



Operations Research

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

Indifference-Zone-Free Selection of the Best

Weiwei Fan, L. Jeff Hong, Barry L. Nelson

To cite this article:

Weiwei Fan, L. Jeff Hong, Barry L. Nelson (2016) Indifference-Zone-Free Selection of the Best. Operations Research

Published online in Articles in Advance 13 Oct 2016

. <http://dx.doi.org/10.1287/opre.2016.1530>

Full terms and conditions of use: <http://pubsonline.informs.org/page/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 © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

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

Indifference-Zone-Free Selection of the Best

Weiwei Fan

Department of Management Science, School of Management, University of Science and Technology of China, Hefei, China, wfan@ustc.edu.cn

L. Jeff Hong

Department of Economics and Finance and Department of Management Sciences, College of Business, City University of Hong Kong, Kowloon Tong, Hong Kong, jeffhong@cityu.edu.hk

Barry L. Nelson

Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208, nelsonb@northwestern.edu

Many procedures have been proposed in the literature to select the simulated alternative with the best mean performance from a finite set of alternatives. Among these procedures, frequentist procedures are typically designed under either the subset-selection (SS) formulation or the indifference-zone (IZ) formulation. Both formulations may encounter problems when the goal is to select the unique best alternative for any configuration of the means. In particular, SS procedures may return a subset that contains more than one alternative, and IZ procedures hinge on the relationship between the chosen IZ parameter and the true mean differences that is unknown to decision makers a priori. In this paper, we propose a new formulation that guarantees to select the unique best alternative with a user-specified probability of correct selection (PCS), as long as the means of alternatives are unique, and we design a class of fully sequential procedures under this formulation. These procedures are parameterized by the PCS value only, and their continuation boundaries are determined based on the Law of the Iterated Logarithm. Furthermore, we show that users can add a stopping criterion to these procedures to convert them into IZ procedures, and we argue that these procedures have several advantages over existing IZ procedures. Lastly, we conduct an extensive numerical study to show the performance of our procedures and compare their performance to existing procedures.

Keywords: selection of the best; indifference-zone-free; fully sequential procedures.

Subject classifications: simulation.

Area of review: Simulation.

History: Received April 2015; revisions received November 2015, March 2016, April 2016; accepted May 2016.

Published online in *Articles in Advance* October 13, 2016.

1. Introduction

Decision makers often face the problem of selecting the best from a finite set of alternatives, where the best refers to the alternative with the largest (or smallest) mean performance. For instance, an inventory manager may need to select an inventory policy to minimize the total mean cost, while a financial investor may want to choose an investment strategy to maximize the mean payoff. In many practical situations, the mean performances of these alternatives are not explicitly available and can only be evaluated by running simulation experiments. This is known as the selection-of-the-best problem. To solve this problem, a decision procedure is often needed to determine the proper sample sizes of all alternatives and to decide which alternative to select. There is a large body of literature on designing such procedures—see Bechhofer et al. (1995), Kim and Nelson (2006a) and Branke et al. (2007) for introductions and overviews.

Among these procedures in the literature, there are frequentist and Bayesian approaches. Frequentist procedures (e.g., Rinott 1978, Kim and Nelson 2001) allocate simulation observations to different alternatives to achieve a prespecified probability of correct selection (PCS) even for the least

favorable configuration of the means. They are typically conservative (i.e., requiring more samples than necessary) for an average case. Bayesian procedures (e.g., Chick and Inoue 2001, Chick and Frazier 2012) often allocate a finite computing budget to different alternatives to either maximize the posterior (Bayesian) PCS or minimize the opportunity cost. They often require fewer observations than frequentist procedures, but typically do not provide a guaranteed (frequentist) PCS. An exception is Frazier (2014) who recently proposed a Bayes-inspired procedure to achieve a prespecified frequentist PCS. In this paper, we take the frequentist viewpoint, and our goal is to design procedures that deliver a user-specified PCS.

The first frequentist selection-of-the-best formulations were designed in the 1950s, when observations were often collected through physical experiments (e.g., agricultural experiments or clinical trials). Each of these experiments may take quite a long time (e.g., weeks to months) to conduct, and it only makes sense to conduct these experiments in batches. Therefore, the stage-wise procedures that select the best at the end of the last stage were prevalent. For these procedures, the critical issue is to determine the appropriate sample size for each alternative so that the best may be

selected. Notice that the means of the best and the second-best alternatives may be arbitrarily close. Therefore, guaranteeing to find the unique best for all configurations of the means could require an arbitrarily large number of observations, which is not practical. To avoid this difficulty, various formulations were proposed to soften the original goal of selecting the unique best alternative. One is the subset-selection (SS) formulation of Gupta (1956, 1965), which guarantees to select a random-sized subset of the alternatives that contains the best (instead of selecting the unique best). The procedures of Panchapakesan et al. (1971) and Sullivan and Wilson (1989) are also designed under this formulation. The other formulation is the indifference-zone (IZ) formulation of Bechhofer (1954), which guarantees to select the unique best alternative when the difference between the means of the best and second-best is at least δ , where $\delta > 0$ is called an IZ parameter. The procedures proposed by Rinott (1978), Clark and Yang (1986), and many others are also designed under this formulation.

Starting in the 1990s, there was a paradigm change in selection-of-the-best problems. Under this new paradigm, observations of alternatives are often obtained through computer simulation experiments. Each of these experiments takes a fraction of a second to a few minutes, and they are often generated sequentially from either a single processor or a small number of parallel processors. This sequential nature of computer simulation experiments motivates the use of sequential procedures that date back to Paulson (1964) under the IZ formulation. The major milestone of this stream of research was the paper of Kim and Nelson (2001) that completely adapts fully sequential procedures to computer simulation experiments. These fully sequential procedures typically specify a continuation region. They then approximate the sum of differences between two competing alternatives as a Brownian motion (BM) process with drift, and make an elimination decision when the process exits the continuation region. Different from stage-wise procedures, fully sequential procedures gradually gather information (i.e., simulation observations) on the unknown differences between the means of competing alternatives and eliminate the inferior alternatives once enough statistical evidence is available. Therefore, they tend to be more efficient (i.e., need fewer observations) than their stage-wise counterparts.

Even though many fully sequential frequentist procedures have been proposed since the work of Kim and Nelson (2001) (e.g., the procedures of Batur and Kim 2006; Hong and Nelson 2005, 2007; Pichitlamken et al. 2006; Hong 2006; Kim and Nelson 2006b; and many others), none of these questioned the necessity of the IZ formulation. Unlike stage-wise procedures, where information on the means of the competing alternatives are unknown, fully sequential procedures gather this information as they progress. This information may free us from having to provide an IZ parameter and help us to select the unique best alternative for any configuration of the means.

In this paper, we propose a new formulation that selects the best alternative with a user-specified PCS, as long as all of the alternatives have unique means. Under this formulation, the means of any pair of alternatives are not equal and can be arbitrarily close. Therefore, the key to designing such a procedure is to construct a continuation region where, with the specified PCS, no elimination decision will be made if the mean difference between these two alternatives is zero, while a correct elimination decision will be made if it is not. To achieve this, we need to construct a continuation region so that a BM with no drift will stay inside while a BM with drift will exit from the correct side. Notice that BM with and without drift differ in the rates at which they grow. In particular, a BM with positive drift approaches infinity at the rate $O(t)$, while a BM without drift approaches infinity at a rate bounded by $O(\sqrt{t \log \log t})$ due to the *Law of the Iterated Logarithm* (Durrett 2010). In light of this, any continuation region, formed by boundaries which grow at a rate between $O(\sqrt{t \log \log t})$ and $O(t)$, can achieve our objective.

To further determine the design parameters of such continuation regions, we need to be able to evaluate the first-exit probability of a BM from the given continuation region. However, it is generally difficult or even impossible to explicitly analyze the first-exit-time distribution of a BM (see Durbin 1985). To solve the problem, we consider two approaches. In the first approach, we obtain the design parameter by numerically solving a one-dimensional stochastic root-finding problem, which can be solved efficiently with the use of an importance sampling scheme. In the second approach, we obtain the design parameter by evaluating the first-exit probability using an asymptotic result, and we prove the validity of the procedure as PCS goes to 1. It is worthwhile noting that the asymptotic regime of letting PCS go to 1 is a classical one and it dates back to Perng (1969) and Dudewicz (1969).

Beyond selecting the unique best alternative, our procedures can also be used as sequential subset-selection procedures. In particular, at any time before the procedures terminate, the surviving subset of alternatives has retained the best alternative with the same probability guarantee.

One shortcoming of our procedures is that they may not stop when two or more alternatives have the same means. To avoid this situation, we propose a stopping criterion for our procedures. The stopping criterion requires the user to specify an error tolerance ε . We show that, with the same PCS, the selected alternative is either the best or within ε of the best. This indicates that our procedures provide both the correct selection guarantee and good selection guarantee, which is the goal of Ni et al. (2014). If one treats ε as an IZ parameter, then our procedures become IZ procedures. However, unlike IZ procedures, the average sample size of our procedures is typically insensitive to the setting of the error tolerance; thus, one may set it very small for practical implementations.

Our research is closely related to the literature on best-arm identification in stochastic multi-armed bandit (MAB)

problems. (Readers interested in the MAB problems may refer to Bubeck and Cesa-Bianchi 2012 for a comprehensive overview). The best-arm identification procedures seek the same goal as our procedures do—i.e., selecting the unique best arm with a predetermined probability. In contrast, our procedures and these best-arm identification procedures differ in at least three aspects. Firstly, they employ different elimination mechanisms. In particular, the best-arm identification procedures (e.g., Karmin et al. 2013, Jamieson and Nowak 2014) sequentially construct a confidence bound for the mean reward of each arm based on the estimated sample mean and make eliminations if there are nonoverlapping confidence bounds. Secondly, the best-arm identification procedures assume that the cumulant generating function of each reward is bounded by a known convex function, while our procedures only require a finite joint moment generating function of rewards in a neighborhood of the zero vector (see Theorem 2), and it need not be known. Lastly, our procedures are able to allow dependence among the rewards of different arms; thus, common random numbers can be used to speed up the selection of the best.

The rest of our paper is organized as follows. In Section 2, we introduce a new selection-of-the-best formulation. In Section 3, we design a class of fully sequential IZ-free procedures under the new formulation when observations from all alternatives are jointly normally distributed. In Section 4, we relax the normality assumption to design the corresponding sequential procedures that are asymptotically valid. In Section 5, we analyze the asymptotic efficiency of our procedures. We discuss a stopping criterion in Section 6. We conduct a comprehensive numerical study to understand our procedures and compare them to existing procedures in Section 7, and conclude in Section 8.

2. A New Selection-of-the-Best Formulation

Suppose there are k (≥ 2) competing alternatives at the beginning of the selection process, and the goal is to select the alternative that has the largest mean performance. For $i = 1, 2, \dots, k$, denote μ_i as the unknown mean performance of alternative i and we evaluate it through computer simulation. Without loss of generality, assume $\mu_1 > \mu_2 > \dots > \mu_k$, implying that the best alternative is unique and alternative 1 is the best. It is worthwhile noting that we make no assumption on the difference between μ_1 and μ_2 because there is often no such knowledge in practice. Further, let $1 - \alpha$ ($0 < \alpha \leq 1 - 1/k$) denote the user-specified PCS. In other words, users target at selecting alternative 1 with probability at least $1 - \alpha$.

