



Operations Research

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

Bonferroni-Free and Indifference-Zone-Flexible Sequential Elimination Procedures for Ranking and Selection

Wenyu Wang, Hong Wan, Xi Chen

To cite this article:

Wenyu Wang, Hong Wan, Xi Chen (2023) Bonferroni-Free and Indifference-Zone-Flexible Sequential Elimination Procedures for Ranking and Selection. *Operations Research*

Published online in *Articles in Advance* 11 Apr 2023

. <https://doi.org/10.1287/opre.2023.2447>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2023, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Methods

Bonferroni-Free and Indifference-Zone-Flexible Sequential Elimination Procedures for Ranking and Selection

Wenyu Wang,^a Hong Wan,^{b,*} Xi Chen^c^aDepartment of Engineering, Eigen Capital LLC, Shenzhen, Guangdong 518031, China; ^bEdward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh, North Carolina 27607; ^cGrado Department of Industrial and Systems Engineering, Virginia Tech, Blacksburg, Virginia 24061

*Corresponding author

Contact: wenyu.wang@eigen.capital (WW); hwan4@ncsu.edu,  <https://orcid.org/0000-0002-2017-4871> (HW); xchen6@vt.edu, <https://orcid.org/0000-0001-7965-9198> (XC)

Received: January 5, 2021

Revised: March 20, 2022; September 5, 2022

Accepted: February 9, 2023

Published Online in Articles in Advance:
April 11, 2023

Area of Review: Simulation

<https://doi.org/10.1287/opre.2023.2447>

Copyright: © 2023 INFORMS

Abstract. This paper proposes two fully sequential procedures for selecting the best system with a guaranteed probability of correct selection (PCS). The main features of the proposed procedures include the following: (1) adopting a Bonferroni-free model that overcomes the conservativeness of the Bonferroni correction and delivers the exact probabilistic guarantee without overshooting; (2) conducting always valid and fully sequential hypothesis tests that enable continuous monitoring of each candidate system and control the type I error rate (or equivalently, PCS) at a prescribed level; and (3) assuming an indifference-zone-flexible formulation, which means that the indifference-zone parameter is not indispensable but could be helpful if provided. We establish statistical validity and asymptotic efficiency for the proposed procedures under normality settings with and without the knowledge of true variances. Numerical studies conducted under various configurations corroborate the theoretical findings and demonstrate the superiority of the proposed procedures.

Funding: W. Wang and H. Wan were supported in part by CollinStar Capital Pty Ltd. X. Chen was supported in part by the National Science Foundation [Grant IIS-1849300 and CAREER CMMI-1846663].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2023.2447>.

Keywords: ranking and selection • indifference-zone-flexible • generalized likelihood ratio statistics

1. Introduction

The ranking and selection (R&S) problem refers to selecting the best system from a finite set of alternatives or systems, where the true performance of each system is unavailable but can be evaluated through sampling or running simulation experiments. R&S problems arise in many practical applications, for example, when Internet companies consider variations of an application to provide the best user experience and when pharmaceutical companies rank treatment options for a target disease, just to name a few.

This paper focuses on the R&S problem of selecting the best system defined as the one with the maximum expected performance. There has been significant interest over the last decades in developing procedures for solving this R&S problem. We refer the interested reader to Kim and Nelson (2006b, 2007), Branke et al. (2007), and Hong et al. (2021) for thorough overviews on existing R&S procedures. Most R&S procedures fall into two categories: fixed-precision and fixed-budget approaches. The fixed-budget techniques aim to efficiently allocate a fixed computing budget among systems to optimize

an objective (e.g., maximize the posterior *probability of correct selection (PCS)*). Within this category, the most widely applied and well-studied algorithms include the optimal computing budget allocation (Chen et al. 2000) and the expected value of information approaches (Chick and Inoue 2001).

In this paper, we focus on solving the R&S problem via a fixed-precision approach, which intends to achieve a prespecified PCS with adaptive sampling. Recent years have witnessed a burgeoning literature on fixed-precision R&S procedures. Typically, fixed-precision approaches assume one of the following three formulations: (a) the subset-selection formulation, (b) the indifference-zone formulation, and (c) the indifference-zone-free formulation. Procedures in the subset-selection formulation guarantee to select a subset of systems that includes the best one with a prescribed high probability (Gupta 1956, 1965); these procedures typically consist of two stages: preliminary evaluation and final elimination. In contrast, procedures in the latter two formulations typically proceed sequentially and terminate after eliminating all systems but one. In particular, procedures in the indifference-zone

formulation guarantee to select the best system with a prespecified *PCS* when an indifference-zone parameter $\delta > 0$ is specified—The mean performance of the best system is greater than that of all other systems by at least δ . Most of the fixed-precision procedures are indifference-zone based. The popular ones include the KN procedure (Kim and Nelson 2001), the KN++ procedure (Kim and Nelson 2006a), the Bayes-inspired indifference zone procedures (Frazier 2014), and so on. Fan et al. (2016) proposed an indifference-zone-free procedure (IZ-free), which can deliver a prescribed *PCS* guarantee without using an indifference-zone parameter. However, IZ-free is more conservative and computationally expensive than the leading indifference-zone procedures if the indifference-zone information is appropriately specified.

There is still an absence of procedures that can bridge the indifference-zone and indifference-zone-free formulations. The indifference-zone procedures are typically (much) more computationally efficient with an appropriate specification of the indifference-zone parameter (Fan et al. 2016). The IZ-free approach is more flexible but cannot use the indifference-zone information even if it is available. Therefore, we are inspired to propose an indifference-zone-flexible formulation, and design corresponding procedures that (1) can deliver the *PCS* guarantee either with or without using the indifference-zone parameter and (2) can take advantage of the parameter if specified.

Another persistent issue that hinders the efficiency of many leading procedures is the use of the Bonferroni adjustment to achieve the prescribed *PCS* guarantee under multiple comparisons (Kim and Nelson 2001, 2006a; Fan and Hong 2014; Fan et al. 2016). That is, dividing the probability of incorrect selection by a factor of $K - 1$, where K denotes the total number of systems under consideration. The consequence is that fixed-precision procedures relying on the Bonferroni correction become computationally expensive and conservative in error rate control when applied for solving large-scale R&S problems (i.e., when K is large). A significant breakthrough is the Bayes-inspired indifference zone procedures (BIZ) developed by Frazier (2014). BIZ can sequentially eliminate inferior systems without resorting to the Bonferroni correction and is found to outperform leading fixed-precision procedures on large-scale R&S problems. It is worth noting that researchers have made continual efforts to tackle the conservativeness of fixed-precision procedures (Wang and Kim 2013; Dieker and Kim 2014, 2021). A notable effort is Dieker and Kim (2021) who propose the DK procedures using a spherical elimination boundary derived from properties of a multidimensional Brownian motion. DK is shown to deliver similar or better performance than BIZ. Nevertheless, both DK and BIZ are tailored for normal distribution and assume the indifference-zone formulation.

To address the previous challenges, we propose new R&S procedures with the following features:

1. **A multiple composite hypothesis testing model.** We model the R&S problem as a multiple composite hypothesis testing problem so that the developed procedures can achieve the desired error rate control without the Bonferroni correction; see Section 2.2 for a more detailed discussion.

2. **Test martingale-based procedures.** We propose a fully sequential hypothesis testing meta procedure whose test statistics and elimination threshold correspond to a test martingale and can be used to continuously monitor and control the type I error rate (or equivalently, *PCS*) at a prescribed level; see Section 2.3 for details.

3. **Procedures in the indifference-zone-flexible formulation.** Specific instantiations of the meta procedure under normality settings are detailed in Sections 3 and 4, where the statistical efficiency and the asymptotic optimality in terms of the average sample size used are verified. These concrete procedures can deliver the prescribed *PCS* guarantee with and without using the indifference-zone parameter.

Our proposed procedures have their roots in sequential analysis. Since the seminal work on the sequential probability ratio test (SPRT; Wald 1945, Wald and Wolfowitz 1948), the last eight decades have witnessed the significant development of sequential likelihood-based hypothesis testing theory and methodologies and their applications in various science and engineering domains. SPRT formulates the sequential testing of two simple hypotheses via the boundary crossing of the likelihood ratio statistic, which is shown to be optimal in terms of minimal expected sample size for achieving prescribed error probabilities. The generalized sequential probability ratio test (GSPRT) provides a natural generalization of the SPRT concerning the problem formulation and the stopping rule, which tackles the sequential testing of two or more composite hypotheses and adopts the generalized likelihood ratio (GLR) test statistic (Li et al. 2014). Over the years, there has been a plethora of work on designing GLR test statistics and proving their asymptotic optimality under various conditions (Lorden 1976; Lai 1977, 1981; Pavlov 1988, 1991; Li et al. 2014). The interested reader is referred to Tartakovsky et al. (2014) for a comprehensive discussion of sequential analysis and GSPRT. In this work, we formulate the R&S problem as a multiple composite hypothesis testing problem, construct adaptive GLR statistics, and provide new R&S procedures; to the best of our knowledge, this work is the first to solve the R&S problem from this new perspective.

The rest of this paper is organized as follows. Section 2 provides a high-level overview of the proposed meta procedure developed according to the multiple composite hypothesis testing model and associated key properties.

Section 3 considers solving the R&S problem of interest under normality with known variances by presenting concrete procedures developed and proving their statistical validity and computational efficiency. Section 4 parallels Section 3 and tackles the R&S problem under normality with unknown variances. Numerical studies are conducted in Section 5. Section 6 concludes the paper.

2. Problem Statement and Methodology Overview

In this section, we formulate the R&S problem of interest as a multiple composite hypothesis testing problem, propose a meta procedure to solve it, and reveal the essential properties of the proposed meta procedure.

2.1. Problem Setup

Let us consider the R&S problem of selecting the best system defined as the one with the maximum expected performance; that is, we are interested in solving the following discrete simulation optimization problem:

$$\arg \max_{i \in [K]} \mu_i^* = \mathbb{E}[X_i], \quad (1)$$

where $[K] = \{1, 2, \dots, K\}$ denotes the set of all candidate systems; and we assume that independent and identically distributed random observations of system i , $\{X_{i,n}, n = 1, 2, \dots\}$, are drawn from a distribution, with $\mu_i^* = \mathbb{E}[X_{i,n}]$ denoting the unknown true mean value of system i , $\forall i \in [K]$. Because the distribution corresponding to system i may be parameterized by unknown parameters other than the mean value, we denote $\theta_i = (\mu_i, \eta_{i,1}, \eta_{i,2}, \dots, \eta_{i,p})$ as the $(p + 1)$ -dimensional parameter vector for system i . For instance, under normality with an unknown variance, $\theta_i = (\mu_i, \sigma_i^2)$, $\forall i \in [K]$. Denote the true parameter vector of system i by θ_i^* , $\forall i \in [K]$. Correspondingly, let us denote $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_K) \in \mathbb{U} \subseteq \mathbb{R}^K$ as the mean vector and $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_K) \in \Theta \subseteq \mathbb{R}^{(p+1)K}$ as the vector comprising all unknown parameters for all K systems, where \mathbb{U} denotes the parameter space of all possible values of $\boldsymbol{\mu}$ and Θ represents the parameter space that contains all possible values of $\boldsymbol{\theta}$. Denote $\boldsymbol{\mu}^* = (\mu_1^*, \mu_2^*, \dots, \mu_K^*)$ as the true mean vector and $\boldsymbol{\theta}^* = (\theta_1^*, \theta_2^*, \dots, \theta_K^*)$ as the true parameter vector.

