, given the historical data $\mathbf{x}_{[t_0]} \equiv (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{t_0})$, decision makers want to find the optimal decisions $\mathbf{u}_t$ with $t \geq t_0$ to minimize the risk-adjusted cost in the planning horizon. For notation simplification, suppose the current time period $t_0 = 0$ and denote the historical data by $\mathbf{x}_{[0]}$. Since it is typically computationally intractable to directly solve the dynamic programming, we build an approximate look-ahead model to guide *real-time decision making*. In this paper, we consider *a generalized two-stage decision model* which has two decision-making points at t_0 and t_1. Specifically, based on the historical data $\mathbf{x}_{[0]}$, we first make decision $\mathbf{u}_0$ at t_0. Then, after collecting the new information from observations $\mathbf{X}_{[1:t_1]} \equiv (\mathbf{X}_1, \ldots, \mathbf{X}_{t_1})$, we further make decisions $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_T$ at t_1 with $t_1 \leq T$. The objective function becomes

$$\min_{\mathbf{u}_0} C_0(\mathbf{S}_0, \mathbf{u}_0) + \rho_{\mathbf{X}_{[1:t_1]}} \left[\min_{\mathbf{u}_1, \ldots, \mathbf{u}_T} \sum_{t=1}^{t_1} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) + \rho_{\mathbf{X}_{[t_1+1:T]}} \left(\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) \right) \right] \tag{3}$$

where $\mathbf{S}_t$ represents physical state variables, i.e., the inventory levels at each stage of supply chains; $C_t(\cdot)$ denotes the cost function in period t for the stochastic system of interest; and r denotes the time discount factor. The state transition $\mathbf{S}_{t+1} = h_t(\mathbf{S}_t, \mathbf{u}_t, \mathbf{X}_{t+1})$ is determined by the system logic, and the system is driven by the posterior predictive distribution with $\mathbf{X}_{t+1} \sim f(\mathbf{x}|\mathbf{x}_{[t]})$. The value of t_1 determines the stage aggregation and T represents the planning horizon. Finding appropriate t_1 and T is our on-going research.

At any time t, the posterior distribution of $\boldsymbol{\Theta}$, denoted by $g_t(\boldsymbol{\theta})$, characterizes our belief about the input model. When time evolves to next time period $t+1$ and the new data $\mathbf{x}_{t+1}$ arrives, our belief is updated by applying Bayes' rule, i.e., $g_{t+1}(\boldsymbol{\theta}) \propto g_t(\boldsymbol{\theta}) \cdot p(\mathbf{x}_{t+1}|\boldsymbol{\theta})$.

For complex stochastic systems, since it is typically impossible to get the analytic form of system risk performance in the planning horizon, we estimate it by simulation. Given a decision policy, the state transition and the decisions $(\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_T)$ are directly obtained by following the logic of the simulated system. To find good real-time decisions, it is critical to *quickly* assess the risk-adjusted cost objective

$$\rho_{\mathbf{X}}[Y(\boldsymbol{\Theta})] = C_0(\mathbf{S}_0, \mathbf{u}_0) + \rho_{\mathbf{X}_{[1:t_1]}} \left[\sum_{t=1}^{t_1} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) + \rho_{\mathbf{X}_{[t_1+1:T]}} \left(\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) \right) \right]. \tag{4}$$

In the literature, the sample average approximation (SAA) is recommended to estimate the risk performance, e.g., CVaR (Shapiro, Dentcheva, and Ruszczynski 2009). The sampling procedure for estimating Objective (4) includes: (1) Generate $\boldsymbol{\theta}^{(i)} \sim g_0(\boldsymbol{\theta})$ and $\mathbf{X}^{(i)}_{[1:t_1]} \sim f(\mathbf{x}_{[1:t_1]}|\boldsymbol{\theta}^{(i)})$ with $i = 1, 2, \ldots, N$; (2) At each $\mathbf{X}^{(i)}_{[1:t_1]}$, update our belief about the input model characterized by $g^{(i)}_{t_1}(\boldsymbol{\theta}) \propto g_0(\boldsymbol{\theta}) \cdot p(\mathbf{X}^{(i)}_{[1:t_1]}|\boldsymbol{\theta})$, and then generate $\boldsymbol{\theta}^{(ij)} \sim g^{(i)}_{t_1}(\boldsymbol{\theta})$ and $\mathbf{X}^{(ij)}_{[t_1+1:T]} \sim f(\mathbf{x}_{[t_1+1:T]}|\boldsymbol{\theta}^{(ij)})$ with $j = 1, 2, \ldots, n_i$, where N and $(n_1, n_2, \ldots, n_N)$ represent the number of replications used to estimate the risk measures $\rho_{\mathbf{X}_{[1:t_1]}}$ and $\rho_{\mathbf{X}_{[t_1+1:T]}}$. Thus, Objective (4) is estimated by

$$\widehat{\rho}_{\mathbf{X}}[Y(\boldsymbol{\Theta})] = C_0(\mathbf{S}_0, \mathbf{u}_0) + \widehat{\rho}_{\mathbf{X}_{[1:t_1]}} \left[\sum_{t=1}^{t_1} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) + \widehat{\rho}_{\mathbf{X}_{[t_1+1:T]}} \left(\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t) \right) \right]$$

where $\widehat{\rho}_{\mathbf{X}_{[1:t_1]}}$ and $\widehat{\rho}_{\mathbf{X}_{[t_1+1:T]}}$ denote the estimates for the risk measures $\rho_{\mathbf{X}_{[1:t_1]}}$ and $\rho_{\mathbf{X}_{[t_1+1:T]}}$. To get a precise estimation over $\rho_{\mathbf{X}}[Y(\boldsymbol{\Theta})]$, it requires thousands of replications for N and $(n_1, n_2, \ldots, n_N)$. Since each simulation run could be computationally expensive, it is important that we can make efficient use of the simulation budget to assess the risk-adjusted cost in (4), which is critical for finding a good policy for

real-time decision making. Notice that since the simulation budget is related to the constraint on the computational time, "a finite simulation budget" typically means that we need to make decisions in a certain amount of time.

