Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Data-Driven Ranking and Selection Under Input Uncertainty

Di Wu, Yuhao Wang, Enlu Zhou

To cite this article:
Di Wu, Yuhao Wang, Enlu Zhou (2022) Data-Driven Ranking and Selection Under Input Uncertainty. Operations Research

Published online in Articles in Advance 18 Oct 2022

. https://doi.org/10.1287/opre.2022.2375

https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–15
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Data-Driven Ranking and Selection Under Input Uncertainty

Di Wu,[a] Yuhao Wang,[b] Enlu Zhou[b,*]

[a] Amazon Web Services, Seattle, Washington 98109; [b] School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332

*Corresponding author

Contact: woody10074026@gmail.com, https://orcid.org/0000-0003-3911-0955 (DW); yuhaowang@isye.gatech.edu, https://orcid.org/0000-0002-9863-273X (YW); enlu.zhou@isye.gatech.edu, https://orcid.org/0000-0001-5399-6508 (EZ)

Received: March 29, 2021
Revised: July 20, 2022; April 10, 2022
Accepted: August 21, 2022
Published Online in Articles in Advance: October 18, 2022

Area of Review: Simulation

https://doi.org/10.1287/opre.2022.2375

Abstract. We consider a simulation-based ranking and selection (R&S) problem with input uncertainty, in which unknown input distributions can be estimated using input data arriving in batches of varying sizes over time. Each time a batch arrives, additional simulations can be run using updated input distribution estimates. The goal is to confidently identify the best design after collecting as few batches as possible. We first introduce a moving average estimator for aggregating simulation outputs generated under heterogeneous input distributions. Then, based on a sequential elimination framework, we devise two major R&S procedures by establishing exact and asymptotic confidence bands for the estimator. We also extend our procedures to the indifference zone setting, which helps save simulation effort for practical usage. Numerical results show the effectiveness and necessity of our procedures in controlling error from input uncertainty. Moreover, the efficiency can be further boosted through optimizing the "drop rate" parameter, which is the proportion of past simulation outputs to discard, of the moving average estimator.

Funding: The authors gratefully acknowledge support by the National Science Foundation Division of Civil, Mechanical and Manufacturing Innovation [Grant CMMI-1453934] and Division of Mathematical Sciences [Grant DMS2053489] and the Air Force Office of Scientific Research [Grants FA9550-19-1-0283 and FA9550-22-1-0244].

Supplemental Material: The electronic companion is available at https://doi.org/10.1287/opre.2022.2375.

Keywords: ranking and selection • optimization via simulation • input uncertainty • fixed confidence • streaming data • online estimation • indifference zone

1. Introduction

Stochastic simulation is used widely for evaluating and optimizing complex systems arising in manufacturing, transportation, and many other domains. Building a stochastic simulation model requires accounting for external random factors (e.g., lead time, traveling time, demand load) that affect the system's performance. A common practice is to model such randomness by probability distributions, from which random samples are generated to simulate real-world scenarios. These distributions are referred to as "input distributions."

In practice, input distributions need to be estimated from finite historical data known as "input data," for which the estimation error results in *input uncertainty* (IU). Aside from this extrinsic uncertainty, the simulation output is also subject to *stochastic uncertainty* (SU), which is the intrinsic uncertainty induced by random samples generated from input distributions. However, these two types of uncertainty differ in nature: whereas SU can be reduced by investing simulation effort, IU is determined by the availability of input data, which is largely beyond a modeler's control. Ignoring IU can be risky when simulation results are used to inform decisions (see, e.g., Zhou and Xie 2015).

In many application problems, input data are collected frequently to reduce the IU of the simulation model. The simulation model, which is often time-consuming to run, is then used to compare different designs/strategies to find the best one with high confidence. Such a procedure is often referred to as ranking and selection (R&S). Data often come on a fast time scale (such as daily), but system designs cannot be changed so frequently; otherwise, it not only incurs excessive cost and labor, but also might cause instability to the system. Hence, simulation becomes extremely useful to test and compare potential designs/strategies before they are implemented in the real system. The following application examples illustrate such a problem setting:

1. Start a ride-sharing service in an airport: In some cities, ride sharing used to be regulated (disallowed) in the airport. When such regulation is lifted, the platform faces a cold-start problem for pricing and matching strategies because of the lack of historical data at the

airport (such as driver and rider arrival rates, pricing and earnings elasticities), which are the inputs to a typical simulation system that the platform uses to choose the pricing and matching strategies. Such simulation can take hours per replication to capture the day/week effect in the market condition. As new data come in daily, the simulation input is updated, and new simulations are run to compare potential pricing and matching strategies until one best strategy is identified and then rolled out to the platform.

2. Supply chain optimization: A multinational retailer makes use of a complex supply chain simulation model to evaluate and compare a few potential inventory policies (including the current policy in use). The model's input distributions capture the uncertainty in production lead time, transit lead time, demand, and so on. As data accrues daily over the selling season, the retailer continuously updates the input models and runs simulations until it identifies the best inventory policy with high confidence and then implements the policy on the real system.

3. Budget R&D capital: A company conducts testing on a set of products with the goal of finding the product with the highest expected net gain within the shortest possible time period. Testing each product is time costly and subject to stochastic error, so multiple replications of testing are needed. As outside investments reveal over time, the company uses this data to update its estimate of the market condition of each product with reduced uncertainty and runs new tests until it narrows down to one best product and then releases it to the market.

These examples motivate us to consider a fixed confidence R&S problem in which new input data arrives over time in batches of possibly varying and random sizes. In each time stage, the number of simulation replications that can be run is limited because of the expensive simulation cost and the time length of each stage. On the one hand, we can use the new data batch to update the input distribution of the simulation model so as to reduce IU. Because of the limited simulation replications at each time stage, it is necessary to aggregate the simulation outputs across different time stages to reduce SU in the performance estimate. On the other hand, simulations are run under a different (updated) input distribution at each stage, and hence, the simulation outputs are correlated and differently distributed across different time stages. This creates a major challenge in designing an R&S procedure for this setting because most classical R&S requires an independent and identically distributed (i.i.d.) condition on the simulation outputs. Further complications lie in how to (i) update the input models and (ii) aggregate simulation outputs generated under heterogenous input models. Whereas there are well-established methods for these two tasks, whether they fit into the R&S framework needs investigation.

To address these challenges in the problem setting, we propose a moving average estimator for system performance to aggregate the simulation outputs across time stages. Intuitively, the moving average estimator drops the obsolete simulation outputs, keeping the more recent outputs that are generated under closer input distributions to the latest one. A parameter, called the drop rate, is used to balance the trade-off between bias (because of keeping the old simulation outputs) and variance (because of the limited number of simulation outputs) in the performance estimate. We then build on the sequential elimination (SE) framework, which was first developed by Even-Dar et al. (2002; 2006), to perform R&S using the moving average estimator. Specifically, by computing the confidence bands to account for both the IU and SU of the moving average estimates, the SE procedures eliminate one (statistically) inferior design each time. We summarize the contributions of this paper as follows:

1. To our knowledge, this paper, along with our earlier conference paper, Wu and Zhou (2019), are the first to consider streaming input data in R&S problems and design a data-driven approach. The simulation outputs generated from different stages are neither independent nor identically distributed, which is a major theoretical challenge that makes our approach drastically different from the classic R&S procedures.

2. We introduce a moving average estimator to aggregate simulation outputs generated under heterogenous input models and substantially extend a sequential elimination framework to handle IU with streaming input data.

3. We design sequential elimination procedures, called SEIU and SEIU-MCB, based on confidence bands established using two different approaches. The SEIU approach relies on exact confidence bands but tends to be conservative; the SEIU-MCB approach relies on asymptotically valid but tighter confidence bands. Specifically, the latter leverages results from "multiple comparison with the best" (MCB) (Chang and Hsu 1992) and an asymptotic normality result that we establish for characterizing the trade-off between IU and SU under streaming data, which is of independent interest and can be viewed as a multistage generalization of a well-known result in Cheng and Holland (1997). We also extend the two SE procedures to the indifference zone setting, named SEIU(IZ) and SEIU-MCB(IZ), which further boosts the procedures for practical usage.

4. The necessity and effectiveness of the procedures are demonstrated numerically through a simple quadratic problem and a more sophisticated production inventory problem. Furthermore, we show that SEIU-MCB can be accelerated by optimizing the drop rate parameter that shows up in the moving average estimator.

We should remark that the aggregation of simulation outputs under different input distributions via the likelihood ratio method is studied by Feng and Staum (2015; 2017), Eckman and Feng (2018), and Feng and Song (2019) under the name "green simulation" and recently by Liu and Zhou (2020) for simulation optimization. Such methods have the potential to be extended to the setting in this paper. We leave the adaptation of these methods to data-driven R&S as an open problem.

1.1. Literature Review

The mixed effect of IU and SU on simulation output is studied in a variety of contexts. In terms of quantifying the impact of IU on a single system design's simulation output, the earliest method at least dates back to Barton and Schruben (1993), followed by many other works, including but not limited to Cheng and Holland (1997), Chick (2001), Zouaoui and Wilson (2003; 2004), Barton et al. (2014), Xie et al. (2014; 2016), Lam and Qian (2022), Lin et al. (2015), Lam and Zhou (2017), Feng and Song (2019), Song and Nelson (2015), and Zhu et al. (2020).

These works assume a fixed batch of input data, but recently Zhou and Liu (2018) and Liu and Zhou (2019) considered streaming input data—the same setting as this paper. However, they take the likelihood ratio method to estimate the performance and a sampling-based method to quantify IU, whereas we develop a moving average estimator and analytically compute confidence bands to quantify IU. We refer the reader to Corlu et al. (2020) for a recent review on the topic of input uncertainty.

There is an abundant and fast-growing R&S literature. Although a comprehensive review is out of this paper's scope, in what follows, we make our best effort to present our work through a broader perspective.

In simulation optimization or ordinal optimization (Ho et al. 1992), R&S arises in the context of identifying the best system design (or a subset of good designs) through noisy simulation outputs. Simulation-based R&S is predominantly studied under the assumption that SU is the only source of uncertainty. Research in this field can roughly be categorized into two problem settings: fixed budget and fixed confidence. In the fixed budget setting, the total number of simulation runs is constrained, and the goal is to maximize the probability of correct selection (PCS) of the best design. In this paper, we focus on the fixed confidence setting, in which the goal is to attain a prespecified PCS using as little simulation effort as possible.

The study on fixed confidence R&S is primarily focused on the IZ formulation, in which the goal is to select the best system design with a target probability when the top two designs differ by at least some value δ in expected performance.

