

RATE-OPTIMAL BUDGET ALLOCATION FOR THE PROBABILITY OF GOOD SELECTION

Taeho Kim¹ and David J. Eckman²

¹Dept. of ISOM, HKUST Business School, Clear Water Bay, Kowloon, HONG KONG

²Dept. of Industrial and Systems Engineering, Texas A&M University, College Station, TX, USA

ABSTRACT

This paper studies the allocation of simulation effort in a ranking-and-selection (R&S) problem with the goal of selecting a system whose performance is within a given tolerance of the best. We apply large-deviations theory to derive an optimal allocation for maximizing the rate at which the so-called probability of good selection (PGS) asymptotically approaches one, assuming that systems' output distributions are known. An interesting property of the optimal allocation is that some good systems may receive a sampling ratio of zero. We demonstrate through numerical experiments that this property leads to serious practical consequences, specifically when designing adaptive R&S algorithms. In particular, we observe that the convergence and even consistency of a simple plug-in algorithm designed for the PGS goal can be negatively impacted. We offer empirical evidence of these challenges and a preliminary exploration of a potential correction.

1 INTRODUCTION

This paper investigates a ranking-and-selection (R&S) problem wherein the goal is to select a system having near-optimal expected performance and each system's performance is estimated via stochastic simulation. In contrast, the predominant goal in the R&S literature is to select the system having the best expected performance. Accordingly, many R&S procedures aim to maximize the probability of identifying the best system, called the probability of correct selection (PCS), given a fixed simulation budget. A consequence of this objective is that PCS-inspired procedures tend to devote too much of the simulation budget trying to distinguish the best system from near-optimal systems when, in practice, a decision maker may be insensitive to slight differences in expected performance, so long as a near-optimal system is ultimately selected. In such cases, a more appropriate goal is to maximize the probability of good selection (PGS), defined as the probability of choosing any system deemed "good." A natural definition of a good system is one whose performance gap relative to the best is less than some user-specified tolerance. Under this relaxed goal, selecting any good system is regarded as a success.

A well-studied problem of both practical and theoretical importance is how best to allocate simulation effort, i.e., replications, across systems. In this paper, we study the sampling allocation that maximizes the rate at which the PGS goes to 1 and identify challenges that arise in designing adaptive procedures that aim to asymptotically achieve the optimal convergence rate. We first exploit large-deviations theory to characterize the exponential convergence rate of the PGS and pose an optimization problem for the sampling allocation that maximizes this convergence rate. Unlike in the PCS case, the optimal allocation for PGS can dictate that we should asymptotically allocate a vanishing fraction of the simulation budget to some good systems. Because of this zero-sampling-ratio property, even small deviations in systems' true performances can dramatically change the optimal allocation. To better understand this phenomenon, we investigate sufficient conditions under which it is optimal to give a sampling ratio of zero to some good systems. We also characterize necessary conditions for the three-system case. We further investigate how the zero-sampling-ratio behavior can inhibit the ability of adaptive algorithms to achieve the optimal sampling allocation in the limit. We shed light on the root causes for this performance degradation through an examination of two bad configurations of sample means and variances.

Past R&S procedures have commonly studied the problem under the indifference-zone (IZ) formulation, which assumes that the difference between the expected performance of the best and second-best systems is greater than some constant specified by the user (Bechhofer 1954). The IZ formulation has been adopted in many R&S procedures to deliver a fixed-confidence guarantee on the PCS (Kim and Nelson 2001; Hong 2006; Kim and Nelson 2006); for a detailed literature review, see Hong et al. (2021) and references therein. Some attention has recently been given to the design of procedures that deliver a fixed-confidence guarantee on the PGS, e.g., the Envelope procedure of Ma and Henderson (2017). The goal of selecting a near-optimal system with high probability also arises in the multi-armed bandit literature (Jourdan et al. 2024), where it is called *probably approximately correct* (PAC) selection. Eckman and Henderson (2021) compare and contrast the PCS and PGS guarantees and discuss their implications for the design of R&S procedures.

