Offline Simulation Online Application: A New Framework of Simulation-Based Decision Making

L. Jeff Hong

School of Data Science and School of Management
Fudan University, Shanghai, China
hong_liu@fudan.edu.cn

Guangxin Jiang*

The Department of Management Science and Engineering, School of Management
Shanghai University, Shanghai, China
gxjiang@shu.edu.cn

July 8, 2019

Abstract

Traditionally, simulation has been used as a tool of design to estimate, compare and optimize the performance of different system designs. It is rarely used in making real-time decisions due to the long computation delay of executing simulation models. However, with the fast growth of computing capability, we have observed more and more works on reusing simulation efforts for repeated experiments with the help of data analytics tools, and the target of these works is to solve real-time decision problems. In this paper, we distill the important features of these works and summarize a new simulation framework, called offline-simulation-online-application (OSOA) framework, which treats simulation as a data generator, applies state-of-the-art analytics tools to build predictive models, and then uses the predictive models for real-time applications. In this paper, we illustrate how to apply the OSOA framework on estimation, ranking and selection and simulation optimization, and provide a prospect of this new framework.

*Corresponding Author

Keywords: Offline-simulation-online-application; simulation analytics; simulation optimization; ranking and selection.

1 Introduction

Simulation is a commonly used tool of design to estimate, compare and optimize the performance of complex systems. For instance, in the area of finance, simulation is used to estimate the prices and sensitivities of complex derivatives and the risk measures of large portfolios (Glasserman 2003 for a survey); in the area of healthcare, simulation is used to compare different medical treatments (Dunnett 1984, Jun and Jacobson 1999); and in the area of transportation, simulation is used to optimize the traffic system designs (Daganzo 1997, Osorio and Bierlaire 2013). These problems are in general offline, that is, the input parameters of the simulation models are known when conducting the simulation experiments and sufficient time is often available to run the experiments before making decisions. As simulation models often take a large amount of time to run, especially when the systems are complex and the solution space is large, traditionally, much of the research attention in the stochastic simulation area is focused on the improvement of the efficiency, for instance, designing variance reduction techniques to speed up estimation (Glasserman et al. 1999, Dingeç and Hörmann 2012, Jiang et al. 2016a), using parallel computing to handle large-scale ranking and selection (R&S) (Luo et al. 2015, Ni et al. 2017), and developing fast search algorithms to conduct simulation optimization (SO) (see Fu 2015 for a survey).

In practice, however, many decisions need to be made in real time (or near real time) based on the parameters observed at the time and the available time is typically not enough to conduct a full-scale simulation study. We call them online problems. For instance, in the area of finance, portfolio risk measures often need to be estimated dynamically based on the current values of the underlying risk factors. As an example, the famous 4:15 report of J. P. Morgan requires consolidating the risks of all trading desks based on the closing values of the underlying factors, available within 15 minutes after the market closes everyday, and the company uses it to decide whether the risk is under control (Jorion 2006). In the area of healthcare, medical treatments sometimes need to be selected almost in real time when the test results of the patients are revealed. The doctors may not have the time or the capability to conduct a simulation study to compare different treatments (Choi et al. 2014). In the area of transportation, self-driving cars need to make the manoeuvre decisions in real time based on the surrounding objects and there is no time to run a simulation study (Santana and Hotz 2016). In these situations, simulation is typically not used due to the long execution time. Instead, simple models based on closed-form approximations or simple rules based on experience are often used. These simple models and simple rules are fast to use but may not lead to the best decisions.

However, with the fast growth of computing capability and with more and more data analytics tools, there are more and more recent works considering how to use complex simulation models and experiments to solve real-time decision problems. In these works, simulation efforts are reused and data analytics tools are applied to build metamodels or to accelerate decision procedures, see Hannah et al. (2010), Nelson (2016), Jiang et al. (2016b), Lin and Nelson (2016), Ouyang and Nelson (2017), Pearce and Branke (2017), Jiang et al. (2019a), Jiang et al. (2019b), Shen et al. (2019), Li et al. (2019), Jin et al. (2019) for instance. In these works, extensive simulation experiments are conducted before the real-time decision problems come, and predictive models or surrogate models are constructed to be used in real-time decision problems. In this paper we summarize these works to a unified framework, which uses the offline simulation results on applications where online decisions need to be made. We call it the *offline-simulation-online-application* (OSOA) framework. The key to linking the offline simulation and online application are *predictive models*.

To understand the importance of predictive models, we first introduce the concept of *covariates*. In many online applications, there are parameters that are only revealed at the spot. The examples include the values of the underlying risk factors in financial risk measurement, the characteristics of the patient in medical treatment selection, and the sizes, locations and speeds of the surrounding objects in self-driving car manoeuvre. These parameters are not exactly known when running simulation experiments, but the possible values (and sometimes ranges) that they may take are typically known. We call these parameters covariates. The

OSOA framework is divided into two stages. *In the offline simulation stage*, we run experiments at different values of the covariates and build predictive models so that the performance of the models may be predicted once the values of the covariates are observed. *In the online application stage*, we use the predictive models in online problems to predict the performance of a decision and to compare and optimize the performance of different decision choices. The predictive models are often closed-form expressions that may be used as analytical formulae of important system performances. For example, the formula of the waiting time for a complex queueing system is usually unavailable, so we can approximate the formula by a predictive model based on a large amount of simulation experiments. Then we can make real-time decisions using the predictive model. In fact, predictive models have been used in the simulation area. In approximate dynamic programming (Powell 2009), for instance, regression models are often used to approximate the value function to overcome the "curse of dimensionality". In this paper, we consider three types of problems, namely "estimation with covariates", "R&S with covariates" and "SO with covariates". In all three types of problems, we emphasize the existence of the covariates to differentiate with the classical offline estimation, R&S and SO problems. Notice that there are some common characteristics for these three types of problems: (i) all these problems contain covariates; (ii) the goals of all these problems are to solve online problems rather than the static design problems; (iii) all these problems need to apply data mining/machine learning tools to simulation data to build predictive models. Therefore, we may unify these problems under the OSOA framework.

In the OSOA framework, the simulation models are essentially data generators. One may use parallel computing environments to generate multiple replications of simulation results on a well-chosen set of values of the covariates. Then, state-of-the-art data mining/analytics tools may be used to construct the predictive models, which may be used for online applications. This view of simulation models as (big) data generators is also in line with the concept of "simulation analytics", recently proposed by Nelson (2016). The main idea of the simulation analytics approach is to treat the simulation model as a generator of multiple replications of system dynamics over time and to apply data analytics tools to mine the data and to estimate conditional statements. Indeed, one may view the OSOA framework as a way of realizing the concept of simulation analytics. Similar to simulation analytics, the OSOA framework gets benefit from the recent development of computer technologies (such as cloud computing and parallel computing) and the recent development of data analytics tools. The new computer technologies allow the users to conduct extensive simulation experiments and to store the data, and the new data analytics tools provide the users a lot of sophisticated predictive models.

Reusing simulation experiments have also been considered by other simulation researchers in recent years. Liu et al. (2010) proposed to run the simulation experiments to construct a good database and use it for future estimations of derivative prices. They call it "simulation on demand". Rosenbaum and Staum (2015) further developed this idea into database Monte Carlo simulation that uses the database to construct control variates to reduce the variances of the estimators. Feng and Staum (2017) proposed the concept of green simulation that uses saved simulation data as a complementary source to new simulation data by employing a change of probability measures. These works are similar to the OSOA framework. However, the difference is that they emphasize on reusing the simulation data on new problems, while our goal is to build predictive models so that we no longer need the simulation data and may solve online problems.

