



INPUT DATA COLLECTION VERSUS SIMULATION: SIMULTANEOUS RESOURCE ALLOCATION

Yuhao Wang
Enlu Zhou

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

ABSTRACT

This paper investigates the problem of ranking and selection under input uncertainty with simultaneous resource allocation. In this problem, two types of resources are sequentially allocated at the same time to collect input data to reduce input uncertainty and run simulations to reduce stochastic uncertainty. We formulate the simultaneous resource allocation problem as a concave optimization problem that aims to maximize the asymptotic probability of correct selection (PCS) through the allocation policy for both input data collection and simulation, based on a moving-average estimator for aggregation of simulation outputs and its asymptotic normality. The two optimal policies are interdependent since they jointly affect the PCS. We derive the optimality equations to characterize the optimal policies and develop a fully sequential algorithm that demonstrates high efficiency through numerical experiments.

1 INTRODUCTION

Stochastic simulations are a common tool for comparing the performance of complex systems or designs in many applications. To assess the performance of a system or design, it is necessary to collect input data for estimating the input distribution, which captures the system's external randomness that affects the system's performance. Based on this estimation, multiple simulation replications are conducted to reduce stochastic error. However, collecting input data can be a challenging and time-consuming task, as there may be multiple data sources with limited availability for different input distributions. Input data collection can also incur high monetary costs. In addition, running simulations can also be a time-consuming process, particularly when the model being simulated is complex, and each replication may take several hours of computing time.

Consider a supply chain optimization problem where the manager wishes to compare various potential inventory policies across multiple products. Due to the large number of products, the various route planning, and the possibly long planning horizon, simulating the supply chain system can be a time-consuming process. In addition, collecting input data for the unknown input distribution (e.g., transit lead time, service time, or product demand) can also be laborious. For instance, to estimate the service rate, the service needs to be deployed to gather the realized service times. The amount of data required for accurate estimation depends on the length of time for collecting such data. Another example of a scenario where input data collection can be time and cost-intensive is a recommendation system. In order to obtain feedback for a certain type of service, the manager must conduct the service to collect actual feedback data. It is crucial to decide on an appropriate data collection policy to avoid redundant efforts and minimize the amount of time and money spent on data collection.

When carrying out input data collection and simulation, one approach is to wait until sufficient input data is collected to obtain a good estimation of the input distributions before proceeding with the simulation. However, given that input data collection and simulation require different resources, conducting both activities simultaneously may be more time-efficient. In this paper, we investigate a simultaneous resource allocation problem, where input data collection and simulation are conducted sequentially at the same time. This procedure involves multiple stages, starting with updating the estimation of input distributions with the collected data from the previous stage and updating the design performance with the simulation output from the previous stage. Based on the current estimates of input distributions and design performance, two optimal resource allocation policies for input data collection and simulation are computed and carried out with two stage-wise budgets. These two policies are interdependent, as they jointly influence the probability of correct selection (PCS).

1.1 Literature Review

The earliest work of ranking and selection (R&S) concerning input uncertainty (IU) assumes a fixed input data set and the simulations are run under fixed input distributions. As a result, no input data collection is involved and the IU is not reduced. One stream of work quantified the impact of IU and give a probability guarantee on the final selection. For instance, Corlu and Biller (2013), Corlu and Biller (2015) considered the subset selection with the existence of IU, where a subset of designs that contains the optimal design with a desired confidence is returned; Song and Nelson (2019) developed asymptotically valid concentration bounds to account for both IU and stochastic uncertainty (SU). Another line of work took a robust approach. Gao et al. (2017), Fan et al. (2020) considered finding the design with the best expected worst-case performance. Kim et al. (2021) also considered the worst-case performance but with a different criterion called the most probable best. Other robust methods include Zhu et al. (2020) and Zhou and Xie (2015), where they used the risk functional value at risk or conditional value at risk as a measure against IU.

While using a fixed input data set can simplify the simulation process by generating independent and identically distributed samples, real-world data often arrives in an online streaming fashion. To address this, researchers have explored running simulations with dynamically updated input distributions, which can reduce input uncertainty (IU). However, there has been limited research on R&S with streaming input data. Two recent studies by Wu et al. (2022) and Wang and Zhou (2022) have focused on the scenario where input data arrive in batches periodically, and the input distribution is updated at the end of each period. Wu et al. (2022) proposed a fully sequential algorithm that returns the optimal design with the desired confidence level, using the fixed confidence formulation. Meanwhile, Wang and Zhou (2022) considered the fixed budget setting and adjusted the optimal budget allocation policy at each period based on the current estimation of the input distribution.

In both Wu et al. (2022) and Wang and Zhou (2022), input data is obtained passively from the environment, and modelers can only adjust their simulation strategies to improve their decision quality. However, in some cases, input data can be actively collected, but at a cost. Wu and Zhou (2017), Xu et al. (2020), and Kim and Song (2022) have followed this approach, considering that input data collection and simulation share a joint budget. Wu and Zhou (2017) formulated joint budget allocation as a two-stage problem, where sufficient input data is collected in the first stage, and the remaining budget is allocated to run simulations in the second stage. Xu et al. (2020) considered scenarios where the cost of data collection is random and compared two cases where the simulation cost is negligible or not compared to the data collection cost. Kim and Song (2022) used the Bayesian posterior to estimate input distributions and the most probable best as their selecting criterion. However, the framework of joint budget allocation has its hidden nature in that simulations start after input data collection, ignoring the possibility that input data collection and simulation can be conducted simultaneously to save time and resources. In this paper, we formulate the resource allocation problem as a simultaneous resource allocation, where input data collection and simulation have their own budgets. We propose an efficient and fully sequential algorithm based on

the asymptotic behavior of PCS, and periodically exchange information to adjust their resource allocation policies. We demonstrate the effectiveness of our algorithm through numerical experiments.