3 A SIMULATION-BASED PREDICTION FRAMEWORK

In this section, given a finite simulation budget, we propose a simulation-based prediction framework to efficiently estimate the risk-adjusted cost $\rho_{\mathbf{X}}[Y(\mathbf{\Theta})]$ in (4), and further deliver a CrI quantifying the simulation estimation uncertainty. Suppose $t_1 \leq T/2$. Notice that the computational cost for estimating $\rho_{\mathbf{X}}[Y(\mathbf{\Theta})]$ mainly comes from estimating the future response at the second decision-making point t_1, denoted by $\mu\left(\mathbf{S}_{t_1}^{(i)}, g_{t_1}^{(i)}(\boldsymbol{\theta})\right) \equiv \rho_{\mathbf{X}_{[t_1+1:T]}}\left[\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t)\right]$, for all realizations $\mathbf{X}_{[1:t_1]}^{(i)}$ with $i = 1, 2, \ldots, N$. Given a decision policy, this risk response only depends on the updated physical state $\mathbf{S}_{t_1}$ and knowledge state or our belief on the input model characterized by the posterior distribution $g_{t_1}(\boldsymbol{\theta})$. Suppose $g_{t_1}(\boldsymbol{\theta})$ can be specified by a finite number of parameters. Motivated by our previous study (Xie, Nelson, and Barton 2014), we introduce *a metamodel-assisted approach* that can efficiently use the simulation budget to reduce the simulation estimation uncertainty. It also provides a posterior distribution of the response surface for $\mu(\mathbf{S}_{t_1}, g_{t_1}(\boldsymbol{\theta}))$ so that we can quantify the simulation estimation uncertainty of the risk-adjusted cost $\rho_{\mathbf{X}}[Y(\mathbf{\Theta})]$.

In this paper, we consider the parametric input model with a conjugate prior; see Section 3.1. Given the historical data $\mathbf{x}_{[0]}$, the input uncertainty can be characterized by a conjugate posterior $g_0(\boldsymbol{\theta})$ specified by parameters, denoted by $\boldsymbol{\beta}_0$. As time evolves to time period t_1, we observe the new data $\mathbf{X}_{[1:t_1]}^{(i)}$ with $i = 1, 2, \ldots, N$. By applying Bayes' rule, the posterior is updated to be $g_{t_1}^{(i)}(\boldsymbol{\theta})$ specified by parameters, denoted by $\boldsymbol{\beta}_{t_1}^{(i)}$. Thus, the uncertainty of updated prediction model $f(\mathbf{x}|\mathbf{x}_{[0]}, \mathbf{X}_{[1:t_1]})$ could be quantified by the samples $\{\boldsymbol{\beta}_{t_1}^{(1)}, \boldsymbol{\beta}_{t_1}^{(2)}, \ldots, \boldsymbol{\beta}_{t_1}^{(N)}\}$. By abusing the notation, we can rewrite the response surface as

$$\mu(\mathbf{S}_{t_1}, \boldsymbol{\beta}_{t_1}) = \rho_{\mathbf{X}_{[t_1+1:T]}}\left[\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t)\right]. \tag{5}$$

To efficiently estimate the response $\mu(\mathbf{S}_{t_1}, \boldsymbol{\beta}_{t_1})$, based on the simulation outputs at a few well-chosen design points, we construct a GP metamodel $M(\cdot)$ characterizing our belief for $\mu(\cdot)$ in Section 3.2. Then, we develop a procedure in Section 3.3 that can deliver a CrI quantifying the simulation estimation uncertainty of $\rho_{\mathbf{X}}[Y(\mathbf{\Theta})]$. Since the evaluation of $\mu(\cdot)$ at each sample $\left(\mathbf{S}_{t_1}^{(i)}, \boldsymbol{\beta}_{t_1}^{(i)}\right)$ through the metamodel is computationally cheap, we can let the number of replications N be large enough so that the finite sampling error introduced by using finite N samples to estimate the outer risk measure $\rho_{\mathbf{X}_{[1:t_1]}}$ is negligible.

3.1 Update the Input Model

When a conjugate prior is used for the parametric input model, the posterior distribution is in the same family as the prior and it is specified by a finite number of parameters. Since many input parametric distributions commonly used in simulation belong to the exponential family, for illustration, we consider the exponential family with a general form, $p(x|\boldsymbol{\theta}) = \eta_1(x)\eta_2(\boldsymbol{\theta})\exp\{\boldsymbol{\theta}^\top T(x)\}$, where $T(x)$ denotes the sufficient statistics. If we start with a conjugate prior, at $t_0 = 0$, our belief on the input model is characterized by the posterior distribution, denoted by $g_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\boldsymbol{\psi}_0, v_0) = f(\boldsymbol{\psi}_0, v_0)\eta_2(\boldsymbol{\theta})^{v_0}\exp(\boldsymbol{\theta}^\top \boldsymbol{\psi}_0)$. It can be specified by the parameters $\boldsymbol{\beta}_0 = (\boldsymbol{\psi}_0, v_0)$ which are functions of prior parameters and the historical data $\mathbf{x}_{[0]}$.