The rest of the paper is organized as follow. We first discuss the OSOA framework in Section 2. In Sections 3 to 5, we consider three types of OSOA problems, estimation with covariates, R&S with covariates and SO with covariates, respectively. We then speculate the future of OSOA research and propose a few problems for future studies in Section 6.

2 The OSOA framework

Recall that in many real-time applications, simple models (often with closed-form expressions) are often used. In finance risk management, for example, closed-form pricing models (e.g., the famous Black-Scholes formula) are typically used despite their known deficiencies. Simulation models are typically more accurate, because they can capture more details. But they are difficult to use in real time due to the long computation time. The OSOA framework fills the gap, and the key is the predictive model (see Figure 1). We use offline simulation to run a large amount of simulation experiments at different combinations of covariate values to build the predictive model that is nearly as accurate as the simulation model itself but has closed-form expressions that may be evaluated in real time. Once the predictive model is built, it acts like the simple models in many real-time applications. It may be used in real time once the values of the covariates are observed. In this way, we combine the benefits of both, the accuracy of simulation models and the speed of simple models. Notice that Figure 1 only shows a general OSOA framework, and there are some other variants. For example, predictive models may be built repeatedly in some problems. In addition, the distinction between offline simulation and online application is blurred in some near real-time problems. We may have time to conduct simulation experiments before making decision, so we can combine the simulation results and the predictive models together to further enhance the decision, and to accumulate the additional data for further use. Overall, the OSOA should be adapted based on the specific problems.

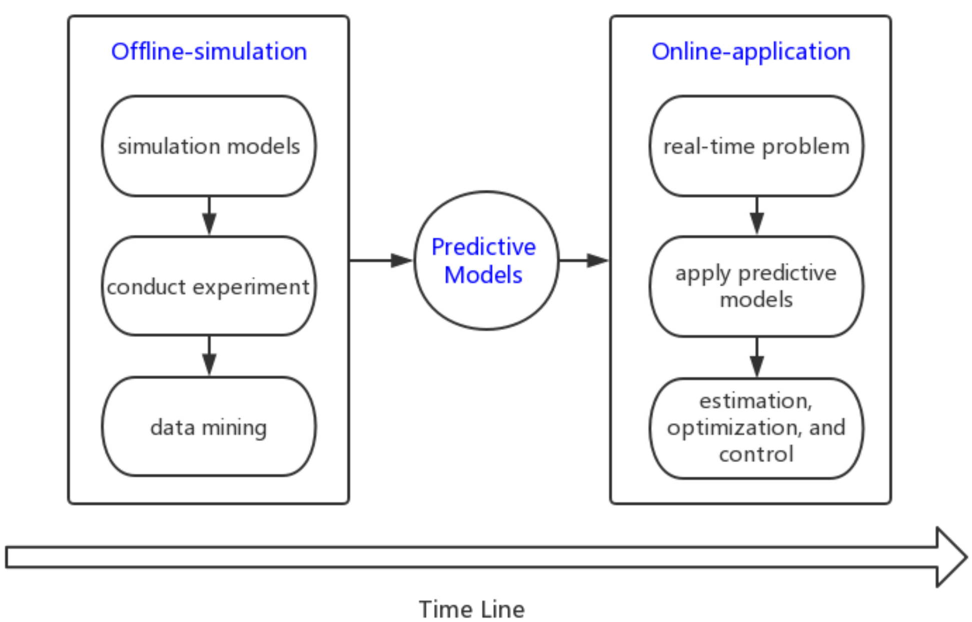

Figure 1: A general OSOA framework

2.1 Offline-Simulation Stage

A simple way to look at the offline-simulation stage is to think it as a surface-fitting stage. The goal of the stage is indeed to fit a surface that may be used to predict once the the covariates are observed. Therefore, many of the state-of-the-art analytics tools may be used (see, for instance, Hastie et al. 2011). The critical issue of surface fitting is typically how to control the bias-variance tradeoff to improve the prediction accuracy. Given the limit on the computation time in most practical applications, the data that may be collected through simulation experiments are often limited. Therefore, the dimension (or at least the effective dimension) of the covariates cannot be too high. Otherwise, we typically do not have the sufficient data to construct an accurate predictive model to make real-time decisions. One may view this as a limitation of the OSOA framework.

In most analytics applications, the data are given and a surface is fitted to achieve the best accuracy, see Hastie et al. (2011) and Tan et al. (2018) for surveys. In the OSOA framework, however, the data collection process is more active and a lot can be done to control the accuracy of the fitted surface. This is often called metamodeling to emphasize the goal is to build a metamodel (i.e., a surface) of a simulation model (see, for instance, Barton and Meckesheimer 2006 for more details on simulation metamodeling). In metamodeling, the critical issue is how to design the experiments, i.e., how to choose the design points (i.e., the values of the covariates) to conduct simulation experiments and the sample size at each design point. Moreover, the experimental design may be conducted in a sequential manner to further improve the prediction accuracy given the same amount of computational budget. These metamodeling tools may certainly be used in the OSOA framework as well.

However, experimental designs in the OSOA framework can sometimes go beyond those in simulation metamodeling problems. The following are some examples encountered in recent works:

1. In simulation metamodeling problems, simulation models are treated as black boxes. In many of the OSOA problems, however, one may have a simple model that may be less accurate than the simulation model but computationally much easier to evaluate. For

instance, in queueing systems, besides the complicated queueing models like the G/G/m queue, one may also have simple models with closed-form expressions of system performances, such as the M/M/m queue. These simple models are typically less accurate locally, but they often capture the global trend of the surface. Therefore, incorporating these simple models in the metamodeling may be beneficial, and it affects the design of experiments. Shen et al. (2018) considered how to incorporate simple models in metamodeling to improve the prediction accuracy and to speed up the optimization process. Specifically, the simple models (also called stylized models), which contain useful information about the shape of the system performance, are incorporated into a stochastic kriging approach by replacing the constant term.

2. In simulation metamodeling problems, only a single metamodel needs to be built, and the accuracy of the metamodel is in the first place. In some of the OSOA problems, however, one may need to build multiple metamodels. Besides the accuracy of the multiple metamodels, we also need to maintain a certain relationship between these metamodels. In derivative pricing, for instance, one needs both the surface of the price and the surfaces of the sensitivities (known as the Greeks). An interesting and also important problem is how to build the metamodels so that they are consistent in differentiation, i.e., the gradient of the price surface matches the values of the sensitivity surfaces. Jiang et al. (2019b) considered this problem and suggested to use stochastic co-kriging to solve the problem.

3. In simulation metamodeling problems, metamodels are evaluated based either on their fitting quality, e.g., the integrated mean squared error, or their prediction accuracy. In the OSOA framework, however, metamodels are the bridge to real-time decision makings. Therefore, the fitting quality may not be the only criterion of metamodeling, and how to best support the decisions may become more important. When comparing two alternatives, for instance, metamodels may need to be more accurate in the regions (of covariates) where the two alternatives have similar performances, but less accurate in regions where the two alternatives are significantly different. We will discuss more of this issue in Sections 4 and 5 when we consider R&S with covariates and SO with

covariates.

Based on these examples, we can see that OSOA problems often bring unique features that allow innovative experimental design and metamodeling. We believe a lot may be done in the offline-simulation stage so that the simulation experiments may be conducted in an efficient way and the metamodeling or machine learning techniques can be applied to build metamodels to better support the online applications.

2.2 Online-Application Stage

Once the predictive models are given, the online applications become quite standard. When the values of the covariates are observed, the predictive models may be evaluated or optimized using very little computational time to support real-time decision makings.