2 PROBLEM STATEMENT

Suppose a set of K designs $\mathcal{I} = \{1, 2, \dots, K\}$ are given. Let $F_i^c(\cdot)$ denote the unknown input distributions which capture the real-world randomness. Let $X_i(F_i^c)$ denote the random performance of design i . The goal is to find the one with the largest expected performance $\mu_i(F_i^c) = \mathbb{E}[X_i(F_i^c)]$. Throughout the paper, we make the assumption that all designs share the same input distribution, that is, $F_i^c = F^c$. Notice this does not limit any practical usage since we can simply include all design-specific input distributions into the common input distribution. We make the following assumption on the commonly shared input distributions.

Assumption 1 (Parametric Input Distribution)

All designs share the same input distributions F^c , which contain S mutually independent input distributions that belong to a parametric family $\{F_\theta(\cdot) | \theta \in \Theta\}$ with known density $f_\theta(\cdot)$ and unknown parameters $\theta^c = (\theta_1^c, \theta_2^c, \dots, \theta_S^c)$.

Assumption 1 is common in the literature of R&S, by which the input distributions F_θ is then a product measure of F_{θ_s} , $s = 1, 2, \dots, S$. That is, $F_\theta = \prod_{s=1}^S F_{\theta_s}$. Accordingly, we use $X_i(\theta)$, $\mu_i(\theta)$ to denote the random and expected performance under input distributions F_θ . Since θ^c is unknown, one needs to collect the input data to get an estimator $\hat{\theta}$, under which simulation is run to get samples $X_i^r(\hat{\theta})$ for design i to estimate the expected performance. We make the following assumption about the input data.

Assumption 2 (unbiased estimator of input parameters)

1. For each θ_s^c , the input data $\zeta_{s,1}, \zeta_{s,2}, \dots$ are independent and identically distributed (i.i.d.) with distribution $F_{\theta_s^c}$.
2. For each θ_s^c , There exists a function D_s such that

$$\hat{\theta}_s^N = \frac{1}{N} \sum_{j=1}^N D_s(\zeta_{s,j})$$

is an unbiased estimator of θ_s^c .

2.1 Sequential Resource Allocation

Suppose input data are actively collected with a cost of c_s per unit for $s = 1, 2, \dots, S$. At the same time, designs are chosen to run simulations to avoid idling computing resources, with a cost of 1 per unit. From this perspective, consider a multi-stage resource allocation procedure. At each stage t , one computes a stage-wise allocation policy $\{m_{i,t}\}$ to allocate a stage-wise budget T_S (part of the total budget) to run the simulation for designs under $\hat{\theta}_t$, which is the estimation of the input distribution at the beginning of stage t , to get samples $X_i^1(\hat{\theta}_t), \dots, X_i^{m_{i,t}}(\hat{\theta}_t)$. At the same time, one also computes a stage-wise allocation policy $\{n_{s,t}\}$ to allocate a stage-wise budget T_I to collect input data and update the input parameter $\hat{\theta}_{t+1}$. The T_S and T_I measure how often we update the input distribution. However, such dynamically updated input distributions bring up a problem of how to aggregate the past simulation outputs, which are generated from different input distributions. We adopt a moving-average estimator which was introduced in Wu et al. (2022). Let $M_{i,t} = \sum_{\ell=1}^t m_{i,\ell}$ be the total budget assigned to design i up to stage t . Let $\eta \in [0, 1)$ be the drop rate and $t_{i,\eta} = \max \tau$ s.t. $M_{i,\tau} \leq M_{i,t}\eta$ be the stage before which the simulation outputs for design i are discarded. The moving-average estimator $\hat{\mu}_{i,t}$ is defined as

$$\hat{\mu}_{i,t} := [M_{i,t} - M_{i,t_{i,\eta}}]^{-1} \sum_{\ell=t_{i,\eta}+1}^t \sum_{r=1}^{m_{i,\ell}} X_i^r(\hat{\theta}_\ell). \quad (1)$$

That is, we only utilize a fraction of $(1 - \eta)$ of the current simulation outputs to estimate the design performance. The drop rate η reflects a trade-off between IU and SU. Setting η small leads to more simulation outputs and reduces SU, while setting η large aggregates simulation outputs generated under recent input distribution estimates, reducing IU. The selection of η will be discussed in Section 3.3.

The analysis of the PCS with finite samples is difficult, even without input uncertainty. Hence, we focus on the asymptotic behavior of PCS. We make the following mild assumption on the input distributions.

Assumption 3 For all $i \in \mathcal{I}$ and $s \in \{1, 2, \dots, S\}$,

- (i) $\Sigma_{D,s} := \text{Cov}(D_s(\zeta_{s,1}))$ exists.
- (ii) $\Sigma(\theta)$ exists and is continuous for all $\theta \in \Theta$, where $\Sigma_{ij}(\theta) = \text{Cov}[X_i(\theta), X_j(\theta)]$.
- (iii) $\mu_i(\cdot)$ is twice continuously differentiable in Θ .

Let $\delta_{ij}(\theta) = \mu_i(\theta) - \mu_j(\theta)$ be the performance difference between design i and j under input parameter θ , $\hat{\delta}_{ij,t} = \hat{\mu}_{i,t} - \hat{\mu}_{j,t}$ be the sample approximation of $\delta_{ij}(\theta^c)$ and $\sigma_i^2(\theta) = \text{Var}(X_i(\theta))$ be the performance variance of design i under θ . The next theorem establishes the asymptotic normality of the moving-average estimator $\hat{\mu}_{i,t}$, which is a variant of Theorem 3 in Wu et al. (2022).