From a methodological viewpoint, our study of the budget allocation problem is closely related to the large-deviations analysis introduced by Glynn and Juneja (2004). In their seminal work, the rate at which the PCS converges to 1 is expressed as a concave function in terms of the fraction of the budget allocated to each system. This function is then optimized to yield the optimal sampling allocation. This formulation has inspired related analyses in other R&S variants, such as R&S with stochastic constraints (Pasupathy et al. 2014), multi-objective R&S (Feldman and Hunter 2018; Applegate et al. 2020), R&S under input uncertainty (Kim et al. 2022), feature-based R&S (Ahn et al. 2021), and contextual R&S (Du et al. 2024; Cakmak et al. 2023), as well as related allocation schemes such as top-two sampling (Russo 2020) and two-moment approximation (Shin et al. 2018). Chen and Ryzhov (2023) propose a fully-sequential procedure called balancing optimal large deviations (BOLD), which aims to satisfy the optimality conditions associated with the aforementioned convex optimization problem.

The optimal sampling allocation and the associated optimality conditions are an important means of justifying the optimality of proposed algorithms, even those that are not derived from large-deviations theory. For instance, the convergence of several myopic procedures to the optimal sampling ratio is established by showing that the optimality conditions are satisfied asymptotically; examples include the modified completed expected improvement (mCEI) algorithm (Chen and Ryzhov 2019), the gradient of CEI (gCEI) procedure (Avci et al. 2023), and dynamic programming approaches (Peng et al. 2018; Li et al. 2022).

The remainder of this paper is organized as follows: Section 2 reviews the large-deviations analysis of the PCS and extends it to the PGS. In Section 3, we discover and discuss a zero-sampling-ratio property exhibited by the optimal sampling allocation for the PGS. Building upon the convergence rate analysis, Section 4 investigates an adaptive algorithm modified for the PGS goal and illuminates underlying issues in its empirical performance. Conclusions and future research directions are presented in Section 5.

2 LARGE-DEVIATIONS ANALYSIS FOR THE PROBABILITY OF GOOD SELECTION

This section analyzes the optimal allocation problem through a large-deviations analysis of the PGS, given that the distributions of all systems' outputs are known. We first review existing results for the PCS.

2.1 Large-Deviations Analysis for the PCS

Assume there are k systems under consideration, where a simulation output from System i is distributed as $N(\mu_i, \sigma_i^2)$, for $i = 1, 2, \dots, k$. Without loss of generality, we assume that systems are indexed in order of their expected performance and that smaller expected performance is better. We further assume that there is a unique optimal system, i.e.,

$$\mu_1 < \mu_2 \leq \dots \leq \mu_k, \quad (1)$$

and that the variances $\sigma_i^2, i = 1, 2, \dots, k$, are known.

In the R&S setting, each μ_i can typically be learned only via the outputs of simulation replications of System i . Let n be a budget of replications and N_i^n be the number of replications allocated to System i for $i = 1, 2, \dots, k$. By definition, $\sum_{i=1}^k N_i^n = n$. We further let $\mu_{i,n}$ denote the sample mean based on N_i^n

simulation outputs sampled from System i . The probability of a correct selection after taking n simulation replications, denoted as PCS_n , is defined as

$$\text{PCS}_n = \text{P}(\arg \min_{1 \leq i \leq k} \mu_{i,n} = \{1\}).$$

Because PCS_n is the probability of a finite intersection of events, it is difficult to obtain an algebraic expression that is directly amenable to optimization. To circumvent this issue, Glynn and Juneja (2004) investigate the large deviations rate (LDR) of PCS_n under a *static allocation regime*. Specifically, for a fixed $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_k) \in \Delta^k := \{\boldsymbol{\alpha} : \mathbf{1}^\top \boldsymbol{\alpha} = 1, \boldsymbol{\alpha} \geq 0\}$, the static allocation of n simulation replications under $\boldsymbol{\alpha}$ satisfies $N_i^n = n\alpha_i$ for all $i = 1, 2, \dots, k$, ignoring issues of integrality. Glynn and Juneja (2004) show that for any static allocation and $\boldsymbol{\alpha} \in \Delta^k$,

$$G_i(\boldsymbol{\alpha}) := \lim_{n \rightarrow \infty} -\frac{1}{n} \log \text{P}(\mu_{i,n} \leq \mu_{1,n}) = \frac{(\mu_i - \mu_1)^2}{2(\sigma_1^2/\alpha_1 + \sigma_i^2/\alpha_i)}$$

for $i = 2, 3, \dots, k$ and

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log(1 - \text{PCS}_n) = \min_{2 \leq i \leq k} G_i(\boldsymbol{\alpha}).$$