To achieve this target, there are two different frequentist formulations in the literature. They are the SS and IZ formulations. As stated in Section 1, both of them may encounter difficulties. In light of this, we propose a new selection-of-the-best formulation in this section. Under the new formulation, procedures guarantee to select the best alternative with

a user-specified PCS value $1 - \alpha$, as long as their means are unique; i.e., $\mu_1 > \mu_2 > \dots > \mu_k$. Equivalently, procedures under the new formulation guarantee to satisfy

$$\mathbb{P}\{\text{select alternative 1}\} \geq 1 - \alpha. \quad (1)$$

Therefore, to use the procedures under the new formulation, one only needs to specify the PCS value $1 - \alpha$. Compared to the IZ formulation, the new formulation frees users from having to specify an IZ parameter, and compared to the SS formulation, it selects the unique best alternative instead of a random subset but can select a subset if procedures are terminated early.

In the following two sections, we design sequential procedures under the new formulation for cases where observations are jointly normally distributed as well as cases where observations are generally distributed, and prove that these procedures can deliver the required statistical guarantee (i.e., Equation (1)). We emphasize that we use sequential procedures not only because they are more efficient but also because stage-wise procedures cannot achieve the necessary statistical guarantee.

3. Procedures with Normally Distributed Observations

Let X_{ir} , $r = 1, 2, \dots$, denote the r th independent observation from alternative i , for $i = 1, 2, \dots, k$. In this section, we consider the case where $(X_{1r}, X_{2r}, \dots, X_{kr})$ ($r = 1, 2, \dots$) are jointly normally distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$ and unknown positive definite covariance matrix. The assumption of possible dependence among observations from different alternatives enables the use of common random numbers (CRNs). The task of this section is to design a sequential procedure that deliver the statistical guarantee stated in (1).

Existing sequential procedures, such as the KN procedure, often decompose the selection-of-the-best problem into pairwise comparisons. In each pairwise comparison, they then approximate the partial-sum difference process between two competing alternatives by a BM with drift and design a continuation region to eliminate the inferior alternative based on the first-exit time of the BM from it. In this paper, we adopt the same mechanism to design a sequential procedure under the new formulation and, consequently, we need to address the following two issues: (1) constructing a Brownian approximation to the partial-sum difference process and (2) designing a proper continuation region.

Under the normality assumption, it is straightforward to construct a Brownian approximation. Consider a pairwise comparison between alternatives i and j , where $i \neq j$. Let $\bar{X}_i(n)$ denote the sample mean calculated from the first n observations from alternative i and define

$$t_{ij}(n) = n/\sigma_{ij}^2 \quad \text{and} \quad (2)$$

$$Z_{ij}(t_{ij}(n)) = t_{ij}(n)[\bar{X}_i(n) - \bar{X}_j(n)],$$

where $\sigma_{ij}^2 = \text{Var}[X_{i1} - X_{j1}]$. Throughout this paper, we define $B(\cdot)$ as a BM with unit variance and no drift, and define $B_\Delta(\cdot)$ as a BM with unit variance and drift Δ . Theorem 1 of Hong (2006) shows that the random sequences $\{Z_{ij}(t_{ij}(n)): n \geq 1\}$ and $\{B_{\mu_i - \mu_j}(t_{ij}(n)): n \geq 1\}$ have the same joint distribution. Therefore, $\{Z_{ij}(t_{ij}(n)): n \geq 1\}$ can be viewed as a BM with drift $\mu_i - \mu_j$ observed at the discrete time points $\{t_{ij}(1), t_{ij}(2), \dots\}$.

Let $(-g_c(t), g_c(t))$ denote the continuation region used in the procedure, where c is a constant determined based on the number of alternatives k , the PCS value $1 - \alpha$, and the common first-stage sample size n_0 . As discussed in Section 1, the continuation region is expected to satisfy that, with a prespecified PCS, a BM with no drift stays inside, while a BM with drift exits from the correct side. Notice that the former grows to infinity at a rate slower than that of the latter. In particular, the *Law of the Iterated Logarithm* (Durrett 2010) states that, the rate of the former is bounded by $O(\sqrt{t \log \log t})$, while the rate of the latter is $O(t)$. The rate difference motivates the use of a boundary function that increases at rates in-between to construct the targeted continuation region. In light of this, boundaries of interest are categorized into the following set:

$$\mathcal{G} = \left\{ g_c(t) \in C^1[0, \infty): \lim_{t \rightarrow \infty} \frac{g_c(t)}{t} = 0 \text{ and } \liminf_{t \rightarrow \infty} \frac{g_c(t)}{\sqrt{t \log \log t}} \in (0, \infty] \right\},$$

where $C^1[0, \infty)$ denotes the set of functions whose first derivatives are continuous on $[0, \infty)$. Furthermore, for the ease of theoretical analysis and practical implementation, we choose $g_c(t)$ from \mathcal{G} such that $g_c(t)$ is increasing in t and $g_c(t)/t$ is decreasing in t for any fixed c . For instance, we can choose $\sqrt{2t \log(c + \log(t+1))}$ ($c \geq 1$), $\sqrt{[c + \log(t+1)](t+1)}$ ($c > 0$), or ct^β ($1/2 < \beta < 1$, $c > 0$). The determination of the design parameter c is deferred to Section 3.1.

3.1. The Procedure

In this subsection, we propose a sequential procedure designed based on the key idea stated above. Particularly, the procedure uses a first stage to estimate the unknown variances and determine a proper continuation region, and then collects observations sequentially to eliminate inferior alternatives until only one is left.

PROCEDURE 1 (SEQUENTIAL PROCEDURE FOR NORMALLY DISTRIBUTED OBSERVATIONS). Setup. Select the PCS $1 - \alpha$ ($0 < \alpha \leq 1 - 1/k$) and a common first-stage sample size $n_0 \geq 2$. Choose a boundary $g_c(t)$ from \mathcal{G} and calculate the constant c that is the unique root to the following equation:

$$\begin{aligned} & \mathbb{E} \left[\mathbb{P} \left\{ B(T) \leq -g_c \left(\frac{n_0 - 1}{Y} T \right) \frac{Y}{n_0 - 1}, T < \infty \mid Y \right\} \right] \\ & = \frac{\alpha}{k - 1}, \end{aligned} \quad (3)$$

where $T = \inf\{t: |B(t)| \geq g_c(((n_0 - 1)/Y)t(Y/(n_0 - 1)))\}$ and Y is a chi-squared distributed random variable with $n_0 - 1$ degrees of freedom that is independent of $B(\cdot)$.

Initialization. Let $I = \{1, 2, \dots, k\}$ be the set of alternatives in contention. For each $i \in I$, simulate n_0 observations $X_{i1}, X_{i2}, \dots, X_{in_0}$ from alternative i and calculate its sample mean as

$$\bar{X}_i(n_0) = \frac{1}{n_0} \sum_{l=1}^{n_0} X_{il}.$$

For any $i, j \in I$ with $i \neq j$, calculate the sample variance of their difference as

$$S_{ij}^2 = \frac{1}{n_0 - 1} \sum_{r=1}^{n_0} [X_{ir} - X_{jr} - (\bar{X}_i(n_0) - \bar{X}_j(n_0))]^2.$$

Set $n = n_0$.

Screening. Let $I^{\text{old}} = I$ and for any $i, j \in I$ with $i \neq j$, let

$$\tau_{ij}(n) = n/S_{ij}^2 \quad \text{and} \quad Z_{ij}(\tau_{ij}(n)) = \tau_{ij}(n)[\bar{X}_i(n) - \bar{X}_j(n)].$$

Let

$$I = I^{\text{old}} \setminus \{i \in I^{\text{old}}: Z_{ij}(\tau_{ij}(n)) \leq -g_c(\tau_{ij}(n)), \text{ for some other } j \in I^{\text{old}}\}.$$

Stopping Rule. If $|I| = 1$, stop and select the alternative whose index is in I as the best. Otherwise, take one additional observation from each alternative $i \in I$, set $n = n + 1$ and go to Screening.

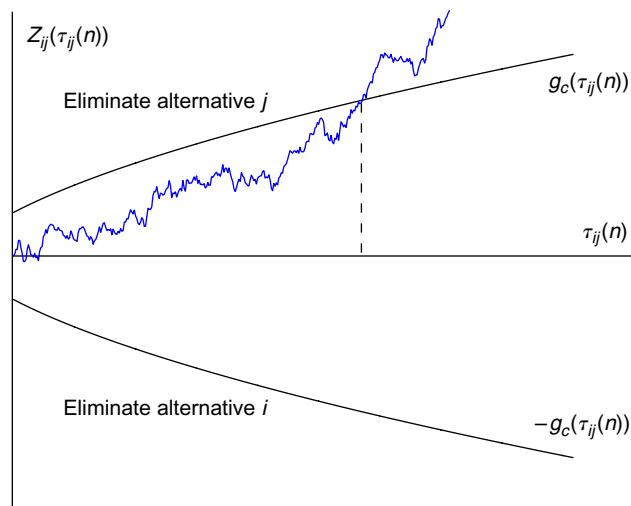
REMARK 1. Procedure 1 can work with many choices of boundaries (i.e., \mathcal{G}). In the numerical studies reported in Section 7, as an instance we use a specific boundary $g_c(t) = \sqrt{[c + \log(t+1)](t+1)}$.

REMARK 2. Please refer to EC.2 in the e-companion (available as supplemental material at <https://doi.org/10.1287/opre.2016.1530>) on how to choose c that satisfies Equation 3.

Procedure 1, like those of Kim and Nelson (2001) and Hong (2006), is essentially a sequential screening procedure and works as follows. It includes all of the alternatives into I at the *Initiation* step and eliminates the inferior alternatives from I over the *Screening* steps until only one alternative is left. At each *Screening* step, it considers all of the possible pairwise comparisons between the surviving alternatives and eliminates the inferior alternatives based on the first-exit time of the partial-sum difference process from the given continuation region. Take the comparison between alternatives i and j , for instance. When the partial-sum difference process $Z_{ij}(\tau_{ij}(\cdot))$ exits the continuation region from below (above), Procedure 1 eliminates alternative $i(j)$; otherwise, it eliminates neither of them and continues collecting observations (see Figure 1). Once an alternative is eliminated, it will not be considered again.

Even though Procedure 1 is designed to select the unique best alternative, users can stop Procedure 1 at any time before termination and it still guarantees to have retained the best alternative. This result is summarized as the following Lemma 1.

Figure 1. (Color online) Continuation region of our sequential procedures.



Note. The x -axis is well defined because $\tau_{ij}(n)$ increases linearly with n .

LEMMA 1. Define N as the number of replications when Procedure 1 terminates and only one alternative is left. Let $I(n)$ for $n \leq N$ denote the set of surviving alternatives in Procedure 1 after n replications. If Procedure 1 delivers the statistical guarantee as stated in Equation (1), then we have, for any $n \in \mathbb{Z}^+$,

$$\mathbb{P}\{1 \in I(n \wedge N)\} \geq 1 - \alpha,$$

where $x \wedge y$ denotes the smaller of x and y .

PROOF. See EC.1.1 in the e-companion. \square

Lemma 1 shows that, if there is a maximal stopping time n , Procedure 1 guarantees with at least probability $1 - \alpha$ to either select the best alternative before n or retain the best alternative in a subset when it is stopped at n . Therefore, Procedure 1 can be utilized as a sequential subset-selection procedure when users have a sampling budget on screening. Besides, Procedure 1 tends to be more efficient than the existing subset-selection procedures (e.g., Gupta 1956) that use equal allocation of observations among alternatives, especially when the number of alternatives is large. This arises because Procedure 1 tends to allocate fewer observations to the significantly inferior alternatives.

3.2. Statistical Validity

In this subsection, we establish the statistical validity of Procedure 1—i.e., it delivers the statistical guarantee stated in Equation (1). To achieve this goal, we need several lemmas.