Since Bechhofer (1954), the IZ formulation has been studied actively by the statistics community. In simulation literature, the milestone work that adapts IZ to simulation-based R&S is Kim and Nelson (2001), in which the proposed KN procedure not only outperforms classic statistical procedures, but also allows the use of common random numbers (CRN), a variance reduction technique, for further enhancement. The KN procedure is extended to steady-state simulation in Goldsman et al. (2002) and Kim and Nelson (2006) and also inspires numerous subsequent works on fully sequential procedures, including Batur and Kim (2006), Hong and Nelson (2005; 2007), Pichitlamken et al. (2006), Hong (2006), and so on. More recently, Frazier (2014) proposes an IZ procedure that is Bayesian in spirit but provides a guarantee on the frequentist PCS; Luo et al. (2015) designs procedures for large-scale R&S problems in parallel computing environments; and Fan et al. (2016) develops procedures that do not require an IZ parameter.

The fixed confidence R&S problem is also studied extensively in the multiarmed bandits literature under the name of best arm identification (BAI). Except for a minor difference in performance criterion, R&S and BAI are essentially the same mathematical problem. Nevertheless, the research in the two fields diverges in several aspects with the most notable difference that BAI focuses more on characterizing complexity results and designing algorithms that match some worst case or problem-specific lower bounds, such as the seminal works of Even-Dar et al. (2002) and Mannor and Tsitsiklis (2004) followed by Gabillon et al. (2012), Karnin et al. (2013), Jamieson and Nowak (2014), and Kaufmann et al. (2016) and culminating in Garivier and Kaufmann (2016).

Aside from these frequentist results, Russo (2016) proposes some simple algorithms from a Bayesian viewpoint.

Compared with the studies on classic R&S and BAI, which have gradually matured over the past few decades, research on R&S under IU is only starting to gain momentum. One stream of work takes a distributionally robust optimization approach by assuming that the true input distribution is contained in a finite set of known distributions (i.e., ambiguity set) with a goal of selecting the design with the best worst case performance over the ambiguity set, such as Gao et al. (2017), Xiao and Gao (2018), Xiao et al. (2020), Fan et al. (2020), and Wu and Zhou (2017). Another line of research aims to screen out as many inferior designs as possible in the presence of IU, such as Corlu and Biller (2013, 2015) and Song and Nelson (2019). Our work has a close connection with Song and Nelson (2019) in that we both assume parametric input distributions and incorporate the MCB framework from Chang and Hsu (1992). However, despite these similarities, our settings

of streaming input data are fundamentally different. In Song and Nelson (2019), the performance estimate only consists of i.i.d. samples simulated under the same input distribution, whereas in this paper, we sequentially update input distributions over time and aggregate past simulation outputs from heterogeneous input distributions. Therefore, the methodologies therein cannot be easily extended to our problem.

All the papers on R&S under IU listed here share one assumption in common: input data are a static data set that does not grow over time. Streaming data have only been considered recently for R&S by Wu and Zhou (2019), which is a preliminary conference version of this paper. Song and Shanbhag (2019) and Liu et al. (2021; 2022) study continuous simulation optimization with streaming data.

The rest of the paper is organized as follows. We formulate the problem in Section 2 and extend the sequential elimination framework in Section 3. We derive the exact and asymptotically valid confidence bands in the SE framework and present the corresponding procedures SEIU and SEIU-MCB in Sections 4 and 5, respectively. We extend the procedures to the IZ setting in Section 6. Finally, we numerically demonstrate the procedures' performance in Section 7 and conclude in Section 8.

2. Problem Formulation

2.1. Review on Fixed Confidence R&S Without IU

We begin with a brief review on the classic fixed confidence R&S problem and lay down some basic notations. Suppose that we are given a set of designs $\mathcal{I} = \{1, \dots, K\}$, and the goal is to find the design with the highest expected performance. A design $i \in \mathcal{I}$ is evaluated through repeated simulation runs, and during each run, we first generate a random sample ξ from a distribution P_i and then compute the output by evaluating a deterministic function of i and ξ.

The input distributions $\{P_i\}$ are used as input to the simulation model to capture various sources of real-world randomness. In classic R&S literature, $\{P_i\}$ are assumed to be known, and the simulation outputs for any design i, denoted by $\{X_{i,r}\}_r$, are i.i.d. Let $\mu_i := \mathbb{E}_{P_i}[X_{i,1}]$ be the true mean performance of design i. Then, μ_i can be estimated by averaging the simulation outputs $\{X_{i,r}\}_r$. Because of finite-sample error, the probability of selecting the best design (PCS) is often adopted to quantify the confidence of selection. In the fixed confidence setting, there is no constraint on the simulation budget, but the PCS is required to exceed a prespecified level. We refer to procedures that guarantee to attain the PCS target as *statistically valid*. The core of the fixed confidence R&S problem is to find statistically valid procedures that terminate after as few simulation runs as possible.

2.2. Fixed Confidence R&S with Streaming Input Data

A limitation of the classic R&S framework is that true input distributions are rarely known in practice; instead, they must be estimated using finite data, which incurs IU that propagates to simulation output. Throughout the paper, we assume that all the designs share the same set of input distributions. Furthermore, to facilitate the characterization of IU, we impose the following structure on the input distributions.

Assumption 1. *The set of input distributions belongs to a known parametric family, which contains S mutually independent distributions with unknown parameters $\{\boldsymbol{\theta}_1^c, \boldsymbol{\theta}_2^c, \dots, \boldsymbol{\theta}_S^c\}$, where $\boldsymbol{\theta}_s^c \in \mathbb{R}^{d_s}$.*

The parametric assumption is common in the literature of R&S. Let $d := \sum_{s=1}^S d_s$ be the total dimension of parameters and $\boldsymbol{\theta} := [\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_S^\top]^\top$ be a vector in $\mathbb{R}^d$ concatenating all parameters. The input distributions then form a product measure $P_{\boldsymbol{\theta}} = \prod_{s=1}^S P_{\boldsymbol{\theta}_s}$. Also, let $\boldsymbol{\theta}^c \in \mathbb{R}^d$ denote the true input parameter. Under Assumption 1, the expected performance of design i becomes a function of θ and is, therefore, denoted $\mu_i(\boldsymbol{\theta})$.

The negative impact of IU on R&S can be seen as follows. For simplicity, assume throughout the paper that $b \in \mathcal{I}$ is the unique true best design, that is, $\mu_b(\boldsymbol{\theta}^c) > \max_{i \neq b} \mu_i(\boldsymbol{\theta}^c)$. When θ^c is unknown and is replaced by a finite-sample estimate $\hat{\boldsymbol{\theta}}$, the best design under $\hat{\boldsymbol{\theta}}$ may be different from b, in which case one cannot correctly identify the true best design even using infinite simulation runs. If input data are given as a static data set that does not grow over time, then our best hope is to gauge the impact of IU and provide a more conservative statistical guarantee (e.g., a lower PCS) because of limited resolution in performance estimation.

Assuming a static input data set is reasonable in applications in which gathering data is expensive, labor-intensive, or time-consuming. Nevertheless, with the advancement in big data technology, additional data can sometimes be collected efficiently and economically. This motivates us to consider an R&S problem with streaming input data, which we elaborate on as follows.

Suppose that, for each input distribution $P_{\boldsymbol{\theta}_s^c}$, the input data consists of i.i.d. samples denoted by $\{\zeta_{s,1}, \zeta_{s,2}, \dots\}$. The input data arrives in batches sequentially in time, and the sample size of batch t is $n_s(t) \geq 1$. Denote by $N_s(t) := \sum_{\ell=1}^t n_s(\ell)$ the total number of data samples up to batch t and let $N_s(0) = 0$. We allow $n_s(t)$ to vary in s and t in order to have the most flexible model for multiple streams of input data with time-varying batch sizes.

The arrival process of input data, thus, divides the simulation process into multiple time stages. At the beginning of stage t, a new batch of input data arrives, which, together with historical data, is used to compute an updated parameter estimate $\hat{\boldsymbol{\theta}}_t$. After that, incremental simulations are run by drawing i.i.d. samples from the updated input distributions $P_{\hat{\boldsymbol{\theta}}_t}$. Similarly, let $m(t) \geq 1$ be the number of incremental simulation runs for each design at stage t and let $M(t) := \sum_{\ell=1}^{t} m(\ell)$ be the total number of runs up until stage t. The simulation batch size $m(t)$ can vary across different stages because the interarrival times between two adjacent data batches could be different. However, here, we require the same batch size of simulation runs for each design because we make a pairwise comparison, which uses the CRN strategy.

Figure 1 illustrates how simulation outputs are generated for a single design, for which we highlight the following observations. First, the parameter estimates $\{\hat{\boldsymbol{\theta}}_t\}$ are correlated because they are computed using the same stream of input data. Second, because $\{\hat{\boldsymbol{\theta}}_t\}$ are random variables that generally take different values, the simulation outputs X are not identically distributed across different stages. Third, conditioned on $\{\hat{\boldsymbol{\theta}}_t\}$, the outputs X are i.i.d. within the same stage and independent across stages, but unconditional independence no longer holds because X are jointly affected by the correlation among $\{\hat{\boldsymbol{\theta}}_t\}$. In sum, the simulation outputs are neither independent nor identically distributed, which sets the problem apart from most of the R&S studies.

The online estimation of an individual design's performance under streaming input data is an important problem in its own right. In this paper, however, we are concerned with designing R&S procedures that can achieve a prespecified PCS upon termination. To make the problem well-defined, we first need to address the following questions:

i. What is the estimator of the input parameter θ^c?

ii. How do we aggregate non-i.i.d. simulation outputs to estimate $\mu_i(\boldsymbol{\theta}^c)$?

Many existing methods can be applied to these two problems, but most of them do not suit our purpose. The key is to find an estimator that allows a tractable decomposition of the correlation among different stages' simulation outputs. This motivates us to make the following assumption.

Assumption 2. *For each input parameter $\boldsymbol{\theta}_s^c$, there exists a function D_s such that*

$$\hat{\boldsymbol{\theta}}_{s,t} := \frac{1}{N_s(t)} \sum_{j=1}^{N_s(t)} D_s(\zeta_{s,j}) \tag{1}$$

is an unbiased estimator of $\boldsymbol{\theta}_s^c$.

Assumption 2 can often be satisfied through reparameterization. Specifically, for a parametric family of distributions with k unknown parameters, the parameters (or some transformation of the parameters, which can be transformed back to the parameters) can often be estimated by the first k moments with $D_s(x)$ being the functions $x, x^2, \ldots, x^k$. For example, if we reparameterize the normal distribution by the first two moments, then $D_s(\zeta) = [\zeta, \zeta^2]^\top$ satisfies the desired property. Let $D_{s,j}$ be shorthand for $D_s(\zeta_{s,j})$. Then, $\{D_{s,j}\}$ are i.i.d. transformed input data samples with $\mathbb{E}[D_{s,1}] = \boldsymbol{\theta}^c$. As we see in Section 5, the additive form of $\hat{\boldsymbol{\theta}}_{s,t}$ plays an important role in designing R&S procedures.

The other question is how to estimate $\mu_i(\boldsymbol{\theta}^c)$ using samples generated under different input distributions.