Theorem 1 (Asymptotic Normality) Suppose $\{n_{s,t}\}$ and $\{m_{i,t}\}$ are uniformly bounded. Furthermore, there exist positive constants $\bar{n}_s$ and $\bar{m}_i$ such that $N_s(t)/t \rightarrow \bar{n}_s$ and $M_i(t)/t \rightarrow \bar{m}_i$ as $t \rightarrow \infty$ almost surely.

$$\sqrt{t} [\hat{\delta}_{ij,t} - \delta_{ij}(\theta^c)] \Rightarrow \mathcal{N}(0, \tilde{\sigma}_{ij}), \quad \text{as } t \rightarrow \infty \quad \text{almost surely,}$$

where $\tilde{\sigma}_{ij}^2 = \lambda_{I,\eta} \sum_{s=1}^S \bar{n}_s^{-1} \partial_{\theta_s} \delta_{ij}(\theta^c)^\top \Sigma_{D,s} \partial_{\theta_s} \delta_{ij}(\theta^c) + \lambda_{S,\eta} \bar{m}_i^{-1} \sigma_i^2(\theta^c) + \lambda_{S,\eta} \bar{m}_j^{-1} \sigma_j^2(\theta^c)$, ∂_{θ_s} is the partial derivative taken with respect to θ_s , $\lambda_{I,\eta} = \left(\frac{2}{1-\eta} + \frac{2\eta \ln \eta}{(1-\eta)^2} \right)$ and $\lambda_{S,\eta} = \frac{1}{1-\eta}$.

We now study the asymptotic behavior of PCS, which is defined as $\text{PCS} := \mathbb{P}(\hat{\mu}_{b,t} \geq \max_{i \neq b} \hat{\mu}_{i,t})$, where $b = \arg \max_{i \in \mathcal{I}} \mu_i(\theta^c)$ is the best design. Notice for $\forall i \neq b$,

$$\mathbb{P}(\hat{\mu}_{b,t} \leq \hat{\mu}_{i,t}) \geq \text{PCS} \geq 1 - \sum_{i \neq b} \mathbb{P}(\hat{\mu}_{b,t} \leq \hat{\mu}_{i,t}).$$

Using the asymptotic normality in Theorem 1, we have approximately

$$\hat{\mu}_{b,t} - \hat{\mu}_{i,t} \sim \mathcal{N} \left(\delta_{bi}(\theta^c), \frac{\tilde{\sigma}_{bi}}{t} \right)$$

For X following a standard normal distribution and $x > 0$,

$$\frac{x}{\sqrt{2\pi(x^2 + 1)}} \exp\left(-\frac{x^2}{2}\right) \leq \mathbb{P}(X < x) \leq \frac{1}{\sqrt{2\pi}x} \exp\left(-\frac{x^2}{2}\right).$$

Hence, with the Gaussian approximation by Theorem 1,

$$\mathbb{P}(\hat{\mu}_{b,t} \leq \hat{\mu}_{i,t}) \geq \frac{\sqrt{t} \delta_{bi}(\theta^c) / \tilde{\sigma}_{bi}(\theta^c)}{\sqrt{2\pi} (\delta_{bi}^2(\theta^c) / \tilde{\sigma}_{bi}^2(\theta^c) t + 1)} \exp\left(-\frac{\delta_{bi}^2(\theta^c)}{2\tilde{\sigma}_{bi}^2} t\right), \quad (2)$$

and

$$\mathbb{P}(\hat{\mu}_{b,t} \leq \hat{\mu}_{i,t}) \leq \frac{\delta_{bi}(\theta^c)}{\sqrt{2\pi t} \tilde{\sigma}_{bi}(\theta^c)} \exp\left(-\frac{\delta_{bi}^2(\theta^c)}{2\tilde{\sigma}_{bi}^2} t\right). \quad (3)$$

With (2) and (3),

$$-\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}(\hat{\mu}_{b,t} \leq \hat{\mu}_{i,t}) = \frac{1}{2} \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2}.$$

Furthermore, we have for PCS

$$-\lim_{t \rightarrow \infty} \frac{1}{t} \log \text{PCS} = \min_{i \neq b} \frac{1}{2} \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2}.$$

In the long run, the budget allocated to a certain design i (or input distribution s) is approximately $t\bar{m}_i$ (or $t\bar{n}_s$). Ignoring the minor issue of $t\bar{m}_i$ and $t\bar{n}_s$ not being integers, we relax $\bar{n}_s$ and $\bar{m}_i$ to be continuous variables. We now aim to solve the following problem,

$$\begin{aligned} \max_{\bar{n}_s, \bar{m}_i \geq 0, z} \quad & z \\ \text{s.t.} \quad & \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2} \geq z \quad \forall i \neq b \\ & \sum_{s=1}^S c_s \bar{n}_s = T_I \\ & \sum_{i=1}^K \bar{m}_i = T_S \end{aligned} \tag{4}$$

To solve the optimization problem (4), we further make the following assumption.

Assumption 4 (Impact of input uncertainty)

For each $1 \leq s \leq S$, there exists $i \neq b$, such that $\partial_{\theta_s} \delta_{bi}(\theta^c) \neq 0$.

Assumption 4 guarantees that, for each input distribution considered, there exists at least one sub-optimal design whose performance is affected by the input distribution in a way that differs from the optimal design. In other words, an inaccurate estimation of the input parameter θ_s^c can lead to difficulty in distinguishing the sub-optimal design i from the optimal design b . Suppose, for instance, that $\delta_{bi}(\theta)$ is a constant for all possible values of θ (in which case $\partial_{\theta_s} \delta_{bi}(\theta^c) = 0$). In such a case, the difference between any two designs would remain the same no matter what the input parameter is, and the input estimation error would have no effect on the selection process. We exclude such cases and focus only on the input distributions that are truly relevant.

The following lemma guarantees Problem (4) is a convex optimization problem, and hence, we can adapt the Karush-Kuhn-Tucker (KKT) condition to derive the optimality conditions.

Lemma 1 Denote by $G_i(\bar{m}_b, \bar{m}_i, \bar{\mathbf{n}}) = \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2}$ the rate function for sub-optimal design i . Suppose Assumption 1 - 4 hold. Then G_i is increasing and concave in $\bar{m}_b, \bar{m}_i$ and $\bar{\mathbf{n}}$ for $\bar{\mathbf{m}}, \bar{\mathbf{n}} \geq 0$.