However, OSOA also brings some new features compared to typical online problems. In the OSOA framework, the predictive models are mainly estimated in the offline-simulation stage using the simulation data collected at a set of design points. As the time goes by, one may also have time to add more experiments in some cases. The following are two examples:

1. In financial risk management, risk needs to be evaluated and monitored continuously during the trading hours. When the market is closed in the evenings and weekends, more simulation experiments may be run given the latest market closing prices. Therefore, over the time, we accumulate simulation data started from different time points and different initial states. How to combine these data to build a predictive model becomes an interesting problem.

2. Notice that the observed values of the covariates are rarely a point in the set of design points. Therefore, bias is often inevitable. In some other applications, e.g., personalized medical treatment, there is often some time for some additional simulation experiments after the covariates are observed, and these new data are typically unbiased observations of the performance given the values of the covariates. It is an interesting problem to consider how to reduce both the bias and the variance of the predictive value by using these additional simulation data.

2.3 Three Types of OSOA Problems

In this paper, we draw mostly from the recent research works to discuss three types of OSOA problems, estimation with covariates, R&S with covariates and SO with covariates. Let $f(\boldsymbol{\theta}, \mathbf{x})$ be a system performance or an objective function, where $\boldsymbol{\theta}$ denotes decision variables and $\mathbf{x}$ denotes covariates. Usually, we cannot obtain $f(\boldsymbol{\theta}, \mathbf{x})$ directly, instead, we can only observe its unbiased random sample $F(\mathbf{Y}, \boldsymbol{\theta}, \mathbf{x})$, where $\mathbf{Y}$ represents the randomness. That is,

$$f(\boldsymbol{\theta}, \mathbf{x}) = \mathbb{E}[F(\mathbf{Y}, \boldsymbol{\theta}, \mathbf{x})].$$

In estimation with covariates, $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ is usually fixed, and we are interested in estimating $f(\mathbf{x}) := f(\boldsymbol{\theta}_0, \mathbf{x})$. In R&S with covariates and SO with covariates, we would like to obtain the optimal solution for different covariates (scenarios)

$$\boldsymbol{\theta}^*(\mathbf{x}) = \arg\max_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta}, \mathbf{x}). \tag{1}$$

For R&S with covariates, $\Theta = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_k\}$ are different alternatives. For SO with covariates, $\Theta \subseteq \mathbb{Z}^d$ (discrete optimization) or $\Theta \subseteq \Re^d$ (continuous optimization), where $\mathbb{Z}^d$ denotes all d-dimensional vectors with integer components.

Notice that, if we regard the covariates $\mathbf{x}$ as random elements (denote as $\mathbf{X}$), then by Jensen's Inequality,

$$\max_{\boldsymbol{\theta} \in \Theta} \mathbb{E}[f(\boldsymbol{\theta}, \mathbf{X}))] \leq \mathbb{E}[\max_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta}, \mathbf{X}))].$$

Therefore, for both R&S with covariates and SO with covariates, the formulation with covariates (i.e., the online problems) on average outperforms the formulation where the covariates are observable after the optimal decision is made (i.e., the offline problems).

In the rest of this paper, we will also demonstrate how to link the offline-simulation stage and the online-application stage, and we will also discuss the experimental design issues and the metamodeling techniques based on the unique features of the applications. Notice that the main difference between the OSOA and the existing simulation approaches is the use of predictive models. In the OSOA framework, the predictive models are constructed in the offline-simulation stage and used in the online-application stage in real time. While most existing simulation approaches focus on how to design a good system/make a good decision without time limit or with a relatively enough time. To make the key idea more clear, we provide a brief description of the offline simulation stage, predictive model, and online application stage for the three types of problems in Table 1.

Table 1: A brief description of OSOA for three types of problems

	Offline simulation	Predictive model	Online application
Estimation with covariates	Estimate performance of system based on different covariate values	Build surface of performance of system w.r.t. the covariates	Evaluate system performance when the covariates are observed
R&S with covariates	Solve R&S based on different covariate values	Build surface of performance of each alternative w.r.t. the covariates	Select best alternative when the covariates are observed
SO with covariates	Solve optimization problems based on different covariate values	Build surface of objective function or optimal solution w.r.t. the covariates	Provide optimal solution when the covariates are observed

3 Estimation with Covariates

The first type of OSOA problems we considered is the estimation with covariates. In this type of problems, we derive the functional relationships among variables of the systems, such as model parameters, simulation outputs, system states, etc. In this section, we discuss the estimation with covariates via two aspects. In the first aspect, we consider the performance estimation, where the classical simulation is usually used. Particularly, we consider how to evaluate the performance of stochastic systems with covariates via the OSOA framework. In the second aspect, we consider the input uncertainty and metamodeling, which are hot topics in the simulation area. Some of the approaches to handle input uncertainty and metamodeling can be included in the OSOA framework.

3.1 Performance estimation with covariates

Simulation is widely used to estimate the performance of stochastic systems. For example, the steady-state waiting time of a queueing system can be estimated via a single long run simulation; the financial risk measures of a portfolio including complex derivatives can be estimated via a nested simulation. However, such estimations are all offline. By the OSOA framework, we can build the predictive models to evaluate the performance of stochastic systems with covariates in real time.

Notice that the performance of stochastic systems is usually a quantity, especially an expectation. In this section, we use the expectation expression of the performance as an example. Let $\mathbf{Y}$ be random variables in a stochastic system, and $F(\cdot)$ be a sample performance function of the system. The performance of the system of interest can be expressed by $\mathbb{E}[F(\mathbf{Y})]$. By applying OSOA, we can estimate the performance with covariates, where we use $F(\mathbf{Y}, \mathbf{x})$ (we drop off $\boldsymbol{\theta}_0$ here) to denote a sample performance function of the system with covariates $\mathbf{x}$. That is, we derive a functional relationship between the system variables and the simulation outputs, which can be expressed by $\mathbb{E}[F(\mathbf{Y}, \mathbf{x})]$ with the expectation with respect to $\mathbf{Y}$. It can be interpreted as when we observe new states or new values of covariates in the system, how will be the performance of the system. Notice that if the covariates are assumed to be random variables, the estimation with covariates turns to be a conditional estimation, i.e., $\mathbb{E}[F(\mathbf{Y})|\mathbf{X} = \mathbf{x}]$, where $\mathbf{X}$ are the covariates in the system with the observed values $\mathbf{x}$. We will provide an example to show the practical applications of performance estimation with covariates.

Jiang et al. (2019a) considered a performance estimation with covariates problem in financial engineering, specifically, estimating the conditional financial risk measures. Financial risk measures are important indicators of risk of financial portfolios and even stability of financial institutions. Simulation studies are often used to estimate portfolio risk measures, and they often produce accurate estimates of risk measures, but very time consuming, especially when portfolios consist of multiple derivative products whose prices also need to be determined by additional simulation efforts (see Hong et al. 2014 for a survey). Nearly all estimation methods proposed in the simulation literature are designed under the offline framework. In practice, however, portfolio risk measures are often needed to be estimated in real time as the current values of the underlying risk factors change; think the famous 4:15 report of J. P. Morgan. Therefore, the OSOA framework can be applied to estimate the conditional risk measures to do online risk monitoring.

The risk measures considered in Jiang et al. (2019a) is the exceedance probability, which expresses as an expectation

$$p = \Pr\left\{L(\mathbf{S}_T) \geq l_0\right\} = \mathbb{E}\left[\mathbf{1}\{L(\mathbf{S}_T) \geq l_0\}\right],$$

where $\{\mathbf{S}_t\}_{t\in[0,T]}$ with $\mathbf{S}_t \triangleq (S_{1,t}, S_{2,t}, \ldots, S_{m,t})$ represents the underlying risk factors in a portfolio, and each coordinate $S_{i,t}, i = 1, 2, \ldots, m$, is driven by a stochastic process. $L(\cdot)$ is the loss function of the portfolio, and l_0 is a preset threshold. Given the initial value of the stochastic processes $\mathbf{S}_0$ and some other parameters in the portfolio, the quantity p can be estimated via a nested simulation.