Figure 1. Illustration of Simulating a Single Design with Streaming Input Data, Where ζ Denotes Input Data, $\hat{\theta}$ Is the Input Parameter Estimate, and X Denotes Simulation Output

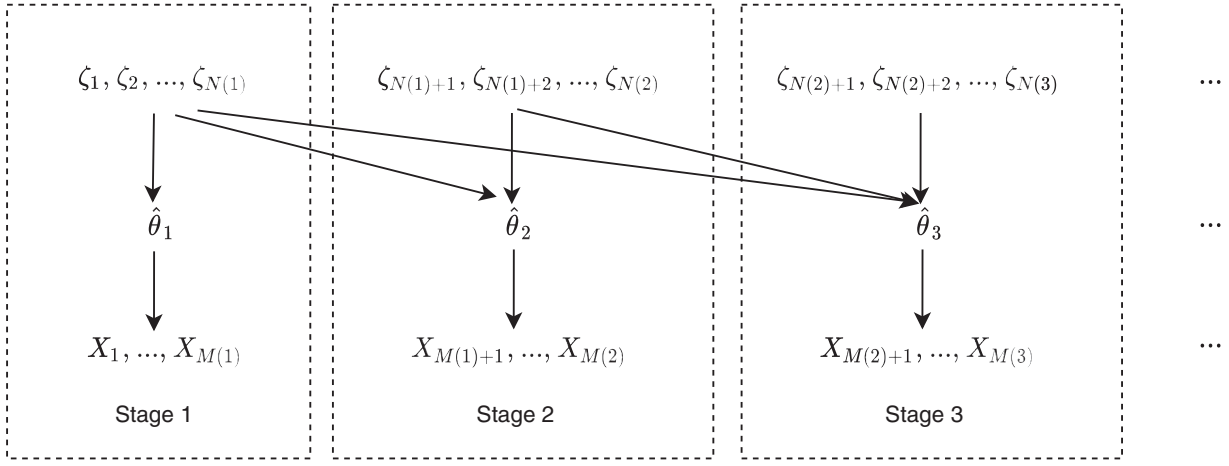

To simplify indexing, let $\{X_{i,r}(\hat{\boldsymbol{\theta}}_t)\}_{r=1}^{m(t)}$ denote the batch of simulation outputs generated at stage t under the input distribution $P_{\hat{\boldsymbol{\theta}}_t}$. We consider the following moving average estimator:

$$\widehat{\mu}_{i,t} := [M(t) - M(t_\eta)]^{-1} \sum_{\ell=t_\eta+1}^{t} \sum_{r=1}^{m(\ell)} X_{i,r}(\hat{\boldsymbol{\theta}}_\ell), \tag{2}$$

where $\eta \in [0,1)$ is called the drop rate, which controls the amount of effective samples used for computation, and $t_\eta := \lfloor \eta t \rfloor$ is the number of stages "discarded." In words, the estimator in (2) only averages the latest $(1-\eta)$ portion of all the simulation output, hence the name "moving average." The idea of throwing away some earlier samples is motivated by the following two extreme cases:

i. $\eta = 0$: Keeping all the outputs helps reduce SU but also retains all the biases $\{\mu_i(\hat{\boldsymbol{\theta}}_t) - \mu_i(\boldsymbol{\theta}^c)\}$ that accumulate over time.

ii. $\eta = 1$: This roughly corresponds to keeping only the latest output $X_{i,m(t)}(\hat{\boldsymbol{\theta}}_t)$, which reduces the bias but also prohibits the estimator from converging to $\mu_i(\boldsymbol{\theta}^c)$.

In Section 5, we explicitly characterize this trade-off between IU and SU through η and show that optimzing η can lead to a lower asymptotic variance than in both of the two extreme cases.

We now formally state the problem as follows. We aim to find the best design among K alternatives via simulation. The designs share the same input distribution with an unknown parameter θ^c, which can be estimated via (1) using input data arriving in batches. The arrival process divides time into multiple stages, and during each stage, we can update the input parameter estimate $\hat{\boldsymbol{\theta}}$ and run incremental simulations under $P_{\hat{\boldsymbol{\theta}}}$. The performance of each design is estimated via a moving average estimator in (2). Given $\alpha \in (0,1)$, the goal is to design a procedure that terminates after a number of stages and outputs the true best design with probability $\geq 1-\alpha$.

On a side note, classic R&S procedures may still be applied to our problem. For instance, we can collect a few batches of input data and apply KN as if the input distribution estimates are accurate. However, statistical validity is no longer guaranteed under the influence of IU. Whereas, in theory, this issue can be overcome with sufficient data, classic procedures have no way of telling when to safely stop collecting data.

3. Sequential Elimination Framework

In fixed confidence R&S, an important idea underlying many sequential procedures is to construct a "continuation region" in which a design can be confidently labeled as suboptimal if its performance estimate exits the region. One example is the triangular region in the KN procedure (see Kim and Nelson 2001). In our problem setting, most of the existing methods are inapplicable because they rely on key assumptions such as i.i.d. and normality. To construct a continuation region in the presence of IU, we resort to an SE framework in Even-Dar et al. (2002, 2006).

In short, an SE procedure sequentially eliminates inferior designs until there is only one left, and the surviving design is output as the estimated best design. Let $\delta_{ij}(\boldsymbol{\theta}) := \mu_i(\boldsymbol{\theta}) - \mu_j(\boldsymbol{\theta})$ and $\hat{\delta}_{ij,t} = \widehat{\mu}_{i,t} - \widehat{\mu}_{j,t}$. A requirement for an SE procedure is a collection of the confidence bands for $\delta_{ij}(\boldsymbol{\theta}^c)$,

$$\mathbb{P}\left\{\bigcap_{t=1}^{\infty}\bigcap_{i<j}\left\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| \leq w_{ij,t}\right\}\right\} \geq 1-\alpha, \tag{3}$$

where $w_{ij,t} \to 0$ as $t \to \infty$. The elimination happens to design i if there exists a design j such that

$$\hat{\delta}_{ij,t} + w_{ij,t} = \widehat{\mu}_{i,t} - \widehat{\mu}_{j,t} + w_{ij,t} < 0. \tag{4}$$

That is if the upper confidence of δ_{ij} is below zero. To see why (4) achieves statistical validity, notice that, if the confidence bands achieve perfect coverage, that is, $\delta_{ij}(\boldsymbol{\theta}^c) \in [\hat{\delta}_{ij,t} - w_{ij,t}, \hat{\delta}_{ij,t} + w_{ij,t}]$ for all i, j, and t, then for any design $i \neq b$ (recall that b is the unique best design),

$$\widehat{\mu}_{b,t} - \widehat{\mu}_{i,t} + w_{bi,t} \geq \mu_b(\boldsymbol{\theta}^c) - \mu_i(\boldsymbol{\theta}^c) > 0,$$

meaning that design b is never eliminated. Meanwhile, because $w_{ij,t} \to 0$ as $t \to \infty$, the procedure terminates in finite time almost surely, and the final output must be design b.

In applying the SE procedure, the core question is how to construct the confidence bands $\{w_{ij,t}\}$. Notice that the false selection happens at any stage when the true optimal design is eliminated. Therefore, we can write the probability of false selection (PFS) as

$$\begin{aligned}
\text{PFS} &= \mathbb{P}(\text{The optimal design } b \text{ is eliminated at some } t) \\
&= \mathbb{P}\left(\bigcup_{t=1}^{\infty}\{\text{The optimal design } b \text{ is eliminated at } t\}\right) \\
&= \sum_{t=1}^{\infty}\mathbb{P}(\text{The optimal design } b \text{ is eliminated at } t) \\
&\leq \sum_{t=1}^{\infty}\mathbb{P}\left(\bigcup_{i<j}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > w_{ij,t}\}\right),
\end{aligned} \tag{5}$$

where the second equality follows from the fact that the events {The optimal design b is eliminated at t}, $t = 1, 2, \ldots$ are disjoint, and the last inequality follows

from

$$\{\text{The optimal design } b \text{ is eliminated at } t\} \subseteq \bigcup_{i<j}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > w_{ij,t}\}.$$

Hence, if we can find $w_{ij,t}$ such that

$$\mathbb{P}\left(\bigcup_{i<j}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > w_{ij,t}\}\right) \le \frac{6\alpha}{\pi^2 t^2},$$

then we have

$$\sum_{t=1}^{\infty}\mathbb{P}\left(\bigcup_{i<j}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > w_{ij,t}\}\right) \le \sum_{t=1}^{\infty}\frac{6\alpha}{\pi^2 t^2} = \alpha. \tag{6}$$

We obtain that PCS is at least $1-\alpha$.

Notice that we can also use a nonpairwise framework that derives the confidence bands $w_{i,t}$ for each μ_i. Nonetheless, the pairwise framework allows for the use of CRN, which often sharpens the comparison between designs. Furthermore, the common input distribution may introduce additional positive correlation that further reduces variance. All things considered, it is promising to achieve $w_{ij,t} < w_{i,t} + w_{j,t}$ and, therefore, faster termination. Therefore, we only focus on the pairwise framework.

The problem is reduced to finding tight confidence bands for $\hat{\delta}_{ij,t}$. In the classic IU-free setting in which $\hat{\delta}_{ij,t}$ are averages of i.i.d. samples, many concentration bounds (e.g., Hoeffding's inequality, Gaussian tail bounds) are readily available. Unfortunately, these bounds do not apply directly to the case with IU. In the upcoming section, we derive IU-compatible confidence bands, on top of which we build SE procedures.

4. SEIU Procedures

We propose procedures based on confidence bands derived through an exact approach. These confidence bands allow us to build idealized SE procedures, which assume full knowledge on several key parameters. The idealized procedures are shown to be statistically valid and are equipped with upper bounds on their expected runtime. We also briefly discuss how to estimate the unknown parameters in practice to end this section.

4.1. Derivation of Exact Confidence Bands

We recall a few notations and concepts before deriving the confidence bands. Denote by $\Theta \subseteq \mathbb{R}^d$ the space in which θ lives. The stream of transformed input data for the sth input distribution, denoted by $\{D_{s,t}\}_t$, are i.i.d. random vectors in $\mathbb{R}^{d_s}$ whose sample mean is an unbiased estimator of $\boldsymbol{\theta}_s^c$. Also recall that a random variable X with mean $\mu = \mathbb{E}[X]$ is called *sub-Gaussian* if there exists $\sigma > 0$ such that $\mathbb{E}[e^{s(X-\mu)}] \le \exp(\sigma^2 s^2/2)$ for all $s \in \mathbb{R}$, in which case we write $X \sim subG(\sigma^2)$, and σ^2 is called a *variance proxy*. In the following assumption, $D_{s,t,j}$ denotes the jth component of the random vector $D_{s,t}$.