Proof. It suffices to show the concavity of the function for $x \in \mathbb{R}_+^n$ with form $f(x) = 1/(\sum_{i=1}^n \frac{a_i}{x_i})$, where $a_i > 0$ for $i = 1, 2, \dots, n$. The proof then follows from Lemma 1 in Wang and Zhou (2022). $\square$

By Lemma 1, the optimal solution for (4) cannot be obtained on the boundary, i.e., all $\bar{n}_s$ and $\bar{m}_i$ should be strictly positive. Otherwise, we have $\tilde{\sigma}_{bi} = \infty$ for some i and $G_i = 0$, since there is either input uncertainty, simulation uncertainty, or both unaddressed.

Theorem 2 Under Assumption 1-4, Problem (4) has an optimal solution satisfying

$$1. \text{ (Local Balance) } \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2} = \frac{\delta_{bj}^2(\theta^c)}{\tilde{\sigma}_{bj}^2} \quad \forall i \neq j \neq b \quad (5)$$

$$2. \text{ (Global Balance) } \bar{m}_b^2 = \sigma_b^2(\theta^c) \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} \quad (6)$$

$$3. \text{ (IU Balance) } \frac{1}{c_s \bar{n}_s^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} g(i, s) = \frac{1}{c_{s'} \bar{n}_{s'}^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} g(i, s') \quad \forall s \neq s' \quad (7)$$

where $g(i, s) = \partial_{\theta_s} \delta_{bi}(\theta^c)^\top \Sigma_{D,s} \partial_{\theta_s} \delta_{bi}(\theta^c)$.

Remark 1. The “Local Balance” equation and the “Global Balance” equation, two of the three optimality equations, were derived in a similar form in previous works Glynn and Juneja (2004), Chen and Ryzhov (2022) without input uncertainty and with passively obtained input data in Wang and Zhou (2022). Following tradition, we adopt these names for the two equations. In Section 3.2, we demonstrate that the “Local Balance” equation can be used to adjust the allocation policy of the simulation budget between two sub-optimal designs, while the “Global Balance” equation can be used to tune the allocation policy of the simulation budget between the optimal design and other designs, as in Chen and Ryzhov (2022) and Wang and Zhou (2022). However, the rate function in the “Local Balance” equation now includes both the allocation policy for the simulation budget and the allocation policy for input data collection since they jointly determine the variance term $\tilde{\sigma}_{bi}^2$ as defined in Theorem 1.

Remark 2. The third optimality equation, the “IU Balance” equation (7), can be used to adjust the allocation policy for input data allocation to reduce IU. Equation (7) suggests that the amount of input data allocated to a particular distribution $F_{\theta_s^c}$ depends on three factors in addition to cost: the simulation effort $\bar{m}_i$, the simulation noise $\sigma_i^2(\theta^c)$ for design $i \neq b$, and its sensitivity $g(i, s)$ to the input parameter θ_s^c . When the function $g(i, s)$ is large, the performance of design i is more sensitive to the input parameter θ_s^c . If at the same time, $\frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)}$ is also large, which means the simulation error of design i is small, more effort should be devoted to reducing IU. To provide a clear example, let us consider a case where the number of independent input distributions equals to the number of designs, i.e., $S = K$. Moreover, each design i is only affected by the i th input distribution $F_{\theta_i^c}$, i.e., $\nabla_{\theta_s} \mu_i(\theta^c) = \begin{cases} = 0, & \text{if } s \neq i \\ > 0, & \text{if } s = i \end{cases}$. Then, we

have $\bar{n}_i \propto \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} g(i, s)$ for $i \neq b$.

Remark 3. It is worth noting that in Wang and Zhou (2022), there exists a different equation called the “Input Balance” equation, which differs from the “IU Balance” equation discussed in this paper. In their scenario, the input data is obtained passively, and the simulation is conducted on a design and fixed input realization pair. Since there is no randomness generated from the input distribution in the simulation output, the “Input Balance” equation is utilized to adjust the allocation policy of the simulation budget toward different input realizations for a fixed design. On the other hand, in this paper, the “IU Balance” equation is utilized to adjust the allocation policy for input data collection, as will be demonstrated in Section 3.2.

Proof. For positive solution $\{\bar{n}_s, \bar{m}_i\}$, apply the KKT conditions and we obtain

$$1 - \sum_{i \neq b} \mu_i = 0 \quad (8)$$

$$c_s \lambda - \sum_{i \neq b} \frac{\mu_i}{2} \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^3} \frac{1}{\bar{n}_s^2} \lambda_{I,\eta} \partial_{\theta_s} \delta_{bi}^\top(\theta^c) \Sigma_{D,s} \partial_{\theta_s} \delta_{bi}(\theta^c) = 0 \quad \forall 1 \leq s \leq S \quad (9)$$

$$\gamma - \frac{\mu_i}{2} \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^3} \frac{1}{\bar{m}_i^2} \lambda_{S,\eta} \sigma_i^2(\theta^c) = 0 \quad \forall i \neq b \quad (10)$$

$$\gamma - \sum_{i \neq b} \frac{\mu_i}{2} \frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^3} \frac{1}{\bar{m}_b^2} \lambda_{S,\eta} \sigma_b^2(\theta^c) = 0 \quad (11)$$

$$\mu_i \left(\frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2} - z \right) = 0 \quad \forall i \neq b \quad (12)$$

For necessity, from (8) we have there exists $i \neq b$ such that $\mu_i > 0$. For this i , from (10) we obtain $\gamma > 0$. Hence for all $i \neq b$, $\mu_i = \frac{2\gamma \tilde{\sigma}_{bi}^3 \bar{m}_i^2}{\lambda_{S,\eta} \sigma_i^2(\theta^c) \delta_{bi}^2(\theta^c)} > 0$. Then, from (12), we have $\frac{\delta_{bi}^2(\theta^c)}{\tilde{\sigma}_{bi}^2} = \frac{\delta_{bj}^2(\theta^c)}{\tilde{\sigma}_{bj}^2} \quad \forall i \neq j \neq b$, which proves (5). Substituting $\mu_i = \frac{2\gamma \tilde{\sigma}_{bi}^3 \bar{m}_i^2}{\lambda_{S,\eta} \sigma_i^2(\theta^c) \delta_{bi}^2(\theta^c)}$ in (11), we obtain $\bar{m}_b^2 = \sigma_b^2(\theta^c) \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)}$, which proves (6). Also, substituting μ_i in (9), we obtain $\forall s \neq s'$,

$$\frac{1}{c_s \bar{n}_s^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} \partial_{\theta_s} \delta_{bi}(\theta^c)^\top \Sigma_{D,s} \partial_{\theta_s} \delta_{bi}(\theta^c) = \frac{1}{c_{s'} \bar{n}_{s'}^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} \partial_{\theta_{s'}} \delta_{bi}(\theta^c)^\top \Sigma_{D,s'} \partial_{\theta_{s'}} \delta_{bi}(\theta^c),$$

which proves (7).