When time evolves to time period t_1, we observe the new realizations $\mathbf{X}_{[1:t_1]}^{(i)}$. The distribution $g_0(\boldsymbol{\theta})$ gives the prior belief of the input model. Then, by applying Bayes' rule, we can have the posterior distribution

$g_{t_1}^{(i)}(\boldsymbol{\theta}) = p\big(\boldsymbol{\theta}|\boldsymbol{\psi}_{t_1}^{(i)}, v_{t_1}^{(i)}\big)$ with $\big(\boldsymbol{\psi}_{t_1}^{(i)}, v_{t_1}^{(i)}\big) = \big(\boldsymbol{\psi}_0 + T(\mathbf{X}_{[1:t_1]}^{(i)}), v_0 + t_1\big)$. Thus, our updated belief on the input model can be specified by parameters $\boldsymbol{\beta}_{t_1}^{(i)} = \big(\boldsymbol{\psi}_{t_1}^{(i)}, v_{t_1}^{(i)}\big)$.

3.2 Construct a GP Metamodeling

In this section, we construct a metamodel for the response surface $\mu(\cdot)$ in Equation (5). Denote the state variables by $\mathbf{z} \equiv (\mathbf{S}_{t_1}, \boldsymbol{\beta}_{t_1})$ with dimension d. In general, we do not have any strong prior information on the parametric form for $\mu(\cdot)$ and only have weak prior knowledge, such as the continuity. That means as the state variables $\mathbf{z}$ and $\mathbf{z}'$ are close to each other, the responses $\mu(\mathbf{z})$ and $\mu(\mathbf{z}')$ tend to be similar. Thus, we construct a spatial dependence metamodel for $\mu(\cdot)$ by using the stochastic kriging (SK) proposed by Ankenman, Nelson, and Staum (2010), which explicitly considers the simulation estimation uncertainty.

SK assumes that the unknown response surface $\mu(\cdot)$ is a realization of GP. At any $\mathbf{z}$, the simulation output Y from the j-th replication is modeled as

$$Y_j(\mathbf{z}) = \mu_0 + W(\mathbf{z}) + \varepsilon_j(\mathbf{z})$$

where μ_0 is a constant parameter, $\varepsilon_j(\mathbf{z})$ represents the simulation estimation error and the mean-zero, second-order stationary GP $W(\mathbf{z})$ accounts for the spatial dependence of the response surface. The spatial covariance is denoted by $\mathrm{Cov}[W(\mathbf{z}), W(\mathbf{z}')] = \tau^2 r(\mathbf{z} - \mathbf{z}')$, where τ^2 denotes the variance and $r(\cdot)$ denotes a correlation function. Based on the study in Xie, Nelson, and Staum (2010), the product-form Gaussian correlation function is used in our empirical study $r(\mathbf{z} - \mathbf{z}'|\boldsymbol{\xi}) = \exp\big[-\sum_{j=1}^{d} \xi_j (z_j - z_j')^2\big]$.

Our prior belief of the response surface $\mu(\cdot)$ is characterized by a GP, $M(\mathbf{z}) \equiv \mu_0 + W(\mathbf{z})$. To reduce the uncertainty of our belief on $\mu(\cdot)$, we run simulations at k well-chosen design points, denoted by $\mathscr{D} \equiv \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_k\}$, and denote the simulation outputs by $\mathbf{Y}_{\mathscr{D}} = (Y_1(\mathbf{z}_i), Y_2(\mathbf{z}_i), \ldots, Y_{m_i}(\mathbf{z}_i))_{i=1}^{k}$, where m_i denotes the replications at $\mathbf{z}_i$. Let the sample means at all design points be $\bar{\mathbf{Y}}_{\mathscr{D}} = (\bar{Y}(\mathbf{z}_1), \bar{Y}(\mathbf{z}_2), \ldots, \bar{Y}(\mathbf{z}_k))^\top$, where $\bar{Y}(\mathbf{z}_i) = \sum_{j=1}^{m_i} Y(\mathbf{z}_i)/m_i$ for $i = 1, 2, \ldots, k$. Without applying common random numbers, the variance of $\bar{\mathbf{Y}}_{\mathscr{D}}$ is a $k \times k$ diagonal matrix $C = \mathrm{diag}\{\sigma_\varepsilon^2(\mathbf{z}_1)/m_1, \sigma_\varepsilon^2(\mathbf{z}_2)/m_2, \ldots, \sigma_\varepsilon^2(\mathbf{z}_k)/m_k\}$, where $\sigma_\varepsilon^2(\mathbf{z}) = \mathrm{Var}[\varepsilon(\mathbf{z})]$.

Let a $N \times d$ matrix $\mathbf{Z}_*$ denote the locations of N prediction points. Let Σ represent the $k \times k$ covariance matrix across all design points and $\Sigma(\mathbf{Z}_*, \cdot)$ be the $k \times N$ covariance matrix between all design points and N prediction points $\mathbf{Z}_*$. If the parameters $(\tau^2, \boldsymbol{\xi}, C)$ are known, given simulation outputs at design points $\mathbf{Y}_{\mathscr{D}}$, the metamodel uncertainty can be characterized by an updated GP

$$M_p(\mathbf{Z}_*) \equiv M(\mathbf{Z}_*)|\mathbf{Y}_{\mathscr{D}} \sim \mathrm{GP}(m_p(\mathbf{Z}_*), \Sigma_p(\mathbf{Z}_*, \mathbf{Z}_*)) \tag{6}$$