In each pairwise comparison, an incorrect-selection event occurs if the better alternative is eliminated, or equivalently, the partial-sum difference process exits the continuation region from the wrong direction. Notice that we only col-

lect observations at discrete time points in the procedure. Hence, the partial-sum process corresponds to observing a BM at discrete time points. The following lemma states that, under very general conditions, the procedure designed for the continuous-time BM provides an upper bound on the probability of incorrect selection (PICS) for the discrete-time BM.

LEMMA 2. Let $B_\Delta(t)$ denote a continuous-time Brownian motion with drift $\Delta > 0$. A discrete-time process is obtained by observing $B_\Delta(t)$ at a (possibly) random, increasing sequence of times $\{t_i: i = 1, 2, \dots\}$ taking values in a given countable set. The value of t_i depends on $B_\Delta(t)$ only through its values in the period $[0, t_{i-1}]$. Define $T_d = \min\{t_i: |B_\Delta(t_i)| \geq g(t_i)\}$ and $T_c = \inf\{t: |B_\Delta(t)| \geq g(t)\}$, where $g(t) \in \mathcal{G}$. Then we have

- (1) $T_c \leq T_d < \infty$, w.p.1,
 - (2) $\mathbb{P}\{B_\Delta(T_d) \leq -g(T_d)\} \leq \mathbb{P}\{B_\Delta(T_c) \leq -g(T_c)\}$
- (Jennison et al. 1980).

PROOF. See EC.1.2 in the e-companion. \square

The drift of the BM refers to the difference of two competing means, which is unknown in our setting. In the following lemma, we show that PICS of the BM with the unknown drift is bounded by its counterpart without drift. This statement is valid for general stochastic processes, not just for BM.

LEMMA 3. Let $Z(\cdot)$ denote a (discrete-time or continuous-time) stochastic process. (Assume that $Z(\cdot)$ has right-continuous sample paths if $Z(\cdot)$ is a continuous-time stochastic process.) Let $Z_\Delta(t) = Z(t) + \Delta t$, where $\Delta > 0$. Further, define T and T_Δ as the stopping times that $Z(\cdot)$ and $Z_\Delta(\cdot)$ first exit the region $(-g(t), g(t))$, where $g(t)$ is a continuous and nonnegative function. Then, we have that

$$\mathbb{P}\{Z_\Delta(T_\Delta) \leq -g(T_\Delta), T_\Delta < \infty\} \leq \mathbb{P}\{Z(T) \leq -g(T), T < \infty\}.$$

PROOF. See EC.1.3 in the e-companion. \square

Now, we can establish the statistical validity of Procedure 1 and summarize it as Theorem 1.

THEOREM 1. Suppose that there are k alternatives and we are interested in selecting the alternative with the largest mean. Assume that $(X_{1r}, X_{2r}, \dots, X_{kr})$, $r = 1, 2, \dots$, are independent and identically distributed (i.i.d.) and jointly normally distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$ and unknown positive definite covariance matrix. If $\mu_1 > \mu_2 > \dots > \mu_k$, Procedure 1 stops in finite time with probability 1 (w.p.1) and selects the best alternative (i.e., alternative 1) with probability at least $1 - \alpha$ (i.e., $\mathbb{P}\{\text{select alternative 1}\} \geq 1 - \alpha$).

REMARK 3. Theorem 1 only requires the covariance matrix to be positive definite. Therefore, it allows the use of CRNs that introduces positive correlation among observations from different alternatives to make the comparisons sharper.

PROOF. For any pair of alternatives i and j with $i \neq j$, since $\mu_i - \mu_j \neq 0$, we see that the elimination between them occurs

in finite time w.p.1 by part (1) of Lemma 2. Similarly, we have all of the pairwise eliminations occur in finite time w.p.1. Therefore, Procedure 1 stops in finite time w.p.1.

When Procedure 1 stops, define an incorrect-selection (ICS) event happens when alternative 1 is eliminated by some other alternative. Then, we have

$$\begin{aligned} \mathbb{P}\{\text{ICS}\} &= \mathbb{P}\left\{\bigcup_{i=2}^k \{\text{alternative 1 is eliminated by alternative } i\}\right\} \\ &\leq \sum_{i=2}^k \mathbb{P}\{\text{alternative 1 is eliminated by alternative } i\}. \quad (4) \end{aligned}$$

The inequality in (4) arises from the Bonferroni inequality. For any $i = 2, 3, \dots, k$, we denote ICS_i as the incorrect-selection event that alternative 1 is eliminated by alternative i . To prove the statistical validity, we view each incorrect-selection event ICS_i separately. In particular, we have

$$\mathbb{P}\{\text{ICS}_i\} = \mathbb{P}\{Z_{1i}(\tau_{1i}(N_i)) \leq -g_c(\tau_{1i}(N_i)), N_i < \infty\},$$

where $N_i = \min\{n \geq n_0: |Z_{1i}(\tau_{1i}(n))| \geq g_c(\tau_{1i}(n)), n \in \mathcal{X}^+\}$. Let $Y = (n_0 - 1)S_{1i}^2/\sigma_{1i}^2$ and it follows that

$$\begin{aligned} \mathbb{P}\{\text{ICS}_i\} &= \mathbb{P}\left\{Z_{1i}(t_{1i}(N_i)) \leq -g_c\left(\frac{\sigma_{1i}^2}{S_{1i}^2}t_{1i}(N_i)\right)\frac{S_{1i}^2}{\sigma_{1i}^2}, N_i < \infty\right\} \\ &= \mathbb{E}\left[\mathbb{P}\left\{Z_{1i}(t_{1i}(N_i)) \leq -g_c\left(\frac{n_0-1}{Y}t_{1i}(N_i)\right)\frac{Y}{n_0-1}, \right. \right. \\ &\quad \left. \left. N_i < \infty \mid Y\right\}\right]. \quad (5) \end{aligned}$$

Under the normality assumption, we have shown that $\{Z_{1i}(t_{1i}(n)): n \in \mathcal{X}\}$ can be viewed as a BM with drift $\mu_1 - \mu_i$ observed at the discrete time points $\{t_{1i}(n): n \geq n_0\}$. By Basu's theorem (Basu 1955), we see that Y follows a chi-squared distribution with $n_0 - 1$ degrees of freedom and is independent of $Z_{1i}(t_{1i}(n_0)) = t_{1i}(n_0)[\bar{X}_1(n_0) - \bar{X}_i(n_0)]$. Further, Y is a function of $X_{11} - X_{i1}, X_{12} - X_{i2}, \dots, X_{1n_0} - X_{in_0}$ and thus is independent of $X_{1, n_0+1} - X_{i, n_0+1}, X_{1, n_0+2} - X_{i, n_0+2}, \dots$. Therefore, Y is independent of $Z_{1i}(t_{1i}(n))$ for any $n \geq n_0$. This implies that, conditioning on Y , $\{Z_{1i}(t_{1i}(n)): n \geq n_0\}$ can still be viewed as a BM with drift $\mu_1 - \mu_i$ observed at the discrete time points $\{t_{1i}(n): n \geq n_0\}$. Then, we get that

$$\begin{aligned} (5) &\leq \mathbb{E}\left[\mathbb{P}\left\{B_{\mu_1 - \mu_i}(T_i) \leq -g_c\left(\frac{n_0-1}{Y}T_i\right)\frac{Y}{n_0-1}, T_i < \infty \mid Y\right\}\right] \\ &\quad \text{by Lemma 2} \\ &\leq \mathbb{E}\left[\mathbb{P}\left\{B(T) \leq -g_c\left(\frac{n_0-1}{Y}T\right)\frac{Y}{n_0-1}, T < \infty \mid Y\right\}\right] \\ &\quad \text{by Lemma 3} \\ &= \frac{\alpha}{k-1}, \end{aligned}$$

where $T_i = \inf\{t: |B_{\mu_1 - \mu_i}(t)| \geq g_c(((n_0 - 1)/Y)t) \cdot (Y/(n_0 - 1))\}$ and $T = \inf\{t: |B(t)| \geq g_c(((n_0 - 1)/Y)t) \cdot (Y/(n_0 - 1))\}$. The last equation holds due to the choice

of the design parameter c (see Equation (3)). Furthermore, combining with (4), we have that $\mathbb{P}\{\text{ICS}\} \leq \alpha$. Therefore, the conclusion of this theorem holds. \square

In Theorem 1, we established the statistical validity of Procedure 1 for the case when all of the means are unique—i.e., $\mu_1 > \mu_2 > \dots > \mu_k$. More generally, the best alternative may be unique, and there might be potential ties among the nonoptimal alternatives—i.e., $\mu_1 > \mu_2 \geq \dots \geq \mu_k$. For this case, we show in the following proposition that Procedure 1 can provide a statistical guarantee that is similar to but weaker than that of Theorem 1.

PROPOSITION 1. *Under the same assumptions as in Theorem 1, if $\mu_1 > \mu_2 \geq \dots \geq \mu_k$, Procedure 1 terminates and selects alternative 1 with probability at least $1 - \alpha$ (i.e., $\mathbb{P}\{\text{terminate and select alternative 1}\} \geq 1 - \alpha$).*

PROOF. For any sample path, part (1) of Lemma 2 implies that Procedure 1 does not terminate in finite time only if all of the remaining alternatives have the same mean. Because the best alternative is unique, it suffices to state that the procedure does not terminate in finite time only if alternative 1 is eliminated. Besides, we have

$$\begin{aligned} \mathbb{P}\{\text{alternative 1 is eliminated}\} &= \mathbb{P}\left\{\bigcup_{i=2}^k \{\text{alternative 1 is eliminated by alternative } i\}\right\} \\ &\leq \sum_{i=2}^k \mathbb{P}\{\text{alternative 1 is eliminated by alternative } i\}. \quad (6) \end{aligned}$$

The inequality above arises from the Bonferroni inequality. Notice that (6) is the same as (4), which is shown to be no larger than α in the proof of Theorem 1. Therefore, we have

$$\begin{aligned} \mathbb{P}\{\text{Procedure 1 does not terminate in finite time}\} &\leq \mathbb{P}\{\text{alternative 1 is eliminated}\} \leq \alpha, \end{aligned}$$

and

$$\begin{aligned} \mathbb{P}\{\text{terminate and select alternative 1}\} &= 1 - \mathbb{P}\{\text{alternative 1 is eliminated}\} \geq 1 - \alpha. \end{aligned}$$

Therefore, the proposition holds. \square

If there are tied means among alternatives other than the best, it is possible that Procedure 1 may not terminate in finite time. Proposition 1 proves that the probability is at most α . Therefore, in practical situations when users are unsure whether the means are unique, the Procedure with a stopping criterion is recommended to use, and the detailed discussion is deferred to Section 6.

4. Procedures with Generally Distributed Observations

In practice, observations are rarely jointly normally distributed and may even follow different distributions across different alternatives. Therefore, we are motivated to relax the normality assumption and extend our procedure to solve general selection-of-the-best problems.

When observations are generally distributed, it appears difficult to design procedures that deliver a finite-sample statistical validity as we do in Section 3. Therefore, we propose a new procedure that delivers an asymptotic statistical validity—i.e., the procedure satisfies Equation (1) in a meaningful limit.

Recall that, to design a sequential procedure, we face two issues: (1) constructing a Brownian approximation to the partial-sum difference process and (2) designing a proper continuation region. For the case of generally distributed observations, neither of these two issues are trivial.