Let $\delta > 0$ denote the indifference-zone parameter—A system is only considered better than all other systems if its mean value is greater than the others' by at least δ . Without loss of generality, assume that there exists a single best system. Consider partitioning the parameter space Θ into $(K + 1)$ disjoint subsets. Specifically, given a prescribed value of δ , let Θ_i denote the set of parameter values such that system i is the best among all systems, that is, $\Theta_i = \{\boldsymbol{\theta} \in \Theta \mid \mu_i \geq \mu_j + \delta, \forall j \neq i, j \in [K]\}$, $\forall i \in [K]$; the indifference parameter set Θ_0 contains all possible values of $\boldsymbol{\theta}$ such that no single system dominates the others by at least δ , that is, $\Theta_0 = \{\boldsymbol{\theta} \in \Theta \mid \forall i \in$

$[K], \exists j \neq i, j \in [K]$ such that $\mu_i < \mu_j + \delta\}$. Therefore, $\{\Theta_0, \Theta_1, \Theta_2, \dots, \Theta_K\}$ satisfy $\Theta = \bigcup_{i=0}^K \Theta_i$ and $\Theta_i \cap \Theta_j = \emptyset$ for $i \neq j, i, j \in \{0\} \cup [K]$. In the degenerate case where the indifference-zone parameter δ is unspecified, we have $\Theta_i = \{\boldsymbol{\theta} \in \Theta \mid \mu_i > \mu_j, \forall j \neq i, j \in [K]\}$ and $\Theta_0 = \{\boldsymbol{\theta} \in \Theta \mid \forall i \in [K], \exists j \neq i, j \in [K]$ such that $\mu_i \leq \mu_j\}$.

2.2. Multiple Composite Hypothesis Testing Model for the R&S Problem

In this section, we provide a formal discussion on formulating the R&S problem of interest as a multiple composite hypothesis testing problem. Let $(\Omega, \mathcal{F}, P_{\boldsymbol{\theta}^*})$ be the probability space on which a sequence of independent random observations of the K systems are defined, where $P_{\boldsymbol{\theta}^*}$ denotes the probability measure in accordance with the true parameter vector $\boldsymbol{\theta}^*$. To ease the notation, we suppress the dependence of the probability measure on $\boldsymbol{\theta}^*$ hereinafter.

Specifically, the R&S problem in (1) is equivalent to the following *multiple composite hypothesis testing problem* with K hypotheses:

$$H_i : \boldsymbol{\theta}^* \in \Theta_i, \quad i \in [K], \quad (2)$$

where $[K] = \{1, 2, \dots, K\}$ and the i th hypothesis states that the true parameter vector $\boldsymbol{\theta}^*$ lies in the parameter subset Θ_i , indicating that system i is the best system. Under the assumption that there exists a single best system (to be made more explicit later), one of the K hypotheses must be true. In other words, the R&S problem in (1) is now equivalent to identifying the one true hypothesis from K of them.

In sharp contrast, the KN family procedures model the R&S problem as a *multiple comparisons problem*, whose hypothesis test formulation was first noted by Hong et al. (2021). Specifically, for each $i \in [K]$, one considers the following test with two hypotheses:

$$H_{0,i} : \mu_i^* \leq \max_{j \neq i} \mu_j^* \quad \text{vs.} \quad H_{1,i} : \mu_i^* > \max_{j \neq i} \mu_j^*, \quad (3)$$

where the pair of hypotheses ($H_{0,i}$ and $H_{1,i}$) focuses on comparisons between system i and all other systems. Model (3) can be rewritten using the notation established in Section 2.1 as follows:

$$H_{0,i} : \boldsymbol{\theta}^* \notin \Theta_i \quad \text{vs.} \quad H_{1,i} : \boldsymbol{\theta}^* \in \Theta_i, \quad i \in [K]. \quad (4)$$

Despite their similar appearance, Models (2) and (4) differ significantly with respect to error rate management. Recall from Section 1 that any fixed-precision procedure intends to achieve a target *PCS* level of, say, $1 - \alpha$ (with $\alpha \in (0, 1)$). Here *PCS* can be expressed as $PCS = P(\text{select system } i^*)$, where i^* denotes the index of the true best system; we assume that there is a single best system, formal statements of which will be provided in the form of assumptions in Sections 3 and 4.

Under the multiple comparisons Model (4), for a procedure to correctly select the true best system i^* ,

it must reject H_{0,i^*} . To this end, existing procedures in the KN family typically proceed sequentially to eliminate $(K - 1)$ systems and conduct error rate control according to the following rationale:

$$\begin{aligned} PCS &= P(\text{reject } H_{0,i^*}) \\ &= 1 - P(\text{fail to reject } H_{0,j} \text{ for some } j \neq i^*, j \in [K]) \end{aligned} \quad (5)$$

$$\geq 1 - \sum_{j \in [K] \setminus \{i^*\}} P(\text{fail to reject } H_{0,j}), \quad (6)$$

where (5) implies that achieving the PCS guarantee for procedures in the KN family is equivalent to controlling the type II error rate when testing $(K - 1)$ pairs of hypotheses given in (4). Furthermore, to control the type I error rate at the prescribed level of α based on (6), these procedures rely on the Bonferroni correction and require

$$P(\text{fail to reject } H_{0,j}) \leq \frac{\alpha}{K-1}, \quad \forall j \neq i^*, j \in [K].$$

In strong contrast, we propose procedures that directly control the error rate achieved for a single hypothesis testing problem, that is, the multiple composite hypothesis testing Model (2), hence avoiding using the Bonferroni correction to allocate the prescribed error probability α over $(K - 1)$ tests. More specifically, the PCS of a procedure under Model (2) can be written as

$$PCS = P(\text{select system } i^*) = P(\text{fail to reject } H_{i^*}).$$

We elaborate on how to achieve a prescribed PCS guarantee in the next section.

2.3. Meta Procedure and Some Essential Properties

In this section, we propose a high-level meta procedure according to Model (2) with the three features mentioned in Section 1. Specific instantiations of this meta procedure along with detailed theoretical guarantees are provided in Sections 3 and 4.

The meta procedure, referred to as $Proc(\Lambda_{i,n}, \mathcal{B})$, is detailed in Algorithm 1. The meta procedure is fully sequential and comprises two major components: (i) the test statistics, $\Lambda_{i,n} \geq 0$ for all $i \in [K]$, where $\Lambda_{i,n}$ denotes the generalized likelihood ratio statistic corresponding to the i th hypothesis in Model (2) and n stands for the iteration index; and (ii) the elimination threshold, \mathcal{B} , to reject the hypotheses corresponding to inferior systems.

Let $\{\mathcal{S}_n \subseteq [K], n \geq 1\}$ denote a sequence of index sets with \mathcal{S}_n comprising systems that have not been eliminated by the end of the n th iteration. The procedure starts with all systems in contention for being the best one on the first iteration, i.e., $\mathcal{S}_1 = [K]$, and keeps eliminating systems sequentially until $|\mathcal{S}_\tau| \leq 1$, where $|\mathcal{A}|$ denotes the cardinality of set \mathcal{A} and τ denotes the index of the iteration on which the procedure terminates. The first n_0 iterations are used to warm up the

procedure, and elimination only takes place on the $(n_0 + 1)$ th iteration and onward. We note that n_0 is a tunable parameter that can impact the computational efficiency of specific instantiations of the meta procedure with no influence on their statistical validity. In addition, it is possible for instantiations of the meta procedure to return an empty candidate set upon termination, that is, $|\mathcal{S}_\tau| = 0$. In practical applications, one may select the system with the highest sample mean when this occurs; however, this does not impact the theoretical guarantees associated with instantiations of the meta procedure to be presented later.

Algorithm 1 (Meta Procedure $Proc(\Lambda_{i,n}, \mathcal{B})$)

Inputs:

- 1: $\mathcal{B} = \alpha$ ▷ Set the elimination threshold to a prescribed error level α
- 2: $[K] = \{1, 2, \dots, K\}$

Procedure:

- 1: set $n \leftarrow 1$;
- 2: set $\mathcal{S}_1 \leftarrow [K]$;
- 3: **repeat**
- 4: generate one observation for each system with its index in \mathcal{S}_n ;
- 5: set $\mathcal{S}_{n+1} \leftarrow \mathcal{S}_n$; ▷ Initialize \mathcal{S}_{n+1} which will be updated later accordingly
- 6: **if** $n \geq n_0 + 1$ **then** ▷ No elimination is conducted in the first n_0 iterations
- 7: **for** $i \in \mathcal{S}_n$ **do**
- 8: calculate the test statistic $\Lambda_{i,n}$;
- 9: **if** $\Lambda_{i,n} < \mathcal{B}$ **then**
- 10: set $\mathcal{S}_{n+1} \leftarrow \mathcal{S}_{n+1} \setminus \{i\}$; ▷ Eliminate system i from \mathcal{S}_{n+1}
- 11: **end if**
- 12: **end for**
- 13: **end if**
- 14: set $n \leftarrow n + 1$;
- 15: **until** $|\mathcal{S}_n| \leq 1$ ▷ Terminate when no more than one system is left in \mathcal{S}_n
- 16: **return** \mathcal{S}_n

Recall that (Ω, \mathcal{F}, P) denotes the probability space on which random observations of all K systems are defined. Let $\mathcal{F}_n \subset \mathcal{F}$ for $n \in \mathbb{Z}^+ = \{1, 2, \dots\}$ be a nondecreasing system of sub- σ -algebras generated by the random observations of all K systems obtained by the end of the n th iteration.

The meta procedure, $Proc(\Lambda_{i,n}, \mathcal{B})$, can be designed to be a valid procedure corresponding to a *sequential test of power one* (Robbins and Siegmund 1974, Siegmund 1976), which refers to those sequential tests that do not reject the true hypothesis in finite time almost surely under P (referred to as almost surely or with probability one for short hereinafter). In the context of R&S, the meta procedure corresponds to a size α sequential test of power one if it admits the following

property:

$$P(\tau_i < \infty) \begin{cases} \leq \alpha, & \text{for } i = i^*; \\ = 1, & \text{for } i \neq i^*, \end{cases} \quad (7)$$

where τ_i denotes the elimination time of system i for each $i \in [K]$ and recall that i^* denotes the true best system's index. That is, the probability of eliminating system i^* is bounded above by α , and hence the target PCS level of $1 - \alpha$ can be achieved.

The rationale behind the proposed meta procedure is to conduct *continuous* monitoring of the test statistics for K hypotheses in Model (2) and reject a hypothesis H_i as soon as its corresponding test statistic $\Lambda_{i,n}$ becomes smaller than \mathcal{B} for the first time. Therefore, requiring the meta procedure to fulfill (7) is equivalent to designing the major components, $\Lambda_{i,n}$ for all $i \in [K]$ and \mathcal{B} , such that

$$P\left(\inf_{n \geq 1} \Lambda_{i,n} \leq \mathcal{B}\right) \begin{cases} \leq \alpha, & \text{if } i = i^*; \\ = 1, & \text{if } i \neq i^*, \end{cases} \quad (8)$$

where i^* denotes the index of the true hypothesis in Model (2).