For sufficiency, it suffices to show KKT conditions are satisfied if the three optimality conditions (5) - (7) are satisfied. Let $i_0 \neq b$ be some fixed sub-optimal design. Let $z = \frac{\delta_{bi_0}^2(\theta^c)}{\tilde{\sigma}_{bi_0}^2}$, $\mu_i = \frac{\tilde{\sigma}_{bi}^3 \bar{m}_i^2}{\sigma_i^2(\theta^c) \delta_{bi}^2(\theta^c)} \left/ \sum_{j \neq b} \frac{\tilde{\sigma}_{bj}^3 \bar{m}_j^2}{\sigma_j^2(\theta^c) \delta_{bj}^2(\theta^c)} \right.$, $\lambda = \frac{\lambda_{I,\eta}}{2c_1 \bar{n}_s^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} \partial_{\theta_s} \delta_{bi}^\top(\theta^c) \Sigma_{D,s} \partial_{\theta_s} \delta_{bi}(\theta^c) \left/ \sum_{j \neq b} \frac{\tilde{\sigma}_{bj}^3 \bar{m}_j^2}{\sigma_j^2(\theta^c) \delta_{bj}^2(\theta^c)} \right.$ and $\gamma = 1 \left/ \sum_{j \neq b} \frac{2\tilde{\sigma}_{bj}^3 \bar{m}_j^2}{\lambda_{S,\eta} \sigma_j^2(\theta^c) \delta_{bj}^2(\theta^c)} \right.$. Then one can verify all the KKT conditions are satisfied. $\square$

3 ALGORITHM

3.1 Parameter Estimation

To design an algorithm, there are several unknown parameters that need to estimate. They include

1. The true input parameter θ^c and its covariance matrix $\Sigma_{D,s}$ for $s = 1, 2, \dots, S$.
2. The true expected performance $\mu_i(\theta^c)$ and variance $\sigma_i^2(\theta^c)$.
3. The gradient $\nabla \mu_i(\theta^c) = (\partial_{\theta_1} \mu_i(\theta^c), \partial_{\theta_2} \mu_i(\theta^c), \dots, \partial_{\theta_S} \mu_i(\theta^c))^\top$.

For θ^c and $\Sigma_{D,s}$, by Assumption 1 and 2, we can use the sample average and sample variance, respectively. Let $\hat{\theta}_{s,t} = \frac{1}{N_{s,t}} \sum_{\ell=1}^{N_{s,t}} D_s(\zeta_{s,\ell})$ and $\hat{\Sigma}_{D,s,t} = \frac{1}{N_{s,t}-1} \sum_{\ell=1}^{N_{s,t}} (D_s(\zeta_{s,\ell}) - \hat{\theta}_{s,t})(D_s(\zeta_{s,\ell}) - \hat{\theta}_{s,t})^\top$. For $\mu_i(\theta^c)$, we use the moving-average estimator defined in (1). Similarly, we estimate $\sigma_i^2(\theta^c)$ as

$$\hat{\sigma}_{i,t}^2 = \frac{1}{M_{i,t} - M_{i,t_i,\eta} - 1} \sum_{\ell=t_i,\eta+1}^t \sum_{r=1}^{m_{i,\ell}} (X_i^r(\hat{\theta}_\ell) - \hat{\mu}_{i,t})^2.$$

Finally, for $\nabla\mu_i(\theta^c)$, one approach is to use important sampling. Suppose $\forall\theta \in \Theta$, f_θ has the same input support Ω . That is, $\forall\xi \in \Omega$, $f_\theta(\xi) > 0$ for all $\theta \in \Theta$. Furthermore, suppose when simulating $X_i^r(\hat{\theta}_\ell)$, one first generates $\xi_{i,\ell}^r \sim f_{\hat{\theta}_\ell}$, then runs the simulation under $\xi_{i,\ell}^r$ and obtains the simulation output $X_i|\xi_{i,\ell}^r = X_i^r(\hat{\theta}_\ell)$. Since

$$\nabla\mu_i(\theta^c) = \nabla\theta \mathbb{E}_{\xi \sim f_{\theta^c}} [X_i|\xi] = \nabla\theta \int_{\xi \in \Omega} f_{\theta^c}(\xi) X_i|\xi d\xi,$$

assuming the interchangeability of the integration and the gradient, we have

$$\nabla\mu_i(\theta^c) = \int_{\xi \in \Omega} \nabla\theta f_{\theta^c}(\xi) X_i|\xi d\xi = \mathbb{E}_{\xi \sim f_\theta} \left[\frac{\nabla\theta f_{\theta^c}(\xi)}{f_\theta(\xi)} X_i|\xi \right], \quad \forall\theta \in \Theta$$

An analogy can be drawn to design the following moving-average important sampling estimator by replacing θ^c and θ with $\hat{\theta}_\ell$ for $\ell = t_{i,\eta} + 1, \dots, t$.

$$\widehat{\nabla}\mu_{i,t} = \frac{1}{M_{i,t} - M_{i,t_{i,\eta}}} \sum_{\ell=t_{i,\eta}+1}^t \sum_{r=1}^{m_{i,\ell}} \frac{\nabla f_{\hat{\theta}_\ell}(\xi_{i,\ell}^r)}{f_{\hat{\theta}_\ell}(\xi_{i,\ell}^r)} X_i^r(\hat{\theta}_\ell).$$

Denote by $\widehat{g}(i, s)$ the estimate of $g(i, s)$ with replacement of the unknown parameters with their estimates. It is worth pointing out that to estimate the unknown parameters, we use all the simulation outputs after stage $t_{i,\eta}$ and do not run any extra simulations for the purpose of efficient sampling.