Consider a pairwise comparison between alternatives 1 and i with $i \neq 1$. As in Equation (2), let $Z_{1i}(\cdot)$ denote the partial-sum difference process on a set of time points $\{t_{1i}(n): n = 1, 2, \dots\}$. Under the normality assumption, we have shown in Section 3 that $\{Z_{1i}(t_{1i}(n)): n = 1, 2, \dots\}$ can be viewed as a BM process observed at discrete time points. However, this will not be true for a general case. Therefore, we consider a standardized version of $\{Z_{1i}(t_{1i}(n)): n = 1, 2, \dots\}$ with a scaling parameter M , that is

$$C_{1i}(M, t_{1i}(s)) := \frac{Z_{1i}(t_{1i}(\lfloor Ms \rfloor)) - t_{1i}(Ms)(\mu_1 - \mu_i)}{\sqrt{M}} \\ = \frac{\sum_{l=1}^{\lfloor Ms \rfloor} (X_{1l} - X_{il}) - Ms(\mu_1 - \mu_i)}{\sqrt{M}\sigma_{1i}^2}, \quad \text{for } s > 0 \quad (7)$$

where $\lfloor x \rfloor$ denotes the largest integer that is no bigger than x . Then we approximate (7) by a BM in the limit using the following Functional Central Limit Theorem (see, e.g., Whitt 2002a, Theorem 4.3.2).

LEMMA 4 (FUNCTIONAL CENTRAL LIMIT THEOREM). *Suppose that $\{X_i: i \geq 1\}$ is a sequence of i.i.d. random variables with mean μ and finite variance σ^2 . Then,*

$$\frac{\sum_{i=1}^{\lfloor nt \rfloor} X_i - nt\mu}{\sqrt{n}\sigma} \xrightarrow{d} B(t), \quad \text{as } n \rightarrow \infty, \text{ in } D[0, \infty),$$

where \xrightarrow{d} denotes convergence in distribution and D is defined as the Skorohod space.

According to Lemma 4 and the self-similarity scaling property of a BM, we have that

$$C_{1i}(M, t_{1i}(s)) \xrightarrow{d} \sigma_{1i}^{-1} B(s) \stackrel{d}{=} B(t_{1i}(s)), \\ \text{as } M \rightarrow \infty, \text{ in } D[0, \infty), \quad (8)$$

where $X \stackrel{d}{=} Y$ denotes that X and Y are identically distributed. Even though $C_{1i}(M, t_{1i}(s))$ is defined for a general distribution, Lemma 4 guarantees its convergence to a known BM in the limit.

We call the parameter M the scaling parameter because it determines how the partial-sum process $Z_{1i}(t_{1i}(n))$ is scaled and how fast the scaled process $C_{1i}(M, t_{1i}(s))$ converges to a BM. To ensure the convergence in (8), we only need M to grow to infinity.

We address the second issue by choosing one boundary $g_c(t)$ from the set \mathcal{G} , similar to that in Section 3. To determine its design parameter c , we need to evaluate the first-exit probability, which shrinks to zero in the asymptotic regime. In this situation, it becomes impractical to find c numerically as in Appendix B. Fortunately, there is a body of literature on asymptotic approximations of this type of first-exit probabilities. For instance, Jennen and Lerche (1981) provide the asymptotic approximations for a list of boundaries. In the following lemma, we present one of these approximations where the boundary is $g_c(t) = \sqrt{[c + \log(t + 1)](t + 1)}$. In this lemma, we need to use the concept of asymptotic equivalence. In particular, we say $f(c)$ is asymptotically equivalent to $g(c)$ with respect to c , denoted as, $f(c) \stackrel{c \rightarrow \infty}{\sim} g(c)$, if $\lim_{c \rightarrow \infty} f(c)/g(c) = 1$.

LEMMA 5 (JENNEN AND LERCHE (1981)). *Let*

$$g_c(t) = \sqrt{[c + \log(t + 1)](t + 1)} \quad (c \geq 0),$$

a nonnegative continuous function, and let $B(t)$ be a Brownian motion without drift. Then $\mathbb{P}\{B(T) \leq -g_c(T), T < \infty\} \stackrel{c \rightarrow \infty}{\sim} \frac{1}{2}e^{-c/2}$, where $T = \inf\{t: |B(t)| \geq g_c(t)\}$.

Lemma 5 provides a closed-form approximation to the first-exit probability when $g_c(t) = \sqrt{(c + \log(t + 1))(t + 1)}$. This allows us to choose the value of c more easily compared to the root-finding problem (3) in Section 3, although this value is only an approximation.

4.1. The Procedure

In this subsection, we propose a sequential procedure for generally distributed observations—i.e., Procedure 2.

PROCEDURE 2 (SEQUENTIAL PROCEDURE FOR GENERALLY DISTRIBUTED OBSERVATIONS). *Setup. Select the PCS $1 - \alpha$ ($0 < \alpha \leq 1 - 1/k$), and a common first-stage sample size $n_0 \geq 2$. Choose the boundary in the form $g_c(t) = \sqrt{[c + \log(t + 1)](t + 1)}$ and set $c = -2 \log(2\alpha/(k - 1))$.*

Initialization. Let $I = \{1, 2, \dots, k\}$ be the set of alternatives in contention. For all $i = 1, 2, \dots, k$, simulate n_0 observations $X_{i1}, X_{i2}, \dots, X_{in_0}$ from alternative i and calculate its sample mean as

$$\bar{X}_i(n_0) = \frac{1}{n_0} \sum_{l=1}^{n_0} X_{il}.$$

For any i, j with $i \neq j$, calculate the sample variance of their difference,

$$S_{ij}^2(n_0) = \frac{1}{n_0 - 1} \sum_{r=1}^{n_0} [X_{ir} - X_{jr} - (\bar{X}_i(n_0) - \bar{X}_j(n_0))]^2.$$

Set $n = n_0$.

Update. If $n > n_0$, calculate the sample mean for each alternative $i \in I$, and for any pair of alternatives $i, j \in I$ with $i \neq j$, update the sample variance $S_{ij}^2(n)$ of their difference.

Screening. Let $I^{\text{old}} = I$ and

$$\tau_{ij}(n) = n/S_{ij}^2(n), \quad \text{and}$$

$$Z_{ij}(\tau_{ij}(n)) = \tau_{ij}(n)[\bar{X}_i(n) - \bar{X}_j(n)].$$

Define

$$I = I^{\text{old}} \setminus \{i \in I^{\text{old}}: Z_{ij}(\tau_{ij}(n)) \leq -g_c(\tau_{ij}(n)), \text{ for some other } j \in I^{\text{old}}\}.$$

Stopping Rule. If $|I| = 1$, stop and select the alternative whose index is in I as the best. Otherwise, take one additional observation $X_{i,n+1}$ from each alternative i for $i \in I$, set $n = n + 1$ and go to Update.

REMARK 4. Similar to Procedure 1, Procedure 2 can also work with many choices of boundaries (i.e., \mathcal{G}). As an instance, Procedure 2 uses a specific boundary $g_c(t) = \sqrt{[c + \log(t + 1)](t + 1)}$, and the following results for Procedure 2 are shown for this specific boundary. Nevertheless, these results can be easily extended to other boundaries in \mathcal{G} .

Procedure 2 differs from Procedure 1 in two aspects. First, Procedure 2 determines the value of c using an asymptotic result, while Procedure 1 does it by solving a stochastic root-finding problem (i.e., Equation (3)). Second, Procedure 2 updates the sample variances sequentially as more and more observations are generated. This enables us to establish the asymptotic validity of Procedure 2 based on the strong consistency of the variance estimators.

4.2. Statistical Validity

Procedure 2 is only parameterized by the PCS value $1 - \alpha$. Therefore, we consider an asymptotic regime, that is, PCS goes to 1 (or $\alpha \rightarrow 0$). Under this asymptotic regime, a procedure is asymptotically valid if its actual probability of incorrect selection (PICS) not only converges to 0 as α converges to 0 but also converges at least as fast as α does; i.e., $\limsup_{\alpha \rightarrow 0} \mathbb{P}\{\text{ICS}\} / \alpha \leq 1$.

To show the asymptotic validity of Procedure 2, we need Lemma 6. It shows that the first-exit time goes to infinity as $\alpha \rightarrow 0$ (or equivalently, $c \rightarrow \infty$) with probability 1 (w.p.1). Based on this lemma, we can prove that the variance estimators at the exit time converge to the true values as $\alpha \rightarrow 0$. This allows us to treat the unknown variances essentially as their true values when analyzing the asymptotic statistical validity of Procedure 2.

LEMMA 6. Let $\{Y_i; i = 1, 2, \dots\}$ denote i.i.d. random variables with $\mathbb{E}[Y_1] = \mu$, and define $Z(n) = \sum_{i=1}^n Y_i$, for $n = 1, 2, \dots$. Let $N_c = \min\{n \in \mathcal{Z}^+: |Z(n)| \geq g_c(n)\}$, where $g_c(t)$ is defined in Procedure 2. Then $N_c \rightarrow \infty$, w.p.1, as $c \rightarrow \infty$.

PROOF. See EC.1.4 in the e-companion. \square

Now we establish the asymptotic validity of Procedure 2 and summarize it in the following theorem.

THEOREM 2. Suppose that there are k alternatives and we are interested in selecting the alternative with the largest mean. Assume that $(X_{1r}, X_{2r}, \dots, X_{kr})$, $r = 1, 2, \dots$, are independent and identically jointly distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$, and their moment generating function exists in a neighborhood of $(0, 0, \dots, 0) \in \mathcal{R}^k$.

Further, let the first-stage sample size n_0 be a function of α that satisfies $n_0 \rightarrow \infty$ as $\alpha \rightarrow 0$. If $\mu_1 > \mu_2 > \dots > \mu_k$, Procedure 2 stops in finite time w.p.1 and satisfies $\limsup_{\alpha \rightarrow 0} \mathbb{P}\{\text{alternative 1 is eliminated}\} / \alpha \leq 1$.

REMARK 5. To ensure the strong consistency of sequentially updated variance estimators, in addition to Lemma 6, we put some technical conditions on the first-stage sample size n_0 , as listed in Theorem 2. Although the theoretical condition facilitates the asymptotic proof, it does not prescribe a specific setting of n_0 in practice.

PROOF. See EC.1.5 in the e-companion. \square

Similar to Proposition 1, we establish Proposition 2 for Procedure 2 in the case when the best alternative is unique but the means are not all unique.

PROPOSITION 2. Under the same assumptions as in Theorem 2, if $\mu_1 > \mu_2 \geq \dots \geq \mu_k$, Procedure 2 terminates and eliminates alternative 1 with probability at most α in an asymptotic regime; i.e.,

$$\limsup_{\alpha \rightarrow 0} \frac{\mathbb{P}\{\text{terminate and eliminate alternative 1}\}}{\alpha} \leq 1.$$

PROOF. See EC.1.8 in the e-companion. \square

5. Asymptotic Efficiency


When choosing a frequentist selection-of-the-best procedure, one not only requires the procedure to deliver the guaranteed PCS, but also wants the procedure to be efficient—i.e., using as few observations as possible to select the best. To understand the efficiency of our procedures, we consider the general case when the observations are jointly generally distributed. In this situation, we focus on our asymptotic procedure (i.e., Procedure 2) and study its average sample size asymptotically.

In Procedure 2, the strong consistency of variance estimators leads to an asymptotic analysis that is similar to the known-variances case. Therefore, for simplicity, we assume that the variances are known in this section. In Section 5.1 we provide an analytic expression for the asymptotic average sample size required by Procedure 2, and in Section 5.2 we compare it to those of existing IZ procedures.

5.1. Asymptotic Average Sample Size

As a building block, we consider a pairwise comparison between alternatives 1 and i with mean difference $\mu_1 - \mu_i$. Lemma 6 implies that the sample size required to detect them grows to infinity in the asymptotic regime. In this case, their partial-sum difference process behaves more and more like a line with slope $\mu_1 - \mu_i$. Therefore, a heuristic method to calculate the average sample size is to find the intersection point between the boundary and the line (see Figure 2). The method is called the mean path approximation. Perng (1969) provides a rigorous proof of this heuristic method for Paulson’s procedure, and the proof can be extended easily

Figure 2. Heuristic mean path approximation method.



to sequential IZ procedures with a triangular continuation region (such as the KN procedure). The result is summarized in Theorem 3.