For the meta procedure to achieve (8), we propose to use the elimination threshold $\mathcal{B} = \alpha$ and the test statistics $\{\Lambda_{i,n}, i \in [K], n \geq 1\}$ that satisfy the following conditions: (i) the sequence of test statistics $\{\Lambda_{i^*,n}^{-1}, n \geq 1\}$ corresponding to system i^* is associated with a test martingale; and (ii) the sequence $\{\Lambda_{i,n}, n \geq 1\}$ corresponding to system i ($\forall i \in [K] \setminus \{i^*\}$) converges to zero with probability one. Formally, a test martingale is defined as a nonnegative martingale $\{Z_n \geq 0, n \geq 1\}$ with $\mathbb{E}[Z_1] = 1$ relative to a sequence of σ -algebras (Shafer et al. 2011); for ease of exposition, consider that the Z_n are adapted to the filtration formed by the \mathcal{F}_n defined earlier in the section. Test martingales yield the maximal inequality property: For any test martingale $\{Z_n, n \geq 1\}$, it holds that $P(\sup_{n \geq 1} Z_n \geq \alpha^{-1}) \leq \alpha$, which follows from Doob's martingale inequality (Revuz and Yor 1999). Therefore, if $\{\Lambda_{i^*,n}^{-1}, n \geq 1\}$ for system i^* is bounded above by a test martingale $\{Z_n, n \geq 1\}$ surely and $\mathcal{B} = \alpha$ is adopted, we have (8) fulfilled, because

$$P\left(\inf_{n \geq 1} \Lambda_{i^*,n} \leq \mathcal{B}\right) = P\left(\sup_{n \geq 1} \frac{1}{\Lambda_{i^*,n}} \geq \frac{1}{\alpha}\right) \leq \alpha. \quad (9)$$

It is worth noting that the meta procedure enables one to monitor the test statistics continuously and “peek” at the test without compromising the type I error (or equivalently, the PCS) guarantee: This feature can be best explained through its connection to always valid p value processes (Johari et al. 2019). An always valid p value process is a sequence of p values $\{p_n, n \geq 1\}$ that can control the type I error rate at level $\alpha \in (0, 1)$ uniformly across the entire sample path under the true hypothesis—such a process satisfies the following

two properties: (1) for any $\alpha \in (0, 1)$ and any $m \geq 1$, $P(p_m \leq \alpha) \leq \alpha$, which is known as the **superuniform** property; and (2) $p_m \geq p_{m+1}$ for any $m \geq 1$ holds surely, which is referred to as the **nonincreasing** property. As noted by Wasserman et al. (2020), unlike “peeking” at the fixed-horizon hypothesis testing that can greatly inflate the type I error rate, the hypothesis testing built on an always valid p value process is able to achieve a prescribed type I error level uniformly across the entire sample path (i.e., across all time steps $m \geq 1$). We are now in a position to point out that $\{\inf_{m \geq \ell \geq 1} \Lambda_{i^*,\ell}, m \geq 1\}$ corresponding to system i^* is an always valid p value process; The second property is satisfied thanks to the definition of the sequence (i.e., $p_m := \inf_{m \geq \ell \geq 1} \Lambda_{i^*,\ell}$) and the first property follows from (9), because

$$P(\text{reject } H_{i^*} \text{ at iteration } m) \leq P\left(\inf_{m \geq \ell \geq 1} \Lambda_{i^*,\ell} \leq \alpha\right) \\ \leq P\left(\inf_{n \geq 1} \Lambda_{i^*,n} \leq \alpha\right) \leq \alpha, \quad \forall m \geq 1.$$

In this section, we provided the meta procedure and proposed to use the elimination threshold $\mathcal{B} = \alpha$ together with the test statistics $\{\Lambda_{i,n}, i \in [K], n \geq 1\}$ satisfying some conditions (without arguing for their existence). In Sections 3 and 4, we will provide specific instantiations of this meta procedure under normality with concrete forms of $\{\Lambda_{i,n}, i \in [K], n \geq 1\}$ provided.

3. Setting of Normality with Known Variances

This section considers solving an R&S problem under the normality setting with known variances. We first present two concrete procedures in the form of the meta procedure provided in Section 3.1 and elaborate on their theoretical properties in Sections 3.2 and 3.3.

To accommodate the normality setting with known variances, the exposition of this section relies on similar notation as established in Section 2.1, but with θ , $\boldsymbol{\theta}$, and Θ replaced by μ , $\boldsymbol{\mu}$, and \mathbb{U} , respectively. For example, $\boldsymbol{\theta}$, the vector comprising the unknown parameters of all K systems, becomes $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_K) \in \mathbb{U}$, where the parameter space Θ reduces to $\mathbb{U} \subseteq \mathbb{R}^K$. The other symbols can be understood analogously when they arise.

Throughout this section, we stipulate the following assumptions.

Assumption 1. *The K systems have distinct true mean values, that is, $\mu_i^* \neq \mu_j^*$, for $i \neq j, i, j \in [K]$. The true mean values of the K systems satisfy that $\mu_{i^*}^* \geq \mu_i^* + \delta$ for all $i \neq i^*, i \in [K]$, where i^* denotes the index of the best system and $\delta > 0$ is the indifference-zone parameter if specified.*

Assumption 2. *The observations of system i , $\{X_{i,n}, n = 1, 2, \dots\}$ are independent and normally distributed with mean μ_i^* and variance σ_i^2 , that is, $\mathcal{N}(\mu_i^*, \sigma_i^2)$, where σ_i^2 is known and μ_i^* is unknown. Furthermore, the observations of different systems are statistically independent, that is,*

$X_{i,n}$ and $X_{j,n'}$ are independent for $n, n' \geq 1$, if $i \neq j$, $i, j \in [K]$.

Assumption 3. The parameter space $\mathbb{U} \subseteq \mathbb{R}^K$ is compact.

Assumption 4. There exists some $r \geq 1$ such that $\sum_{m=1}^{\infty} \mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)} [\rho(\boldsymbol{\mu}^*, \hat{\boldsymbol{\mu}}_m)]^r < \infty$, where $\hat{\boldsymbol{\mu}}_m$ represents the maximum likelihood estimate (MLE) of $\boldsymbol{\mu}^*$ obtained using observations from the first m iterations, and $\rho(\boldsymbol{\mu}, \boldsymbol{\mu}')$ denotes the Kullback-Leibler (KL) divergence from $f(\cdot|\boldsymbol{\mu}')$ to $f(\cdot|\boldsymbol{\mu})$ that can be written as

$$\rho(\boldsymbol{\mu}, \boldsymbol{\mu}') = \mathbb{E}_{f(\cdot|\boldsymbol{\mu})} \left[\log \left(\frac{f(\cdot|\boldsymbol{\mu})}{f(\cdot|\boldsymbol{\mu}')} \right) \right],$$

with $\mathbb{E}_{f(\cdot|\boldsymbol{\mu})}$ denoting the expectation taken with respect to $f(\cdot|\boldsymbol{\mu})$. Here, $f(\cdot|\boldsymbol{\mu}) := \prod_{i=1}^K g_i(\cdot|\mu_i)$ denotes the joint density function of observations from all K systems, and $g_i(\cdot|\mu_i)$ denotes the normal density function corresponding to system i with the mean value being specified as μ_i , $\forall i \in [K]$.

Some remarks follow from Assumptions 1–4. Assumption 1 is stipulated such that a single best system exists and its true mean value is higher than those of all other systems by at least a margin of the indifference-zone parameter if specified. Assumption 2 requires no dependencies among systems so that the eliminations of two distinct systems are statistically independent; it can be replaced by a less restrictive conditional-independent condition to allow the use of common random numbers. Assumptions 2 and 3 imply that the density functions are of normal distributions with a compact parameter space. Assumption 4 stipulates that $\mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)} [\rho(\boldsymbol{\mu}^*, \hat{\boldsymbol{\mu}}_m)]^r$ converges to zero sufficiently fast for some $r \geq 1$. This condition holds for many typical distributions in the multivariate exponential family including Gaussian, Poisson, Bernoulli, and so on (e.g., see condition (5.84) in Tartakovsky et al. (2014)).

3.1. Two Concrete Procedures

In this section, we first provide two concrete procedures in the form of the meta procedure given in Algorithm 1, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, which are suitable for the normality setting with known variances, then we reveal the relationship between them.

3.1.1. Procedure $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. The test statistic used in $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ is a generalized likelihood ratio (GLR) statistic, denoted by $\Lambda_{i,n}^{(1)}$, which is defined as follows:

$$\Lambda_{i,n}^{(1)} = \frac{\sup_{\boldsymbol{\mu} \in \mathbb{U}_i} L_n(\boldsymbol{\mu})}{\pi_n}, \quad n \geq n_0 + 1, \quad \forall i \in [K], \quad (10)$$

where $\mathbb{U}_i = \{\boldsymbol{\mu} \in \mathbb{U} | \mu_i \geq \mu_j + \delta, \forall j \neq i, j \in [K]\}$ if the indifference-zone parameter $\delta > 0$ is specified; otherwise, $\mathbb{U}_i = \{\boldsymbol{\mu} \in \mathbb{U} | \mu_i > \mu_j, \forall j \neq i, j \in [K]\}$. Notice that $\Lambda_{i,n}^{(1)}$ in (10) is defined for $n \geq n_0 + 1$, because $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ does not start any system elimination using the test statistic until the $(n_0 + 1)$ th iteration (recall Algorithm 1).

We may set $\Lambda_{i,n}^{(1)} = 1$ for $1 \leq n \leq n_0$ for completeness; however, this is irrelevant to the implementation and the statistical validity of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$.

In (10), $L_n(\boldsymbol{\mu})$ denotes the likelihood of all observations collected up to the n th iteration ($n \geq 1$) for all K systems, which is given by

$$L_n(\boldsymbol{\mu}) = \prod_{m=1}^n \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \mu_i). \quad (11)$$

On the other hand, π_n in (10) is defined as

$$\begin{aligned} \pi_n &= L_{n_0}(\hat{\boldsymbol{\mu}}_{n_0}) \cdot \prod_{m=n_0+1}^n \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,m-1}), \\ &= \prod_{m=1}^{n_0} \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,n_0}) \cdot \prod_{m=n_0+1}^n \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,m-1}), \\ & \quad n \geq n_0 + 1, \end{aligned} \quad (12)$$

which consists of two components. The first term, $L_{n_0}(\hat{\boldsymbol{\mu}}_{n_0})$, denotes the likelihood of obtaining the observations collected during the warm-up period (i.e., the first n_0 iterations) calculated using the MLE of $\boldsymbol{\mu}^*$ attained by the end of the warm-up period, $\hat{\boldsymbol{\mu}}_{n_0} = (\hat{\mu}_{1,n_0}, \hat{\mu}_{2,n_0}, \dots, \hat{\mu}_{K,n_0})$. The second term, $\prod_{m=n_0+1}^n \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,m-1})$, denotes the adaptive maximum likelihood of obtaining the individual observations after the warm-up period using the corresponding most updated MLEs. That is, we use $\hat{\mu}_{i,m-1}$ —the MLE of μ_i^* attained based on the observations of system i ($i \in \mathcal{S}_m$) collected up to the $(m-1)$ th iteration—in calculating the likelihood of seeing the m th observation for system i , $m = n_0 + 1, \dots, n$.

We next elaborate on the calculation of $\Lambda_{i,n}^{(1)}$. Under Assumptions 1–4, the density function $g_j(\cdot|\mu_j)$ in (11) and (12) follows as

$$g_j(X_{j,m} | \mu_j) = \frac{1}{\sqrt{2\pi\sigma_j^2}} \exp\left(-\frac{(X_{j,m} - \mu_j)^2}{2\sigma_j^2}\right). \quad (13)$$

Denote $\tau_j^{(1)}$ as the stopping time for system j , that is, the iteration index that system j gets eliminated by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$, $\forall j \in [K]$; we have $\hat{\mu}_{j,n}$, the MLE of μ_j^* obtained based on the observations of system j collected up to the n th iteration, given by

$$\hat{\mu}_{j,n} := \arg \sup_{\mu \in \mathbb{U}} L_n(\boldsymbol{\mu}) = \frac{\sum_{m=1}^n \mathbf{1}\{\tau_j^{(1)} \geq m\} X_{j,m}}{\sum_{m=1}^n \mathbf{1}\{\tau_j^{(1)} \geq m\}}, \quad \forall j \in [K].$$

Finally, the following result details the calculation of $\sup_{\boldsymbol{\mu} \in \mathbb{U}_i} L_n(\boldsymbol{\mu})$ used in (10), and its proof is given in Online Appendix EC.2.2.

Theorem 1. Under Assumptions 1–4, denote $\hat{\boldsymbol{\mu}}_n^{\mathbb{U}_i} := \arg \sup_{\boldsymbol{\mu} \in \mathbb{U}_i} L_n(\boldsymbol{\mu})$ as the constrained MLE of the mean vector $\boldsymbol{\mu}$ —obtained over the parameter space \mathbb{U}_i ($\forall i \in [K]$)—by the end of the n th iteration. Then $\hat{\boldsymbol{\mu}}_n^{\mathbb{U}_i} = (\hat{\mu}_{1,n}^{\mathbb{U}_i}, \hat{\mu}_{2,n}^{\mathbb{U}_i}, \dots, \hat{\mu}_{K,n}^{\mathbb{U}_i})$ can

be obtained as follows:

$$\hat{\mu}_{j,n}^{\mathbb{U}_i} = \begin{cases} \hat{\mu}_{\cdot,n}^{\mathbb{U}_i} + \delta, & \text{if } j = i; \\ \hat{\mu}_{\cdot,n}^{\mathbb{U}_i}, & \text{if } \hat{\mu}_{j,n} > \hat{\mu}_{i,n} - \delta; \\ \hat{\mu}_{j,n}^{\mathbb{U}_i}, & \text{if } \hat{\mu}_{j,n} \leq \hat{\mu}_{i,n} - \delta, \end{cases} \quad (14)$$

where $\hat{\mu}_{\cdot,n}^{\mathbb{U}_i}$ denotes the average of all observations from those systems whose MLEs are higher than the MLE of system i minus δ :

$$\hat{\mu}_{\cdot,n}^{\mathbb{U}_i} = \frac{\sum_{m=1}^n \sum_{j \in \mathcal{S}_m} \mathbf{1}\{\hat{\mu}_{j,n} \geq \hat{\mu}_{i,n} - \delta\} X_{j,m} \sigma_j^{-2} - n \delta \sigma_i^{-2}}{\sum_{m=1}^n \sum_{j \in \mathcal{S}_m} \mathbf{1}\{\hat{\mu}_{j,n} \geq \hat{\mu}_{i,n} - \delta\} \sigma_j^{-2}}, \quad (15)$$

with $\delta > 0$ denoting the indifference-zone parameter if specified and $\delta = 0$ otherwise.