In order to monitor risk online, the conditional exceedance probability is required, i.e.,

$$p(\mathbf{s}) = \mathbb{E}\left[\mathbf{1}\{L(\mathbf{S}_T) \geq l_0\} | \mathbf{S}_t = \mathbf{s}\right],$$

where $\mathbf{S}_t = \mathbf{s}$ means observing a realization of underlying risk factors at time $t \in (0, T)$. Jiang et al. (2019a) proposed to use logistic models as the predictive model in OSOA to estimate the conditional exceedance probability. The logistic regression model is given by

$$\log\left(\frac{p(\mathbf{S}_t)}{1 - p(\mathbf{S}_t)}\right) = \boldsymbol{\beta}(t)^\top \mathbf{X}(\mathbf{S}_t),$$

where $\mathbf{X}(\cdot) : \Re^m \to \Re^d$ denotes a basis function mapping to make the model more flexible, and $\boldsymbol{\beta}(t) = (\beta_1(t), \ldots, \beta_d(t))^\top$ is the vector of coefficients. When the portfolio is formed, a thorough nested simulation study is typically conducted to analyze the risk profile of the portfolio. After the simulation study, we have n simulated sample paths of the underlying risk factors $\mathbf{S}_1(t), \mathbf{S}_2(t), \ldots, \mathbf{S}_n(t), 0 \leq t \leq T$, and we can also easily calculate the portfolio loss at time T based on the simulated realizations of the underlyings, denoted them as

$L_1(T), L_2(T), \ldots, L_n(T)$. Then, the coefficient $\boldsymbol{\beta}(t)$ can be estimated via maximum likelihood estimation. Jiang et al. (2019a) showed that the conditional exceedance probability estimator is consistent and has a central limit theorem under some mild conditions.

Furthermore, the risk managers sometimes need to know whether the portfolio is safe or not in real time, that is, the online risk classification problem (see Lucas and Klaassen 2006 for example). In this problem, a threshold α is given, and if the exceedance probability $p(\mathbf{x})$ is less than or equal to α, the portfolio is in the safe zone. Otherwise, the portfolio is in the dangerous zone. In practice, $p(\mathbf{x})$ is unknown and $\hat{p}(\mathbf{x})$ is its estimate. Let I and $\hat{I}$ denote the safe/dangerous indicators under the true and estimated probabilities, respectively, i.e.,

$$I = \begin{cases} 1 & \text{if } p(\mathbf{x}) \leq \alpha \\ 0 & \text{if } p(\mathbf{x})) > \alpha \end{cases}, \qquad \hat{I} = \begin{cases} 1 & \text{if } \hat{p}(\mathbf{x}) \leq \alpha \\ 0 & \text{if } \hat{p}(\mathbf{x}) > \alpha \end{cases}.$$

Then, $\hat{I} \neq I$ denotes a misclassification. Jiang et al. (2019a) showed that the probability of misclassification, i.e., $\Pr\{\hat{I} \neq I\}$, converges exponentially fast as $n \to \infty$, under some mild conditions. In addition, by introducing the *lasso*, *perturbation technique*, performance of the conditional exceedance probability estimator and classifier can be enhanced. Specifically, Jiang et al. (2019a) provided two ways to utilize the *additional simulation* data. One is combing the estimators derived from the additional data set and the original data set separately, and the other is to combine these two data sets together. One may refer to Jiang et al. (2019a) for more details.

The above example introduces the application of the OSOA framework in financial engineering. In other areas, the OSOA framework can be also applied to evaluate performances of stochastic systems with covariates. For example in queueing, one may be interested in the waiting time of a customer to a service system conditional on a particular event occurring at a particular time. Lin and Nelson (2016) proposed to retain the simulation samples in offline, and apply k nearest neighborhood method to build a predictive model for online application. Ouyang and Nelson (2017) considered a two-step method to predict the probability that the queueing system state belongs to a certain subset. Li et al. (2019) proposed to use stochastic kriging prediction model in the offline-simulation stage and studied the convergence rate of the prediction error with the number of covariate points sampled. By applying OSOA framework, the traditional performance estimation problem can be extended to the online setting, which makes the performance more practical.

3.2 Input Uncertainty and Metamodeling

Estimation with covariates may have a general scope including input uncertainty and metamodeling. Input uncertainty refers to the uncertainty that is derived from using estimated input models (parameters), which are the probability distributions used to drive the simulation, and they are specified from real-world data or the subjective experience; think interarrival times and service times in queueing; drifts and volatilities of underlying assets in option pricing. There are various methods proposed to quantify the uncertainty due to input models, and the methods can be basically divided into frequentist and Bayesian approaches, see Barton (2012) and Lam (2016) for surveys. For some recent works, interested readers may see Xie et al. (2014), Song and Nelson (2015), Barton et al. (2018), etc. We use the same notation in Xie et al. (2014), and let $\mathbf{z}_m$ denote the real-world data, where m is the data size. For fixed input model parameter $\boldsymbol{\theta}$, $\mu(\boldsymbol{\theta})$ is the true simulation mean response.

In frequentist approaches, a point estimate of the input model parameters $\hat{\boldsymbol{\theta}}$ is estimated from $\mathbf{z}_m$. Since these data are only the realizations of the true input models distributions, the uncertainty is quantified by its sampling distribution. The input uncertainty is then propagated to the simulation output. If $\mu(\boldsymbol{\theta})$ is known, the input uncertainty can be assessed easily via some functional transformations, and the confidence interval (CI) of simulation outputs can be constructed accounting for the input uncertainty. However, $\mu(\cdot)$ is usually unknown, so the direct simulation or metamodeling is needed to approximate $\mu(\cdot)$, and the input uncertainty is difficult to be assessed in the simulation output.

In Bayesian approaches, a prior distribution of input model parameters $\pi_\Theta(\boldsymbol{\theta})$ is first assumed, and the data $\mathbf{z}_m$ are applied to update the posterior distribution via $p_\Theta(\boldsymbol{\theta}|\mathbf{z}_m) \propto \pi_\Theta(\boldsymbol{\theta}) \cdot p_{\mathbf{z}_m}(\mathbf{z}_m|\boldsymbol{\theta})$, where $p_{\mathbf{z}_m}$ is assumed likelihood function of $\mathbf{z}_m$ given $\boldsymbol{\theta}$. Similar to frequentist approaches, if $\mu(\boldsymbol{\theta})$ is known, the propagation of input uncertainty can be tracked and CI of simulation outputs can be constructed to assess the risk of input uncertainty. If $\mu(\boldsymbol{\theta})$ is unknown, direct simulation or metamodeling is also needed to learn the response surface of $\mu(\boldsymbol{\theta})$. For example, Xie et al. (2014) proposed a Gaussian process model to learn $\mu(\boldsymbol{\theta})$, and use it in the posterior distribution of simulation output response.

Notice that, no matter in frequentist approaches or Bayesian approaches, if the surface of $\mu(\boldsymbol{\theta})$ is unknown, we need to learn it from the simulation data, which has the same spirit with the OSOA framework, that is, we learn the the true simulation mean response offline via some metamodeling tools, like Gaussian process regression (see Ankenman et al. 2010), and use the learned response surface to assess input uncertainty online. As we have mentioned in Section 2, metamodeling tools are essential in the offline-simulation stage, and using a proper metamodeling method will bring a significant benefit in solving the online problem.