THEOREM 3 (PERNG (1969)). *Suppose that there are k alternatives and we are interested in selecting the alternative with the largest mean. Assume that $(X_{1r}, X_{2r}, \dots, X_{kr})$, $r = 1, 2, \dots$, are independent and identically jointly distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$ and a positive definite covariance matrix. Further, assume $\mu_1 - \delta \geq \mu_2 \geq \dots \geq \mu_k$, where $\delta > 0$ is the IZ parameter. Let $(-a_{ij} + \lambda_{ij}t, a_{ij} - \lambda_{ij}t)$ ($0 < \lambda_{ij} < \delta$, $0 \leq t \leq a_{ij}/\lambda_{ij}$) be the triangular region used for a pairwise comparison between alternatives i and j . Define N_i as the sample size required to distinguish μ_1 and μ_i , for $i = 2, 3, \dots, k$, and define $1 - \alpha$ as the probability of correct selection. Then,*

$$\mathbb{E}[N_i] \stackrel{\alpha \rightarrow 0}{\sim} \frac{a_{1i}\sigma_{1i}^2}{\mu_1 - \mu_i + \lambda_{1i}}, \quad \text{for all } i = 2, 3, \dots, k,$$

where $\sigma_{1i}^2 = \text{Var}[X_{11} - X_{i1}]$.

In particular, the KN procedure often chooses $\lambda_{ij} = \delta/2$ and $a_{ij} = -(1/\delta) \log(2\alpha/(k-1))$. Based on Theorem 3, it is straightforward to derive a theoretical expression for the asymptotic average sample size of the KN procedure; i.e.,

$$\mathbb{E}_{KN}[N_i] \stackrel{\alpha \rightarrow 0}{\sim} \frac{f^2 \sigma_{1i}^2}{(\mu_1 - \mu_i + \delta/2)\delta}, \quad \text{for } i = 2, 3, \dots, k, \quad (9)$$

where $f^2 = -\log(2\alpha/(k-1))$.

However, a rigorous proof has not been given so far for the sequential procedures with a general boundary, such as the one used in Procedure 2. In the next theorem, we accomplish the task and show that this heuristic method is also valid for Procedure 2 under mild conditions. The proof is included in the appendix.

THEOREM 4. *Suppose that there are k alternatives and we are interested in selecting the alternative with the largest mean. Assume that $(X_{1r}, X_{2r}, \dots, X_{kr})$, $r = 1, 2, \dots$, are independent and identically jointly distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$, and their moment generating*

function exists in the neighborhood of $(0, 0, \dots, 0) \in \mathcal{R}^k$. Let $(-g_c(t), g_c(t))$ denote the continuation region in Procedure 2, where $g_c(t) = \sqrt{[c + \log(t+1)](t+1)}$. Define N_i as the sample size required to distinguish μ_1 and μ_i , for $i = 2, 3, \dots, k$, and define $1 - \alpha$ as the probability of correct selection. If $\mu_1 > \mu_2 > \dots > \mu_k$, then, in Procedure 2, we have

$$\mathbb{E}[N_i] \stackrel{\alpha \rightarrow 0}{\sim} \frac{c\sigma_{1i}^2}{(\mu_1 - \mu_i)^2}, \quad \text{for all } i = 2, 3, \dots, k, \quad (10)$$

where $\sigma_{1i}^2 = \text{Var}[X_{11} - X_{i1}]$.

REMARK 6. Although the theorem is proven when the boundary function is $g_c(t) = \sqrt{[c + \log(t+1)](t+1)}$, the proof can be easily extended to other boundaries.

PROOF. See EC.1.6 in the e-companion. \square


Theorem 4 provides a theoretical foundation to evaluate the asymptotic average sample size using the mean path approximation. For the comparison between alternatives 1 and i , it shows that the asymptotic average sample size is inversely proportional to $(\mu_1 - \mu_i)^2$ —i.e., the square of their true mean difference.

5.2. Comparisons Between Procedure 2 and IZ Procedures

In this subsection, we compare the asymptotic average sample sizes of Procedure 2 to the KN⁺⁺ procedure. There are two reasons why we choose the KN⁺⁺ procedure as the benchmark. First, the KN⁺⁺ procedure is the most similar procedure to Procedure 2. Both allow unknown and unequal variances and CRNs, and both break the comparisons into pairwise comparisons and approximate the partial sums of differences by BMs. Second, the KN⁺⁺ procedure, as the most efficient one in the family of the KN procedures, has been well studied and compared in the literature (Branke et al. 2007, Wang and Kim 2013) and implemented by commercial simulation software. Notice that it is not our intention to argue that the KN⁺⁺ procedure is the most efficient procedure in the literature. Indeed, it is not. When there are a large number of alternatives with equal unknown variance, the BIZ procedure of Frazier (2014) is more efficient because it avoids the use of Bonferroni inequality to break the comparisons into pairs, which causes inefficiency for the KN⁺⁺ procedure and our procedures as well. Therefore, when the number of alternatives is large, we also expect the BIZ procedure to outperform ours. However, the objective of this comparison is to show that the avoidance of setting IZ parameters may give an advantage to our procedures. Even though the comparison is done between the KN⁺⁺ procedure and Procedure 2, we believe that the insights obtained here hold robustly for all IZ procedures.

Since both the KN⁺⁺ procedure and Procedure 2 decompose the selection-of-the-best problem into pairwise comparisons, we compare their asymptotic average sample sizes for the case of two alternatives (i.e., $k = 2$). Let Δ denote the difference of their means and set the variance of their

Figure 3. (Color online) Ratios of asymptotic average sample sizes using the KN^{++} procedure and Procedure 2 (with the boundary $g_c(t) = \sqrt{[c + \log(t+1)](t+1)}$) when Δ/δ varies.



difference to 1. It is easy to calculate the asymptotic average sample size required by Procedure 2 based on Theorem 4. Meanwhile, given any specified IZ parameter δ , the asymptotic average sample size of the KN^{++} procedure can be also evaluated according to Theorem 3.

To compare the asymptotic average sample sizes, we vary Δ/δ , which indicates how the chosen IZ parameter relates to the true mean difference. In Figure 3, we plot the ratios of sample sizes of the KN^{++} procedure and Procedure 2 with boundary $\sqrt{[c + \log(t+1)](t+1)}$ for different values of Δ/δ . From Figure 3, we find that the KN^{++} Procedure needs about one-fifth of the samples of Procedure 2 when $\Delta/\delta = 1$. However, when Δ/δ is above 5, Procedure 2 begins to out-perform the KN^{++} procedure. Particularly, as Δ/δ becomes larger and larger than 5, Procedure 2 requires fewer and fewer samples than the KN^{++} procedure.

Even though Procedure 2 appears to require more samples when Δ/δ is smaller than 5, the disadvantage may vanish when the number of alternatives is large. In this situation, there are often many alternatives whose means are significantly smaller than the best. This may result in large values of Δ/δ for many pairwise comparisons and thus increase the total number of observations required by the KN^{++} procedure relative to Procedure 2.

To close this section, we summarize the asymptotic average samples sizes required by Procedure 2, the KN^{++} procedure as well as Rinott's procedure, and list them in Table 1. For simplicity, Table 1 considers a pairwise comparison where Δ and σ^2 are the mean and variance of their difference. We have two findings from the table, and they are consistent with the findings from Figure 3. First, both the KN^{++} and Rinott's procedures are sensitive to the choice of δ . In particular, the asymptotic sample sizes of the KN^{++} and Rinott's procedures are inversely proportional to δ and its square. Second, the asymptotic sample size of Procedure 2 is immune to the choice of δ because Procedure 2 is

Table 1. Asymptotic average sample sizes of the Rinott's procedure, the KN^{++} procedure and Procedure 2 to select from two random variables whose difference is of mean Δ and variance σ^2 .

Procedures	Rinott's procedure	KN^{++} procedure	Procedure 2
Asymptotic sample size	$\frac{h^2 \sigma^2}{\delta^2}$	$\frac{f^2 \sigma^2}{(\Delta - \delta/2)\delta}$	$\frac{c \sigma^2}{\Delta^2}$

Note. h^2 is the Rinott's constant (see Rinott 1978).

IZ-free. Therefore, Procedure 2 is relatively most efficient if δ is chosen too conservatively.

6. A Stopping Criterion

One drawback of Procedures 1 and 2 is that the continuation regions are open on the right-hand side, which may potentially cause the procedures to run a very long time before they stop or maybe never stop. As a remedy, part (1) of Lemma 2 shows that these procedures can stop in finite time w.p.1 in each pairwise comparison as long as the two alternatives have different means.

However, there might be two or more alternatives that have the same means. Then, there is a positive probability that our procedures may not terminate (seen from Propositions 1 and 2), resulting in an infinite sample size. To avoid this problem, we design truncated procedures, which force the original procedures to stop when a stopping criterion is satisfied, even if there is more than one alternative in contention.

Beyond avoiding an infinite sample size, users may have practical reasons for designing such stopping criteria. For instance, they may have a practically significant difference below which difference in mean performances is not meaningful. This could arise if the resolution of the simulation model itself is limited (e.g., nearest \$10,000), or if the precision of the software itself is finite (e.g., five decimal digits is the default in Matlab). We refer to this limit as the error tolerance ε . Therefore, the revised goal is to select an alternative that is within ε of the best, and to design a stopping criterion that terminates once this goal is achieved.

Now, we illustrate how to design a stopping criterion that corresponds to the error tolerance ε . Denote I_ε as the set of alternatives whose means are within ε of the best, and the goal is to select an alternative that is in I_ε . Intuitively, we should choose a termination time $T^*(\varepsilon)$ such that only alternatives in I_ε can survive at $T^*(\varepsilon)$. In particular, we design $T^*(\varepsilon)$ as follows:

Stopping Criterion. Determine a termination time $T^*(\varepsilon)$ that is a root to $T\varepsilon - g_c(T) = 0$.

Basically, the choice of $T^*(\varepsilon)$ above is driven by the mean path approximation (see Theorem 4). This method shows that the asymptotic average time required to eliminate all of the alternatives outside I_ε is bounded by the value which is a

root of $T\varepsilon - g_c(T) = 0$. In light of this, we determine $T^*(\varepsilon)$ as stated in *Stopping Criterion* above.

Provided the stopping criterion, we design a truncated procedure—i.e., Procedure 3—that is the same as Procedure 2 except adding a stopping rule as follows: if $|I| > 1$ and $\tau_{ij}(n) \geq T^*(\varepsilon)$ for all $i, j \in I$ with $i \neq j$, let $i^* = \arg \max_{i \in I} \bar{X}_i(n)$ and return i^* . In the following theorem, we show that Procedure 3 achieves the goal of selecting an alternative within ε of the best, with at least the user-specified PCS in the limit.

THEOREM 5. Assume that $(X_{1r}, X_{2r}, \dots, X_{kr})$, $r = 1, 2, \dots$, are independent and identically jointly distributed with unknown mean vector $(\mu_1, \mu_2, \dots, \mu_k)$ and their moment generating function exists in a neighborhood of $(0, 0, \dots, 0) \in \mathbb{R}^k$. Further, let the first-stage sample size n_0 be a function of α that satisfies $n_0 \rightarrow \infty$ as $\alpha \rightarrow 0$. Without loss of the generality, assume that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_k$. Let ε denote the tolerance error and $1 - \alpha$ denote the user-specified PCS. Let i^* be the alternative selected by Procedure 3. Then, we have, $\limsup_{\alpha \rightarrow 0} \mathbb{P}\{\mu_{i^*} < \mu_1 - \varepsilon\} / \alpha \leq 1$.