where $m_p(\cdot)$ is the minimum mean squared error (MSE) linear predictor

$$m_p(\mathbf{Z}_*) = \widehat{\mu}_0 \cdot \mathbf{1}_{N\times 1} + \Sigma(\mathbf{Z}_*, \cdot)^\top [\Sigma + C]^{-1} (\bar{\mathbf{Y}}_{\mathscr{D}} - \widehat{\mu}_0 \cdot \mathbf{1}_{k\times 1}), \tag{7}$$

and the corresponding variance covariance matrix is

$$\begin{aligned} \Sigma_p(\mathbf{Z}_*, \mathbf{Z}_*) \;=\; & \Sigma(\mathbf{Z}_*, \mathbf{Z}_*) - \Sigma(\mathbf{Z}_*, \cdot)^\top [\Sigma + C]^{-1} \Sigma(\mathbf{Z}_*, \cdot) \\ & + k\eta^\top [\mathbf{1}_{k\times 1}^\top (\Sigma + C)^{-1} \mathbf{1}_{k\times 1} \mathbf{1}_{k\times 1}^\top (\Sigma + C)^{-1} \mathbf{1}_{k\times 1}]^{-1} \eta \end{aligned} \tag{8}$$

where $\widehat{\mu}_0 = [\mathbf{1}_{k\times 1}^\top (\Sigma + C)^{-1} \mathbf{1}_{k\times 1}]^{-1} \mathbf{1}_{k\times 1}^\top (\Sigma + C)^{-1} \bar{\mathbf{Y}}_{\mathscr{D}}$ and $\eta = \mathbf{1}_{1\times N} - \mathbf{1}_{k\times 1}^\top (\Sigma + C)^{-1} \Sigma(\mathbf{Z}_*, \cdot)$ (Ankenman, Nelson, and Staum 2010). Since in reality the spatial correlation parameters τ^2 and $\boldsymbol{\xi}$ are unknown, the maximum likelihood estimators are typically used for prediction.

Then, we consider the system performance estimate for $\rho_{\mathbf{X}_{[t_1+1:T]}}\big[\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t, \mathbf{u}_t)\big]$ with $\rho_{\mathbf{X}} = \mathrm{CVaR}_\alpha$ and its estimation variance at each design point. Specifically, at the design point $\mathbf{z}_i$, there are n_i simulation runs, denoted by $\tilde{L}_{i\ell} \equiv \sum_{t=t_1+1}^{T} r^t C_t\big(\mathbf{S}_t^{(i\ell)}, \mathbf{u}_t^{(i\ell)}\big)$ with $\ell = 1, 2, \ldots, n_i$. We can divide n_i runs into m_i batches with $N_i = n_i/m_i$ being an integer. Based on samples in each batch, we can have a CVaR_α estimate or

the simulation output $Y_j(\mathbf{z}_i)$ with $j = 1,2,\ldots,m_i$. After that, we can calculate the sample mean $\bar{Y}(\mathbf{z}_i)$ and also estimate its variance. However, based on the study in Chen and Kim (2014), we can minimize the mean squared error (MSE) of CVaR_α estimation by letting $N_i = n_i$. Thus, in this paper, we use all samples $\{\tilde{L}_{i\ell}\}_{\ell=1}^{n_i}$ to estimate the response $\text{CVaR}_\alpha\left[\sum_{t=t_1+1}^{T} r^t C_t(\mathbf{S}_t,\mathbf{u}_t)\right]$, and the simulation output becomes

$$Y(\mathbf{z}_i) = \tilde{v}_i - \frac{1}{n_i\alpha}\sum_{\ell=1}^{n_i}[\tilde{v}_i - \tilde{L}_{i\ell}]^+ \tag{9}$$

where $\tilde{v}_i = \inf\{x : \tilde{F}_i(x) \geq \alpha\}$ and $\tilde{F}_i(x) = \frac{1}{n_i}\sum_{\ell=1}^{n_i} \mathrm{I}\{\tilde{L}_{i\ell} \leq x\}$ with $\mathrm{I}(\cdot)$ denoting an indicator function. The variance of $Y(\mathbf{z}_i)$ could be estimated by the bootstrap approach (Cheung and Lee 2005). Accounting for the bias introduced by using finite n_i to estimate the risk measure $\rho_{\mathbf{X}_{[t_1+1:T]}}$ is in our future research.

3.3 Procedure for the Metamodel-Assisted Approach

In this section, we describe the procedure for our metamodel-assisted approach and it includes the main steps as follows. We first find a design space E that can cover the most likely prediction points for $(\mathbf{S}_{t_1}^{(i)}, \boldsymbol{\beta}_{t_1}^{(i)})$. To have all prediction points located close to design points, we use a space-filling design, the orthogonal max-min Latin Hypercube Design (LHD), to generate k design points evenly covering E (Liu and Staum 2010; Barton, Nelson, and Xie 2010). Then, run the simulations at design points and construct a GP $M_p(\cdot)$ for the risk response surface $\mu(\mathbf{S}_{t_1}, \boldsymbol{\beta}_{t_1})$. After that, we use $m_p(\cdot)$ to get a point estimate (PE) for $\rho_X[Y(\boldsymbol{\Theta})]$ in Step 3, and further use the posterior distribution of $\mu(\cdot)$ to generate a $(1-\alpha_0)100\%$ percentile CrI quantifying the simulation estimation uncertainty in Step 4.