3.1.2. Procedure $Proc(\Lambda_{i,n}^{(2)}, \alpha)$. The test statistic for system i used in $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, $\Lambda_{i,n}^{(2)}$, is defined as follows:

$$\Lambda_{i,n}^{(2)} = \min_{j \neq i, j \in [K]} \Lambda_{i,j,n}^{(2)}, \quad n \geq n_0 + 1, \quad \forall i \in [K],$$

where $\Lambda_{i,j,n}^{(2)}$ is the minimum of all pairwise statistics $\Lambda_{i,j,n}^{(2)}$ s that are defined as

$$\Lambda_{i,j,n}^{(2)} = \frac{\sup_{\boldsymbol{\mu} \in \mathbb{U}_{ij}} L_n(\boldsymbol{\mu})}{\pi_n}, \quad n \geq n_0 + 1, \quad j \neq i, \quad j \in [K]. \quad (16)$$

The definitions of $\Lambda_{i,n}^{(2)}$ and $\Lambda_{i,j,n}^{(2)}$ are specified for $n \geq n_0 + 1$ due to the same reason as given for defining $\Lambda_{i,n}^{(1)}$ in (10) for $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. We may set $\Lambda_{i,n}^{(2)} = 1$ for $1 \leq n \leq n_0$ for completeness, but this impacts neither the implementation nor the statistical validity of $Proc(\Lambda_{i,n}^{(2)}, \alpha)$.

In (16), $L_n(\boldsymbol{\mu})$ and π_n are, respectively, as given in (11) and (12), with $L_n(\boldsymbol{\mu})$ being maximized over the parameter space \mathbb{U}_{ij} , where $\mathbb{U}_{ij} = \{\boldsymbol{\mu} \in \mathbb{U} \mid \mu_i \geq \mu_j + \delta\}$ if $\delta > 0$ is specified; otherwise, $\mathbb{U}_{ij} = \{\boldsymbol{\mu} \in \mathbb{U} \mid \mu_i > \mu_j\}$. Notice that $\Lambda_{i,j,n}^{(2)}$ essentially gives the likelihood that system i is better than system j , and the test statistic $\Lambda_{i,n}^{(2)}$ is the minimax statistic for the likelihood that system i dominates all other systems. The calculation of $\Lambda_{i,n}^{(2)}$ can be carried out in the same vein as that of $\Lambda_{i,n}^{(1)}$ in (10); hence, we omit the details for the sake of brevity.

We close this section with the following result that sheds light on the relationship between the two test statistics $\Lambda_{i,n}^{(1)}$ and $\Lambda_{i,n}^{(2)}$ used in the two concrete procedures. The proof is deferred to Online Appendix EC.2.3.

Theorem 2. Under Assumptions 1–4, $\Lambda_{i,n}^{(2)} \geq \Lambda_{i,n}^{(1)}$ holds true surely for system i , $\forall i \in [K]$.

Theorem 2 states that, seeing the same data, $\Lambda_{i,n}^{(2)}$ is bounded below by $\Lambda_{i,n}^{(1)}$. Because an elimination occurs when a test statistic drops below the threshold \mathcal{B} , Theorem 2 indicates that, with the same observations seen by these two concrete procedures, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ always terminates earlier than $Proc(\Lambda_{i,n}^{(2)}, \alpha)$. Hence, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ can be regarded as a conservative version of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$.

This result becomes handy when we analyze the statistical validity and the computational efficiency of these two concrete procedures in Sections 3.2 and 3.3; for example, see Theorems 3 and 4 and their proofs for details.

3.2. Statistical Validity

In this section, we show that the two concrete procedures given in Section 3.1—respectively using the proposed test statistics, $\Lambda_{i,n}^{(1)}$ and $\Lambda_{i,n}^{(2)}$, together with the elimination threshold $\mathcal{B} = \alpha$ —can indeed deliver the prescribed PCS guarantee.

We first show that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ terminate in finite time with probability 1; the proof is deferred to Online Appendix EC.2.4.

Theorem 3. Under Assumptions 1–4, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ terminate in finite time with probability one.

Let $PCS^{(1)}$ and $PCS^{(2)}$ denote the probabilities of correct selection corresponding to $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, respectively. We next show that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ can deliver the prescribed PCS guarantee.

Theorem 4. Under Assumptions 1–4, upon termination, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ select the best system with probability at least $1 - \alpha$, that is,

$$PCS^{(1)} \geq 1 - \alpha, \quad PCS^{(2)} \geq 1 - \alpha.$$

Proof of Theorem 4. We consider $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ first and show $PCS^{(1)} \geq 1 - \alpha$. Recall that $\boldsymbol{\mu}^* = (\mu_1^*, \mu_2^*, \dots, \mu_K^*)$ denotes the true mean vector for the K systems. Without loss of generality, assume that system 1 is the best system (i.e., $i^* = 1$). Recall the multiple composite hypothesis testing Model (2); in this case, the true hypothesis among the K hypotheses considered is $H_1 : \boldsymbol{\mu}^* \in \mathbb{U}_1$.

Let $\tau_i^{(1)} = \min\{n \mid \Lambda_{i,n}^{(1)} \leq \alpha\}$ denote the elimination time of system i by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ for any $i \in [K]$ and let $\tau^{(1)}$ denote the stopping time of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. The probability of correct selection of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ can be bounded below as follows:

$$\begin{aligned} PCS^{(1)} &= P(Proc(\Lambda_{i,n}^{(1)}, \alpha) \text{ terminates infinite time} \\ &\quad \text{and system 1 is not eliminated}) \\ &= P((\tau_1^{(1)} > \tau^{(1)}) \cap (\tau^{(1)} < \infty)) \\ &= P(\tau_1^{(1)} > \tau^{(1)}) \\ &\quad + P(\tau^{(1)} < \infty) - P((\tau_1^{(1)} > \tau^{(1)}) \cup (\tau^{(1)} < \infty)) \\ &\geq P(\tau_1^{(1)} > \tau^{(1)}) + P(\tau^{(1)} < \infty) - 1. \end{aligned} \quad (17)$$

Recall Theorem 3 that states that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ terminates in finite time with probability one, that is, $P(\tau^{(1)} < \infty) = 1$; it follows from (17) that

$$PCS^{(1)} \geq P(\tau_1^{(1)} > \tau^{(1)}) \geq 1 - P(\tau_1^{(1)} < \infty). \quad (18)$$

By definition, $\tau_1^{(1)} = \min\{n \mid \Lambda_{1,n}^{(1)} \leq \alpha\}$; recall the definition of $\Lambda_{1,n}^{(1)}$ given by (10) and the assumption that the

true best system is system 1, that is, $\boldsymbol{\mu}^* \in \mathbb{U}_1$, we have

$$\Lambda_{1,n}^{(1)} = \frac{\sup_{\boldsymbol{\mu} \in \mathbb{U}_1} L_n(\boldsymbol{\mu})}{\pi_n} \geq \frac{L_n(\boldsymbol{\mu}^*)}{\pi_n}. \quad (19)$$

Define $\Gamma_n = \pi_n/L_n(\boldsymbol{\mu}^*)$. It follows that $\{\Gamma_n, n \geq n_0 + 1\}$ is a nonnegative martingale. Specifically,

$$\begin{aligned} \Gamma_n &= \frac{\pi_{n-1}}{L_{n-1}(\boldsymbol{\mu}^*)} \cdot \prod_{i \in \mathcal{S}_n} \frac{g_i(X_{i,n} | \hat{\mu}_{i,n-1})}{g_i(X_{i,n} | \mu_i^*)} \\ &= \Gamma_{n-1} \cdot \prod_{i \in \mathcal{S}_n} \frac{g_i(X_{i,n} | \hat{\mu}_{i,n-1})}{g_i(X_{i,n} | \mu_i^*)} \geq 0. \end{aligned}$$

Hence, it follows that

$$\begin{aligned} \mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)}[\Gamma_n | \Gamma_{n-1}] &= \Gamma_{n-1} \times \mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)} \left(\prod_{i \in \mathcal{S}_n} \frac{g_i(\cdot | \hat{\mu}_{i,n-1})}{g_i(\cdot | \mu_i^*)} \right) \\ &= \Gamma_{n-1} \int \prod_{i \in \mathcal{S}_n} \frac{g_i(\cdot | \hat{\mu}_{i,n-1})}{g_i(\cdot | \mu_i^*)} \\ &\quad \cdot \prod_{m=1}^n \prod_{i \in \mathcal{S}_m} g_i(\cdot | \mu_i^*) dX_{i,m} \\ &= \Gamma_{n-1}, \quad \text{for } n \geq n_0 + 2. \end{aligned}$$

In addition, we have $\mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)}(\Gamma_n) = 1$, because

$$\begin{aligned} \mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)}[\Gamma_n] &= \int \Gamma_n \cdot \prod_{m=1}^n \prod_{i \in \mathcal{S}_m} g_i(\cdot | \mu_i^*) dX_{i,m} \\ &= \int \prod_{m=1}^{n_0} \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,n_0}) \\ &\quad \cdot \prod_{m=n_0+1}^n \prod_{i \in \mathcal{S}_m} g_i(X_{i,m} | \hat{\mu}_{i,m-1}) dX_{i,m} = 1, \end{aligned}$$

where the last step holds because g_i s are valid density functions.

Given that $\{\Gamma_n, n \geq n_0 + 1\}$ is a nonnegative martingale with respect to $\{\mathcal{F}_n, n \geq n_0 + 1\}$ satisfying $\mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)}(\Gamma_n) = 1$ (hence a test martingale), it follows immediately from Doob's martingale inequality (Revuz and Yor 1999) that

$$\mathbb{P} \left(\max_{n_0+1 \leq n \leq N} \Gamma_n \geq \frac{1}{\alpha} \right) \leq \alpha \cdot \mathbb{E}_{f(\cdot|\boldsymbol{\mu}^*)}(\Gamma_N) = \alpha, \quad \forall N \geq n_0 + 1.$$

Based on (19) and the definition of Γ_n , we have $\Lambda_{1,n}^{(1)} \geq 1/\Gamma_n$ for $n \geq n_0 + 1$; furthermore,

$$\mathbb{P} \left(\min_{n_0+1 \leq n \leq N} \Lambda_{1,n}^{(1)} \leq \alpha \right) \leq \mathbb{P} \left(\max_{n_0+1 \leq n \leq N} \Gamma_n \geq \frac{1}{\alpha} \right) \leq \alpha, \quad \forall N \geq n_0 + 1.$$

Given the inessential definition that $\Lambda_{1,n}^{(1)} = 1$ for $1 \leq n \leq n_0$, it follows that

$$\begin{aligned} \mathbb{P}(\tau_1^{(1)} < N) &= \mathbb{P} \left(\min_{1 \leq n \leq N} \Lambda_{1,n}^{(1)} \leq \alpha \right) = \mathbb{P} \left(\min_{n_0+1 \leq n \leq N} \Lambda_{1,n}^{(1)} \leq \alpha \right) \\ &\leq \alpha, \quad \forall N \geq n_0 + 1. \end{aligned} \quad (20)$$

In light of (18), we have $\text{PCS}^{(1)} \geq 1 - \mathbb{P}(\tau_1^{(1)} < \infty) \geq 1 - \alpha$.