Estimation with covariates is an important problem. In many applications, the conditional expected functions are critical to online decision makings. Besides, it is also the building block of other types of methods, e.g., approximate dynamic programming (Powell 2011). There are many research issues worthy of studying in estimation with covariates. For example, we may consider how to deal with the misspecification of models (both simulation models and predictive models), and how to combine the additional simulation data with the original data to improve the accuracy of estimations more efficiently.

4 Ranking and Selection with Covariates

In this section, we consider the second type of OSOA problems, ranking and selection (R&S) with covariates. R&S is a kind of discrete simulation optimization problem, which identifies the best one from a finite set of competing alternatives. The problem can be written by

$$i^* = \arg\max_{i \in \mathcal{S}} \{\mu_i = \mathbb{E}[Y_i]\},$$

where $\mathcal{S} = \{1, \ldots, k\}$ is a finite set of alternatives indexed by i. For each alternative i, its performance μ_i cannot be computed exactly, but can be estimated using output from a stochastic simulation represented by Y_i. For example in inventory management, $\mathcal{S}$ may be different inventory policies, μ_i is the expected average profit of each policy, and Y_i is the realization of each simulation on the ith policy. The aim of the management is to find the policy with the largest expected profit. For R&S procedures and applications, see Kim and Nelson (2006), Kim and Nelson (2007), and Hong et al. (2015) for surveys.

In classical R&S procedures, all the model parameters are preset, and the performances of all the alternatives are usually unconditional. However, in some applications, some model parameters cannot be determined in advance, and can only be observed online, that is, there are covariates in the stochastic systems. At the same time, the decision should be made immediately. For example, online advertisements or promotions need to be pushed to consumers based on their purchasing behaviors and personal information like income and address. After obtaining the information (covariates), the chosen advertisement should be presented to the consumers on their screens immediately. In this setting, the performance of an alternative is no longer a quantity, and varies as a function of the covariates.

Shen et al. (2019) proposed new procedures to solve the R&S with covariates under the OSOA framework. Distinct from the classical setting, they assumed that the performance of each alternative depends on $\mathbf{X} = (X_1, \ldots, X_d)^\top \in \mathbb{X} \subseteq \Re^d$. $\mu_i(\mathbf{X})$ denotes the mean performance of alternative i. For each $i = 1, \ldots, k$ and $l = 1, 2, \ldots,$ $Y_{il}(\mathbf{X})$ denotes the lth sample from alternative i. The goal is to select the alternative with the largest mean performance conditionally on $\mathbf{X} = \mathbf{x}$,[1]

$$i^*(\mathbf{x}) = \arg\max_{i \in \mathcal{S}} \{\mu_i(\mathbf{X}) | \mathbf{X} = \mathbf{x}\}.$$

Notice that when $\mathbf{X}$ changes, the previous best alternative is not necessarily the current best alternative, so it needs to know the surfaces of performances of alternatives with respect to $\mathbf{X}$. Shen et al. (2019) assumed a linear model in which $\mu_i(\mathbf{X})$ is linear in $\mathbf{X}$ and $Y_{i,l}(\mathbf{X})$ is

[1] R&S with covariates can be formulated as (1). Here we use another formulation in line with the classical R&S.

unbiased, that is, for each $i = 1, \ldots, k$ and $l = 1, 2, \ldots$, conditionally on $\mathbf{X} = \mathbf{x}$,

$$\begin{aligned} \mu_i(\mathbf{x}) &= \mathbf{x}^\top \boldsymbol{\beta}_i, \\ Y_{il}(\mathbf{x}) &= \mu_i(\mathbf{x}) + \epsilon_{il}(\mathbf{x}), \end{aligned}$$

where $\boldsymbol{\beta}_i \in \Re^d$ is a vector of unknown coefficients and $\epsilon_{il}(\mathbf{x}) \sim \mathcal{N}(0, \sigma_i^2(\mathbf{x}))$ is the sampling error and is independent of $\epsilon_{i'l'}(\mathbf{x}')$ for any $(i, l, \mathbf{x}) \neq (i', l', \mathbf{x}')$.

The OSOA framework is applied in Shen et al. (2019). In the offline-simulation stage, they first sample each alternative at some values of the covariates, and use the data to estimate the coefficient $\boldsymbol{\beta}$. In the online-application stage, the estimates of $\hat{\mu}_i(\mathbf{x}) = \mathbf{x}^\top \hat{\boldsymbol{\beta}}_i, i = 1, \ldots, k$ are applied to solve online decision making problems. More specifically, in the offline-simulation stage, they use the indifference-zone formulation, and defined the conditional probability of correct selection (PCS) as

$$\mathrm{PCS}(\mathbf{x}) = \mathbb{P}\left(\mu_{i^*(\mathbf{X})}(\mathbf{X}) - \mu_{\hat{i}^*(\mathbf{X})}(\mathbf{X}) < \delta | \mathbf{X} = \mathbf{x}\right),$$

where δ is the indifference-zone parameter. In addition, the fixed design setting is assumed, i.e., the design points $\mathbf{x}_1, \ldots, \mathbf{x}_m$ are chosen properly and fixed, and that alternative i can be sampled at $\mathbf{x}_j$ repeatedly arbitrarily many times, for each $i = 1, \ldots, k$ and $j = 1, \ldots, m$. Let $\mathcal{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_m)^\top \in \Re^{m \times d}$ and $\mathbf{Y}_{il} = (Y_{il}(\mathbf{x}_1), \ldots, Y_{il}(\mathbf{x}_m))^\top$. Then they proposed several two-stage R&S procedures under the homoscedastic ($\sigma^2(\mathbf{x}_i) = \sigma^2$ for all $\mathbf{x}_i, i = 1, \ldots, \mathbf{x}_m$) and heteroscedastic ($\sigma^2(\mathbf{x}_i)$ is different on $\mathbf{x}_i, i = 1, \ldots, m$) sampling errors setting, respectively. The basic procedure is described as follow:

- In the first stage sampling, take n_0 independent samples of each alternative i at each design point, and estimate the coefficient $\hat{\boldsymbol{\beta}}_i(n_0)$ and the variance of each alternative

$$S_i^2 = \frac{1}{n_0 m - d} \sum_{l=1}^{n_0} \left(\mathbf{Y}_{il} - \mathcal{X}\hat{\boldsymbol{\beta}}_i(n_0)\right)^\top \left(\mathbf{Y}_{il} - \mathcal{X}\hat{\boldsymbol{\beta}}_i(n_0)\right).$$

- In the second-stage sampling, compute the total sample size $N_i = \max\{\lceil h^2 S_i^2/\delta^2 \rceil, n_0\}$ for each i, where h is a constant determined in advance based on different PCS criteria and different sampling error settings (homoscedastic or heteroscedastic), and $\lceil a \rceil$ denotes the smallest integer no less than a.

- In the selection stage, for each alternative i, compute the overall estimate of its unknown coefficients

$$\hat{\boldsymbol{\beta}}_i = \frac{1}{N_i}\left(\mathcal{X}^\top \mathcal{X}\right)^{-1} \mathcal{X}^\top \sum_{l=1}^{N_i} \mathbf{Y}_{il},$$

  and return $\hat{i^*}(\mathbf{x}) = \arg\max_{i\in\mathcal{S}}\{\mathbf{x}^T\hat{\boldsymbol{\beta}}_i\}$.

Shen et al. (2019) proved the statistical validities under different settings, and pointed out that research on complex models to capture the relationship between the response of an alternative and the covariates is a good potential directions for future work.