0. Provide the current physical and knowledge states $(\mathbf{S}_0, \boldsymbol{\beta}_0)$.
1. Identify a design space E of state variables $\mathbf{z}$ at the second decision-making point, which is the smallest cube covering a large percentage of prediction points $\mathscr{P} = \{\mathbf{z}_1^*, \mathbf{z}_2^*, \ldots, \mathbf{z}_N^*\}$, where $\mathbf{z}_i^* = (\mathbf{S}_{t_1}^{(i)}, \boldsymbol{\beta}_{t_1}^{(i)})$.
2. Use the orthogonal max-min LHD to generate k design points $\mathscr{D} = \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_k\}$ evenly covering the design region E. At each $\mathbf{z}_i$, run n_i simulation runs to obtain $\tilde{L}_{i\ell}$ with $i = 1,2,\ldots,k$ and $\ell = 1,2,\ldots,n_i$. Then, obtain the simulation output $Y(\mathbf{z}_i)$ by Equation (9) and estimate the variance $\sigma_\varepsilon^2(\mathbf{z}_i)$ by the bootstrapping. Fit the metamodel by using Equations (7) and (8).
3. Use $m_p(\cdot)$ to estimate the responses at the prediction points $\mathscr{P}$. Then, calculate the PE for risk-adjusted cost, $\widehat{\rho}_{\mathbf{X}}[Y(\boldsymbol{\Theta})] = C_0(\mathbf{S}_0,\mathbf{u}_0) + \widehat{v} - \frac{1}{N\alpha}\sum_{i=1}^{N}[\widehat{v} - \widehat{L}_i]^+$, where $\widehat{L}_i = \sum_{t=1}^{t_1} r^t C_t(\mathbf{S}_t^{(i)}, \mathbf{u}_t^{(i)}) + m_p(\mathbf{z}_i^*)$ is generated by simulations driven by $\mathbf{X}_{[1:t_1]}^{(i)}$ with $i = 1,2,\ldots,N$, $\widehat{v} = \inf\{x : \widehat{F}(x) \geq \alpha\}$ and $\widehat{F}(x) = \frac{1}{N}\sum_{i=1}^{N}\mathrm{I}\left\{\widehat{L}_i \leq x\right\}$. The costs $C_t(\mathbf{S}_t,\mathbf{u}_t)$ with $t = 0,1,\ldots,t_1$ are obtained by simulation.
4. Return a CrI quantifying the simulation estimation uncertainty for $\rho_X[Y(\boldsymbol{\Theta})]$. Generate B sample paths of $M_p(\cdot)$ by using Equations (6) – (8). For each sample path, estimate the responses $\mu(\cdot)$ at the prediction points $\mathscr{P}$, and then by following the similar procedure in Step 3, obtain an estimate of the risk-adjusted cost objective, denoted by U_j with $j = 1,2,\ldots,B$. Record the $(1-\alpha_0)100\%$ percentile CrI, $\left[U_{(\lceil B\frac{\alpha_0}{2}\rceil)}, U_{(\lceil B(1-\frac{\alpha_0}{2})\rceil)}\right]$, where $U_{(1)} \leq U_{(2)} \leq \cdots \leq U_{(B)}$ are the sorted values.

4 EMPIRICAL STUDY

In this section, a single-item multi-period inventory example with a daily review (R,Q) policy is used to study the performance of our approach. Thus, the time unit is in terms of days. Suppose the supply lead time is $L = 2$. At the beginning of time period t, we first receive the order u_{t-L} and then observe the demand X_t. If the stock level $S_{t-1} + u_{t-L}$ is greater than the demand, we have the holding cost $h(S_{t-1} + u_{t-L} - X_t)$, where S_{t-1} denotes the inventory level in the last time period and the holding cost per unit is $h = 1.5$. If X_t exceeds the stock level, we subcontract the remaining demand to the third party

at price $p = 5$ per unit. After that, we update the inventory level $S_t = \max\{S_{t-1} + u_{t-L} - X_t, 0\}$. If the inventory position (equal to the sum of stock on hand and outstanding orders) drops to or below the reorder point R, we place an order with $u_t = Q$; Otherwise, set $u_t = 0$. Thus, the total cost occurs in time period t is $C_t(S_t, u_t) = cu_t + hS_t + p\max\{X_t - S_{t-1} - u_{t-L}, 0\}$, where the ordering cost per unit is $c = 1.2$. The demands in each time period follow an exponential distribution, $X_t \overset{i.i.d.}{\sim} \exp(\theta^c)$ with rate $\theta^c = 1/300$, and the system starts with an initial inventory 100. Since high-tech products tend to have a short life time, for simplification, we set the discount factor $r = 1$. For any given policy specified by (R, Q), we want to quickly and efficiently estimate the risk-adjusted cost $\rho_X[Y(\Theta)]$ in Equation (4) with $\rho = \text{CVaR}_{0.95}$. For illustration, suppose $R = 100$ and $Q = 200$.

To assess the performance of our approach, suppose the input distribution family is known with the unknown parameter θ^c estimated from real-world data. We start with a noninformative conjugate prior Gamma$(0, 0.5)$. At the current time period t_0, the unknown parameter is estimated by m valid historical data, denoted by $\mathbf{x}_{[t_0]} = (x_{t_0-m+1}, x_{t_0-m+2}, \ldots, x_{t_0})$. Thus, our belief on θ^c is characterized by the posterior distribution $p(\theta|\mathbf{x}_{[t_0]})$ and it is Gamma(a, b) with $(a, b) = (m, 0.5 + \sum_{t=t_0-m+1}^{t_0} x_t)$. For notation simplification, set $t_0 = 0$. For the generalized two-stage decision model in Equation (4), at the second decision-making point t_1, the knowledge state is characterized by the updated posterior distribution, Gamma$(a + t_1, b + \sum_{t=1}^{t_1} X_t)$, which is specified by $\beta_{t_1} = \sum_{t=1}^{t_1} X_t$.