3.2 Balancing Approach: A Fully Sequential Procedure

In this section, we propose a fully sequential algorithm that uses a so-called "Balancing" approach. "Balancing" means reducing the gap between two sides of the optimality equations. To see how this works, first notice in (5)-(7), if we multiply an optimal allocation policy $\{\bar{n}_s, \bar{m}_i\}$ by any constant C , then the three optimality equations remain valid. Then, at stage t , We can substitute $\bar{n}_s$ and $\bar{m}_i$ with $N_{s,t}$ and $M_{i,t}$ in (5)-(7), respectively. For input data selection, we can collect input data for input distribution F_{θ_s} , where s maximizes

$$\frac{1}{c_s N_{s,t}^2} \sum_{i \neq b} \frac{M_{i,t}^2}{\sigma_i^2(\theta^c)} g(i, s).$$

By doing so, we decrease the maximal value of $\frac{1}{c_s \bar{n}_s^2} \sum_{i \neq b} \frac{\bar{m}_i^2}{\sigma_i^2(\theta^c)} g(i, s)$ and thus reduce the biggest gap in equations (7). Similarly for design selection, if $M_{b,t}^2 < \sigma_b^2(\theta^c) \sum_{i \neq b} \frac{M_{i,t}^2}{\sigma_i^2(\theta^c)}$, we simulate for design b . Otherwise, we pick a sub-optimal design i that minimizes $\frac{\delta_{bi}^2(\theta^c)}{\bar{\sigma}_{bi}^2}$, where $\bar{m}_i$ and $\bar{n}_s$ are substituted with $M_{i,t}$ and $N_{s,t}$, respectively. This procedure is shown in Algorithm 1.

3.3 Boosting through η

In the previous section, we derive the asymptotic optimality conditions by maximizing the rate functions $\min_{i \neq b} G_i$ through the (asymptotic) budget allocation policy $\{\bar{n}_s, \bar{m}_i\}$. Although the rate function is also influenced by the drop rate η , we cannot give a relatively simple expression of the optimal solution. However, given an allocation policy $\{\bar{n}_s, \bar{m}_i\}$, we can adjust the drop rate for the sake of increasing $\min_{i \neq b} G_i$. For a fixed $i \neq b$, maximizing G_i over η is equivalent to minimizing

$$\tilde{\sigma}_{bi} \propto \frac{\kappa}{1-\eta} + \frac{\eta \ln \eta}{(1-\eta)^2},$$

Algorithm 1 Simultaneous Resources Allocation

1: **Input:** Number of designs K , number of input distributions S , input data collection cost c_s , $s = 1, \dots, S$, stage-wise budget T_I and T_S , initial budget n_0 and m_0 , drop rate η , maximal stage T .

2: **Initialize** Collect n_0 input data for each input distribution and run m_0 simulation replications for each design. Estimate $\hat{\theta}_0$, $\hat{\Sigma}_{D,s,0}$, $\hat{\mu}_{i,0}$, $\hat{\sigma}_{i,0}$ and $\hat{\nabla}\mu_{i,0}$ as in Section 3.1. $\hat{b} = \arg \max_i \hat{\mu}_{i,0}$, $N_{s,0} = n_0$, $\forall s$, $M_{i,0} = m_0$, $\forall i$. Set $t = 0$.

3: **for** $t = 1 : T$ **do**

4: $M_{i,t} = M_{i,t-1}$, $N_{s,t} = N_{s,t-1}$, $\forall 1 \leq i \leq K$, $1 \leq s \leq S$.

5: Do the following two **WHILE** loops for **Data Collection** and **Simulation** simultaneously.

6: **while** $\sum_{s=1}^S c_s N_{s,t} < t \times T_I$ **do**

7: $s^* = \arg \max_s \frac{1}{c_s N_{s,t}} \sum_{i \neq \hat{b}} \frac{M_{i,t}^2}{\hat{\sigma}_{i,t}^2} \hat{g}(i, s)$.

8: Collect 1 input data for distribution $F_{\theta_{s^*}}$, $N_{s^*,t} = N_{s^*,t} + 1$. Update $\hat{\theta}_s$ and $\hat{\Sigma}_{D,s,t}$.

9: **end while**

10: **while** $\sum_{i=1}^K M_{i,t} < t \times T_S$ **do**

11: **if** $M_{b,t}^2 - \hat{\sigma}_{b,t}^2 \sum_{i \neq \hat{b}} \frac{M_{i,t}^2}{\hat{\sigma}_{i,t}^2} < 0$ **then**

12: Simulate 1 sample for design $\hat{b}$. $M_{b,t} = M_{b,t} + 1$. Update $\hat{\mu}_{b,t}$, $\hat{\sigma}_{b,t}$ and $\hat{\nabla}\mu_{b,t}$.

13: **else**

14: Choose $i^* = \arg \min_{i \neq \hat{b}} \frac{(\hat{\mu}_{b,t} - \hat{\mu}_{i,t})^2}{\frac{\lambda_{I,\eta}}{\lambda_{S,\eta}} \sum_{s=1}^S \frac{\hat{g}(i,s)}{N_{s,t}} + \frac{\hat{\sigma}_{i,t}^2}{M_{i,t}} + \frac{\hat{\sigma}_{b,t}^2}{M_{b,t}}}$. Simulate 1 sample for design i^* . Update $\hat{\mu}_{i^*,t}$, $\hat{\sigma}_{i^*,t}$ and $\hat{\nabla}\mu_{i^*,t}$.

15: **end if**

16: $\hat{b} = \arg \max_i \hat{\mu}_{i,t}$.