Assumption 3.

i. *Θ is a compact subset of $\mathbb{R}^d$.*

ii. *For any s, t, and j, $D_{s,t,j} \sim subG(v_s^2)$ for some $v_s > 0$ and $\{D_{s,t,j}\}$ are i.i.d.*

iii. *For any $\boldsymbol{\theta} \in \Theta$ and $i \in \mathcal{I}$, $X_i(\boldsymbol{\theta}) \sim subG(\bar{\sigma}_i^2)$ for some $\bar{\sigma}_i > 0$ independent of θ.*

iv. *For any design i, the performance function μ_i is Lipschitz continuous with respect to input parameter θ. More specifically, there exists $\bar{L}_i > 0$ such that*

$$|\mu_i(\boldsymbol{\theta}_1) - \mu_i(\boldsymbol{\theta}_2)| \le \bar{L}_i \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_1, \quad \forall \boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta,$$

where $\|\cdot\|_1$ denotes the ℓ^1-norm.

In Assumption 3, the major condition is the compactness of Θ, which can be satisfied in most real-world applications in which input distributions are supported on a bounded set (e.g., number of visitors on a given day, customer waiting time, etc.). With compactness, $\bar{\sigma}_i^2$ can often be taken as the maximum variance proxy over Θ, and the Lipschitz continuity of μ_i becomes a local property that usually holds in practice. Also, the boundedness of input data and simulation outputs implies their sub-Gaussianity. Furthermore, Assumption 3 implies that there exists $\bar{\sigma}_{ij} > 0$ such that $X_i(\boldsymbol{\theta}) - X_j(\boldsymbol{\theta}) \sim subG(\bar{\sigma}_{ij}^2)$ for all $\boldsymbol{\theta} \in \Theta$ and there exists $\bar{L}_{ij} > 0$ such that $\delta_{ij}(\cdot) = \mu_i(\cdot) - \mu_j(\cdot)$ is $\bar{L}_{ij}$-Lipschitz continuous on Θ.

We now motivate exact confidence bands for the moving average estimator $\hat{\delta}_{i,t}$. The goal is to find bounds $B_t(\cdot)$ such that, for any fixed $x > 0$, we have

$$\mathbb{P}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > x\} \le B_t(x),$$

where $B_t(x) \to 0$ as $t \to \infty$. To gain intuition for an IU-compatible bound, consider a simplified setting in which $\boldsymbol{\theta}^c \in \mathbb{R}$ and there is no SU, that is,

$$\hat{\delta}_{ij,t} = \frac{1}{t}\sum_{\ell=1}^{t}\delta_{ij}(\hat{\boldsymbol{\theta}}_\ell). \tag{7}$$

In this case, error only comes from the biases $\{\delta_{ij}(\hat{\boldsymbol{\theta}}_\ell) - \delta ij(\boldsymbol{\theta}^c)\}$. Then, for any $x > 0$,

$$\mathbb{P}\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > x\} \le \mathbb{P}\left\{\frac{1}{t}\sum_{\ell=1}^{t}|\delta_{ij}(\hat{\boldsymbol{\theta}}_\ell) - \delta_{ij}(\boldsymbol{\theta}^c)| > x\right\} \le \mathbb{P}\left\{\sum_{\ell=1}^{t}|\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^c| > tx/\bar{L}_{ij}\right\} \tag{8}$$

by the Lipschitz continuity of δ_{ij}. To further decompose the last term in (8), notice that, for any nonnegative sequence $\{\omega_\ell\}$ with $\sum_{\ell=1}^t \omega_\ell = 1$, we have

$$\mathbb{P}\left\{\sum_{\ell=1}^t |\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^c| > tx/\bar{L}_{ij}\right\} \le \mathbb{P}\left\{\bigcup_{\ell=1}^t \left\{|\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^c| > \omega_\ell tx/\bar{L}_{ij}\right\}\right\} \le \sum_{\ell=1}^t \mathbb{P}\left\{|\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^c| > \omega_\ell tx/\bar{L}_{ij}\right\}.$$

The question is whether we can choose $\{\omega_\ell\}$ such that the resulting bound converges to zero as $t \to \infty$. Note that taking equal weights $\omega_\ell = 1/t$ is not helpful because the resulting bound is bounded from below by a positive constant:

$$\sum_{\ell=1}^t \mathbb{P}\left\{|\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^c| > x/\bar{L}_{ij}\right\} \ge \mathbb{P}\left\{|\hat{\boldsymbol{\theta}}_1 - \boldsymbol{\theta}^c| > x/\bar{L}_{ij}\right\} = \mathbb{P}\left\{|D_1 - \boldsymbol{\theta}^c| > x/\bar{L}_{ij}\right\}.$$

Instead, an intuitive way to tighten this bound is to allow larger deviations (assign a greater ω_ℓ) for earlier stages and smaller deviations for later stages. This idea is made concrete by the following proposition.

Proposition 1. *Let Assumption 3 hold. Then, for any design i and any $x, y > 0$,*

$$\mathbb{P}\left\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > x + y\right\} \le 2\exp\left\{-\frac{[M(t) - M(t_\eta)]x^2}{2\bar{\sigma}_{ij}^2}\right\} + 2(t - t_\eta)\Sigma_{s=1}^S d_s \exp\left\{-\frac{[M(t) - M(t_\eta)]^2 y^2}{2d^2\bar{L}_{ij}^2 v_s^2 \gamma_{s,\eta}^2(t)}\right\}, \tag{9}$$

where $\gamma_{s,\eta}(t) := \sum_{\ell=t_\eta+1}^t m(\ell)/\sqrt{N_s(\ell)}$.

Remark 1. The term $\gamma_{s,\eta}$ loosely characterizes the rate at which the IU-induced error accumulates over time: the $1/\sqrt{N_s(\ell)}$ term reflects how fast IU (i.e., estimation error in $\hat{\boldsymbol{\theta}}_\ell$) diminishes as input data grows, but the error is multiplied by a factor of $m(\ell)$ because, during each stage, we perform $m(\ell)$ simulation runs under the inaccurate input distribution $P_{\hat{\boldsymbol{\theta}}_\ell}$.

Remark 2. In the proof of Proposition 1, the choice of $\omega_\ell \propto m(\ell)/\sqrt{N_s(\ell)}$ echoes our interpretation of $\gamma_{s,\eta}$, and it reweighs IU at different stages to yield a useful upper bound. It is also worth noting that, despite the multiplicative factor $(t - t_\eta)$ in (9), the bound still converges to zero exponentially fast as $t \to \infty$.

Proposition 1 provides IU-compatible bounds for a pair of designs. When it comes to designing SE procedures, however, it is more advantageous to use the following simultaneous bounds for all pairs of designs, which exploits the fact that all designs share the same input distribution and is, therefore, tighter than directly applying Bonferroni's inequality to (9).

Lemma 1. *Let Assumption 3 hold. Then, for any $x_{ij}, y_{ij} > 0$,*

$$\mathbb{P}\left\{\bigcup_{i<j}\left\{|\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)| > x_{ij} + y_{ij}\right\}\right\} \le \sum_{i<j} \exp\left\{-\frac{[M(t) - M(t_\eta)]x_{ij}^2}{2\bar{\sigma}_{ij}^2}\right\} + 2(t - t_\eta)\Sigma_{s=1}^S d_s \max_{i<j}\left\{\exp\left\{-\frac{[M(t) - M(t_\eta)]^2 y_{ij}^2}{2v_s^2 d^2 \bar{L}_{ij}^2 \gamma_{s,\eta}^2(t)}\right\}\right\}, \tag{10}$$

where $\gamma_{s,\eta}(t) := \sum_{\ell=t_\eta+1}^t m(\ell)/\sqrt{N_s(\ell)}$.

4.2. The SEIU Procedures

Based on Lemma 1, we design the SEIU procedure based on exact confidence band derivation, in which SEIU stands for sequential elimination under input uncertainty. The procedure relies on the key parameters $v, \bar{\sigma}_{ij}$, and $\bar{L}_{ij}$, which are assumed to be known for the moment.

Procedure: SEIU

- Initialization. $\alpha \in (0,1), \eta \in (0,1), \mathcal{S} = \mathcal{I}, t = 1$.
- Step 1. At stage t, update the estimate of θ^c using new input data and run additional simulations. For each pair of designs $i > j$ in $\mathcal{S}$, compute $\hat{\delta}_{ij,t} = \widehat{\mu}_{i,t} - \widehat{\mu}_{j,t}$ using (2).
- Step 2. Compute the confidence bands $w_{ij,t} = u_{ij,t} + v_{ij,t}$ for each pair of designs $i < j$ using

$$u_{ij,t} = 2\bar{\sigma}_{ij}\sqrt{\frac{\ln\left(\sqrt{\frac{K(K-1)\pi^2}{3\alpha}}t\right)}{M(t) - M(t_\eta)}}, \tag{11}$$

$$v_{ij,t} = \max_{1\le s\le S}\left\{\frac{dv_s\bar{L}_{ij}\gamma_{s,\eta}(t)}{M(t) - M(t_\eta)}\sqrt{2\ln\left(\frac{2Sd_s\pi^2(1-\eta)t^3}{3\alpha}\right)}\right\}, \tag{12}$$

where $\gamma_{s,\eta}(t) := \sum_{\ell=t_\eta+1}^t m(\ell)/\sqrt{N_s(\ell)}$.

- Step 3. For each design $i \in \mathcal{I}$, if

$$\min_{i>j}\left\{\hat{\delta}_{ij,t} + w_{ij,t}\right\} < 0, \tag{13}$$

then set $\mathcal{S} \leftarrow \mathcal{S} \setminus \{i\}$. Go to Output if $|\mathcal{S}| = 1$; otherwise, set $t \leftarrow t + 1$ and go to step 1.
- Output. Output the only design in $\mathcal{S}$.

In the SEIU procedure presented here, $u_{i,t}$ and $v_{i,t}$ are chosen to ensure that PCS $\geq 1-\alpha$. Specifically, by substituting x and y with u and v in Lemma 1, we have

$$\mathbb{P}\left\{\bigcup_{i<j}\left\{|\hat{\delta}_{ij,t}-\delta_{ij}(\boldsymbol{\theta}^c)|>u_{ij,t}+v_{ij,t}\right\}\right\}\leq\frac{6\alpha}{\pi^2t^2},$$

that is, the probability of any design's estimate at any stage exiting its confidence band does not exceed $\frac{6\alpha}{\pi^2t^2}$, and as a result, the probability of false selection does not exceed α as shown in (5) and (6). To run SEIU, a user first needs to specify the target PCS level $1-\alpha$ and a drop rate η for the moving average estimator $\widehat{\mu}_{i,t}$. Then, the procedure collects input data and runs simulations alternately and removes design(s) from the survival set $\mathcal{S}$ whenever the elimination criterion (13) is met. Finally, the last surviving design is output as the selected best.

The following theorem shows the statistical validity of the procedure SEIU and characterizes the number of stages needed for the procedure to terminate, denoted as $\tilde{\tau}$, assuming full knowledge of $\nu_s, \bar{\sigma}_{ij}$ and $\bar{L}_{ij}$.