REMARK 7. We refer to “correct selection” as selecting the best alternative and refer to “good selection” as selecting any alternative within ε of the best. (The definitions are consistent with those in Ni et al. 2014.) Theorem 5 states that Procedure 3 provides a good selection guarantee when $\mu_1 \leq \mu_2 + \varepsilon$. Furthermore, when $\mu_1 > \mu_2 + \varepsilon$, Procedure 3 provides a correct selection guarantee because $\{\mu_{i^*} \geq \mu_1 - \varepsilon\}$ implies $i^* = 1$. Therefore, Procedure 3 can provide both correct selection and good selection guarantees.

PROOF. See EC.1.7 in the e-companion. \square

If one treats ε as an IZ parameter, Procedure 3 turns into an IZ procedure because the selected alternative is within ε of the best (shown by Theorem 5). Nevertheless, we want to emphasize that the concept of an error tolerance is different from an IZ parameter. When the difference between two alternatives is larger than ε , the average sample size required by the existing IZ procedures depends critically on the choice of ε . In other words, a conservative ε may lead to excessive samples. In contrast, a conservative ε will in general not affect the required sample size as our procedures are expected to stop long before they reach the maximal sample size (due to Theorem 4). Another way to see this is that the continuation regions of our procedures are not a function of ε . Therefore, an error tolerance can be set very conservatively.

7. Numerical Experiments

In this section, we test the performance of our procedures through extensive numerical experiments and compare them with existing procedures. Firstly, we establish the small-sample performance of Procedures 1 and 2. Secondly, we consider the selection-of-the-best problem with various numbers of alternatives. In this experiment, we compare the

efficiency of Procedure 2 with the KN^{++} procedure under three different configurations of variances. Thirdly, we test Procedure 3 that incorporates an error tolerance and compare it with IZ procedures under different configurations of means. Lastly, we test our procedures by addressing a realistic problem.

For simplicity, CRNs are not used in the experiments because using CRNs to increase their efficiency is not our focus.

7.1. Small-Sample Performance

To get rid of the effect of using the Bonferroni inequality, we consider the case of only two alternatives. Assume that their observations are normally distributed with mean (μ_1, μ_2) and variance $(5, 5)$. Further, assume alternative 1 is the better alternative which we would like to select. In this experiment, we implement Procedures 1 and 2 to select alternative 1 with a desired PCS 0.95. For different combinations of (μ_1, μ_2) and the first-stage sample size n_0 , we report the approximated values of c , the estimated PCS, and the average sample size with 95% confidence interval across 10,000 microreplications in Table 2.

From Table 2, we obtain several findings. First, both procedures deliver a larger PCS than the desired PCS, and their delivered PCS decreases slightly as the true mean difference Δ decreases from 0.20 to 0.05. This arises because both procedures are designed for the case when the mean difference is zero. Second, the average sample size of Procedure 1 depends on the setting of n_0 . In particular, for a given desired PCS, a smaller n_0 leads to a larger sample size. This occurs because Procedure 1 explicitly accounts for the $n_0 - 1$ degrees of freedom of the variance estimator through using different values of c . Third, Procedure 1 often needs more samples than Procedure 2 does, but their difference tends to be smaller as n_0 increases. Specially, as n_0 increases to 40, Procedure 1 needs fewer observations than Procedure 2 does. To understand this, as n_0 increases, the randomness coming from variance estimators fades away, and the corresponding value of c for Procedure 1 tends to be smaller and even less conservative (or smaller) than that of Procedure 2, which is obtained by asymptotic approximation.

Table 2 shows that both procedures can be used to select the best in this experiment. Therefore, users may face the choice between Procedures 1 and 2, and our suggestions are as follows. When observations are normally distributed, Procedure 1 would be preferred unless it is significantly less efficient than Procedure 2, because Procedure 1 guarantees PCS in small samples. Otherwise, Procedure 2 would be preferred because it guarantees PCS although in an asymptotic regime while Procedure 1 can not. Besides, Procedure 2 is more easily used because of its closed-form expressions for the design parameter c and achieves a higher efficiency because it updates variance estimators sequentially.

Table 2. The estimated PCS (Est. PCS) and the average sample sizes (Avg. SSize) with 95% confidence intervals for Procedures 1 and 2 when observations are normally distributed.

$\mu_1 - \mu_2$	n_0	Procedure 1			Procedure 2		
		c	Est. PCS	Avg. SSize	c	Est. PCS	Avg. SSize
0.05	10	7.358	0.97	$(1.149 \pm 0.019) \times 10^5$	4.605	0.96	$(8.126 \pm 0.125) \times 10^4$
	15	5.870	0.97	$(9.905 \pm 0.160) \times 10^4$		0.97	$(8.298 \pm 0.127) \times 10^4$
	20	5.163	0.97	$(9.181 \pm 0.148) \times 10^4$		0.97	$(8.393 \pm 0.127) \times 10^4$
	25	5.040	0.97	$(9.124 \pm 0.142) \times 10^4$		0.97	$(8.518 \pm 0.126) \times 10^4$
	30	4.671	0.97	$(8.710 \pm 0.138) \times 10^4$		0.98	$(8.536 \pm 0.126) \times 10^4$
	35	4.514	0.97	$(8.509 \pm 0.136) \times 10^4$		0.98	$(8.500 \pm 0.125) \times 10^4$
	40	4.415	0.98	$(8.411 \pm 0.133) \times 10^4$		0.98	$(8.492 \pm 0.125) \times 10^4$
0.20	10	7.358	0.99	$(5.849 \pm 0.097) \times 10^3$	4.605	0.98	$(3.952 \pm 0.067) \times 10^3$
	15	5.870	0.99	$(4.971 \pm 0.083) \times 10^3$		0.99	$(4.020 \pm 0.065) \times 10^3$
	20	5.163	0.99	$(4.526 \pm 0.078) \times 10^3$		0.99	$(4.070 \pm 0.067) \times 10^3$
	25	5.040	0.99	$(4.427 \pm 0.074) \times 10^3$		0.99	$(4.057 \pm 0.066) \times 10^3$
	30	4.671	0.99	$(4.173 \pm 0.070) \times 10^3$		0.99	$(4.120 \pm 0.066) \times 10^3$
	35	4.514	0.99	$(4.181 \pm 0.070) \times 10^3$		0.99	$(4.108 \pm 0.067) \times 10^3$
	40	4.415	0.99	$(4.077 \pm 0.068) \times 10^3$		0.99	$(4.124 \pm 0.060) \times 10^3$

7.2. Large-Scale Selection of the Best

In practice, the means of alternatives are often spread out when the number of alternatives is large. In light of this, we consider a monotone decreasing configuration of means $(\mu_1, \mu_2, \dots, \mu_k)$ where $\mu_i = 1.5 - 0.5i$, and an equal-variance configuration where $\sigma_i^2 = 10$ for all i . Our target is to select alternative 1 with a desired PCS 0.95. To show the performance of our procedure to solve large-scale selections of the best, we vary the number of alternatives as $k = 20, 50, 100, 500$.

To solve this problem, we use Procedure 2 for simplicity of the simulation study because its design parameter c is more easily determined. As a comparison, we implement the KN procedure by choosing an IZ parameter δ . It is worthwhile emphasizing that δ is chosen to enable the implementation of the KN procedure, and we do not refer it to the smallest difference users care to detect in this experiment. Further, to make a fair comparison, we compare Procedure 2 with the KN^{++} procedure because both of them update variance estimators sequentially. In this experiment, we set n_0 to 10.

To compare Procedure 2 with the KN^{++} procedure, we consider different settings of δ where $(\mu_1 - \mu_2)/\delta = 1/4, 1/2, 1, 2, 4, 8$. In Table 3, we report the estimated PCS and the average sample sizes with 95% confidence intervals based on 1,000 macroreplications of each procedure. From the table, we obtain the following three findings. Firstly, when users select δ that is larger than $\mu_1 - \mu_2$ in the KN^{++} procedure, they may suffer from a lower PCS than they require, seen from the third and fourth columns. Secondly, for any fixed number of alternatives (see each row in Table 3), Procedure 2 demands fewer observations than the KN^{++} procedure, when δ is much smaller than the true difference $\mu_1 - \mu_2$. Furthermore, the gap of their sample sizes becomes larger as δ becomes smaller. Thirdly, from the second and fifth columns, we can conclude that, although Procedure 2

needs more observations than the KN^{++} procedure when δ is exactly $\mu_1 - \mu_2$, the relative difference between their average sample sizes tends to be narrower as the number of alternatives k increases. This implies the average sample size required by Procedure 2 grows more slowly with k than the KN^{++} procedure.

Furthermore, Procedure 2 is shown to be asymptotically valid when the observations are generally distributed. To show its small-sample performance for this general case, we repeat the experiment above using the same parameters except that all of the observations are exponentially distributed. The results are reported in Table 4 and they are similar to that of the normally distributed case.

Last, we further investigate the efficiency of Procedure 2 and the KN^{++} Procedure in equal, increasing, and decreasing configurations of variances where $\sigma_i^2 = 10, 10 \times (0.95 + 0.05i)$, and $10/(0.95 + 0.05i)$ for $i = 1, 2, \dots, k$, respectively. To avoid reporting the similar results, in the following Table 5, we only report the results when observations are normally distributed. From this table, we find that the same conclusions seen from Tables 3 and 4 are valid.

7.3. Comparisons with IZ Procedures

In this subsection, we consider a selection-of-the-best problem with an error tolerance ε and solve this problem using Procedure 3. If one treats ε as an IZ parameter, IZ procedures can also be used to solve this problem. In this experiment, we choose the (two-stage) Rinott's procedure and the (fully sequential) KN^{++} procedures as representatives of IZ procedures, and compare Procedure 3 with them in terms of the efficiency.

We consider two different configurations of means. The first one is the slippage configuration (SC) where $\mu_1 = 0.5, \mu_2 = \mu_3 = \dots = \mu_{50} = 0$. The second one is the monotone decreasing configuration of means (MDM) where $\mu_i = 1 - 0.5i$ for $i = 1, 2, \dots, 50$. For both configurations

Table 3. The estimated PCS and average sample sizes with 95% confidence interval for Procedure 2 and the KN⁺⁺ procedure under equal configuration of variances when all observations are normally distributed.

<i>k</i>	Procedure 2	KN ⁺⁺ procedure: $(\mu_1 - \mu_2)/\delta$					
		1/4	1/2	1	2	4	8
20	0.99	0.70	0.91	1.00	1.00	1.00	1.00
	2.816×10^3	2.639×10^2	5.503×10^2	1.371×10^3	3.247×10^3	7.045×10^3	1.470×10^4
	$\pm 0.082 \times 10^3$	$\pm 0.020 \times 10^2$	$\pm 0.064 \times 10^2$	$\pm 0.016 \times 10^3$	$\pm 0.029 \times 10^3$	$\pm 0.050 \times 10^3$	$\pm 0.007 \times 10^4$
50	1.00	0.72	0.94	0.99	1.00	1.00	1.00
	3.588×10^3	5.904×10^2	9.619×10^2	2.014×10^3	4.454×10^3	9.779×10^3	2.070×10^4
	$\pm 0.089 \times 10^3$	$\pm 0.023 \times 10^2$	$\pm 0.075 \times 10^2$	$\pm 0.018 \times 10^3$	$\pm 0.033 \times 10^3$	$\pm 0.054 \times 10^3$	$\pm 0.008 \times 10^4$
100	1.00	0.76	0.95	0.99	1.00	1.00	1.00
	4.388×10^3	1.112×10^3	1.540×10^3	2.710×10^3	5.506×10^3	1.179×10^4	2.514×10^4
	$\pm 0.089 \times 10^3$	$\pm 0.003 \times 10^3$	$\pm 0.008 \times 10^3$	$\pm 0.020 \times 10^3$	$\pm 0.034 \times 10^3$	$\pm 0.006 \times 10^4$	$\pm 0.009 \times 10^4$
500	0.99	0.79	0.97	1.00	1.00	1.00	1.00
	9.138×10^3	5.164×10^3	5.721×10^3	7.217×10^3	1.077×10^4	1.877×10^4	3.628×10^4
	$\pm 0.102 \times 10^3$	$\pm 0.003 \times 10^3$	$\pm 0.009 \times 10^3$	$\pm 0.021 \times 10^3$	$\pm 0.004 \times 10^4$	$\pm 0.006 \times 10^4$	$\pm 0.009 \times 10^4$

Table 4. The estimated PCS and the average sample sizes with 95% confidence intervals of Procedure 2 and the KN⁺⁺ procedure when all observations are exponentially distributed.