Next, we consider $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$ and show $\text{PCS}^{(2)} \geq 1 - \alpha$. Let $\tau_i^{(2)} = \min\{n | \Lambda_{i,n}^{(2)} \leq \alpha\}$ denote the elimination time of system i ($\forall i \in [K]$) by $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$ and let $\tau^{(2)}$ denote the stopping time of $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$.

With an analogous argument as given in (17), we arrive at the following lower bound for $\text{PCS}^{(2)}$:

$$\begin{aligned} \text{PCS}^{(2)} &= \mathbb{P}(\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha) \text{ terminates in finite time} \\ &\quad \text{and system 1 is not eliminated}) \\ &\geq \mathbb{P}(\tau_1^{(2)} > \tau^{(2)}) + \mathbb{P}(\tau^{(2)} < \infty) - 1. \end{aligned}$$

In light of Theorem 3, we have $\mathbb{P}(\tau^{(2)} < \infty) = 1$. Hence,

$$\text{PCS}^{(2)} \geq \mathbb{P}(\tau_1^{(2)} > \tau^{(2)}) \geq 1 - \mathbb{P}(\tau_1^{(2)} < \infty). \quad (21)$$

Thanks to Theorem 2 and the inessential definition that $\Lambda_{1,n}^{(2)} = 1$ for $1 \leq n \leq n_0$, we have for any $N \geq n_0 + 1$,

$$\begin{aligned} \mathbb{P}(\tau_1^{(2)} < N) &= \mathbb{P} \left(\min_{1 \leq n \leq N} \Lambda_{1,n}^{(2)} \leq \alpha \right) = \mathbb{P} \left(\min_{n_0+1 \leq n \leq N} \Lambda_{1,n}^{(2)} \leq \alpha \right) \\ &\leq \mathbb{P} \left(\min_{n_0+1 \leq n \leq N} \Lambda_{1,n}^{(1)} \leq \alpha \right) \leq \alpha, \end{aligned} \quad (22)$$

where the last inequality follows from (20). Finally, $\text{PCS}^{(2)} \geq 1 - \alpha$ follows from the last inequality and (21) by letting $N \rightarrow \infty$. \square

Theorem 4 ensures that the two concrete procedures select the best system with probability at least $1 - \alpha$ when they terminate. The following result further reveals that if $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$) is stopped before the termination criterion is met, that is, $|\mathcal{S}_n| > 1$, it still holds that the subset of systems returned by $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively, $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$) contains the best system with the prescribed probabilistic guarantee. The proof of the following result is deferred to Online Appendix EC.2.5.

Corollary 1. *Suppose that Assumptions 1–4 are satisfied. While implementing $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively, $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$), the set \mathcal{S}_n contains the best system with probability at least $1 - \alpha$ for any $n \geq 1$.*

3.3. Asymptotic Efficiency

In this section, we study the asymptotic efficiency of the proposed procedures. Without loss of generality, assume that system 1 is the best system, that is, $\mu_1^* > \mu_i^*, \forall i \in [K] \setminus \{1\}$, or equivalently, $\boldsymbol{\mu}^* \in \mathbb{U}_1$. Recall that $\tau_i^{(1)}$ (respectively, $\tau_i^{(2)}$) denotes the number of iterations used by $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively, $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$) to eliminate system $i, \forall i \in [K] \setminus \{1\}$. Let $\mathbb{U}_{i,1} = \{\boldsymbol{\mu} \in \mathbb{U} | \mu_i \geq \mu_1 + \delta\}$ denote the set of parameter values such that system i is better than system 1, if $\delta > 0$ is specified; otherwise, let $\mathbb{U}_{i,1} = \{\boldsymbol{\mu} \in \mathbb{U} | \mu_i > \mu_1\}, \forall i \in [K] \setminus \{1\}$. We first study the asymptotic limit of $\mathbb{E}[\tau_i^{(1)}]$ for $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ in relation to that of $\mathbb{E}[\tau_i^{(2)}]$ for $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$ as the error probability α approaches zero.

Theorem 5. Suppose that Assumptions 1–4 are satisfied. Under the settings of $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ and $\text{Proc}(\Lambda_{i,n}^{(2)}, \alpha)$, as $\alpha \rightarrow 0$, the respective expected sample sizes used by the two procedures to eliminate system i satisfy

$$\mathbb{E}_{f(\cdot|\mu^*)}[\tau_i^{(1)}] \leq \mathbb{E}_{f(\cdot|\mu^*)}[\tau_i^{(2)}] \sim \frac{\log(1/\alpha)}{\inf_{\mu \in \mathbb{U}_{i,1}} \rho(\mu^*, \mu)} (1 + o(1)),$$

$$\forall \mu^* \in \mathbb{U}_1, i \in [K] \setminus \{1\},$$
(23)

where $\rho(\mu^*, \mu)$ is the KL divergence from $f(\cdot|\mu)$ to $f(\cdot|\mu^*)$ and $a \sim b$ denotes the asymptotic equivalence of a and b .

The proof of Theorem 5 is provided in Online Appendix EC.2.6. Theorem 5 can be intuitively understood as follows. The numerator, $\log(1/\alpha)$, denotes the amount of “information” to make an elimination decision. The denominator, $\inf_{\mu \in \mathbb{U}_{i,1}} \rho(\mu^*, \mu)$, represents the expected amount of “information” regarding the difference between μ and μ^* obtained from observations collected on one iteration. The expected sample size to reject system i hence is roughly given by $\log(1/\alpha) / \inf_{\mu \in \mathbb{U}_{i,1}} \rho(\mu^*, \mu)$.

In light of Theorem 5, we have the following asymptotic efficiency result for the normality setting with a known equal variance. Proof of Corollary 2 is provided in Online Appendix EC.2.7.

Corollary 2. Suppose that the assumptions of Theorem 5 are satisfied and $X_{i,n} \sim \mathcal{N}(\mu_i^*, \sigma^2)$, $\forall i \in [K]$. As $\alpha \rightarrow 0$, the respective expected sample sizes used by the two proposed procedures to eliminate system i satisfy

$$\mathbb{E}_{f(\cdot|\mu^*)}[\tau_i^{(1)}] \leq \mathbb{E}_{f(\cdot|\mu^*)}[\tau_i^{(2)}] \sim \frac{-4 \log(\alpha) \cdot \sigma^2}{(\mu_1^* - \mu_i^* + \delta)^2}, \quad \forall i \in [K] \setminus \{1\},$$

where $\delta > 0$ denotes the indifference-zone parameter specified and $\delta = 0$ is used when this parameter is unspecified.

We close this section with a comparison of the asymptotic expected sample sizes of the KN++ procedure, the IZ-free procedure, and $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ under the normality setting with a known equal variance. The asymptotic expected sample sizes used for eliminating system i ($\forall i \in [K] \setminus \{1\}$) by all three procedures can be given as functions of the true difference in the mean values of the best system and system i ($\Delta := \mu_1^* - \mu_i^*$), the indifference-zone parameter (δ), and the known variance (σ^2), as shown in Table 1. Two interesting observations can be made. First and foremost, when δ is unspecified, $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ adopts $\delta = 0$ and the corresponding expected sample size is of the same order as that of IZ-free in terms of the true difference Δ ; when δ is specified, $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ can use and benefit from this information. Second, it is known that a proper specification of δ is required for KN++ to work appropriately and the optimal choice is $\delta = \Delta$, because any other values would

Table 1. Asymptotic Expected Sample Sizes of KN++ Procedure, IZ-Free Procedure, and $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$ as $\alpha \rightarrow 0$

KN++	IZ-free	$\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$
$\frac{2\gamma\sigma^2}{(\Delta - \delta/2)\delta}$	$\frac{2c\sigma^2}{\Delta^2}$	$\frac{-4 \log(\alpha) \cdot \sigma^2}{(\Delta + \delta)^2}$

Notes. For KN++, $\gamma = \log((k-1)/(2\alpha))$; for IZ-free, c is a parameter whose value can be determined numerically. Results for KN++ and IZ-free are available in Fan et al. (2016).

result in a larger expected sample size. Such an issue does not exist for $\text{Proc}(\Lambda_{i,n}^{(1)}, \alpha)$, however.

4. Setting of Normality with Unknown Variances

This section considers solving an R&S problem under the normality setting with unknown variances. The R&S problem is as described in Section 2.1 with notation already established there.

Throughout this section, the following assumptions are stipulated.

Assumption 5. The K systems have distinct true mean values, that is, $\mu_i^* \neq \mu_j^*$, for $i \neq j, i, j \in [K]$. The true mean values of the K systems satisfy that $\mu_{i^*}^* \geq \mu_i^* + \delta$ for all $i \neq i^*, i \in [K]$, where i^* denotes the index of the best system and $\delta > 0$ is the indifference-zone parameter if specified.

Assumption 6. The observations of system i , $\{X_{i,n}, n = 1, 2, \dots\}$ are independent and normally distributed with mean μ_i^* and variance σ_i^{2*} , that is, $\mathcal{N}(\mu_i^*, \sigma_i^{2*})$, with $\theta_i^* = (\mu_i^*, \sigma_i^{2*})$ being unknown. Furthermore, the observations of different systems are statistically independent, that is, $X_{i,n}$ and $X_{j,n'}$ are independent for $n, n' \geq 1$, if $i \neq j, i, j \in [K]$.

Assumption 7. The parameter space $\Theta \subseteq \mathbb{R}^{2K}$ is compact.

Assumption 8. There exists some $r \geq 1$, such that $\sum_{m=1}^{\infty} \mathbb{E}_{f(\cdot|\theta^*)}[\rho(\theta^*, \hat{\theta}_m)]^r < \infty$, where $\hat{\theta}_m$ represents the MLE of θ^* obtained using observations from the first m iterations, and $\rho(\theta, \theta')$ denotes the KL divergence from $f(\cdot|\theta')$ to $f(\cdot|\theta)$, which can be written as

$$\rho(\theta, \theta') = \mathbb{E}_{f(\cdot|\theta)} \left[\log \left(\frac{f(\cdot|\theta)}{f(\cdot|\theta')} \right) \right],$$

with $\mathbb{E}_{f(\cdot|\theta)}$ denoting the taken expectation with respect to $f(\cdot|\theta)$. Here, $f(\cdot|\theta) := \prod_{i=1}^K g_i(\cdot|\theta_i)$ denotes the joint density function of observations from all K systems, and $g_i(\cdot|\theta_i)$ denotes the normal density function corresponding to system i with the parameter vector given by $\theta_i = (\mu_i, \sigma_i^2)$, $\forall i \in [K]$.

Assumptions 5–8 stipulated in this section are in parallel with Assumptions 1–4 in Section 3, with the parameter space and the parameter sets of interest in this section being expanded to accommodate the unknown variances. In fact, with the knowledge of the variances, Assumptions 5–8 imply Assumptions 1–4.

4.1. Two Concrete Procedures

As in Section 3.1, we elaborate on two concrete procedures, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, which are suitable for the normality setting with unknown variances, and provide some insight into their practical implementation.