17: **end while**

18: **end for**

19: **Output:** $\hat{b} = \arg \max_i \hat{\mu}_{i,T}$.

where $\propto$ means proportional to and $\kappa = 1 + \frac{\bar{m}_b^{-1} \sigma_b^2(\theta^c) + \bar{m}_i^{-1} \sigma_i^2(\theta^c)}{2 \sum_{s=1}^S \bar{n}_s^{-1} \partial_{\theta_s} \delta_{ij}(\theta^c)^\top \Sigma_{D,s} \partial_{\theta_s} \delta_{ij}(\theta^c)}$. The following lemma which shows we can adjust η somehow by optimizing a convex function.

Lemma 2 (Wu et al. (2022)) Let $h(\eta) := \frac{\kappa}{1-\eta} + \frac{\eta \ln \eta}{(1-\eta)^2}$, $\kappa \geq 1$. $h(\cdot)$ is strictly convex in $\eta \in (0, 1)$.

Thanks to the convexity property established in Lemma 2, determining the optimal value of η to maximize G_i is a straightforward task. In Algorithm 1, we incorporate this insight by adjusting η at the end of each stage. Specifically, at the conclusion of stage t , we identify the sub-optimal design i^* for which G_{i^*} is the smallest among all designs except for the baseline b . Next, we employ a numerical optimization method such as gradient descent to solve for the value of η that maximizes G_{i^*} . For more details on this process, see Algorithm 2.

4 CONSISTENCY OF ALGORITHM 1

Lemma 3 Assume there exists $\bar{x}, \bar{F} > 0$, such that $|X_i(\theta)| \leq \bar{x}$ almost surely and $\mathbb{E} \left[\left(\frac{\nabla_{\theta} f_{\theta}(\xi)}{f_{\theta}(\xi)} \right)^2 \right] \leq \bar{F} < \infty$ for all $\theta \in \Theta$ and $1 \leq i \leq K$. Suppose $N_{s,t} \rightarrow \infty$ and $M_{i,t} \rightarrow \infty$ as $t \rightarrow \infty$ almost surely for all $1 \leq s \leq S, 1 \leq i \leq K$, then (a) $\hat{\mu}_{i,t} \rightarrow \mu_i(\theta^c)$, (b) $\hat{\sigma}_{i,t}^2 \rightarrow \sigma_i^2(\theta^c)$ and (c) $\hat{\nabla}\mu_{i,t} \rightarrow \nabla_{\theta} \mu_i(\theta^c)$ as $t \rightarrow \infty$ almost surely.

Algorithm 2 Boosting through η

-
- 1: Run Algorithm 1 until Line 17
 - 2: Let $i^* = \arg \min_{i \neq b} \frac{(\hat{\mu}_{b,t} - \hat{\mu}_{i,t})^2}{\frac{\lambda_{I,\eta}}{\lambda_{S,\eta}} \sum_{s=1}^S \frac{\hat{g}(i,s)}{N_{s,t}} + \frac{\hat{\sigma}_{i,t}^2}{M_{i,t}} + \frac{\hat{\sigma}_{b,t}^2}{M_{b,t}}}$.
 - 3: Solve $\eta = \arg \min_{\eta} \lambda_{I,\eta} \sum_{s=1}^S \frac{\hat{g}(i,s)}{N_{s,t}} + \lambda_{S,\eta} (\frac{\hat{\sigma}_{i,t}^2}{M_{i,t}} + \frac{\hat{\sigma}_{b,t}^2}{M_{b,t}})$.
 - 4: Update $\hat{\mu}_{i,t}$, $\hat{\sigma}_{i,t}^2$ and $\widehat{\nabla} \mu_{i,t}$ for $i = 1, \dots, K$.
 - 5: Continue running Algorithm 1.
-

Proof. Due to the page limit, we omit the proof details. The idea is to partition the estimation error into two parts, accounting for SU and IU separately. Specifically, to prove (a), rewrite

$$\begin{aligned} \hat{\mu}_{i,t} - \mu_i(\theta^c) &= \frac{1}{M_{i,t} - M_{i,t,\eta}} \sum_{\tau=t_{i,\eta}+1}^t \sum_{\ell=1}^{m_{i,\tau}} [X_{i,\ell}(\hat{\theta}_\tau) - \mu_i(\hat{\theta}_\tau)] \\ &\quad + \frac{1}{M_{i,t} - M_{i,t,\eta}} \sum_{\tau=t_{i,\eta}+1}^t m_{i,\tau} [\mu_i(\hat{\theta}_\tau) - \mu_i(\theta^c)]. \end{aligned}$$

By the Strong Law of Large Number for Martingale difference sequence (e.g., see Csörgő (1968)), we can prove the first term converges to zero almost surely. The convergence for the second term holds due to the convergence of $\hat{\theta}_t$ and the fact that μ_i is continuous in θ . We can prove (b) and (c) in a similar way. $\square$

With Lemma 3, we can prove the consistency of Algorithm 1 by showing $N_{s,t} \rightarrow \infty$ and $M_{i,t} \rightarrow \infty$ as $t \rightarrow \infty$, which is stated in the following Theorem 3. The proof is omitted due to page limit.

Theorem 3 (Consistency) Algorithm 1 selects the optimal design almost surely as $T \rightarrow \infty$.

5 NUMERICAL EXPERIMENT

In this section, we carry out some numerical experiments to test the efficiency of the proposed algorithms.

5.1 Comparison Baselines

1. Simultaneous Resource Allocation (SRA). The proposed Algorithm 1.
2. Simultaneous Resource Allocation with adaptive η (SRA- η). The proposed Algorithm 2.
3. Equal Allocation (EA). Equally allocated the simulation budget to all designs and the budget for input data collection to all input distributions.
4. Equal Allocation + OCBA (EA_OCBA). Equally allocated the budget for input data collection to all input distributions, and implement OCBA for simulation budget allocation, where all the simulation outputs are treated as i.i.d. data.
5. Joint Budget Allocation (JBA). The two-stage joint budget allocation procedure in Wu and Zhou (2017), where the input data collection and simulation share a common budget and the simulation is conducted after input data collection. To implement the algorithm in our setting, set the stage-wise budget $T_I = T_S$ after re-scaling c_s and the total joint budget is $T \cdot T_I$.