To examine the robustness of our approach and study the effects of t_1 and T on the objective $\rho_X[Y(\Theta)]$, we consider different settings with $m = 10, 20, 100$, $T = 10, 20$ and $t_1 = 3, 5$. Since there is no analytical form for $\rho_X[Y(\Theta)]$ in Equation (4), we do a side experiment to estimate the true value of risk-adjusted cost by using the Monte Carlo approach with a large enough simulation budget. We run N replications to estimate the outer risk measure. Since we do not have any prior knowledge about $\mu(S_{t_1}, \beta_{t_1}) = \rho_{X_{[t_1+1:T]}}[\sum_{t=t_1+1}^{T} r^t C_t(S_t, u_t)]$, we set equal replications to each sample or prediction point $(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ with $i = 1, 2, \ldots, N$. That means $n_1 = n_2 \cdots = n_N = n$.

To determine the appropriate replications N and n to estimate the true risk-adjusted cost $\rho_X[Y(\Theta)]$, we select $N = 10^4$ and $n = 10^4$ as the benchmark, and then run 10 macro-replications to test different settings. Since there is larger input uncertainty when $m = 10$, we only consider $m = 10$ in the side experiment. In each macro-replication, we first generate m historical observations, run simulations to estimate $\rho_X[Y(\Theta)]$ and record the relative error. Then, the results of maximum error in the unit % obtained from 10 macro-replications are provided in Table 1. We observe that the relative error is more sensitive to the value of N than n. Since the setting $N = 10^4$ and $n = 10^3$ can achieve a relative error less than 0.3%, to balance precision and computational cost, we choose it to estimate the true cost objective value.

Table 1: The maximum absolute relative difference compared to the estimate obtained with $N = 10^4$ and $n = 10^4$ (in the unit %).

N	1000			5000			10000	
n	1000	5000	10000	1000	5000	10000	1000	5000
$t_1 = 3, T = 10$	1.5	1.6	1.6	1.3	1.4	1.4	0.19	0.08
$t_1 = 5, T = 10$	1.9	1.9	1.9	1.1	1.1	1.0	0.17	0.07
$t_1 = 3, T = 20$	1.4	1.2	1.1	0.9	0.9	0.8	0.29	0.07
$t_1 = 5, T = 20$	1.4	1.2	1.2	0.9	0.9	0.8	0.29	0.08

4.1 Comparison of Metamodel-Assisted and SAA Approaches

For a given simulation budget, we compare our metamodel-assisted approach with the direct SAA approach. We first study the statistical properties of the PE of $\rho_X[Y(\Theta)]$ obtained by both approaches. Table 2 records the mean and standard error (SE) of the relative error of PEs in the unit %. They are based on 100

macro-replications. In each macro-replication, we first generate m historical data from F^c. Then, the metamodel-assisted and direct SAA approaches are used to estimate the risk-adjusted cost. We record the relative errors for both approaches. For the metamodel-assisted approach, when we construct the metamodel for $\mu(S_{t_1}, \beta_{t_1})$, we choose $k = 20$ design points by following the '$10d$' rule recommended in Jones, Schonlau, and Welch (1998). At each design point, we assign 100 replications to estimate the inner risk measure $\mu(S_{t_1}, u_{t_1}) = \rho_{X_{[t_1+1:T]}}[\sum_{t=t_1+1}^{T} r^t C_t(S_t, u_t)]$. Based on the simulation results at k design points, we can construct the fitted response surface $m_p(S_{t_1}, \beta_{t_1})$ by Equation (7) and use it to estimate the risk-adjusted cost objective. Since it is computationally cheap to assess the inner level risk measure by using the metamodel, we set $N = 1000$ when we assess $\rho_X[Y(\Theta)]$. Considering that the computational cost mainly comes from assessing the inner risk measure, for simplification, we assign an equal simulation budget to estimate the inner level risk measure for both metamodel-assisted and SAA approaches. Thus, for the SAA approach, since the relative error is more sensitive to the value of N than n, we set $N = 100$ and $n = 20$ so that we can have one observation falling in the 5% right tail part of $\sum_{t=t_1+1}^{T} r^t C_t(S_t, u_t)$. Table 2 demonstrates that the metamodel-assisted approach provides much smaller estimation error and it can efficiently use the simulation budget to estimate the risk-adjusted cost $\rho_X[Y(\Theta)]$. Further, we observe that as t_1 increases, the estimation errors from both approaches tend to decrease.

Table 2: The relative error of PEs and the relative width of CrIs obtained by the metamodel-assisted and SAA approaches (in the unit %).

$t_1 = 3, T = 10$	metamodel-assisted				SAA			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	5.5	0.4	11.3	7.6	13.0	1.0	30.1	8.9
$m = 20$	4.9	0.3	7.2	3.6	9.0	0.6	26.5	7.1
$m = 100$	3.0	0.2	3.9	1.5	7.3	0.5	21.1	1.5
$t_1 = 5, T = 10$	metamodel-assisted				SAA			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	4.8	0.4	7.3	5.5	10.0	0.7	25.6	7.3
$m = 20$	4.1	0.3	5.6	3.2	7.0	0.5	21.6	5.5
$m = 100$	3.2	0.2	3.0	1.0	5.2	0.4	17.5	4.5