4.1.1. Procedure $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. To tackle the R&S problem under normality with unknown variances, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ adopts a GLR statistic, denoted by $\Lambda_{i,n}^{(1)}$, which is defined as follows:

$$\Lambda_{i,n}^{(1)} = \frac{\sup_{\theta \in \Theta_i} L_n(\theta)}{\pi_n}, \quad n \geq n_0 + 1, \quad \forall i \in [K]. \quad (24)$$

As noted in Section 3.1, we may set $\Lambda_{i,n}^{(1)} = 1$ for $1 \leq n \leq n_0$ for completeness; however, this is irrelevant to the implementation and the statistical validity of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. In (24), $L_n(\theta)$ denotes the likelihood of all observations collected up to the n th iteration ($n \geq 1$) for all K systems, which is given by

$$L_n(\theta) = \prod_{m=1}^n \prod_{i \in S_m} g_i(X_{i,m} | \theta_i); \quad (25)$$

and π_n in (24) is given by

$$\begin{aligned} \pi_n &= L_{n_0}(\hat{\theta}_{n_0}) \cdot \prod_{m=n_0+1}^n \prod_{i \in S_m} g_i(X_{i,m} | \hat{\theta}_{i,m-1}) \\ &= \prod_{m=1}^{n_0} \prod_{i \in S_m} g_i(X_{i,m} | \hat{\theta}_{i,n_0}) \cdot \prod_{m=n_0+1}^n \prod_{i \in S_m} g_i(X_{i,m} | \hat{\theta}_{i,m-1}), \end{aligned} \quad (26)$$

which denotes the adaptive maximum likelihood of obtaining the observations collected up to the n th iteration for $n \geq n_0 + 1$. The first term in (26), $L_{n_0}(\hat{\theta}_{n_0})$, denotes the likelihood of obtaining the observations collected during the warm-up period calculated using the MLE of θ^* attained by the end of the warm-up period, $\hat{\theta}_{n_0} = (\hat{\theta}_{1,n_0}, \hat{\theta}_{2,n_0}, \dots, \hat{\theta}_{K,n_0}) := \arg \sup_{\theta \in \Theta} L_{n_0}(\theta)$. The second term, $\prod_{m=n_0+1}^n \prod_{i \in S_m} g_i(X_{i,m} | \hat{\theta}_{i,m-1})$, denotes the adaptive maximum likelihood of obtaining the individual observations after the warm-up period using the corresponding most updated MLEs, where $\hat{\theta}_{i,m-1}$ denotes the MLE of θ_i^* attained based on the observations of system i ($i \in S_m$) collected up to the $(m - 1)$ th iteration, $m = n_0 + 1, \dots, n$.

The optimization problem involved in the expression of $\Lambda_{i,n}^{(1)}$ in (24) is solvable under normality with unknown variances; nevertheless, it may not admit a closed-form solution. Because implementing $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ requires calculating $\Lambda_{i,n}^{(1)}$ for $\forall i \in [K]$ on every iteration, in case a closed-form solution does not exist, the computational cost of implementing $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ can be prohibitive.

4.1.2. Procedure $Proc(\Lambda_{i,n}^{(2)}, \alpha)$. The test statistic used in $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ is constructed based on pairwise system comparisons as follows:

$$\Lambda_{i,n}^{(2)} = \min_{j \neq i, j \in [K]} \Lambda_{i,j,n}^{(2)}, \quad n \geq n_0 + 1, \quad \forall i \in [K],$$

which denotes the minimum of all pairwise statistics $\Lambda_{i,j,n}^{(2)}$ and $\Lambda_{i,j,n}^{(2)}$ denotes the generalized likelihood ratio of the event that system i is better than system j , that is,

$$\Lambda_{i,j,n}^{(2)} = \frac{\sup_{\theta \in \Theta_{ij}} L_n(\theta)}{\pi_n}, \quad n \geq n_0 + 1, \quad j \neq i, \quad j \in [K],$$

where $L_n(\theta)$ and π_n are, respectively, given in (25) and (26). Here, $L_n(\theta)$ is maximized over the parameter set Θ_{ij} , where $\Theta_{ij} = \{\theta \in \Theta | \mu_i \geq \mu_j + \delta\}$ when $\delta > 0$ is specified; otherwise, $\Theta_{ij} = \{\theta \in \Theta | \mu_i > \mu_j\}$.

Similarly as in the normality setting with known variances, we can show that the following relationship between the test statistics adopted in the two concrete procedures hold in the normality setting with unknown variances. The proof is provided in Online Appendix EC.3.2.

Theorem 6. Under Assumptions 5–8, $\Lambda_{i,n}^{(2)} \geq \Lambda_{i,n}^{(1)}$ holds true surely for system i , $\forall i \in [K]$.

We next elaborate on the calculation of $\Lambda_{i,n}^{(2)}$ in $Proc(\Lambda_{i,n}^{(2)}, \alpha)$. Let $\{X_{i,n}, n \geq 1\}$ be a sample of independent and normally distributed observations of system i with unknown mean μ_i^* and unknown variance σ_i^{2*} , $\forall i \in [K]$. Under Assumptions 5–8, $L_n(\theta)$ in (25) can be written as

$$L_n(\theta) = \prod_{m=1}^n \prod_{l \in S_m} g_l(X_{l,m} | \mu_l, \sigma_l^2),$$

where

$$g_l(X_{l,m} | \mu_l, \sigma_l^2) = \frac{1}{\sqrt{2\pi\sigma_l^2}} \exp\left(-\frac{(X_{l,m} - \mu_l)^2}{2\sigma_l^2}\right), \quad \forall l \in [K];$$

recall that S_n contains those systems that have not been eliminated by the n th iteration, and $\hat{\theta}_n = (\hat{\theta}_{1,n}, \hat{\theta}_{2,n}, \dots, \hat{\theta}_{K,n}) := \arg \sup_{\theta \in \Theta} L_n(\theta)$ denote the MLE of θ^* obtained at the n th iteration, with $\hat{\theta}_{l,n} = (\hat{\mu}_{l,n}, \hat{\sigma}_{l,n}^2)$ denoting the vector of MLEs of the mean μ_l^* and the variance σ_l^{2*} for all $l \in [K]$. Denote $\tau_l^{(2)}$ as the elimination time of system l by $Proc(\Lambda_{i,n}^{(2)}, \alpha)$; then $\hat{\mu}_{l,n}$ and $\hat{\sigma}_{l,n}^2$ can be, respectively, given as

$$\begin{aligned} \hat{\mu}_{l,n} &= \frac{\sum_{m=1}^n \mathbf{1}\{\tau_l^{(2)} \geq m\} X_{l,m}}{\sum_{m=1}^n \mathbf{1}\{\tau_l^{(2)} \geq m\}}, \\ \hat{\sigma}_{l,n}^2 &= \frac{\sum_{m=1}^n \mathbf{1}\{\tau_l^{(2)} \geq m\} (X_{l,m} - \hat{\mu}_{l,n})^2}{\sum_{m=1}^n \mathbf{1}\{\tau_l^{(2)} \geq m\}}, \quad \forall l \in [K]. \end{aligned}$$

Let $\tilde{\theta}_n = (\tilde{\theta}_{1,n}, \tilde{\theta}_{2,n}, \dots, \tilde{\theta}_{K,n}) := \arg \sup_{\theta \in \Theta_{ij}} L_n(\theta)$ denote the constrained MLE of θ^* attained at the n th iteration, where $\tilde{\theta}_{l,n} = (\tilde{\mu}_{l,n}, \tilde{\sigma}_{l,n}^2)$ denotes the vector of constrained MLEs of the mean μ_l^* and the variance σ_l^{2*} over the parameter set Θ_{ij} for all $l \in [K]$. It is easy to see that the constrained MLE $\tilde{\theta}_{l,n}$ and the MLE $\hat{\theta}_{l,n}$ for system l are identical, $\forall l \in [K] \setminus \{i, j\}$. Therefore, under Assumption 6, $\Lambda_{i,j,n}^{(2)}$ can be written explicitly as

$$\Lambda_{i,j,n}^{(2)} = \pi_n^{-1} \prod_{m=1}^n \prod_{l \in \mathcal{S}_n \setminus \{i, j\}} g_l(X_{l,m} | \hat{\mu}_{l,n}, \hat{\sigma}_{l,n}^2) \cdot \prod_{m=1}^n g_i(X_{i,m} | \tilde{\mu}_{i,n}, \tilde{\sigma}_{i,n}^2) \cdot g_j(X_{j,m} | \tilde{\mu}_{j,n}, \tilde{\sigma}_{j,n}^2), n \geq n_0 + 1.$$

It is worth noting that an analytical solution for $\tilde{\theta}_n = \arg \sup_{\theta \in \Theta_{ij}} L_n(\theta)$ (in particular, a closed-form expression for $\tilde{\theta}_{i,n} = (\tilde{\mu}_{i,n}, \tilde{\sigma}_{i,n}^2)$ and $\tilde{\theta}_{j,n} = (\tilde{\mu}_{j,n}, \tilde{\sigma}_{j,n}^2)$) does not exist under normality with unknown variances. Iterative algorithms are typically recommended to solve such an optimization problem which is nonconvex. Hence, we approximate the constrained MLEs by repeating the following two steps through iterations: (1) updating $(\tilde{\mu}_{i,n}, \tilde{\mu}_{j,n})$ while holding $(\tilde{\sigma}_{i,n}^2, \tilde{\sigma}_{j,n}^2)$ fixed, and (2) updating $(\tilde{\sigma}_{i,n}^2, \tilde{\sigma}_{j,n}^2)$ while holding $(\tilde{\mu}_{i,n}, \tilde{\mu}_{j,n})$ fixed. To simplify the calculation of $\Lambda_{i,j,n}^{(2)}$ in implementation, we approximate $\tilde{\mu}_{i,n}, \tilde{\mu}_{j,n}, \tilde{\sigma}_{i,n}^2$, and $\tilde{\sigma}_{j,n}^2$ as follows:

$$(\tilde{\mu}_{i,n}, \tilde{\mu}_{j,n}) \approx \begin{cases} (\hat{\mu}_{i,n}, \hat{\mu}_{j,n}) & \text{if } \hat{\mu}_{i,n} \geq \hat{\mu}_{j,n} + \delta; \\ ((\hat{\mu}_{i,n} + \hat{\mu}_{j,n} + \delta)/2, (\hat{\mu}_{i,n} + \hat{\mu}_{j,n} - \delta)/2) & \text{otherwise,} \end{cases}$$

$$(\tilde{\sigma}_{i,n}^2, \tilde{\sigma}_{j,n}^2) \approx \left(\frac{1}{n_i} \sum_{m=1}^{n_i} (X_{i,m} - \tilde{\mu}_{i,n})^2, \frac{1}{n_j} \sum_{m=1}^{n_j} (X_{j,m} - \tilde{\mu}_{j,n})^2 \right),$$

where $\delta > 0$ denotes the indifference-zone parameter if specified and $\delta = 0$ is used when this parameter is unspecified, $n_i = \min\{n, \tau_i\}$, and $n_j = \min\{n, \tau_j\}$, $\forall i \neq j, i, j \in [K]$. These expressions are essentially the results obtained upon running one iteration of the aforementioned two-step approximation.

In view of the relative ease of calculating $\Lambda_{i,n}^{(2)}$ compared with $\Lambda_{i,n}^{(1)}$, we remark that $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ is more applicable under the normality setting with unknown variances. However, in other settings where $\Lambda_{i,n}^{(1)}$ admits a closed-form expression, we would recommend implementing $Proc(\Lambda_{i,n}^{(1)}, \alpha)$, which can be less conservative than $Proc(\Lambda_{i,n}^{(2)}, \alpha)$.

4.2. Statistical Validity

In this section, we provide an in-depth study of the theoretical properties of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ under normality with unknown variances.

The first result establishes that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ terminate in finite time with probability

one, whose proof is provided in Online Appendix EC.3.3.

Theorem 7. Under Assumptions 5–8, both $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ terminate in finite time with probability one.

Given that the two procedures terminate in finite time with probability one, we further show that they both deliver the prescribed PCS guarantee. The proof of the following result is deferred to Online Appendix EC.3.4.

Theorem 8. Under Assumptions 5–8, both $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ select the best system with probability at least $1 - \alpha$, that is,

$$PCS^{(1)} \geq 1 - \alpha, \quad PCS^{(2)} \geq 1 - \alpha. \quad (27)$$

Theorem 8 establishes that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ select the best system with probability at least $1 - \alpha$ when they terminate. Nonetheless, if $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ are stopped before the termination criterion is met, that is, $|\mathcal{S}_n| > 1$, we have that the set \mathcal{S}_n contains the best system with the prescribed probabilistic guarantee. The proof of the following result is provided in Online Appendix EC.3.5.

Corollary 3. Suppose that Assumptions 5–8 are satisfied. While implementing $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$), the set \mathcal{S}_n contains the best system with probability at least $1 - \alpha$ for any $n \geq 1$.