<i>k</i>	Procedure 2	KN ⁺⁺ procedure: $(\mu_1 - \mu_2)/\delta$					
		1/4	1/2	1	2	4	8
20	0.99	0.71	0.90	0.99	1.00	1.00	1.00
	2.902×10^3	2.740×10^2	5.470×10^2	1.344×10^3	3.143×10^3	6.927×10^3	1.463×10^4
	$\pm 0.084 \times 10^3$	$\pm 0.022 \times 10^2$	$\pm 0.067 \times 10^2$	$\pm 0.015 \times 10^3$	$\pm 0.028 \times 10^3$	$\pm 0.046 \times 10^3$	$\pm 0.007 \times 10^4$
50	0.99	0.74	0.95	1.00	1.00	1.00	1.00
	3.610×10^3	5.972×10^2	9.430×10^2	1.932×10^3	4.215×10^3	9.316×10^3	2.020×10^4
	$\pm 0.088 \times 10^3$	$\pm 0.026 \times 10^2$	$\pm 0.074 \times 10^2$	$\pm 0.017 \times 10^3$	$\pm 0.030 \times 10^3$	$\pm 0.052 \times 10^3$	$\pm 0.008 \times 10^4$
100	1.00	0.74	0.95	1.00	1.00	1.00	1.00
	4.396×10^3	1.119×10^3	1.525×10^3	2.627×10^3	5.306×10^3	1.124×10^4	2.414×10^4
	$\pm 0.092 \times 10^3$	$\pm 0.003 \times 10^3$	$\pm 0.008 \times 10^3$	$\pm 0.018 \times 10^3$	$\pm 0.031 \times 10^3$	$\pm 0.006 \times 10^4$	$\pm 0.009 \times 10^4$
500	1.00	0.79	1.00	1.00	1.00	1.00	1.00
	9.095×10^3	5.167×10^3	5.699×10^3	7.114×10^3	1.051×10^4	1.806×10^4	3.483×10^4
	$\pm 0.104 \times 10^3$	$\pm 0.003 \times 10^3$	$\pm 0.010 \times 10^3$	$\pm 0.018 \times 10^3$	$\pm 0.004 \times 10^4$	$\pm 0.006 \times 10^4$	$\pm 0.009 \times 10^4$

of means, we use the equal-variance configuration where $\sigma_i^2 = 10$, for all $i = 1, 2, \dots, 50$. For each configuration, we select the best using three different procedures and report their average sample sizes based on 1,000 macroreplications of each procedure in Table 6. We set PCS to 0.95 and n_0 to 10.

From Table 6, we obtain the following insights. Firstly, the last column shows that the sample size required by Procedure 3 appears stable as the ratio $(\mu_1 - \mu_2)/\varepsilon$ varies under MDM. It follows from that, under MDM, Procedure 3 tends to terminate early before reaching the stopping criterion. Secondly, Procedure 3 has a wider confidence interval under SC than the other two procedures. This occurs because Procedure 3 either stops long before reaching the stopping criterion or stops when the stopping criterion is met. This implies that Procedure 3 effectively avoids infinite samples. Thirdly, from the first two columns, we find that the sample size required by the Rinott's procedure increases nearly 4 times as $(\mu_1 - \mu_2)/\varepsilon$ increases from 1 to 2, and increases nearly 25 times as $(\mu_1 - \mu_2)/\varepsilon$ increases from 1 to 5. Similar results hold for the KN⁺⁺ procedure, although it increases

more slowly. Therefore, the Rinott's procedure is the most sensitive to the choice of ε , then the KN⁺⁺ procedure, and Procedure 3 is the least sensitive to the choice of ε . The results in Table 6 are consistent with the theoretical results listed in Table 1. Lastly, similar results appear for both the normally and exponentially distributed cases.

7.4. Ambulance Allocation in Emergency Medical Services System

An Emergency Medical Services (EMS) system is an important component of public safety that attempts to allocate scarce resources as critical events occur. In the EMS system, one critical resource-allocation challenge is deploying several ambulances to best serve requested emergency calls—for instance, keeping response time small. The response time for a call is the elapsed time from when the call is received to when an ambulance arrives at the scene. Here, we measure the performance as the fraction of requested calls received that have response times of eight minutes or less. In this experiment, we consider an example similar to that in Ni et al. (2012).

Table 5. The estimated PCS and average sample sizes with 95% confidence interval for Procedure 2 and the KN⁺⁺ procedure under different configurations of variances when all observations are normally distributed.

k	Variances	Procedure 2	KN ⁺⁺ procedure: $(\mu_1 - \mu_2)/\delta$			
			1	2	4	8
20	Decreasing	0.99	0.99	1.00	1.00	1.00
		2.780×10^3 $\pm 0.081 \times 10^3$	1.326×10^3 $\pm 0.016 \times 10^3$	3.092×10^3 $\pm 0.028 \times 10^3$	6.720×10^3 $\pm 0.049 \times 10^3$	1.411×10^4 $\pm 0.007 \times 10^4$
	Equal	0.99	1.00	1.00	1.00	1.00
		2.816×10^3 $\pm 0.082 \times 10^3$	1.371×10^3 $\pm 0.016 \times 10^3$	3.247×10^3 $\pm 0.029 \times 10^3$	7.045×10^3 $\pm 0.050 \times 10^3$	1.470×10^4 $\pm 0.007 \times 10^4$
	Increasing	0.99	0.99	1.00	1.00	1.00
		2.916×10^3 $\pm 0.085 \times 10^3$	1.473×10^3 $\pm 0.017 \times 10^3$	3.423×10^3 $\pm 0.031 \times 10^3$	7.437×10^3 $\pm 0.049 \times 10^3$	1.536×10^4 $\pm 0.008 \times 10^4$
50	Decreasing	0.99	1.00	1.00	1.00	1.00
		3.551×10^3 $\pm 0.086 \times 10^3$	1.881×10^3 $\pm 0.017 \times 10^3$	4.120×10^3 $\pm 0.033 \times 10^3$	8.970×10^3 $\pm 0.052 \times 10^3$	1.910×10^4 $\pm 0.008 \times 10^4$
	Equal	1.00	0.99	1.00	1.00	1.00
		3.588×10^3 $\pm 0.089 \times 10^3$	2.014×10^3 $\pm 0.018 \times 10^3$	4.454×10^3 $\pm 0.033 \times 10^3$	9.779×10^3 $\pm 0.054 \times 10^3$	2.070×10^4 $\pm 0.008 \times 10^4$
	Increasing	0.99	0.99	1.00	1.00	1.00
		3.635×10^3 $\pm 0.090 \times 10^3$	2.147×10^3 $\pm 0.018 \times 10^3$	4.958×10^3 $\pm 0.034 \times 10^3$	1.085×10^4 $\pm 0.005 \times 10^4$	2.268×10^4 $\pm 0.002 \times 10^4$
100	Decreasing	1.00	1.00	1.00	1.00	1.00
		4.351×10^3 $\pm 0.093 \times 10^3$	2.592×10^3 $\pm 0.019 \times 10^3$	5.061×10^3 $\pm 0.033 \times 10^3$	1.060×10^4 $\pm 0.006 \times 10^4$	2.238×10^4 $\pm 0.009 \times 10^4$
	Equal	1.00	0.99	1.00	1.00	1.00
		4.388×10^3 $\pm 0.089 \times 10^3$	2.710×10^3 $\pm 0.020 \times 10^3$	5.506×10^3 $\pm 0.034 \times 10^3$	1.179×10^4 $\pm 0.006 \times 10^4$	2.514×10^4 $\pm 0.009 \times 10^4$
	Increasing	0.99	0.99	1.00	1.00	1.00
		4.444×10^3 $\pm 0.093 \times 10^3$	2.934×10^3 $\pm 0.020 \times 10^3$	6.350×10^3 $\pm 0.036 \times 10^3$	1.412×10^4 $\pm 0.006 \times 10^4$	3.000×10^4 $\pm 0.009 \times 10^4$

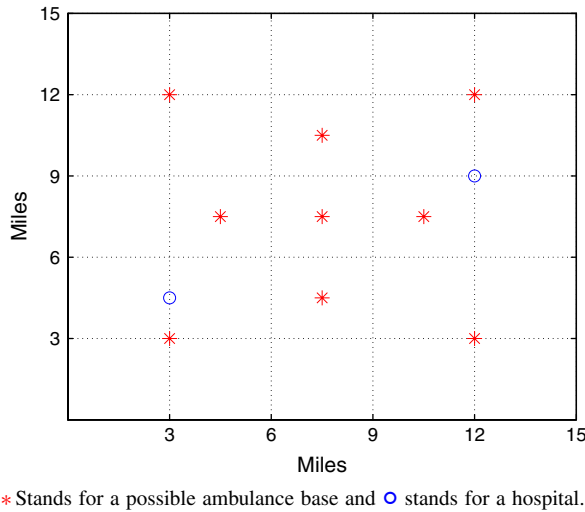
Consider a city of size 15 miles by 15 miles. Suppose there are four ambulances with traveling speed 24 miles per hour, nine ambulance bases, and two hospitals in the city. Distances from point to point are measured by the Manhattan metric. See Figure 4 for a map of this city. When a requested

call arrives, the EMS system will dispatch the call to the nearest free ambulance. After the dispatched ambulance arrives at the scene, the time spent there is exponentially distributed with mean 12 minutes before carrying the patient to the nearest hospital. Once the patient reaches the hospital, the

Table 6. Average sample sizes with 95% confidence intervals for the Rinott’s procedure, the KN⁺⁺ procedure, and Procedure 3.