Theorem 1. *Let Assumption 3 hold, and suppose that ν_s, $\bar{\Sigma}_{ij}$, and $\bar{L}_{ij}$ are known. Then, SEIU selects the best design with probability at least $1-\alpha$. Furthermore,*

$$\mathbb{E}[\tilde{\tau}]\leq\max_{i\neq b}\min_{j\neq i}\tilde{\tau}_{ij}+\alpha, \tag{14}$$

where

$$\tilde{\tau}_{ij}=\inf\{t|w_{ij,t}<\delta_{ij}(\boldsymbol{\theta}^c)\}. \tag{15}$$

The intuition behind the bound on $\mathbb{E}[\tilde{\tau}]$ is that, if all confidence bands achieve perfect coverage, then they eventually become narrow enough for us to distinguish between any two designs. Hence, $\tilde{\tau}_{ij}$ is the worst case time it takes for a design j to eliminate design i, which is infinity if j is inferior to i. By taking the minimum of $\tilde{\tau}_{ij}$ over all designs j other than i, we obtain the longest survival time for i. Thus, the total runtime does not exceed the maximum survival time of all suboptimal designs. Although $\tilde{\tau}_{ij}$ does not have a simple closed form, it clearly depends on two factors: (i) the widths of the confidence bands $\{w_{ij,t}\}$ and (ii) the difference between means, that is, $\delta_{ij}(\boldsymbol{\theta}^c)$.

As in many elimination-based procedures, SEIU can be prestopped at any stage to output a subset $\mathcal{S}$ that contains the true best with probability at least $1-\alpha$. Moreover, the following corollary provides a selection at any stage with an IZ-type guarantee.

Corollary 1. *Let Assumption 3 hold and suppose that $\nu_s, \bar{\sigma}_{ij}$, and $\bar{L}_{ij}$ are known. Following the SEIU procedure, at any stage t, with the remaining set of designs $\mathcal{S}$, the design with the largest moving average mean, that is, the design $i^*=\arg\max_{i\in\mathcal{S}_t}\widehat{\mu}_{i,t}$, is ϵ_t-optimal with probability $1-\alpha$, that is, with probability $1-\alpha$, $\mu_{i^*}(\boldsymbol{\theta}^c)\geq\max_{1\leq i\leq K}\mu_i(\boldsymbol{\theta}^c)-\epsilon_t$, where $\epsilon_t=\max_{j\in\mathcal{S},j\neq i^*}w_{i^*j,t}$.*

Corollary 1 implies that, before the procedure eliminates all the inferior designs, one can also make the selection at any intermediate stage t by choosing the remaining design that has the largest sample mean. Such a selection gives an ϵ_t-optimal design with confidence level $1-\alpha$. The value of ϵ_t depends on the confidence bands $w_{ij,t}$, which is the sum of $\{u_{ij,t}\}$ and $\{v_{ij,t}\}$ that account for SU and IU, respectively.

To gain more insight, consider a special case in which $m(\ell)\equiv m_0$ and $n_s(\ell)\equiv n_0$, that is, a constant batch size of data and simulation over different stages. Then, (11) and (12) can be simplified as

$$u_{ij,t}=2\frac{\bar{\sigma}_{ij}}{\sqrt{m_0}}\sqrt{\frac{\ln\left(\sqrt{\frac{K(K-1)\pi^2}{3\alpha}}t\right)}{t-t_\eta}}, \tag{16}$$

$$v_{ij,t}\approx\max_{1\leq s\leq S}\left\{\frac{2d\nu_s\bar{L}_i}{\sqrt{n_0}(1+\sqrt{\eta})}\sqrt{\frac{2\ln\left(\frac{2Sd_s\pi^2(t-t_\eta)^3}{3(1-\eta)^2\alpha}\right)}{t}}\right\}. \tag{17}$$

Thus, $u_{i,t}$ and $v_{i,t}$ are both of order $O(\sqrt{\ln t/t})$, which is common in SE procedures. Moreover, it can be seen from (16) that the width of $u_{i,t}$ primarily depends on the magnitude of $\bar{\sigma}_{ij}/\sqrt{m_0}$, that is, the discounted variance proxy after averaging m_0 output samples at each stage; clearly, a smaller $\bar{\sigma}_{ij}$ and a higher m_0 yield narrower bands. The drop rate η also matters: notice that $u_{ij,t}\to\infty$ as $\eta\to 1$, meaning that the confidence band needs to expand in order to cover inflated SU with more samples thrown away.

Also clear from (17) is that greater values of $\nu_s/\sqrt{n_0}$ and $\bar{L}_{ij}$, which capture the extent of IU and δ_{ij}'s sensitivity to IU, respectively, result in a wider $v_{ij,t}$. It can further be checked that $v_{ij,t}\to 0$ as $\eta\to 1$, which agrees with the intuition that the more stages we discard, the less we suffer from the biases $\{\delta_{ij}(\hat{\boldsymbol{\theta}}_t)-\delta_{ij}(\boldsymbol{\theta}^c)\}$ from previous stages.

In the SEIU procedure, we require full knowledge of the parameters $\bar{\sigma}_{i,j}, \nu_s, \bar{L}_{i,j}$, which is often unknown and needs to be estimated in practice. We propose some heuristic methods and discussions for estimating these unknown parameters, which can be found in the electronic companion, Section EC.1.

5. SEIU-MCB Procedures

In this section, we propose a more efficient SEIU-MCB procedure by deriving asymptotically valid confidence bands, which utilizes the result of MCB proposed by

Chang and Hsu (1992). MCB is extended by Song and Nelson (2019) to construct confidence bands accounting for IU, in which they exploit the jointly asymptotic normality of pairwise differences of performance estimators (referring to $\{\hat{\delta}_{ij,t}\}_{j\neq i}$) that consist of samples under the current estimated input distribution. We extend this jointly asymptotic normality to our moving average estimator. The major challenge of establishing the asymptotic normality result for our moving average estimator is that the estimator aggregates simulation outputs from different stages, which disallows the application of the central limit theorem by a direct partition of the total error into two parts caused, respectively, by IU and SU. Instead, we adopt the Lindeberg–Feller theorem (see proposition 2.2.7 in Van der Vaart 2000) to prove the asymptotic normality of our moving average estimator.

Compared with SEIU, the SEIU-MCB procedure requires much less restrictive assumptions and achieves higher efficiency. Furthermore, the MCB framework helps to avoid the usage of Boole's inequality across designs, and as a result, the asymptotic normality result yields much tighter confidence bands to control the cumulative error across stages. In addition, our result explicitly characterizes how η affects the trade-off between IU and SU and shows that the elimination between any two designs can be potentially boosted by optimizing η.

5.1. MCB

The MCB result is shown as the following theorem.

Theorem 2. *Let $\widehat{\mu}_{i,t}$ be the estimator of μ_i for $i = 1, 2, \ldots, K$, $x^+ = max(x, 0)$ and $x^- = -min(x, 0)$. If, for each fixed i,*

$$P\left\{\widehat{\mu}_{i,t} - \mu_i - (\widehat{\mu}_{j,t} - \mu_j) \geq -w_{ij,t} \textbf{ for } i \neq j\right\} = 1 - \alpha, \quad (18)$$

then we can make the joint probability statement

$$\mathbb{P}\left\{\mu_i - \max_{j\neq i} \mu_j \in [W_i^-, W_i^+] \quad \text{for } i = 1, 2, \ldots K\right\} = 1 - \alpha,$$

where W_i^- and W_i^+ can be computed as

$$W_i^+ = \left(\min_{j\neq i}[\widehat{\mu}_{i,t} - \widehat{\mu}_{j,t} + w_{ij,t}]\right)^+, \mathbb{G} = \{i : W_i^+ > 0\}$$

$$W_i^- = \begin{cases} 0 & \text{if } \mathbb{G} = \{i\} \\ -\left(\min_{j\neq i}[\widehat{\mu}_{i,t} - \widehat{\mu}_{j,t} - w_{ji,t}]\right)^- & \text{otherwise.} \end{cases} \quad (19)$$

Theorem 2 says that, if we are provided with the simultaneous pairwise confidence bands $w_{ij,t}$ for each i, then we can get a simultaneous confidence band for i that measures how the design i is compared with the remaining best. If the upper bound W_i^+ is zero, then we eliminate design i because it is inferior to the remaining best design. If W_i^- is zero, then we directly choose design i as the optimal design.

Now, the key to the problem remains as how to compute the simultaneous pairwise confidence band $w_{ij,t}$, which depends on the joint distribution of

$$\Delta_{i,t} \triangleq (\hat{\delta}_{i1,t} - \delta_{i1}(\boldsymbol{\theta}^c), \ldots, \hat{\delta}_{ii-1,t} - \delta_{ii-1}(\boldsymbol{\theta}^c), \hat{\delta}_{ii+1,t} - \delta_{ii+1}(\boldsymbol{\theta}^c), \ldots, \hat{\delta}_{iK,t} - \delta_{iK}(\boldsymbol{\theta}^c)).$$

However, usually the joint distribution is hardly computable, and even so, the computation can be expensive. Hence, instead of finding the exact joint distribution, we show the asymptotic joint distribution in the following Section 5.2.

5.2. Asymptotic Normality

In Section 4, a critical assumption underpinning the SEIU framework is the compactness of input distribution's parameter space Θ, which plays an important role in rigorously bounding the variation of simulation output in the presence of IU. We now relax this condition and shift to the following assumption.

Assumption 4. *For all $i \in \mathcal{I}$ and $s \in \{1, 2, \ldots, S\}$,*

i. *$\Sigma_{D,s} := Cov(D_{s,1})$ exists.*

ii. *$\Sigma(\boldsymbol{\theta})$ exists and is continuous for all $\boldsymbol{\theta} \in \Theta$, where $\Sigma_{ij}(\boldsymbol{\theta}) = Cov[X_i(\boldsymbol{\theta}), X_j(\boldsymbol{\theta})]$.*

iii. *$\mu_i(\cdot)$ is twice continuously differentiable in Θ.*

iv. *$\{n_s(t)\}$ and $\{m(t)\}$ are uniformly bounded. Furthermore, there exist positive constants $\bar{n}_s$ and $\bar{m}$ such that $N_s(t)/t \to \bar{n}_s$ and $M(t)/t \to \bar{m}$ as $t \to \infty$.*

In Assumption 4, (i) and (ii), the existence of input data's covariance and simulation output's variance is far less stringent than the sub-Gaussianity conditions in Assumption 3 and holds in most real-world applications. The smoothness of $\mu_i(\cdot)$ in (iii) is also a reasonable assumption for many parametric families of $P_{\boldsymbol{\theta}}$. Moreover, (iv) only requires the limit of input data and simulation batch sizes to exist in the Cesàro sense, and the uniform boundedness of $\{n_s(t)\}$ and $\{m(t)\}$ is guaranteed in practice. With Assumption 4, we establish the following asymptotic result for $\Delta_{i,t}$.