5.2 Simulation Example

Service Comparison with random return. Consider one wants to compare K types of services within a time period τ . Assume customers arrive as a Poisson process with unknown arrival rate $\theta_{i,1}^c$. For each customer serviced, a return r_i is obtained, which follows a normal distribution with unknown mean $\theta_{i,2}^c$

and variance σ^2 . The return and arrival time are independent. The goal is then to find service b such that $b = \arg \max_i \mathbb{E}[D_i r_i]$, where $D_i \sim \text{Poi}(\theta_{i,1}^c \tau)$ and $r_i \sim \mathcal{N}(\theta_{i,2}^c, \sigma^2)$.

5.3 Results

For implementation details, we set the number of designs $K = 10$, time period $\tau = 1$, customer arrival rate $\theta_{i,1}^c = 0.5 \cdot (i + 1)$, average reward $\theta_{i,2}^c = \frac{K}{2} - |i - \frac{K}{2}|$ and reward variance $\sigma^2 = 1$. The maximal stages $T = 500$, stage-wise simulation budget $T_S = 10$, stage-wise budget for input data collection $T_I = 10$, input data collection cost $c_s = 2, \forall s$. Figure 1 shows the Average PCS obtained of different algorithms at each stage by running all algorithms for 200 times.




Figure 1: Performance comparison with different drop rates.

1. The results shown in Figure 1 indicate that both SRA and SRA- η perform well in both scenarios, with SRA- η performing the best. On the other hand, EA and EA_OCBA perform much worse than other algorithms, highlighting the importance of considering both IU and SU.
2. In terms of comparing the fully sequential procedures, SRA and SRA- η , with the two-stage procedure JBA, JBA reaches a high final PCS but has a relatively low intermediate PCS due to not running any simulations before collecting all input data. In contrast, SRA and SRA- η maintain high intermediate PCS and reach high final PCS.
3. Regarding the impact of the drop rate η , the experiment shows that for a drop rate of $\eta = 0.1$, SRA reaches a similar final PCS as JBA, which is lower than the PCS obtained by SRA- η . For $\eta = 0.25$, SRA outperforms JBA and performs more similarly to SRA- η . During the experiment, we found that the optimal choice of η for SRA is between 0.25 and 0.3.
4. It is worth noting that while SRA- η automatically adjusts η to obtain higher PCS, it is more computationally expensive. At each stage, the adjustment of η requires solving a convex optimization problem. Furthermore, Figure 1 indicates that SRA- η reaches the lowest PCS at very early stages due to large estimation errors for any unknown parameter, which may lead to a bad choice of η and further enlarge the estimation error.

REFERENCES

- Chen, Y., and I. O. Ryzhov. 2022. “Balancing Optimal Large Deviations in Sequential Selection”. *Management Science* 69(6):3457–3473.
- Corlu, C. G., and B. Biller. 2013. “A Subset Selection Procedure under Input Parameter Uncertainty”. In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 463–473. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

- Corlu, C. G., and B. Biller. 2015. "Subset Selection for Simulations Accounting for Input Uncertainty". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, V. W. Chan, I.-C. Moon, T. M. Roeder, C. Macal, and M. D. Rossetti, 437–446. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Csörgő, M. 1968. "On the Strong Law of Large Numbers and the Central Limit Theorem for Martingales". *Transactions of the American Mathematical Society* 131(1):259–275.
- Fan, W., L. J. Hong, and X. Zhang. 2020. "Distributionally Robust Selection of the Best". *Management Science* 66(1):190–208.
- Gao, S., H. Xiao, E. Zhou, and W. Chen. 2017. "Robust Ranking and selection with Optimal Computing Budget Allocation". *Automatica* 81:30–36.
- Glynn, P., and S. Juneja. 2004. "A Large Deviations Perspective on Ordinal Optimization". In *Proceedings of the 2004 Winter Simulation Conference, 2004.*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 585–593. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Kim, K.-K., T. Kim, and E. Song. 2021. "Selection of the Most Probable Best under Input Uncertainty". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Kim, T., and E. Song. 2022. "Optimizing Input Data Acquisition for Ranking and Selection: A View Through the Most Probable Best". In *Proceedings of the 2022 Winter Simulation Conference*, edited by B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C. Corlu, L. Lee, E. Chew, T. Roeder, and P. Lendermann, 2258–2269. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Song, E., and B. L. Nelson. 2019. "Input-output Uncertainty Comparisons for Discrete Optimization via Simulation". *Operations Research* 67(2):562–576.
- Wang, Y., and E. Zhou. 2022. "Data-driven Optimal Computing Budget Allocation under Input Uncertainty". *arXiv preprint arXiv:2209.11809*.
- Wu, D., Y. Wang, and E. Zhou. 2022. "Data-Driven Ranking and Selection Under Input Uncertainty". *Operations Research* 0:1–15.
- Wu, D., and E. Zhou. 2017. "Ranking and Selection under Input Uncertainty: A Budget Allocation Formulation". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan, G. Z. A. D'Ambrogio, N. Mustafee, G. Wainer, and E. Page, 2245–2256. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Xu, J., Z. Zheng, and P. W. Glynn. 2020. "Joint Resource Allocation for Input Data Collection and Simulation". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 2126–2137. 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.
- Zhu, H., T. Liu, and E. Zhou. 2020. "Risk Quantification in Stochastic Simulation under Input Uncertainty". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 30(1):1–24.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support by the Air Force Office of Scientific Research under Grants FA9550-19-1-0283 and FA9550-22-1-0244 and the National Science Foundation under Grant NSF-DMS2053489.

AUTHOR BIOGRAPHIES

YUHAO WANG is a Ph.D. student at the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. He received his B.S. degree from the Department of Mathematics at Nanjing University, China, in 2021. His research interests include simulation and stochastic optimization and reinforcement learning. His email address is yuhaowang@gatech.edu.

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