Pearce and Branke (2017) also considered R&S with covariates. However, the formulation in their setting is different. They assumed the distribution of the covariate $\mathbf{X}$ as $\mathbb{P}(\mathbf{x})$, and the aim is to find a mapping $i^*(\mathbf{x})$ that maximizes the overall expected performance across all the possibilities of the covariate $\mathbf{X}$, i.e., $\int_\Theta \mu_{i(\mathbf{x})}(\mathbf{x})\mathbb{P}(\mathbf{x})dx$. In addition, the problem is solved under the Bayesian framework. More specifically, the performance surface $\mu_i(\mathbf{x})$ is predicted by Gaussian process regression model, and the sampling policies for determining the sampling point and the chosen alternative based on the expected improvement. To approximate the expected improvement, three numerically integrating approaches are proposed.

R&S with covariates is a new problem under the R&S category. For each alternative, its performance is not a constant but a surface (function) of covariates, so the aim of R&S with covariates is not to find the best alternative, but to provide a policy (a mapping from the covariates to alternatives) that tell us which alternative is the best when covariates are observed. The complexity of the performance surfaces of alternatives and the large dimensionality of the covariates make R&S with covariate more difficult than the classical setting generally. The OSOA framework may be a potential way to overcome the difficulties. By applying data analytics tools (e.g., linear regression and Gaussian regression in the previous two works), the surfaces of alternatives performances can be learned in the offline-simulation stage. In the online-application stage, the learned surfaces can be used to make decisions immediately. There are some other R&S procedures such as knowledge gradient (Frazier et al. 2008, Ryzhov et al. 2012), optimal computing budget allocation (Chen et al. 2000, Chen et al. 2010, Xiao et al. 2017, Xiao and Gao 2018), and we may consider how to adapt these procedures under the OSOA framework to solve R&S with covariates.

To close this section, we introduce another type of problems with covariates, that is, multi-armed bandit (MAB) with covariates. MAB and R&S were proposed at almost the same time, and MAB is an important class of sequential decision making problems in the fields of operations research, statistics and machine learning; see, for instance, Bubeck and Cesa-Bianchi (2012) for a review. Different from the R&S to identify the best alternative[2], the aim of MAB is to make decisions that maximizes the cumulative rewards (roughly speaking, when sampling an alternative, there is a reward for selecting this alternative). The offline-simulation stage and the online-application (decision) stage are coupled at the same time in MAB. MAB with covariates (also known as contextual MAB, bandits with side information, etc.) has been studied widely in recent years. Parametric models (see Goldenshluger and Zeevi 2009, Wang et al. 2005) and nonparametric models (see Perchet and Rigollet 2013, Slivkins 2014) have been considered in the literature of MAB with covariates. The aim of MAB with covariate is different from that of R&S with covariate. The former aims to obtain the optimal discounted cumulative rewards, and does not care whether the true surfaces are learned. However, the latter aims to estimate the surfaces with respect to covariates so that it can select the best alternative or good alternatives with statistically significance.

5 Simulation Optimization with Covariates

In this section, we consider the third type of OSOA problems, simulation optimization (SO) with covariates. Distinct from Section 4, we focus on continuous problems. Recall that the classical SO problem is to solve

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}\in\Theta} \mathbb{E}[F(\boldsymbol{\theta}, \mathbf{Y})], \tag{2}$$

[2]Notice that the problem "best arm identification in MAB" is similar to R&S, but the stochastic performance (reward) is not normally distributed in this problem, and the way of analysis and the procedures of selection are also different from R&S. See Audibert et al. (2010) and Kaufmann et al. (2016) for more details.

where $\Theta \subseteq \Re^d$, $\mathbf{Y}$ are random variables in the stochastic system, and F is the sample performance function of the stochastic system. There are many methods proposed to solve this problem, such as stochastic approximation (see Kushner and Yin 2003 for a survey), and sample average approximation (see Shapiro 2003 for a survey).

In some applications, there may be covariates in the objective function. For example, in wind farms, the managers need to determine the amount of the electricity production an hour in advance based on the knowledge on current observed informations such as the time of day, the time of year, wind speed, electricity price, etc; see Hannah et al. (2010). In the urban and air transportation, the optimal control policies should change with the observed states in the transportation system and other information like weather and time. Let $\mathbf{x} \in \mathcal{X}$ be the covariates, and the SO with the covariate can be formulated by

$$\boldsymbol{\theta}^*(\mathbf{x}) = \arg\min_{\boldsymbol{\theta}\in\Theta} \mathbb{E}[F(\boldsymbol{\theta}, \mathbf{Y}, \mathbf{x})], \tag{3}$$

where the distributions of $\mathbf{Y}$ may be affected by $\mathbf{x}$. Different from (2), the aim of (3) is to find the surface of the optimal solution with respect to covariates $\mathbf{x}$.

Hannah et al. (2010) considered a convex stochastic optimization problem with covariates, which is similar to SO with covariates in this paper. However, in their setting, the covariates are uncontrollable and observed online, and the realizations of the random variables $\mathbf{Y}_i$ are determined only after observing the covariates $\mathbf{x}_i$ and the action $\boldsymbol{\theta}_i$. They proposed two approaches to solve the problem.

The first approach is the function-based. Let $\{\mathbf{x}_i, \mathbf{Y}_i\}_{i=1}^n$ be a set of n observations, and $\{w_n(\mathbf{x}, \mathbf{x}_i)\}_{i=1}^n$ be a set of weights such that $\sum_{i=1}^n w_n(\mathbf{x}, \mathbf{x}_i) = 1$. Similar to SAA in approximating the objective function, they developed a locally weighted average approximation, that is,

$$\bar{F}(\boldsymbol{\theta}|\mathbf{x}) = \sum_{i=1}^n w_n(\mathbf{x}, \mathbf{x}_i) F(\boldsymbol{\theta}, \mathbf{Y}_i, \mathbf{x}_i).$$

Then the estimated optimal solution is given by

$$\hat{\boldsymbol{\theta}}^*(\mathbf{x}) = \arg\min_{\boldsymbol{\theta}\in\Theta} \bar{F}(\boldsymbol{\theta}|\mathbf{x}).$$

Notice that, since $\mathbf{x}$ may affect the distributions of $\mathbf{Y}$, $\mathbf{Y}_i, i = 1, 2, \ldots, n$, are not necessarily i.i.d. So it requires a weighted average approximation. The weights are like the likelihoods in importance sampling that make the approximation be consistent. Hannah et al. (2010) proposed to use kernel weights and Dirichlet process weights in the approximation, and proved that $\hat{\boldsymbol{\theta}}^*(\mathbf{x})$ converges to the true optimal solution $\boldsymbol{\theta}^*(\mathbf{x})$ under some mild conditions.

The second approach is gradient-based. In this approach, the sample performance function $F(\boldsymbol{\theta}_i, \mathbf{Y}_i, \mathbf{x}_i)$ cannot be observed, but only a gradient estimate at $\boldsymbol{\theta}_i$ instead, i.e.,

$$\hat{g}(\boldsymbol{\theta}_i, \mathbf{Y}_i, \mathbf{x}_i) = \nabla_{\boldsymbol{\theta}} F(\boldsymbol{\theta}_i, \mathbf{Y}_i, \mathbf{x}_i).$$

In addition, the objective function $\mathbb{E}[F(\boldsymbol{\theta}, \mathbf{Y}, \mathbf{x})]$ is assumed to have some special structures, and approximated by

$$\bar{F}_n(\boldsymbol{\theta}|\mathbf{x}) = \sum_{k=1}^{d} f_n^k(\theta^k|\mathbf{x}),$$

where θ^k is the kth component of $\boldsymbol{\theta}$. $f_n^k(y|\mathbf{x})$ is a univariate, piecewise linear function in y, and it is convex for every $\mathbf{x} \in \mathcal{X}$. Based on this special structure and the observations $\{\boldsymbol{\theta}_i, \mathbf{x}_i, \hat{g}(\boldsymbol{\theta}_i, \mathbf{Y}_i, \mathbf{x}_i)\}_{i=1}^n$, the approximation $\bar{F}_n(\boldsymbol{\theta}|\mathbf{x})$ can be updated sequentially. Since $f_n^k(\theta^k|\mathbf{x})$ is piecewise linear, the slopes of the linear functions can be determined based on the gradient estimates $\hat{g}(\boldsymbol{\theta}_i, \mathbf{Y}_i, \mathbf{x}_i)$. However, $\mathbf{x}_i, i = 1, 2, \ldots, n$, may be different from a new value $\mathbf{x}$, so, similar to function-based approach, kernel weights and Dirichlet process weights are introduced in the procedure to estimate the slopes on $\mathbf{x}$. Details refers to Algorithm 2 in Hannah et al. (2010). In that paper, convergence of gradient-based approach was also studied under some mild conditions.