Theorem 3. *Let Assumption 4 hold. Then, for any $\eta \in [0, 1)$ and any design i,*

$$\sqrt{t}\Delta_{i,t} \Rightarrow \mathcal{N}(0, \Sigma_{i,\infty}), \quad as \quad t \to \infty, \quad (20)$$

where $\Rightarrow$ means convergence in distribution, $\mathcal{N}$ denotes the normal distribution, and

$$\Sigma_{i,\infty}(j, j') := \lambda_{I,\eta} \nabla\delta_{ij}(\boldsymbol{\theta}^c)^\top \bar{\Sigma}_D \nabla\delta_{ij'}(\boldsymbol{\theta}^c) + \lambda_{S,\eta} \bar{m}^{-1} \boldsymbol{Cov}\left(X_i(\boldsymbol{\theta}^c) - X_j(\boldsymbol{\theta}^c), X_i(\boldsymbol{\theta}^c) - X_{j'}(\boldsymbol{\theta}^c)\right), \quad (21)$$

in which $\bar{\Sigma}_D := diag(\bar{n}_1^{-1}\Sigma_{D,1}, \ldots, \bar{n}_S^{-1}\Sigma_{D,S})$ *and*

$$\lambda_{I,\eta} := \frac{2}{1-\eta} + \frac{2\eta \ln \eta}{(1-\eta)^2}, \quad \lambda_{S,\eta} := \frac{1}{1-\eta}. \tag{22}$$

Theorem 3 shows that the asymptotic covariance matrix of $\Delta_{i,t}$ is the weighted sum of two components,

$$\Sigma_{i,IU}(j,j') := \nabla\delta_{ij}(\boldsymbol{\theta}^c)^\top \bar{\Sigma}_D \nabla\delta_{ij'}(\boldsymbol{\theta}^c),$$
$$\Sigma_{i,SU}(j,j') := \bar{m}^{-1}\mathbf{Cov}\Big(X_i(\boldsymbol{\theta}^c) - X_j(\boldsymbol{\theta}^c), X_i(\boldsymbol{\theta}^c) - X_{j'}(\boldsymbol{\theta}^c)\Big), \tag{23}$$

which quantify the variance induced by IU and SU, respectively. Furthermore, if we let $\partial_{\boldsymbol{\theta}_s}\delta_{ij}(\boldsymbol{\theta}^c)$ denote the partial derivative $\partial\delta_{ij}(\boldsymbol{\theta}^c)/\partial\boldsymbol{\theta}_s$, then we can decompose $\Sigma_{i,IU}$ as

$$\Sigma_{i,IU}(j,j') = \sum_{s=1}^{S} \partial_{\boldsymbol{\theta}_s}\delta_{ij}(\boldsymbol{\theta}^c)^\top \Sigma_{D,s}\partial_{\boldsymbol{\theta}_s}\delta_{ij'}(\boldsymbol{\theta}^c)/\bar{n}_s,$$

which attributes the variance to each individual input distribution. The $O(1/\bar{n}_s)$ and $O(1/\bar{m})$ convergence rate in (21) are standard. However, the weights $\lambda_{I,\eta}$ and $\lambda_{S,\eta}$ are less common and, therefore, deserve a closer look.

The expressions of $\lambda_{I,\eta}$ and $\lambda_{S,\eta}$ are shaped by two major factors: (i) the structure imposed on $\hat{\boldsymbol{\theta}}_t$ (see Assumption 2) and (ii) the choice of estimator $\widehat{\mu}_{i,t}$. In particular, the sample average structure of $\hat{\boldsymbol{\theta}}_t$ makes asymptotic analysis more tractable because it allows the correlated biases $\{\delta_{ij}(\hat{\boldsymbol{\theta}}_t) - \delta_{ij}(\boldsymbol{\theta}^c)\}$ to be decoupled into a weighted sum of i.i.d. samples $D_{s,t}$. To get a sense of the limiting behavior of $\lambda_{I,\eta}$ and $\lambda_{S,\eta}$, we investigate two extreme cases:

Case 1. If $\eta \to 1$, then $\lambda_{I,\eta} \to 1$ and $\lambda_{S,\eta} \to \infty$, which coincides with the intuition that IU can be reduced by dropping samples, but there will be fewer simulation outputs to average out SU.

Case 2. If $\eta = 0$, then $\lambda_{I,\eta} = 2$ and $\lambda_{S,\eta} = 1$, that is,

$$\sqrt{t}\Big[\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)\Big] \Rightarrow \mathcal{N}(0, 2\Sigma_{i,IU} + \Sigma_{i,SU}), \quad \text{as } t \to \infty. \tag{24}$$

To put (24) in perspective, we cite an asymptotic normality result from Wu and Zhou (2017) for the case of static input data. Suppose that there is only a single batch of input data, which we use to compute the parameter estimate $\hat{\boldsymbol{\theta}}_t$ and then run simulations under $P_{\hat{\boldsymbol{\theta}}_t}$. Let $\tilde{\mu}_{i,t} = \frac{1}{M_i(t)}\sum_{r=1}^{M_i(t)} X_{i,r}(\hat{\boldsymbol{\theta}}_t)$ be the estimator for $\mu_i(\boldsymbol{\theta}^c)$. Under some mild conditions, we have

$$\sqrt{t}\big[\tilde{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c)\big] \Rightarrow \mathcal{N}(0, \Sigma_{i,IU} + \Sigma_{i,SU}), \quad \text{as} \quad t \to \infty. \tag{25}$$

By comparing (24) with (25), it can be seen that the weight of $\Sigma_{i,IU}$ is doubled in the case of streaming input data, whereas the weight of $\Sigma_{i,SU}$ remains the same, and the inflation of $\Sigma_{i,IU}$ is due to estimation error accumulated over multiple stages.

With Theorem 3, we approximate the joint distribution of $\Delta_{i,t}$ as

$$\Delta_{i,t} \overset{D}{\approx} \mathcal{N}(0, \Sigma_{i,\infty}/t).$$

Then, $\{w_{ij,t}, j \neq i\}$ is a $(K-1)$-multidimensional quantile. When given the confidence level $1-\alpha$, the choice of $K-1$-dimensional quantile is not unique because the degree of freedom is $K-1$. We can reduce the degree of freedom to one by letting

$$\mathcal{P}\Big(\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c) \geq -w_{ij,t}\Big)$$
$$= \mathcal{P}\Big(\hat{\delta}_{ij',t} - \delta_{ij'}(\boldsymbol{\theta}^c) \geq -w_{ij',t}\Big) \quad \forall j \neq j'.$$

That is, all the confidence bands result in the same probability coverage. This can be done easily under the asymptotic normality of $\hat{\delta}_{ij,t}$. To be specific, for each i, we find the smallest $w_{i,t}$, such that

$$\mathcal{P}\Big(\hat{\delta}_{ij,t} - \delta_{ij}(\boldsymbol{\theta}^c) \geq -w_{i,t}\sigma_{ij,t}, \quad \forall j \neq i\Big) \geq 1-\alpha,$$

where $\sigma_{ij,t}$ refers to the variance of $\hat{\delta}_{ij,t}$ and we assume $\hat{\delta}_{ij,t}$ approximately follows a normal distribution centered at $\delta_{ij}(\boldsymbol{\theta}^c)$. We present the SEIU-MCB procedure in the next section.

5.3. SEIU-MCB Procedure

With the elimination and selection rule stated in Section 5.1, we have the SEIU-MCB procedure as follows.

SEIU-MCB

- Initialization. $\alpha \in (0,1), \eta \in (0,1), t = 1, \mathcal{S} = \mathcal{I}$.
- Step 1. At stage t, update the estimate of θ^c using new input data and run additional simulations using CRN. For each system $i \in \mathcal{S}$, compute $\widehat{\mu}_{i,t}$ using (2).
- Step 2. Compute the simultaneous pairwise confidence bands $w_{ij,t}$ using (18) and (20) with confidence level set as $\frac{6\alpha}{\pi^2 t^2}$. Then, compute W_i^-, W_i^+ using (19).
- Step 3. For each i, if $W_i^- = 0$, then $\mathcal{S} = \{i\}$; else if $W_i^+ = 0$, $\mathcal{S} \leftarrow \mathcal{S}\backslash\{i\}$. If $|\mathcal{S}| = 1$, then go to Output. Increment t by one and go to step 1.
- Output. Output the only design in $\mathcal{S}$ as the optimal system.

As we can see from the procedure, compared with the SEIU in which we have the sub-Gaussian and Lipschitz continuous assumption, SEIU-MCB requires much less knowledge of those parameter estimates. This is attributed to the general asymptotic normality result established in Section 5.2, by which we can compute the confidence bands easily.

Moreover, from Section 5.2, we can see the drop rate balances the impact of input uncertainty and simulation

uncertainty characterized by the asymptotic covariance matrix. Detailed discussions on minimizing such impact through η are presented in Section EC.2.

6. Extension to the IZ Setting

In Sections 4 and 5, the SEIU and SEIU-MCB procedures are designed to proceed until there is only one design left in the remaining candidate set because of our goal to identify the unique optimal design. However, in practice, selecting one of the "near-optimal" designs is often sufficient, which is the motivation behind the indifference zone R&S procedures. To be more specific, an IZ R&S procedure aims to find an ϵ-optimal design with a given $\epsilon > 0$, that is, a design i^* such that $\mu_{i^*} \geq \max_{1\leq i\leq K}\mu_i - \epsilon$, where the ϵ is often referred to as the IZ preference. Such an ϵ compromise can sometimes dramatically save simulation effort, especially when the optimal performance $\mu_b(\boldsymbol{\theta}^c)$ is close to the second optimal performance $\mu_{b'}(\boldsymbol{\theta}^c)$. Inspired by Corollary 1, we can also extend our procedures to the IZ setting and terminate the procedure when either of these cases happens: (i) there is only one design remaining or (ii) $\epsilon_t < \epsilon$. Denote by SEIU(IZ) and SEIU-MCB(IZ), the procedures with the IZ setting. We add an additional step 4 in both procedures to verify whether an ϵ-optimal is found.

- Step 4 in SEIU(IZ). Let $i^* = \arg\max_{i\in\mathcal{S}_t}\widehat{\mu}_{i,t}$. If

$$\max_{j\in\mathcal{S},j\neq i^*} w_{i^*j,t} < \epsilon,$$

then set $\mathcal{S} \leftarrow \{i^*\}$ and go to Output; otherwise, set $t \leftarrow t+1$ and go to step 1.

- Step 4 in SEIU-MCB(IZ). Let $\mathcal{I} = \{i \in \mathcal{S} | -W_i^- < \epsilon\}$. If $|\mathcal{I}| \geq 1$, let $\hat{i}^* = \arg\max_{i\in\mathcal{I}}\widehat{\mu}_{i,t}$, set $\mathcal{S} \leftarrow \{\hat{i}^*\}$, and go to Output. Otherwise, increment t by one and go to step 1.

With similar intuition as Theorem 1, the following corollary provides the statistical validity of SEIU(IZ) and characterization of the number of stages needed to terminate the algorithm.