Then, we study the statistical properties of the 95% percentile CrIs obtained by using the metamodel-assisted and SAA approaches to quantify the simulation estimation uncertainty. Let $B = 1000$. Table 2 records the mean and standard deviation (SD) of the CrI width divided by the true risk-adjusted cost $\rho_X[Y(\Theta)]$. The results are based on 100 macro-replications. In each macro-replication, we first generate m real-world data from F^c. For the metamodel-assisted approach, let $N = 1000$ and we build the metamodel for $\mu(S_{t_1}, \beta_{t_1})$ by using $k = 20$ and $n = 100$. Then, we can construct the CrI by following the procedure described in Section 3.3. For the SAA approach, we first generate $N = 100$ prediction points, $(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ with $i = 1, 2, \ldots, N$. At each sample $(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$, we estimate the response $\mu(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ by using $n = 20$ replications, denoted by $\widehat{\mu}(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$, and then obtain the estimation variance by the bootstrap method (Cheung and Lee 2005), denoted by $\widehat{\sigma}^2(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$. Without using common random numbers, the simulation estimation uncertainty at different prediction points is independent. Following the similar idea with that in Xie, Li, and Zhang (2016), we characterize the simulation estimation uncertainty of the response $\mu(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ with the posterior distribution $\mathcal{N}(\widehat{\mu}(S_{t_1}^{(i)}, \beta_{t_1}^{(i)}), \widehat{\sigma}^2(S_{t_1}^{(i)}, \beta_{t_1}^{(i)}))$, which holds asymptotically. Then, we draw a multivariate sample at N prediction points from $\mathcal{N}(\widehat{\boldsymbol{\mu}}, \widehat{\Sigma})$, where the $N \times 1$ vector $\widehat{\boldsymbol{\mu}}$ with element

$\widehat{\mu}_i = \widehat{\mu}(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ and the $N \times N$ diagonal matrix $\widehat{\Sigma}$ with $\widehat{\Sigma}_{ii} = \widehat{\sigma}^2(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ for $i = 1, 2, \ldots, N$. After that, we can obtain an estimate of the risk-adjusted cost $\rho_X[Y(\Theta)]$, denoted by $\tilde{U}_b$. By repeating this procedure B times, we have $\{\tilde{U}_1, \tilde{U}_2, \ldots, \tilde{U}_B\}$. Record the 95% percentile CrI, $\left[\tilde{U}_{(\lceil B\frac{\alpha_0}{2} \rceil)}, \tilde{U}_{(\lceil B(1-\frac{\alpha_0}{2}) \rceil)}\right]$, where $\tilde{U}_{(1)} \leq \tilde{U}_{(2)} \leq \cdots \leq \tilde{U}_{(B)}$ are the sorted values. Notice that these CrIs only account for the simulation uncertainty on estimating $\mu(S_{t_1}, \beta_{t_1})$. They do not consider the finite sampling uncertainty introduced by using finite N to estimate the outer risk measure, which could be large for the SAA approach. Table 2 indicates that the metamodel-assisted approach can efficiently use the simulation budget to reduce the simulation estimation uncertainty.

4.2 Study the Effects of N, t_1 and T

In this section, we use the metamodel-assisted approach to investigate the effects of the number of prediction points N, the second decision point t_1, and the length of planning horizon T on the estimation of risk-adjusted cost $\rho_X[Y(\Theta)]$. We first study the impact of sample size N. As N increases, the finite sampling error for estimating the outer risk measure tends to decrease. Since the SAA approach needs to run simulations at each prediction point $(S_{t_1}^{(i)}, \beta_{t_1}^{(i)})$ with $i = 1, 2, \ldots, N$ to estimate the corresponding the inner risk measure, the computational cost increases dramatically as N becomes large. Differing with SAA, the metamodel-assisted approach separates the estimation of inner and outer risk measures and it allows us to make predictions for a large number of prediction points. Thus, we use the metamodel-assisted approach to study the impacts of the choice of N on the estimation of outer risk measure. Table 3 provides the results on the relative error of PE and the relative CrI width of the risk-adjusted cost when $N = 1000, 2000$. They indicate that $N = 2000$ gives narrower CrI width, while the relative errors of PEs obtained by $N = 1000, 2000$ are close to each other. That means the finite sampling uncertainty is not ignorable when $N = 1000$. Notice that as the required N increases, the advantage of the metamodel-assisted approach becomes more obvious.

In addition, Table 3 indicates that as t_1 increases, both the relative estimation errors of PE and the relative CrI width decrease. By the asymptotic consistency, as t_1 increases, the posterior $g_{t_1}(\theta)$ becomes more concentrated around θ^c and the design space is smaller. Given the same simulation budget, we can estimate the inner risk measure better.

Table 3: The relative error of PEs and the relative width of CrIs obtained by the metamodel-assisted approach when $N = 1000, 2000$ (in the unit %).

$t_1 = 3, T = 10$	$N = 1000$				$N = 2000$			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	5.5	0.4	11.3	7.6	5.6	0.4	8.1	6.0
$m = 20$	4.9	0.3	7.2	3.6	4.0	0.4	5.9	3.4
$m = 100$	3.0	0.2	3.9	1.5	3.0	0.2	3.4	1.4
$t_1 = 5, T = 10$	$N = 1000$				$N = 2000$			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	4.8	0.4	7.3	5.5	5.2	0.4	6.1	4.9
$m = 20$	4.1	0.3	5.6	3.2	3.8	0.3	5.2	2.9
$m = 100$	3.2	0.2	3.0	1.0	3.0	0.2	2.6	1.0

Then, we study the impact of the planning horizon T. While it is computationally cheaper to estimate the risk-adjusted cost when we use a smaller T value, the larger T gives better approximation to the infinite planning horizon. We study cases with $T = 10, 20$ and the results are shown in Tables 4. The relative

errors of PEs and relative CrI width do not change significantly. Since the value of $\rho_X[Y(\Theta)]$ increases as T increases, the absolute error and the CrI width become larger.

Table 4: The relative error of PEs and the relative width of CrIs obtained by the metamodel-assisted approach when $T = 10, 20$ (in the unit %).