Similar to R&S with covariate, we need to learn the optimal solution surface in SO with covariates, which makes the problem more challenging. Distinct from the setting of Hannah et al. (2010), the covariates may be controllable by us in some problems, and we can design the covariates to obtain a good optimal solution surface. In addition, in the offline-simulation stage, we can both build predictive models for objective functions and for optimal solutions, so we may derive fundamentally different ways to solve this problem.

6 Conclusion and Future Research

Due to the famous Moore's law, the computing power has grown tremendously in past fifty years, and it has profoundly changed every aspect of human life. In the simulation area, the growing computing power has also changed the ways that simulation models are built and simulation experiments are conducted and the applications that simulation may be used on. On one hand, the availability of large-scale cheap and easily accessible computing power makes a large amount of offline simulation experiments possible; on the other hand, simulation experiments still take minutes to hours to execute and, therefore, cannot be used in making many online decisions, where time windows for making decision are commonly in the order of seconds to minutes, sometimes even milliseconds. To fill the gap, we summarize the OSOA framework in this paper so that simulation becomes a viable tool for real-time decision makings.

Based on our research experience, we propose some further research directions about OSOA:

- *Building new metamodeling techniques and theories in OSOA.* As we have introduced in Section 2.1, the goal of metamodeling problems in the OSOA framework is different from that of classical metamodeling problems. In classical metamodeling problems, fitting accuracy is the most important criterion. However, in the OSOA framework, the metamodel (fitted surface) is used to make decisions, so the accuracy is not the only criterion. How to best support the decisions is the most important goal of the metamodeling in the OSOA framework. Therefore, building new metamodeling techniques and new theories under the OSOA framework are good research directions.

- *Treating OSOA as a feedback loop.* In the OSOA framework, the metamodels are often used repeatedly (with different observed values of the covariates). Real data are often observed after the decisions are made. Based on the real data and the simulated data, the simulation models need to be updated to capture the features of the stochastic system, e.g., better calibrating the model parameters and reducing model misspecifications. Therefore, how to handle these feedback loops becomes another interesting research direction.

- *Dealing with high dimensional covariates.* In many practical problems, the dimension of the covariates may be very high. For example in healthcare, a lot of information, such as the demographic information and gene information, may be collected for a patient. However, some diseases are only affected by a few of factors. That is, the effective dimension is low. In the OSOA framework, since we may repeat the simulation experiments and receive feedbacks regularly, we may utilize these information and some machine learning techniques to identify the effective dimensions. How to deal with high dimensional covariates and how to reduce dimensionality are interesting research directions in the OSOA framework.

- *Embedding SO algorithms in OSOA.* There are many efficient SO algorithms proposed in last two decades. An interesting research direction is to embed these algorithms to solve SO problems with covariates. Notice that, in the OSOA framework, the problems of SO with covariates may have both high dimensional decision variables and high dimensional covariates. When designing the SO algorithms, these features need to be considered.

Acknolwedgement

The research reported in this paper is partially supported by Hong Kong Research Grants Council Grant [GRF 16203214, 11504017], National Natural Science Foundation of China [Grants 71801148], and Shanghai Young Eastern Scholar Program [N.60-D129-18-202].

References

Ankenman, B., Nelson, B. L., and Staum, J. (2010). Stochastic kriging for simulation metamodeling. *Operations Research*, 58:371–382.

Audibert, J., Bubeck, S., and Munos, R. (2010). Best arm identification in multi-armed bandits. In *Proceedings of the Conference on Learning Theory*, pages 41–53.

Barton, R., Lam, H., and Song, E. (2018). Revisiting direct bootstrap resampling for input model uncertainty. In Rabe, M., Juan, A. A., Mustafee, N., Skoogh, A., Jain, S., and Johansson, B., editors, *Proceedings of the 2018 Winter Simulation Conference*, pages 1635–1645.

Barton, R. R. (2012). Tutorial: Input uncertainty in output analysis. In Laroque, C., Himmelspach, J., Pasupathy, R., Rose, O., and Uhrmacher, A., editors, *Proceedings of the 2012 Winter Simulation Conference*, pages 67–78, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R. and Meckesheimer, M. (2006). Metamodel-based simulation optimization. In Henderson, S. G. and Nelson, B. L., editors, *Handbooks in Operations Research and Management Science*, volume 13, chapter 18, pages 535–574. Springer, Amsterdam.

Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122.

Chen, C. H., Lin, J., Yucesan, E., and Chick, S. E. (2000). Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Journal of Discrete Event Dynamic Systems: Theory and Applications*, 10:251–270.

Chen, C. H., Yucesan, E., Dai, L., and Chen, H. C. (2010). Efficient computation of optimal budget allocation for discrete event simulation experiment. *IIE Transactions*, 42:251–270.

Choi, S. E., Perzan, K. E., Tramontano, A. C., Kong, C. Y., and Hur, C. (2014). Statins and aspirin for chemoprevention in barretts esophagus: Results of a cost-effectiveness analysis. *Cancer Prevention Researh*, 7:341–350.

Daganzo, C. F. (1997). *Fundamentals of Transportation and Traffic Operations*. Emerald Group Publishing Limited, Bingley, UK.

Dingeç, K. D. and Hörmann, W. (2012). A general control variate method for option pricing under Lévy processes. *European Journal of Operational Research*, 221:368–377.

Dunnett, C. W. (1984). Selection of the best treatment in comparison to a control with an application to a medical trial. In Santner, T. J. and Tamhane, A. C., editors, *Design of Experiments: Ranking and Selection*, pages 21–34. CRC Press, New York.

Feng, M. and Staum, J. (2017). Green simulation: Reusing the output of repeated experiments. *ACM Transactions on Modeling and Computer Simulation*, 37(4):Article 23.

Frazier, P. I., Powell, W. B., and Dayanik, S. (2008). A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47:2410–2439.

Fu, M. C. (2015). Stochastic gradient estimation. In Fu, M. C., editor, *Handbooks of Simulation Optimization*, chapter 5, pages 105–147. Springer, New York.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer, New York.

Glasserman, P., Heidelberger, P., and Shahabuddin, P. (1999). Asymptotically optimal importance sampling and stratification for pricing path-dependent options. *Mathematical Finance*, 9:117–152.

Goldenshluger, A. and Zeevi, A. (2009). Woodroofes one-armed bandit problem revisited. *The Annals of Applied Probability*, 19(4):1603–1633.

Hannah, L. A., Powell, W. B., and Blei, D. M. (2010). Nonparametric density estimation for stochastic optimization with an observable state variable. In Lafferty, J. D., Williams, C. K. I., Shawe-Taylor, J., Zemel, R. S., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 820–828. Curran Associates, Inc.