Corollary 2. *Let Assumption 3 hold and suppose that $v_s, \bar{\sigma}_{ij}$, and $\bar{L}_{ij}$ are known. Then, SEIU(IZ) selects an ϵ-optimal design with probability at least $1-\alpha$. Furthermore, let τ denote the number of stages it takes to terminate. Then,*

$$\tau \leq \inf\left\{t \middle| \max_{1\leq i<j\leq K} w_{ij,t} < \epsilon\right\} \quad a.s. \tag{26}$$

Corollary 2 shows that an advantage of using the IZ criterion is that we can characterize the number of stages needed to terminate in a stronger, almost sure sense. The intuition is that, with the IZ criterion, when the confidence band is small enough, we always either identify the best design or reach the specified accuracy (ϵ-optimal). Again, although we do not have a closed form of the upper bound on τ, it clearly depends on the width of the confidence bands, the true difference $\delta_{ij}(\boldsymbol{\theta}^c)$, and the IZ preference ϵ.

Another advantage of our IZ procedures compared with the well-known KN procedure (see Kim and Nelson 2001) is that our procedures compute the confidence bands independently of IZ preference and terminate when either the best selection is made or the IZ preference is satisfied. In contrast, the KN procedure computes the confidence bands by considering a so-called "slippage configuration" in which the best design has expected performance exactly δ better than the second best. If the actual difference is bigger than δ, the KN procedure computes a larger confidence band and makes redundant simulations. In our procedures, if we set δ smaller than the true difference of the expected performance between the best and second best designs, it does not affect the performance of the procedures because they act at least as good as the SEIU and SEIU-MCB procedures.

7. Numerical Results

The numerical study mainly consists of three parts: (i) comparing our procedures with the KN++ procedure, which does not consider input uncertainty; (ii) investigating the performance of all our proposed procedures; (iii) optimizing η to boost the efficiency of SEIU-MCB.

7.1. Test Problems

Two problems are used for numerical testing. One is to minimize a quadratic objective function with a single source of IU, and the other is a more complex production inventory optimization problem from Hong (2009) with multiple independent sources of IU. A detailed description of the two problems can be found in Sections EC.3.1 and EC.3.2. In all subsequent experiments, we consider the following settings unless otherwise stated:

i. Equal batch size: $n_s(t) = n, m(t) = m$ for some fixed $n, m \in \mathbb{Z}^+$.

ii. Random batch size: $n_s(t)$ and $m(t)$ are i.i.d. samples drawn from the uniform distribution on the lattice $\{\bar{k}, 2\bar{k}, \ldots, 5\bar{k}\}$ for some fixed $\bar{k}$.

More specifically, in the case of unequal batch size, the random samples of $n_s(t)$ and $m(t)$ are independent of input data and simulation outputs.

7.2. Experiment Results

In the following sections, we first use the simple quadratic example to show the necessity of accounting for input uncertainty by comparing the KN++ procedure with our two IZ-type procedures: SEIU(IZ) and SEIU-MCB(IZ). We then demonstrate the applicability and effectiveness of all our proposed procedures on the more general production inventory example, in which

we have multiple input distributions and allow different batch sizes for both input data and simulation budget. Finally, we show that the performance of SEIU-MCB can be further boosted by optimizing η.

7.2.1. Necessity of Considering IU. We first use the quadratic problem with equal batch size and $\mathcal{I} = \{\boldsymbol{\theta}^c + 0.3 \cdot i : i \in [-5, 5] \cap \mathcal{Z}\}$, to test the three procedures with the IZ setting: SEIU(IZ), SEIU-MCB(IZ), and KN++. The KN++ procedure (Kim and Nelson 2006) is an extension of the KN procedure that allows nonnormal observation data and updating the variance estimator. The KN++ procedure uses average batch mean to estimate the variance and updates the variance estimator as more simulation outputs are collected, and hence, it can be directly applied to our example with equal batch size. Specifically, it is applied by ignoring the existence of input uncertainty and carrying out simulations in each stage conditioned on the current estimated input parameter. Table 1 shows the empirical PCS and average termination stage along with their 95% confidence intervals based on 100 independent runs of the three procedures with different values of the IZ parameter δ and different input data batch sizes. The empirical PCS is computed as the percentage of algorithm runs that correctly select the true optimal design. The target PCS is set as 90%. Some observations can be made as follows:

1. The empirical PCS of the KN++ procedure is less than the target PCS 0.9 across all settings. The true difference between the best design and the second best is 0.16. Hence, with $\delta = 0.1$, it is sufficient to identify the best design. Furthermore, by examining the estimated input parameters when incorrect selection or elimination happens, we find that the KN++ procedure makes the false selection mainly when procedures terminate in earlier stages, in which input parameters are estimated more inaccurately. This implies the necessity of controlling the input error.

2. The SEIU-MCB(IZ) achieves the target PCS when $\delta = 0.1$ and shoots a higher PCS than KN++ in all settings with mostly earlier termination. This implies the high efficiency of the SEIU-MCB procedure even when it takes IU into consideration.

3. Comparing SEIU(IZ) and SEIU-MCB(IZ), the latter procedure requires much less simulation effort and achieves more accurate empirical PCS. It shows that the usage of asymptotic normality and MCB in SEIU-MCB significantly relieves the conservativeness in the SEIU procedure. The empirical PCS of SEIU(IZ) always overshoots the target PCS, revealing the conservativeness of the procedure. The conservativeness is mainly a result of the Bonferroni inequality used in deriving the confidence bands in the SEIU procedure.

7.2.2. Inventory Production Problem: A More General Setting. In this section, we show the applicability and effectiveness of all our proposed procedures using an inventory production example, in which we have multiple unknown input distributions with random batch sizes for both input and simulation data. We set the maximal production amount $R^* = 4$, holding cost $c_H = 0.5$, and backlog cost $c_B = 1$. We consider two scenarios with different input dimension S and input parameter θ^c: $S = 2$, $\boldsymbol{\theta}^c = [4,5]^\top$, and $S = 4$, $\boldsymbol{\theta}^c = [4,5,3,3]^\top$. The optimal order-up-to quantity is selected among $\mathcal{I} = \{1, 2, \ldots, 10\}$. For IZ procedures, we set the IZ preference $\delta = 0.3$. Table 2 shows the empirical PCS and average termination stage along with their 95% confidence intervals

Table 1. PCS and Termination Stage for Quadratic Example (Equal Batch Size)

$\eta = 0.5, m_s = n_s = 10$			
δ	SEIU(IZ)	SEIU-MCB(IZ)	KN++(IZ)
0.3	98% ± 1.7% 77 ± 6.2	82% ± 5% 17 ± 1.9	75% ± 8% 15 ± 1.5
0.2	99% ± 0.6% 111 ± 6.7	88% ± 4% 18 ± 1.4	83% ± 7% 21 ± 1.3
0.1	100% ± 0% 521 ± 31	90% ± 2% 25 ± 0.9	85% ± 6% 34 ± 5.4
$\eta = 0.5, m_s = 10, \delta = 0.2$			
n_s	SEIU(IZ)	SEIU-MCB(IZ)	KN++(IZ)
5	99% ± 0.6% 170 ± 9.7	87% ± 5% 25 ± 1.6	74% ± 5% 24 ± 1.5
10	99% ± 0.6% 111 ± 6.7	88% ± 4% 18 ± 1.4	83% ± 7% 21 ± 1.3
15	97% ± 1.8% 89 ± 5.0	89% ± 4% 16 ± 0.9	86% ± 3% 18 ± 1.0

Table 2. PCS and Termination Stage for Production Inventory Problem with Random Batch Size ($\eta = 0.5$)

$\theta^c = [4,5]^\top (S = 2)$				
$\bar{k}$	SEIU	SEIU(IZ)	SEIU-MCB	SEIU-MCB(IZ)
10	100% ± 0% 364 ± 16	100% ± 0% 205 ± 13	90% ± 4% 44 ± 4.7	81% ± 6% 17 ± 1.2
20	100% ± 0% 156 ± 7.9	100% ± 4% 90 ± 5.5	94% ± 3% 25 ± 2.4	87% ± 5% 11 ± 0.6
30	100% ± 0% 102 ± 5.7	100% ± 4% 54 ± 3.8	96% ± 2% 19 ± 1.4	90% ± 4% 7.5 ± 0.4
$\theta^c = [4,5,3,3]^\top (S = 4)$				
$\bar{k}$	SEIU	SEIU(IZ)	SEIU-MCB	SEIU-MCB(IZ)
10	100% ± 0% 1137 ± 152	100% ± 0% 528 ± 84	93% ± 3% 60 ± 9.6	90% ± 4% 30 ± 3.1
20	100% ± 0% 522 ± 102	100% ± 4% 235 ± 58	95% ± 2% 42 ± 6	92% ± 3% 18 ± 1.6
30	100% ± 0% 301 ± 33	100% ± 4% 139 ± 30	96% ± 2% 36 ± 4.1	90% ± 4% 14 ± 1.2

based on 100 independent runs of each procedure. We summarize our observations as follows:

1. The SEIU and SEIU(IZ) procedures obtain the 100% empirical PCS across all settings, overshooting the target PCS. The improved procedures SEIU-MCB and SEIU-MCB(IZ) significantly relieve the conservativeness and achieve an empirical PCS much closer to the target and terminating in a much smaller number of stages.

2. Comparing the two scenarios in which the number of unknown input distributions $S = 2$ and $S = 4$, the actual difference between the best design and the second best is similar, but all four procedures tend to terminate in later stages when $S = 4$ than $S = 2$. Specifically, the termination stages of SEIU and SEIU(IZ) increase roughly by a factor of three, whereas those of SEIU-MCB and SEIU-MCB(IZ) increase by a factor of less than two. Their empirical PCS also increase and overshoot the target PCS slightly when $S = 4$. This implies that a higher dimension of input uncertainty can cause a larger impact on both the PCS and the termination stage, and SEIU-MCB and SEIU-MCB(IZ) are more robust to this impact.

7.2.3. Optimizing the Drop Rate η. Finally, we explore boosting the performance of SEIU-MCB by optimizing η. The experiment setting and result can be found in Section EC3.3. The result shows the optimized η improves both the PCS and the termination stage.

8. Conclusion and Future Directions

We study ranking and selection under input uncertainty in which input data arrive in time-varying batches over time. A moving average estimator is proposed to aggregate the simulation outputs under different input models across time stages because of the limited simulation budget at each stage. We further propose the SEIU and SEIU-MCB procedures by, respectively, deriving the exact and asymptotic confidence bands in a substantially extended sequential elimination framework. We also derive the corresponding SEIU(IZ) and SEIU-MCB(IZ) procedures to extend our algorithms to the indifference zone setting. We analyze the impact of the drop rate of the moving average estimator on the procedures. Numerical results show the necessity of our proposed procedures under input uncertainty and the statistical validity to achieve the target PCS. The SEIU-MCB procedure is highly efficient for practical use and can be accelerated by optimizing the drop rate.

There are several future research directions. First, other methods can be used to update the input model and aggregate simulation outputs over time stages. Second, other R&S procedures can be extended to the setting of streaming input data, especially when (asymptotic) normality of the aggregated performance estimate could be established. Third, other types of performance measure, such as quantiles, can be considered in data-driven ranking and selection problems.