$t_1 = 3$	$T = 10$				$T = 20$			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	5.5	0.4	11.3	7.6	5.2	0.4	10.9	8.0
$m = 20$	4.9	0.3	7.2	3.6	4.3	0.3	8.9	5.6
$m = 100$	3.0	0.2	3.9	1.5	2.5	0.2	4.2	2.0
$t_1 = 5$	$T = 10$				$T = 20$			
	error of PE		CrI width		error of PE		CrI width	
	mean	SE	mean	SD	mean	SE	mean	SD
$m = 10$	4.8	0.4	7.3	5.5	5.4	0.4	7.1	6.6
$m = 20$	4.1	0.3	5.6	3.2	3.7	0.4	6.2	4.9
$m = 100$	3.2	0.2	3.0	1.0	2.5	0.2	3.8	1.6

5 CONCLUSIONS

There exist various challenges when we make real-time decisions for supply chain risk management in the bio-pharmaceutical manufacturing industry. In this paper, a generalized two-stage decision model is studied and we propose a simulation-based prediction framework. Since the risk-adjusted cost objective in the decision model involves the nested risk measures, it could be computationally prohibitive to assess the system performance by using the classical SAA approach. Considering that the computational cost mainly comes from estimating the inner risk measure, a metamodel-assisted approach is introduced. It can efficiently use a finite simulation budget to estimate the risk-adjusted cost objective and further deliver a CrI quantifying the simulation estimation uncertainty, which is critically important to find good real-time decisions. An empirical study on an inventory example demonstrates clear advantages of the metamodel-assisted approach compared to the SAA approach. Then, our approach is further used to study the effects of the settings of two-stage decision model on the cost objective.

REFERENCES

Ankenman, B. E., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58:371–382.

Barton, R., B. L. Nelson, and W. Xie. 2010. "A Framework for Input Uncertainty Analysis". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yucesan, 1189 – 1198. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Beale, E. 1955. "On Minimizing A Convex Function Subject to Linear Inequalities". *Journal of the Royal Statistical Society Series B* 17:173–184.

Chen, F., M. Xu, and Z. G. Zhang. 2009. "A Risk-Averse Newsvendor Model Under the CVaR Criterion". *Operations Research* 57 (4): 1040–1044.

Chen, X., and K. Kim. 2014. "Stochastic Kriging with Biased Sample Estimates". *ACM Transactions on Modeling and Computer Simulation* 24:8/1–8/23.

Chen, X., M. Sim, D. Simchi-Levi, and P. Sun. 2007. "Risk Aversion in Inventory Management". *Operations Research* 55:828 – 842.

Cheung, K. Y., and S. M. S. Lee. 2005. "Variance Estimation for Sample Quantiles Using the *m* Out of *n* Bootstrap". *Annals of the Institute of Statistical Mathematics* 57 (2): 279–290.

Delage, E., and Y. Ye. 2010. "Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems". *Operations Research* 58 (3): 595–612.

Dentzig, G. 1955. "Linear Programming under Uncertainty". *Management Science* 1:197–206.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 2004. *Bayesian Data Analysis*. 2nd ed. New York: Taylor and Francis Group, LLC.

Gotoh, J., and Y. Takano. 2007. "Newsvendor Solutions via Conditional Value-At-Risk Minimization". *European Journal of Operational Research* 179:80–96.

Hu, J., W. Xie, and B. Xu. 2016. "Data-Driven Risk-Averse Stochastic Optimization Integrating Input and Systemic Uncertainties". Working Paper.

Jones, D., M. Schonlau, and W. Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions". *Journal of Global Optimization* 13:455–492.

Kaminsky, P., and Y. Wang. 2015. "Analytical Models for Biopharmaceutical Operations and Supply Chain Management: A Survey of Research Literature". *Pharmaceutical Bioprocessing* 3 (1): 61–73.

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

Nelson, B. L. 2013. *Foundations and Methods of Stochastic Simulation: A First Course*. Springer.

Qian, P., Z. Wang, and Z. Wen. 2015. "A Composite Risk Measure Framework for Decision Making under Uncertainty". Submitted.

Scarf, H., K. Arrow, and S. Karlin. 1958. "A Min-Max Solution of an Inventory Problem". *Studies in the Mathematical Theory of Inventory and Production* 10:201–209.

Shapiro, A., D. Dentcheva, and A. Ruszczynski. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia: SIAM.

Shapiro, A., and T. Homem-de Mello. 1998. "A Simulation-Based Approach to Two-Stage Stochastic Programming with Recourse". *Mathematical Programming* 81:301–325.

Xie, W., C. Li, and P. Zhang. 2016. "Statistical Uncertainty Quantification in Stochastic Simulation via A Bayesian Nonparametric Hierarchical Framework". *Under Review*.

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.

Xie, W., B. L. Nelson, and J. Staum. 2010. "The Influence of Correlation Functions on Stochastic Kriging Metamodels". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yucesan, 1067–1078. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Zhou, E., and W. Xie. 2015. "Simulation Optimization when Facing Input Uncertainty". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3714–3724. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

WEI XIE is an assistant professor in the Department of Industrial and Systems Engineering at Rensselaer Polytechnic Institute. She holds a Ph.D. degree in the Industrial Engineering and Management Sciences from Northwestern University. Her research interests are in computer simulation, risk management and data analytics. Her email address is xiew3@rpi.edu and her web page is http://homepages.rpi.edu/~xiew3/.

YUAN YI is a Ph.D. student in the Department of Industrial and Systems Engineering at the Rensselaer Polytechnic Institute. His research interests are in the interplay of simulation, supply chain and and risk management. His email address is yiy2@rpi.edu.