$(\mu_1 - \mu_2)/\varepsilon$	SC			MDM		
	Rinott	KN ⁺⁺	Procedure 3	Rinott	KN ⁺⁺	Procedure 3
Normally distributed observations						
1	6.071×10^4 $\pm 0.035 \times 10^4$	1.576×10^4 $\pm 0.061 \times 10^4$	4.781×10^4 $\pm 0.025 \times 10^4$	6.045×10^4 $\pm 0.025 \times 10^4$	2.002×10^3 $\pm 0.018 \times 10^3$	3.119×10^3 $\pm 0.053 \times 10^3$
2	2.421×10^5 $\pm 0.010 \times 10^5$	4.022×10^4 $\pm 0.042 \times 10^4$	6.430×10^4 $\pm 0.217 \times 10^4$	2.426×10^5 $\pm 0.010 \times 10^5$	4.425×10^3 $\pm 0.033 \times 10^3$	3.633×10^3 $\pm 0.091 \times 10^3$
5	1.515×10^6 $\pm 0.006 \times 10^6$	1.129×10^5 $\pm 0.008 \times 10^5$	8.857×10^4 $\pm 1.354 \times 10^4$	1.518×10^6 $\pm 0.006 \times 10^6$	1.254×10^4 $\pm 0.006 \times 10^4$	3.615×10^3 $\pm 0.086 \times 10^3$
Exponentially distributed observations						
1	6.071×10^4 $\pm 0.046 \times 10^4$	1.602×10^4 $\pm 0.015 \times 10^4$	4.879×10^4 $\pm 0.058 \times 10^4$	6.046×10^4 $\pm 0.048 \times 10^4$	1.933×10^3 $\pm 0.017 \times 10^3$	3.117×10^3 $\pm 0.050 \times 10^3$
2	2.440×10^5 $\pm 0.019 \times 10^5$	4.020×10^4 $\pm 0.035 \times 10^4$	6.351×10^4 $\pm 0.148 \times 10^4$	2.426×10^5 $\pm 0.019 \times 10^5$	4.223×10^3 $\pm 0.030 \times 10^3$	3.583×10^3 $\pm 0.087 \times 10^3$
5	1.518×10^6 $\pm 0.012 \times 10^6$	1.131×10^5 $\pm 0.007 \times 10^5$	7.685×10^4 $\pm 0.985 \times 10^4$	1.516×10^6 $\pm 0.012 \times 10^6$	1.194×10^4 $\pm 0.006 \times 10^4$	3.646×10^3 $\pm 0.089 \times 10^3$

Figure 4. (Color online) Map for the city.



ambulance travels back to its home base to wait for the next call. We assume that the call location is uniformly distributed in the city. Our target is allocating these four ambulances to possible bases to maximize the performance—i.e., the portion of calls with response time no more than eight minutes. In this experiment, we assume that more than one ambulance can be allocated to the same base.

To select the optimal policy in this example, we are faced with at least three challenges. First, there are 495 alternative policies in this problem, which is not a small number for the classic selection procedures. Second, the observations from each alternative policy are generally distributed and with unknown variances. Last and most important, the ideal case is to set the IZ parameter to zero and select the best policy in this problem because when human life is concerned, no difference is so small that we are indifferent. In other words, when the IZ parameter is set as 0.01, it implies that $0.01 \times 100\%$ of the citizens are in a dangerous situation because it may be fatal if their calls are not handled in time. To circumvent these challenges, we use Procedure 2 that is IZ-free. In our experiment, we set the PCS to 0.99 and replicate the selection process 1,000 times. The result shows that the optimal policy is selected as the best with probability 1.000

and the average total sample size required is 2.10×10^4 , seen from Table 7.

One may argue that the IZ procedures (e.g., the KN^{++} procedure) can also be utilized to select the optimal policy by choosing a conservative IZ parameter. To compare the performances of Procedure 3 and the KN^{++} procedure, we run the following experiments when ε is set as 0.01 and 0.001, respectively; the results are listed in Table 7. Via Monte Carlo simulation, we evaluated the performance of Policies 1, 2, 3, and 4 as 0.8189, 0.8162, 0.8151, and 0.8144 because they are all selected as the best in our experiments, and Policy 1 is the optimal policy.

From Table 7, we obtain the following conclusions. First, both procedures select a policy within δ of the best with a probability greater than 0.99. Second, when δ is 0.01, Procedure 3 achieves a larger PCS, although it costs about 1.5 times more samples than the KN^{++} procedure. As δ decreases to 0.001, the sample size required by the KN^{++} procedure increases faster than that of Procedure 3, implying that the KN^{++} procedure is more sensitive to the setting of δ .

8. Conclusion

In this paper, we propose a new frequentist selection-of-the-best formulation which selects the best alternative with a user-specified probability of correct selection whenever the alternatives have unique means. Under this formulation, we design a class of sequential procedures based on the Law of the Iterated Logarithm. Among these procedures, the design parameters are obtained by solving numerically a one-dimensional root-finding problem or using asymptotic approximating results. We call them Procedures 1 and 2, and they are shown to be statistically valid in a finite regime and an asymptotic regime, respectively. The merit of Procedures 1 and 2 is that they free users from having to specify an IZ parameter. In addition, we add a stopping criterion to turn them into an IZ procedure (Procedure 3). The numerical results show that our procedures are less sensitive to the setting of IZ parameter and more efficient than the KN procedure when there are a large number of alternatives with many significantly inferior ones.

Table 7. Comparisons of our procedures and the KN^{++} procedure under different settings of the IZ parameter.

δ	Good selection	Procedure 3		KN^{++} procedure	
		Selection prob.	Avg. SSize	Selection prob.	Avg. SSize
0.01	Policy 1	1.00		0.93	
	Policy 2	0.00	1.859×10^4	0.04	7.943×10^3
	Policy 3	0.00	$\pm 0.013 \times 10^4$	0.02	$\pm 0.027 \times 10^3$
	Policy 4	0.00		0.01	
0.001	Policy 1	1.00	2.102×10^4	1.00	1.107×10^5
			$\pm 0.022 \times 10^4$		$\pm 0.003 \times 10^5$
0	Policy 1	1.00	2.097×10^4	NA	NA
			$\pm 0.021 \times 10^4$		

Supplemental Material

Supplemental material to this paper is available at <https://doi.org/10.1287/opre.2016.1530>.

Acknowledgments

The authors would like to thank Professor Shane G. Henderson of Cornell University for his helpful discussions. The authors would also like to thank the associate editor and two anonymous referees for their insightful and invaluable comments that have significantly improved this paper. A preliminary version of this paper (Fan and Hong 2014) was published in the Proceedings of the 2014 Winter Simulation Conference. This research was partially supported by the Hong Kong Research Grants Council [GRF 613012 and GRF 16203214].

References

- Basu D (1955) On statistics independent of a complete sufficient statistic. *Sankhyā: The Indian J. Statist.* 35:377–380.
- Batur D, Kim SH (2006) Fully sequential selection procedures with a parabolic boundary. *IIE Trans.* 38(9):749–764.
- Bechhofer RE (1954) A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Ann. Math. Statist.* 16–39.
- Bechhofer RE, Santner TJ, Goldsman DM (1995) *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons* (John Wiley & Sons, New York).
- Branke J, Chick SE, Schmidt C (2007) Selecting a selection procedure. *Management Sci.* 53(12):1916–1932.
- Bubeck S, Cesa-Bianchi N (2012) Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Foundations Trends Machine Learn.* 5(1):1–122.
- Chick SE, Inoue K (2001) New two-stage and sequential procedures for selecting the best simulated system. *Oper. Res.* 49(5):732–743.
- Chick SE, Frazier PI (2012) Sequential sampling with economics of selection procedures. *Management Sci.* 58(3):550–569.
- Clark GM, Yang W (1986) A Bonferroni selection procedure when using common random numbers with unknown variances. *Proc. 1986 Winter Simulation Conf. (WSC)*, 313–315.
- Dudewicz EJ (1969) An approximation to the sample size in selection problems. *Ann. Math. Statist.* 40(2):492–497.
- Durbin J (1985) The first-passage density of a continuous gaussian process to a general boundary. *J. Appl. Probab.* 99–122.
- Durrett R (2010) *Probability: Theory and Examples* (Cambridge University Press, New York).
- Fan W, Hong LJ (2014) A frequentist selection-of-the-best procedure without indifference zone. Buckley SJ, Miller JA, eds. *Proc. 2014 Winter Simulation Conf. (WSC)* (IEEE, Piscataway, NJ), 3737–3748.
- 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. Ph.D. thesis, University of North Carolina, Chapel Hill.
- Gupta SS (1965) On some multiple decision (selection and ranking) rules. *Technometrics* 7(2):225–245.
- Hong LJ (2006) Fully sequential indifference-zone selection procedures with variance-dependent sampling. *Naval Res. Logist.* 53(5):464–476.
- Hong LJ, Nelson BL (2005) The tradeoff between sampling and switching: New sequential procedures for indifference-zone selection. *IIE Trans.* 37(7):623–634.
- Hong LJ, Nelson BL (2007) Selecting the best system when systems are revealed sequentially. *IIE Trans.* 39(7):723–734.
- Jamieson K, Nowak R (2014) Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. *48th Annual Conf. Inform. Sci. Systems (CISS)* (IEEE, Piscataway, NJ), 1–6.
- Jennen C, Lerche HR (1981) First exit densities of brownian motion through one-sided moving boundaries. *Zeitschrift Wahrscheinlichkeitstheorie Verwandte Gebiete* 55(2):133–148.
- Jennison C, Johnston IM, Turnbull BW (1980) Asymptotically optimal procedures for sequential adaptive selection of the best of several normal means. Technical Report, Department of ORIE, Cornell University, Ithaca, New York.
- Karnin Z, Tomer K, Oren S (2013) Almost optimal exploration in multi-armed bandits. *Proc. 30th Internat. Conf. Machine Learn. (ICML)* (JMLR.org), 1238–1246.
- 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) Selecting the best system. *Handbooks in Operations Research and Management Science: Simulation*, Vol. 13 (Elsevier, Amsterdam), 501–534.
- Kim SH, Nelson BL (2006b) On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Oper. Res.* 54(3):475–488.
- Ni EC, Henderson SG, Hunter SR (2014) A comparison of two parallel ranking and selection procedures. Buckley SJ, Miller JA, eds. *Proc. 2014 Winter Simulation Conf. (WSC)* (IEEE, Piscataway, NJ), 3761–3772.
- Ni EC, Hunter SR, Henderson SG, Topaloglu H (2012) Exploring bounds on ambulance deployment policy performance. Rose O, Uhrmacher AM, eds. *Proc. 2012 Winter Simulation Conf. (WSC)* (IEEE, Piscataway, NJ), 1–12.
- Panchapakesan S, Gupta SS, Yackel J (1971) On a subset selection procedure for the most probable event in a multinomial distribution. *Statist. Decision Theory Related Topics* 275–298.
- Paulson E (1964) A sequential procedure for selecting the population with the largest mean from k normal populations. *Ann. Math. Statist.* 35(1):174–180.
- Perng SK (1969) A comparison of the asymptotic expected sample sizes of two sequential procedures for ranking problem. *Ann. Math. Statist.* 40(6):2198–2202.
- Pichitlamken J, Nelson BL, Hong LJ (2006) A sequential procedure for neighborhood selection-of-the-best in optimization via simulation. *Eur. J. Oper. Res.* 173(1):283–298.
- Rinott Y (1978) On two-stage selection procedures and related probability-inequalities. *Comm. Statist. Theory Methods* 7(8):799–811.
- Sullivan DW, Wilson JR (1989) Restricted subset selection procedures for simulation. *Oper. Res.* 37(1):52–71.
- Wang H, Kim SH (2013) Reducing the conservativeness of fully sequential indifference-zone procedures. *IEEE Trans. Automatic Control* 58(6):1613–1619.
- Whitt W (2002a) *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues* (Springer, New York).

Weimei Fan is an assistant professor in the School of Management at University of Science and Technology of China. Her research interests include stochastic simulation, input uncertainty, and robust simulation.

L. Jeff Hong is the Chair Professor of Management Sciences in the College of Business at City University of Hong Kong. His research interests include stochastic simulation, stochastic optimization, business analytics, and financial risk management.

Barry L. Nelson is the Walter P. Murphy Professor of Industrial Engineering and Management Sciences at Northwestern University. His research addresses statistical issues in the design and analysis of stochastic computer simulation experiments, including metamodeling, multivariate input modeling, simulation optimization, input uncertainty quantification, and variance reduction. He is a Fellow of INFORMS and the Institute of Industrial Engineers (IIE).