Acknowledgments

The authors thank the associate editor and anonymous reviewers for their careful reading of the paper and very constructive comments that led to a substantially improved paper.

References

Barton RR, Schruben LW (1993) Uniform and bootstrap resampling of empirical distributions. Evans GW, Mollaghasemi M, Russell EC, Biles WE, eds. *Proc. 1993 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 503–508.

Barton RR, Nelson BL, Xie W (2014) Quantifying input uncertainty via simulation confidence intervals. *INFORMS J. Comput.* 26(1):74–87.

Batur D, Kim SH (2006) Fully sequential selection procedures with a parabolic boundary. *IIE Trans.* 38(9):749–764.

Bechhofer RE (1954) A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Ann. Math. Statist.* 25(1):16–39.

Chang JY, Hsu JC (1992) Optimal designs for multiple comparisons with the best. *J. Statist. Planning Inference* 30(1):45–62.

Cheng RC, Holland W (1997) Sensitivity of computer simulation experiments to errors in input data. *J. Statist. Comput. Simulation* 57(1–4):219–241.

Chick SE (2001) Input distribution selection for simulation experiments: Accounting for input uncertainty. *Oper. Res.* 49(5):744–758.

Corlu CG, Biller B (2013) A subset selection procedure under input parameter uncertainty. Pasupathy R, Kim S-H, Tolk A, Hill R, Kuhl ME, eds. *Proc. 2013 Winter Simulation* (IEEE, Piscataway, NJ), 463–473.

Corlu CG, Biller B (2015) Subset selection for simulations accounting for input uncertainty. Yilmaz L, Chan WKV, Moon I, Roeder TMK, Macal C, Rossetti MD, eds. *Proc. 2015 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 437–446.

Corlu CG, Akcay A, Xie W (2020) Stochastic simulation under input uncertainty: A review. *Oper. Res. Perspect.* 7:100162.

Eckman DJ, Feng MB (2018) Green simulation optimization using likelihood ratio estimators. Rabe M, Juan AA, Mustafee N, Skoogh A, Jain S, Johansson B, eds. *Proc. 2018 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 2049–2060.

Even-Dar E, Mannor S, Mansour Y (2002) PAC bounds for multi-armed bandit and Markov decision processes. Kivinen J, Sloan RH, eds. *Internat. Conf. Comput. Learn. Theory* (Springer, Berlin), 255–270.

Even-Dar E, Mannor S, Mansour Y (2006) Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *J. Machine Learn. Res.* 7:1079–1105.

Fan W, Hong LJ, Nelson BL (2016) Indifference-zone-free selection of the best. *Oper. Res.* 64(6):1499–1514.

Fan W, Hong LJ, Zhang X (2020) Distributionally robust selection of the best. *Management Sci.* 66(1):190–208.

Feng BM, Song E (2019) Efficient input uncertainty quantification via green simulation using sample path likelihood ratios. Mustafee N, Bae K-HG, Lazarova-Molnar S, Rabe M, Szabo C, Haas P, Son Y-J, eds. *Proc. 2019 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3693–3704.

Feng M, Staum J (2015) Green simulation designs for repeated experiments. Yilmaz L, Chan WKV, Moon I, Roeder TMK, Macal C, Rossetti MD, eds. *Proc. 2015 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 403–413.

Feng M, Staum J (2017) Green simulation: Reusing the output of repeated experiments. *ACM Trans. Model. Comput. Simulation* 27(4):1–28.

Frazier PI (2014) A fully sequential elimination procedure for indifference-zone ranking and selection with tight bounds on probability of correct selection. *Oper. Res.* 62(4):926–942.

Fu MC, ed. (2015) *Handbook of Simulation Optimization*, vol. 216 (Springer, Berlin).

Gabillon V, Ghavamzadeh M, Lazaric A (2012) Best arm identification: A unified approach to fixed budget and fixed confidence. Pereira F, Burges CJ, Bottou L, Weinberger KQ, eds. *Adv. Neural Inform. Processing Systems*, vol. 25 (MIT Press, Cambridge, MA), 3212–3220.

Gao S, Xiao H, Zhou E, Chen W (2017) Robust ranking and selection with optimal computing budget allocation. *Automatica* 81:30–36.

Garivier A, Kaufmann E (2016) Optimal best arm identification with fixed confidence. Feldman V, Rakhlin A, Shamir O, eds. *Conf. Learn. Theory* (PMLR, New York), 998–1027.

Goldsman D, Kim SH, Marshall WS, Nelson BL (2002) Ranking and selection for steady-state simulation: Procedures and perspectives. *INFORMS J. Comput.* 14(1):2–19.

Ho YC, Sreenivas R, Vakili P (1992) Ordinal optimization of DEDS. *Discrete Event Dynamic Systems* 2(1):61–88.

Hong LJ (2006) Fully sequential indifference-zone selection procedures with variance-dependent sampling. *Naval Res. Logist.* 53(5):464–476.

Hong LJ (2009) Estimating quantile sensitivities. *Oper. Res.* 57(1): 118–130.

Hong LJ, Nelson BL (2005) The tradeoff between sampling and switching: New sequential procedures for indifference-zone selection. *IIE Trans.* 37(7):623–634.

Hong LJ, Nelson BL (2007) Selecting the best system when systems are revealed sequentially. *IIE Trans.* 39(7):723–734.

Jamieson K, Nowak R (2014) Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. Czachórski T, Gelenbe E, Lent R, eds. *48th Annual Conf. Inform. Sci. Systems* (IEEE, Piscataway, NJ), 1–6.

Karnin Z, Koren T, Somekh O (2013) Almost optimal exploration in multi-armed bandits. Dasgupta S, McAllester D, eds. *Internat. Conf. Machine Learn.* (PMLR, New York), 1238–1246.

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

Kim SH, Nelson BL (2001) A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Model. Comput. Simulation* 11(3):251–273.

Kim SH, Nelson BL (2006) On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Oper. Res.* 54(3):475–488.

Lam H, Qian H (2022) Subsampling to enhance efficiency in input uncertainty quantification. *Oper. Res.* 70(3):1891–1913.

Lam H, Zhou E (2017) The empirical likelihood approach to quantifying uncertainty in sample average approximation. *Oper. Res. Lett.* 45(4):301–307.

Lin Y, Song E, Nelson BL (2015) Single-experiment input uncertainty. *J. Simulation* 9(3):249–259.

Liu T, Zhou E (2019) Online quantification of input model uncertainty by two-layer importance sampling. Preprint, submitted December 24, https://arxiv.org/abs/1912.11172.

Liu T, Zhou E (2020) Simulation optimization by reusing past replications: Don't be afraid of dependence. Bae K-H, Feng B, Kim S, Lazarova-Molnar S, Zheng Z, Roeder T, Thiesing R, eds. *Proc. 2020 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 2923–2934.

Liu T, Lin Y, Zhou E (2021) A Bayesian approach to online simulation optimization with streaming input data. Kim S, Feng B, Smith K, Masoud S, Zheng Z, Szabo C, Loper M, eds. *Proc. 2021 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1–12.

Liu T, Lin Y, Zhou E (2022) Bayesian stochastic gradient descent for stochastic optimization with streaming input data. Preprint, submitted February 15, https://arxiv.org/abs/2202.07581.

Luo J, Hong LJ, Nelson BL, Wu Y (2015) Fully sequential procedures for large-scale ranking-and-selection problems in parallel computing environments. *Oper. Res.* 63(5):1177–1194.

Mannor S, Tsitsiklis JN (2004) The sample complexity of exploration in the multi-armed bandit problem. *J. Machine Learn. Res.* 5:623–648.

Pichitlamken J, Nelson BL, Hong LJ (2006) A sequential procedure for neighborhood selection-of-the-best in optimization via simulation. *Eur. J. Oper. Res.* 173(1):283–298.

Russo D (2016) Simple Bayesian algorithms for best arm identification. Feldman V, Rakhlin A, Shamir O, eds. *Conf. Learn. Theory* (PMLR, New York), 1417–1418.

Song E, Nelson BL (2015) Quickly assessing contributions to input uncertainty. *IIE Trans.* 47(9):893–909.

Song E, Nelson BL (2019) Input–output uncertainty comparisons for discrete optimization via simulation. *Oper. Res.* 67(2): 562–576.

Song E, Shanbhag UV (2019) Stochastic approximation for simulation optimization under input uncertainty with streaming data. Mustafee N, Bae K-HG, Lazarova-Molnar S, Rabe M, Szabo C, Haas P, Son Y-J, eds. *Proc. 2019 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3597–3608.

Van der Vaart AW (2000) *Asymptotic Statistics*, vol. 3 (Cambridge University Press, Cambridge, UK).

Wu D, Zhou E (2017) Ranking and selection under input uncertainty: A budget allocation formulation. Chan WKV, D'Ambrogio A, Zacharewicz G, Mustafee N, Wainer G, Page E, eds. *Proc. 2017 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 2245–2256.

Wu D, Zhou E (2019) Fixed confidence ranking and selection under input uncertainty. Mustafee N, Bae K-HG, Lazarova-Molnar S, Rabe M, Szabo C, Haas P, Son Y-J, eds. *Proc. 2019 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3717–3727.

Xiao H, Gao S (2018) Simulation budget allocation for selecting the top-m designs with input uncertainty. *IEEE Trans. Automatic Control* 63(9):3127–3134.

Xiao H, Gao F, Lee LH (2020) Optimal computing budget allocation for complete ranking with input uncertainty. *IISE Trans.* 52(5):489–499.

Xie W, Nelson BL, Barton RR (2014) A Bayesian framework for quantifying uncertainty in stochastic simulation. *Oper. Res.* 62(6):1439–1452.

Xie W, Nelson BL, Barton RR (2016) Multivariate input uncertainty in output analysis for stochastic simulation. *ACM Trans. Model. Comput. Simulation* 27(1):1–22.

Zhou E, Liu T (2018) Online quantification of input uncertainty for parametric models. *Proc. 2018 Winter Simulation Conf.*, 1587–1598.

Zhou E, Xie W (2015) Simulation optimization when facing input uncertainty. Yilmaz L, Chan WKV, Moon I, Roeder TMK, Macal C, Rossetti MD, eds. *Proc. 2015 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3714–3724.

Zhu H, Liu T, Zhou E (2020) Risk quantification in stochastic simulation under input uncertainty. *ACM Trans. Model. Comput. Simulation* 30(1):1–24.

Zouaoui F, Wilson JR (2003) Accounting for parameter uncertainty in simulation input modeling. *IIE Trans.* 35(9):781–792.

Zouaoui F, Wilson JR (2004) Accounting for input-model and input-parameter uncertainties in simulation. *IIE Trans.* 36(11): 1135–1151.

Di Wu is an applied scientist in Amazon Web Services. He received his PhD degree in operations research from the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology in 2019.

Yuhao Wang is a PhD student in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. His research interests include simulation and stochastic optimization.

Enlu Zhou is an associate professor in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. Her research interests include simulation optimization, stochastic optimization, and stochastic control.