Hastie, T., Tibshirani, R., and Friedman, J. (2011). *The Elements of Statistical Learning*. Springer-Verlag, New York, second edition.

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

Hong, L. J., Nelson, B. L., and Xu, J. (2015). Discrete optimization via simulation. In Fu, M., editor, *Handbook of Simulation Optimization*, chapter 2, pages 9–44. Springer, New York.

Jiang, G., Fu, M. C., and Xu, C. (2016a). On sample average approximation algorithms for determining the optimal importance sampling parameters in pricing financial derivatives on Lévy processes. *Operations Research Letters*, 44(1):44–49.

Jiang, G., Hong, L. J., and Nelson, B. L. (2016b). A simulation analytics approach to dynamic risk monitoring. In Roeder, T. M. K., Frazier, P. I., Szechtman, R., Zhou, E., Huschka, T., and Chick, S. E., editors, *Proceedings of the 2016 Winter Simulation Conference*, pages 437–447, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Jiang, G., Hong, L. J., and Nelson, B. L. (2019a). Online risk monitoring using offline simulation. *INFORMS Journal on Computing*. Forthcoming.

Jiang, G., Hong, L. J., and Shen, H. (2019b). Constructing consistent surfaces of price and greeks for financial derivatives. Working paper.

Jin, X., Li, H., and Lee, L. H. (2019). Optimal budget allocation in simulation analytics. Working paper.

Jorion, P. (2006). *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, New York, third edition.

Jun, J. B. and Jacobson, S. H. (1999). Application of discrete-event simulation in health care clinics: A survey. *Journal of the Operational Research Society*, 50:109–123.

Kaufmann, E., Cappé, O., and Garivier, A. (2016). On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42.

Kim, S.-H. and Nelson, B. L. (2006). Selecting the best system. In S. G. Henderson, B. L. N., editor, *Handbooks in Operations Research and Management Science*, volume 13, chapter 17, pages 501–534. North-Holland Publishing, Amsterdam.

Kim, S.-H. and Nelson, B. L. (2007). Recent advances in ranking and selection. In Henderson, S. G., Biller, B., Hsieh, M.-H., Shortle, J., Tew, J. D., and Barton, R. R., editors, *Proceedings of the 2007 Winter Simulation Conference*, pages 162–172, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Kushner, H. and Yin, G. G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, New York.

Lam, H. (2016). Advanced tutorial: Input uncertainty and robust analysis in stochastic simulation. In Roeder, T. M. K., Frazier, P. I., Szechtman, R., Zhou, E., Huschka, T., and Chick, S. E., editors, *Proceedings of the 2016 Winter Simulation Conference*, pages 178–192.

Li, C., Gao, S., and Du, J. (2019). Convergence rates of stochastic kriging-assisted simulation with covariates. Working paper.

Lin, Y. and Nelson, B. L. (2016). Simulation analytics for virtual statistics via k nearest neighbors. In Roeder, T. M. K., Frazier, P. I., Szechtman, R., Zhou, E., Huschka, T., and Chick, S. E., editors, *Proceedings of the 2016 Winter Simulation Conference*, pages 448–459, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Liu, M., Nelson, B. L., and Staum, J. (2010). Simulation on demand for pricing many securities. In Johansson, B., Jain, S., Montoya-Torres, J., Hugan, J., and Yucesan, E., editors, *Proceedings of the 2010 Winter Simulation Conference*, pages 2782–2789, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Liu, M. and Staum, J. (2010). Stochastic kriging for efficient nested simulation of expected shortfall. *Journal of Risk*, 12(3):3–27.

Lucas, A. and Klaassen, P. (2006). Discrete versus continuous state switching models for portfolio credit risk. *Journal of Banking & Finance*, 30:23–35.

Luo, J., Hong, L. J., and Nelson, B. L. (2015). Fully sequential procedures for large-scale ranking-and-selection problems in parallel computing environments. *Operations Research*, 63:1177–1194.

Nelson, B. L. (2016). 'Some tactical problems in digital simulation' for the next ten years. *Journal of Simulation*, 10(1):2–11.

Ni, E. C., Ciocan, D. F., and Henderson, S. G. (2017). Efficient ranking and selection in high performance computing environments. *Operations Research*, 65:821–836.

Osorio, C. and Bierlaire, M. (2013). A simulation-based optimization framework for urban transportation problems. *Operations Research*, 1:1333–1345.

Ouyang, H. and Nelson, B. L. (2017). Simulation-based predictive analytics for dynamic queueing systems. In Chan, W. K. V., DAmbrogio, A., Zacharewicz, G., Mustafee, N., Wainer, G., and Page, E., editors, *Proceedings of the 2017 Winter Simulation Conference*, pages 1716–1727, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Pearce, M. and Branke, J. (2017). Efficient expected improvement estimation for continuous multiple ranking and selection. In Chan, W. K. V., DAmbrogio, A., Zacharewicz, G., Mustafee, N., Wainer, G., and Page, E., editors, *Proceedings of the 2017 Winter Simulation Conference*, pages 2161–2172, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Perchet, V. and Rigollet, P. (2013). The multi-armed bandit problem with covariates. *The Annals of Statistics*, 41(2):693–721.

Powell, W. B. (2009). *Approximate dynamic programming: solving the curses of dimensionality*. Wiley, Hoboken, New Jersey, second edition.

Powell, W. B. (2011). *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, Hoboken, New Jewsey.

Rosenbaum, I. and Staum, J. (2015). Database Monte Carlo for simulation on demand. In Yilmaz, L., Chan, W. K. V., Moon, I., Roeder, T. M. K., Macal, C., and Rossetti, M. D., editors, *Proceedings of the 2015 Winter Simulation Conference*, pages 679–688.

Ryzhov, I. O., Powell, W. B., and Frazier, P. I. (2012). The knowledge-gradient algorithm for a general class of online learning problems. *Operations Research*, 60:180–195.

Santana, E. and Hotz, G. (2016). Learning a driving simulator. Working paper.

Shapiro, A. (2003). Monte Carlo sampling methods. In Ruszczynski, A. and Shapiro, A., editors, *Handbooks in Operations Research and Management Science*, volume 10, chapter 6, pages 353–425. Elsevier, New York.

Shen, H., Hong, L. J., and Zhang, X. (2018). Enhancing stochastic kriging for queueing simulation with stylized models. *IISE Transactions*, 50:943–958.

Shen, H., Hong, L. J., and Zhang, X. (2019). Ranking and selection with covariates for personalized decision making. Working paper.

Slivkins, A. (2014). Contextual bandits with similarity information. *Journal of Machine Learning Research*, 15:2533–2568.

Song, E. and Nelson, B. L. (2015). Quickly assessing contributions to input uncertainty. *IIE Transactions*, 47:893–909.

Tan, P. N., Steinbach, M., Karpatne, A., and Kumar, V. (2018). *Introduction to Data Mining*. Pearson, New York, second edition.

Wang, C., Kulkarni, S. R., and Poor, H. V. (2005). Bandit problems with side observations. *IEEE Transactions on Automatic Control*, 50(3):338–355.

Xiao, H. and Gao, S. (2018). Simulation budget allocation for selecting the top-m designs with input uncertainty. *IEEE Transactions on Automatic Control*, 63:3127–3134.

Xiao, H., Gao, S., and Lee, L. H. (2017). Simulation budget allocation for simultaneously selecting the best and worst subsets. *Automatica*, 84:177–127.

Xie, W., Nelson, B. L., and Barton, R. R. (2014). A Bayesian framework for quantifying uncertainty in stochastic simulation. *Operations Research*, 62(2):1439–1452.