4.3. Asymptotic Efficiency

In this section, we investigate the asymptotic efficiency. Without loss of generality, assume that system 1 is the best system, that is, $\mu_1^* > \mu_i^*, \forall i \in [K] \setminus \{1\}$, or equivalently, $\theta \in \Theta_1$. Recall that $\tau_i^{(1)}$ (respectively, $\tau_i^{(2)}$) denotes the number of iterations used by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ (respectively, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$) to eliminate system i , $\forall i \in [K] \setminus \{1\}$. Recall that $\Theta_{i,1}$ denotes the set of parameter values such that system i is better than system 1, $\forall i \in [K] \setminus \{1\}$. The following result reveals the asymptotic limit of $\mathbb{E}[\tau_i^{(1)}]$ for $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and that of $\mathbb{E}[\tau_i^{(2)}]$ for $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ as the error probability α tends to zero. The proof is provided in Online Appendix EC.3.6.

Theorem 9. Suppose that Assumptions 5–8 are satisfied. Under the settings of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, as $\alpha \rightarrow 0$, the respective expected sample sizes used by the two procedures to eliminate system i satisfy

$$\mathbb{E}_{f(\cdot|\theta^*)}[\tau_i^{(1)}] \leq \mathbb{E}_{f(\cdot|\theta^*)}[\tau_i^{(2)}] \sim \frac{\log(1/\alpha)}{\inf_{\theta \in \Theta_{i,1}} \rho(\theta^*, \theta)} (1 + o(1)),$$

$$\forall \theta^* \in \Theta_1, i \in [K] \setminus \{1\}.$$

5. Numerical Evaluations

In this section, we conduct numerical experiments to evaluate the performance of the proposed procedures.

In the first study, we evaluate the statistical validity and the computational efficiency of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ for solving an R&S problem involving two systems. In the second study, we assess the performance of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ in comparison with two state-of-the-art R&S procedures in selecting the best system under different indifference-zone configurations with up to 500 systems. The third study focuses on evaluating $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ in comparison with two leading R&S procedures for solving an R&S problem in which all systems are subject to unknown heterogeneous variances. It is worth noting that we recommend working with the logarithms of the likelihoods, the test statistics, and the elimination threshold to enhance the proposed procedures' numerical stability.

5.1. Study 1: Statistical Validity and Computational Efficiency of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$

This study focuses on evaluating the statistical validity and computational efficiency of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ proposed for solving R&S problems under normality with unknown mean values and known variances (recall Section 3.1).

Consider comparing $K = 2$ systems with statistically independent observations. The observations of system i are independent and normally distributed with mean μ_i^* and variance 1, $i = 1, 2$. For different combinations of the target PCS level (or equivalently, the target error probability level α) and the true mean value difference ($\Delta = \mu_1^* - \mu_2^*$), we perform the experiment for 10,000 independent macro-replications and record the estimated PCS and the total sample size consumed by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. In all experiments, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ is implemented without a specification of the indifference-zone parameter, and it starts with $n_0 = 2$ iterations to collect observations from each system without conducting elimination.

Table 2 summarizes the average total sample sizes (\overline{SSize}) with 95% confidence interval and the estimated PCSs (\widehat{PCS}) obtained under different combinations of $\Delta = \mu_1^* - \mu_2^*$ and the target PCS level considered. Two interesting observations are worth mentioning. First, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ tends to overshoot the target PCS level

when the gap between the mean values of the two systems, Δ , is large. However, as Δ decreases, \widehat{PCS} drops steadily to the target PCS level. The reasoning behind is that the larger Δ , the more likely the GLR statistic of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ overshoots the elimination threshold, and hence the higher \widehat{PCS} obtained. As Δ decreases, \widehat{PCS} drops and approaches the target PCS level, which confirms the statistical validity of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$. Second, as Δ decreases by a factor of two, \overline{SSize} increases approximately by a factor of four, echoing Corollary 2, which states that the average total sample size asymptotically approaches the limit $-4 \log \alpha / \Delta^2$ as $\alpha \rightarrow 0$.

Figure 1 plots the average total sample size obtained across 10,000 macro-replications and the analytical sample size given by Corollary 2 against $\Delta = \mu_1^* - \mu_2^*$ for each target PCS level considered. Despite some discrepancy between the numerical and the analytical results, we observe that the average total sample size used by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ approaches the theoretical lower bound as Δ decreases. Two possible explanations for this observation can be given. On the one hand, the GLR statistic of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ in the numerical implementation is a discrete-time stochastic process, but is approximated as a continuous-time stochastic process in the theoretical analysis. On the other hand, the analytical result stated in Corollary 2 assumes the knowledge of μ_1^* and μ_2^* , whereas this information is unavailable from the numerical implementation perspective. We conclude that the observations made in this study corroborate the theoretical results regarding the statistical validity and computational efficiency of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ obtained in Section 3.

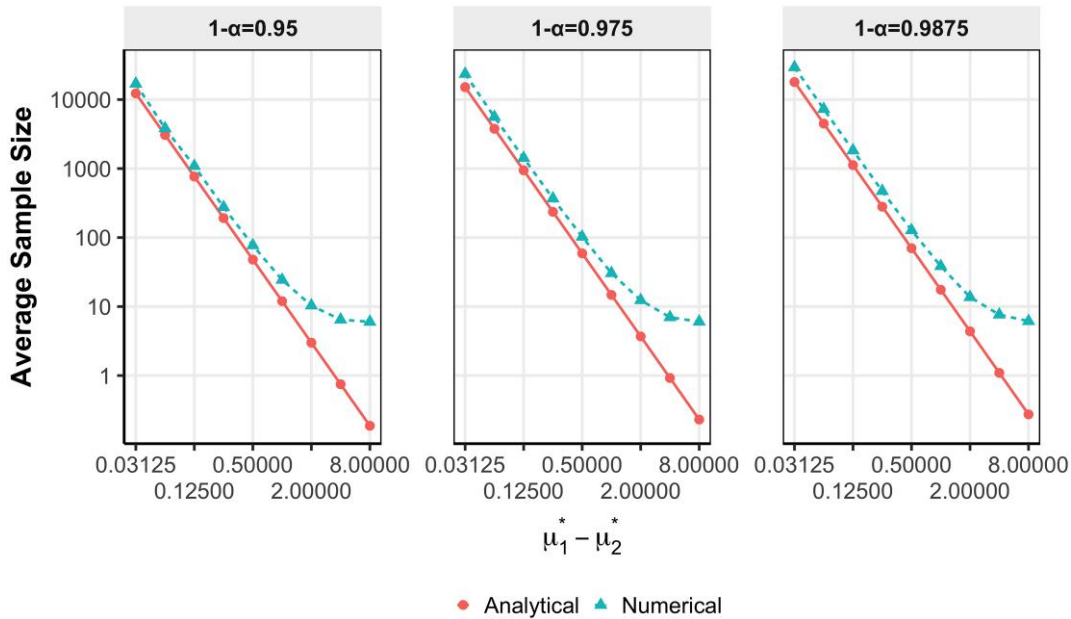
5.2. Study 2: Selecting the Best System Under Different Indifference-Zone Parameter Specifications

This study examines the performance of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ under different indifference-zone parameter settings in comparison with two benchmarking methods, the KN++ procedure and the IZ-free procedure, for solving an R&S problem under normality with known variances (recall Section 3.1). To facilitate comparison, identical

Table 2. Study 1: Estimated PCS (\widehat{PCS}) and Average Total Sample Size (\overline{SSize}) with 95% Confidence Interval for $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ Under Normality with Unknown Means and Variance 1, Where $\alpha \in \{0.05, 0.025, 0.0125\}$

Δ	$1 - \alpha = 0.95$		0.975		0.9875	
	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}
2^3	6.004 ± 0.00	1.000	6.002 ± 0.012	1.000	6.112 ± 0.032	1.000
2^2	6.92 ± 0.012	1.000	7.10 ± 0.13	1.000	7.58 ± 0.19	1.000
2^1	9.98 ± 0.25	0.999	11.90 ± 0.31	1.000	14.38 ± 0.41	1.000
2^0	25.79 ± 1.08	0.997	30.95 ± 1.22	0.998	37.26 ± 1.57	0.998
2^{-1}	89.86 ± 4.67	0.997	109.41 ± 5.30	0.998	132.17 ± 5.78	0.999
2^{-2}	333.01 ± 17.47	0.995	421.91 ± 21.22	0.996	513.53 ± 23.23	0.998
2^{-3}	$1,227 \pm 72$	0.979	$1,538 \pm 79$	0.988	$1,933 \pm 88$	0.996
2^{-4}	$4,523 \pm 279$	0.970	$5,919 \pm 333$	0.979	$7,583 \pm 376$	0.989
2^{-5}	$16,939 \pm 1,109$	0.951	$23,499 \pm 1,297$	0.975	$29,366 \pm 1,482$	0.988

Figure 1. (Color online) Study 1: Log-Log Plot of the Average Total Sample Size (y Axis) Against the True Mean Value Difference (x Axis), Displayed Using Logarithmic Scales on Both Axes



Note. Solid dots denote the analytical average total sample sizes given by Corollary 2, and solid triangles denote the average total sample sizes reported in Table 2.

experimental settings as considered in Fan and Hong (2014) and Fan et al. (2016) are adopted here. It is worth noting that a major difference between $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and the two benchmarking procedures is that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ does not require the Bonferroni correction.

Consider an R&S problem involving K systems with statistically independent observations. The observations of system i are independent and normally distributed with mean $\mu_i^* = 1.5 - 0.5i$ and variance $\sigma_i^2 = 10$ for $\forall i \in [K]$. We consider varying the number of systems $K \in \{20, 50, 100, 500\}$ and the specification of the indifference-zone parameter $\delta \in \{2^{-4}, 2^{-3}, 2^{-2}, 2^{-1}, 0\}$ with zero, indicating that δ is unspecified. Notice that $\delta = 0.5$ corresponds to the true gap between the best system and the second best: $\Delta = 0.5$. With the target PCS level set to $1 - \alpha = 0.95$, under each combination of K and δ , we perform the experiment for 1,000 independent macro-replications and record the estimated PCS obtained and the total sample sizes consumed by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ and the two benchmarking methods. The implementation of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ under each combination of δ and K starts with $n_0 = 10$ iterations to collect observations from each system without conducting elimination.

Table 3 summarizes the estimated PCSs (\widehat{PCS}) and the average total sample sizes (\widehat{SSize}) with 95% confidence interval obtained under the different combinations of δ and K considered. We note that the results corresponding to the cases where δ is unspecified do not exist for the KN++ procedure, because it applies

only when $\delta > 0$ is provided. Conversely, the results corresponding to the cases where $\delta > 0$ assumes different specifications do not exist for the IZ-free procedure, as it does not admit an indifference-zone parameter.

Several important observations can be made from Table 3. First and foremost, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ outperforms the two benchmarking procedures in terms of the average total sample sizes used under all experimental settings. Specifically, compared with KN++, the average total sample size used by $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ grows much more slowly as the indifference-zone parameter δ approaches zero. The results suggest that the computational efficiency of KN++ hinges heavily on the specification of δ , which declines drastically when $\delta \ll \Delta$ (recall $\Delta = 0.5$ in this study), whereas that of $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ is much less affected. On the other hand, when δ is unspecified, $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ uses much smaller average total sample sizes compared with IZ-free to achieve the target PCS level. Second, in terms of the \widehat{PCS} obtained, we find that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ is less conservative than KN++ and IZ-free because both procedures deliver \widehat{PCS} much higher than the target level 0.95 as a result of using the Bonferroni correction. We conclude that $Proc(\Lambda_{i,n}^{(1)}, \alpha)$ is a statistically valid and computationally efficient R&S procedure with and without specifying the indifference-zone parameter. Indeed, it helps fill the gap between procedures developed under the indifference-zone formulation and the indifference-zone-free formulation.

Table 3. Study 2: Estimated PCS (\widehat{PCS}) and Average Total Sample Size (\overline{SSize}) with 95% Confidence Interval Obtained from 1,000 Independent Macro-Replications for Implementing $Proc(\Lambda_{i,n}^{(1)}, \alpha)$, KN++ Procedure, and IZ-Free Procedure with $1 - \alpha = 0.95$

K	δ	δ specified				δ unspecified			
		KN++		$Proc(\Lambda_{i,n}^{(1)}, \alpha)$		IZ-free		$Proc(\Lambda_{i,n}^{(1)}, \alpha)$	
		\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}
20	2^{-1}	1,371 ± 16	0.99	490 ± 20	0.95	2,816 ± 82	0.99	1,156 ± 133	0.98
	2^{-2}	3,247 ± 29	1.00	594 ± 52	0.98				
	2^{-3}	7,045 ± 50	1.00	757 ± 95	0.95				
	2^{-4}	14,700 ± 70	1.00	830 ± 93	0.98				
50	2^{-1}	2,014 ± 18	1.00	782 ± 23	0.95	3,588 ± 89	0.99	1,471 ± 153	0.99
	2^{-2}	4,454 ± 33	1.00	866 ± 44	0.95				
	2^{-3}	9,779 ± 54	1.00	1,069 ± 76	0.97				
	2^{-4}	20,700 ± 80	1.00	1,327 ± 130	0.97				
100	2^{-1}	2,710 ± 20	0.99	1,359 ± 20	0.95	4,388 ± 89	0.99	1,955 ± 159	0.98
	2^{-2}	5,506 ± 34	1.00	1,428 ± 49	0.98				
	2^{-3}	11,790 ± 60	1.00	1,691 ± 80	0.98				
	2^{-4}	25,140 ± 90	1.00	1,850 ± 116	0.97				
500	2^{-1}	7,217 ± 21	0.99	5,683 ± 21	0.96	9,138 ± 102	1.00	6,335 ± 143	0.98
	2^{-2}	10,770 ± 40	1.00	5,873 ± 53	0.97				
	2^{-3}	18,770 ± 60	1.00	6,050 ± 87	0.97				
	2^{-4}	36,280 ± 90	1.00	6,193 ± 106	0.99				

Note. Results for KN++ and IZ-free are cited from Fan et al. (2016).

5.3. Study 3: Selecting the Best System Under Unknown Heterogeneous Variances

In this study, we investigate the performance of $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ in comparison with the KN++ procedure and the IZ-free procedure for solving an R&S problem under the normality setting with unknown, heterogeneous variances (recall Section 4.1).

Consider solving an R&S problem involving K systems with the target PCS level set to $1 - \alpha = 0.95$ and the indifference-zone parameter δ being either 0.5 or unspecified. The random observations of system i ($\forall i \in [K]$) are independent and normally distributed with mean $\mu_i^* = 1.5 - 0.5i$ and variance σ_i^{2*} , which is specified next. Notice that $\delta = 0.5$ corresponds to the true gap between the best system and the second best, $\Delta = 0.5$. Three configurations of variances for all $i \in [K]$ are considered, respectively, equal ($\sigma_i^{2*} = 10$), monotonically increasing ($\sigma_i^{2*} = 10 \times (0.95 + 0.05i)$), and monotonically decreasing ($\sigma_i^{2*} = 10 / (0.95 + 0.05i)$). We consider varying $K \in \{20, 50, 100\}$. For each combination of K , δ , and the configuration of variances, we perform the experiment for 1,000 independent macro-replications and record the estimated PCS obtained and the total sample sizes consumed by $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ and the two benchmarking methods. Every evaluation of $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ under each experimental setting uses $n_0 = 10$ iterations to collect observations from all systems without conducting elimination.

Table 4 summarizes the estimated PCSs (\widehat{PCS}) and the average total sample sizes (\overline{SSize}) with 95% confidence interval obtained under different experimental settings, from which we have the following observations. First, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ demonstrates its statistical validity

across all experimental settings and in particular it delivers robust performance under different configurations of variances. Second, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ outperforms the two benchmarking procedures in terms of computational efficiency under all indifference-zone parameter settings considered. Specifically, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ uses smaller total sample sizes on average than KN++, even when δ is specified to be $\Delta = 0.5$ that brings out the best in KN++. Last but not least, in terms of the \widehat{PCS} obtained, $Proc(\Lambda_{i,n}^{(2)}, \alpha)$ is less conservative than KN++ and IZ-free. Although the target PCS level is 0.95, both KN++ and IZ-free overshoot the target by a sizable margin and consume much larger total sample sizes on average.

6. Conclusion

In this paper, we proposed two new fully sequential procedures for R&S. Unlike many other leading methods, the new procedures do not use the Bonferroni correction, can continuously monitor and control the type I error rate at a prescribed level, and are indifference-zone-flexible. Under the settings of normality with and without the knowledge of true variances, the proposed procedures are shown to be statistically valid and asymptotically efficient and provide a prescribed PCS guarantee. Numerical studies corroborate the theoretical results and demonstrate the superiority of the proposed procedures under various configurations.

There are several future research directions to pursue, including, but not limited to (1) developing generalized sequential likelihood ratio test-based procedures that can achieve a guaranteed probability of good selection (this direction is highly relevant to tackling large-scale

Table 4. Study 3: Estimated PCS (\widehat{PCS}) and Average Total Sample Size (\overline{SSize}) with 95% Confidence Interval Obtained from 1,000 Independent Macro-Replications for Implementing $Proc(\Lambda_{i,n}^{(2)}, \alpha)$, KN++ Procedure, and IZ-Free Procedure with $1 - \alpha = 0.95$

K	Variances	$\delta = 0.5$				δ unspecified			
		KN++		$Proc(\Lambda_{i,n}^{(2)}, \alpha)$		IZ-free		$Proc(\Lambda_{i,n}^{(2)}, \alpha)$	
		\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}	\overline{SSize}	\widehat{PCS}
20	Decreasing	1,326 ± 16	0.99	806 ± 32	0.97	2,780 ± 81	0.99	2,302 ± 78	0.99
	Equal	1,371 ± 16	0.99	888 ± 35	0.96	2,816 ± 82	0.99	2,486 ± 86	0.99
	Increasing	1,473 ± 17	0.99	985 ± 41	0.95	2,916 ± 85	0.99	2,559 ± 88	0.98
50	Decreasing	1,881 ± 17	0.99	1,179 ± 31	0.97	3,551 ± 86	0.99	2,840 ± 97	0.99
	Equal	2,014 ± 18	1.00	1,360 ± 37	0.97	3,588 ± 89	1.00	2,938 ± 99	0.99
	Increasing	2,147 ± 18	0.99	1,601 ± 39	0.96	3,635 ± 90	0.99	3,158 ± 116	0.97
100	Decreasing	2,592 ± 19	1.00	2,213 ± 32	0.97	4,351 ± 93	1.00	3,657 ± 102	0.99
	Equal	2,710 ± 20	1.00	2,402 ± 37	0.97	4,388 ± 89	1.00	3,792 ± 106	0.98
	Increasing	2,934 ± 20	0.99	2,679 ± 38	0.95	4,444 ± 93	0.99	4,005 ± 112	0.98

Note. Results for KN++ and IZ-free are cited from Fan et al. (2016).

R&S problems); (2) extending the procedures to solve R&S problems with unknown underlying distributions; (3) establishing guidelines for setting the parameter n_0 and studying its impact on the proposed procedures' performance; and (4) conducting comprehensive comparisons with leading Bonferroni-free R&S procedures such as BIZ (Frazier 2014) and DK (Dieker and Kim 2021).

Acknowledgments

The authors thank three anonymous reviewers and the associate editor for carefully reading the paper and for suggestions that significantly improved the paper.

References

Branke J, Chick SE, Schmidt C (2007) Selecting a selection procedure. *Management Sci.* 53(12):1916–1932.
 Chen CH, Lin J, Yücesan E, Chick SE (2000) Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems* 10(3):251–270.
 Chick SE, Inoue K (2001) New two-stage and sequential procedures for selecting the best simulated system. *Oper. Res.* 49(5):732–743.
 Dieker AB, Kim SH (2014) A fully sequential procedure for known and equal variances based on multivariate Brownian motion. Tolk A, Diallo SY, Ryzhov IO, Yilmaz L, Buckley S, Miller JA, eds. *Proc. Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3749–3760.
 Dieker AB, Kim SH (2021) Efficient fully sequential indifference-zone procedures using properties of multidimensional Brownian motion exiting a sphere. Preprint, submitted April 18, <https://arxiv.org/abs/2104.08784>.
 Fan W, Hong LJ (2014) A frequentist selection-of-the-best procedure without indifference zone. Tolk A, Diallo SY, Ryzhov IO, Yilmaz L, Buckley S, Miller JA, eds. *Proc. Winter Simulation Conf.* (IEEE, Piscataway, NJ), 3737–3748.
 Fan W, Hong LJ, Nelson BL (2016) Indifference-zone-free selection of the best. *Oper. Res.* 64(6):1499–1514.
 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.
 Gupta SS (1956) *On a Decision Rule for a Problem in Ranking Means* (University of North Carolina at Chapel Hill, Chapel Hill).
 Gupta SS (1965) On some multiple decision (selection and ranking) rules. *Technometrics* 7(2):225–245.

Hong LJ, Fan W, Luo J (2021) Review on ranking and selection: A new perspective. *Frontiers Engrg. Management* 8(3):321–343.
 Johari R, Koomen P, Pekelis L, Walsh D (2019) Always valid inference: Continuous monitoring of A/B tests. *Oper. Res.* 70(3):1806–1821.
 Kim SH, Nelson BL (2001) A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Modeling Comput. Simulation* 11(3):251–273.
 Kim SH, Nelson BL (2006a) On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Oper. Res.* 54(3):475–488.
 Kim SH, Nelson BL (2006b) Selecting the best system. Henderson SG, Nelson BL, eds. *Handbooks in Operations Research and Management Science*, vol. 13 of Simulation (Elsevier, New York), 501–534.
 Kim SH, Nelson BL (2007) Recent advances in ranking and selection. Henderson SG, Biller B, Hsieh MH, Shortle J, Tew JD, Barton RR, eds. *Proc. Winter Simulation Conf.* (IEEE, Piscataway, NJ), 162–172.
 Lai TL (1977) Power-one tests based on sample sums. *Ann. Statist.* 5(5):866–880.
 Lai TL (1981) Asymptotic optimality of invariant sequential probability ratio tests. *Ann. Statist.* 9(2):318–333.
 Li X, Liu J, Ying Z (2014) Generalized sequential probability ratio test for separate families of hypotheses. *Sequential Anal.* 33(4):539–563.
 Lorden G (1976) 2-SPRT's and the modified Kiefer-Weiss problem of minimizing an expected sample size. *Ann. Statist.* 4(2):281–291.
 Pavlov I (1988) A sequential procedure for testing many composite hypotheses. *Theory Probability Appl.* 32(1):138–142.
 Pavlov I (1991) Sequential procedure of testing composite hypotheses with applications to the Kiefer-Weiss problem. *Theory Probability Appl.* 35(2):280–292.
 Revuz D, Yor M (1999) *Continuous Martingales and Brownian Motion* (Springer-Verlag, Berlin).
 Robbins H, Siegmund D (1974) The expected sample size of some tests of power one. *Ann. Statist.* 2(3):415–436.
 Shafer G, Shen A, Vereshchagin N, Vovk V (2011) Test martingales, Bayes factors and p -values. *Statist. Sci.* 26(1):84–101.
 Siegmund D (1976) Importance sampling in the Monte Carlo study of sequential tests. *Ann. Statist.* 4(4):673–684.
 Tartakovsky A, Nikiforov I, Basseville M (2014) *Sequential Analysis: Hypothesis Testing and Changepoint Detection* (CRC Press, Boca Raton, FL).
 Wald A (1945) Sequential method of sampling for deciding between two courses of action. *J. Amer. Statist. Assoc.* 40(231):277–306.
 Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *Ann. Math. Statist.* 19(3):326–339.

Wang H, Kim SH (2013) Reducing the conservativeness of fully sequential indifference-zone procedures. *IEEE Trans. Automated Control* 58(6):1613–1619.

Wasserman L, Ramdas A, Balakrishnan S (2020) Universal inference. *Proc. National Acad. Sci. USA* 117(29):16880–16890.

Wenyu Wang is a research scientist at Eigen Capital LLC.

Hong Wan is an associate professor in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina

State University. Her research interests include statistical experimental design and analysis of simulation experiments, learning-based simulation, and blockchain system analysis and simulation.

Xi Chen is an associate professor in the Grado Department of Industrial and Systems Engineering at Virginia Tech. Her research addresses statistical issues in the design and analysis of stochastic computer simulation experiments, including metamodeling, input screening, sensitivity analysis, and simulation optimization.