Intuitively speaking, $G_i(\boldsymbol{\alpha})$ represents the LDR of the probability that the sample mean of System i falsely outperforms that of System 1 under the static allocation $\boldsymbol{\alpha}$. The LDR of the probability of false selection (PFS), i.e., $1 - \text{PCS}$, is then expressed as a minimum of the LDRs from the pairwise comparisons between Systems 1 and i , $i \neq 1$. In other words, the exponential convergence rate of the PFS is determined by the pairwise comparison having the slowest convergence rate.

The rate-optimal sampling allocation for the PCS, say $\boldsymbol{\alpha}^{\text{PCS}}$, is

$$\boldsymbol{\alpha}^{\text{PCS}} := \arg \max_{\boldsymbol{\alpha} \in \Delta^k} \min_{2 \leq i \leq k} G_i(\boldsymbol{\alpha}). \tag{2}$$

The optimization problem in (2) is a convex program in $\boldsymbol{\alpha}$ with necessary optimality conditions for (2) derived from the Karush-Kuhn-Tucker (KKT) conditions, as stipulated in Theorem 1.

Theorem 1 (Glynn and Juneja (2004)) For all $i = 1, 2, \dots, k$, we have $\alpha_i^{\text{PCS}} > 0$. Furthermore, $\boldsymbol{\alpha}^{\text{PCS}}$ is the optimal solution to (2) if and only if $\boldsymbol{\alpha}^{\text{PCS}}$ satisfies

- (i) $G_2(\boldsymbol{\alpha}^{\text{PCS}}) = G_3(\boldsymbol{\alpha}^{\text{PCS}}) = \dots = G_k(\boldsymbol{\alpha}^{\text{PCS}})$;
- (ii) $\sum_{j=2}^k \frac{\partial G_j(\boldsymbol{\alpha})/\partial \alpha_1}{\partial G_j(\boldsymbol{\alpha})/\partial \alpha_j} \Big|_{\boldsymbol{\alpha}=\boldsymbol{\alpha}^{\text{PCS}}} = 1$, or equivalently, $(\alpha_1^{\text{PCS}}/\sigma_1)^2 = \sum_{j=2}^k (\alpha_j^{\text{PCS}}/\sigma_j)^2$.

Condition (i) implies that, under $\boldsymbol{\alpha}^{\text{PCS}}$, the probabilities of making errors when comparing Systems 1 and i converge to zero at the same exponential rate. That is, all pairwise comparisons are equally important. The second condition states the balance condition between the derivatives of the LDRs of suboptimal solutions. It is worth noting that Theorem 1 gives necessary and sufficient conditions that guarantee optimality. The aforementioned works of Peng et al. (2018), Chen and Ryzhov (2019), Li et al. (2022), and Avci et al. (2023) exploit the necessity and sufficiency of the conditions to show that the empirical allocations under their respective algorithms converge to $\boldsymbol{\alpha}^{\text{PCS}}$ as $n \rightarrow \infty$.

2.2 Large-Deviations Analysis for the PGS

We extend the previous analysis to the setting of maximizing the PGS. Suppose that the decision maker is indifferent to selecting any system whose expected performance is within some tolerance $\delta > 0$ of the best. In this case, unlike in (1), the best system need not necessarily be unique. We instead assume that the configuration of expected performances satisfies

$$\mu_1 \leq \mu_2 \leq \dots \leq \mu_\ell \leq \mu_1 + \delta < \mu_{\ell+1} \leq \dots \leq \mu_k \text{ for some } 1 \leq \ell < k,$$

where ℓ represents the number of good systems, i.e., $\ell = |\{j : \mu_j \leq \mu_1 + \delta\}|$. This paper focuses on the nontrivial case when $1 \leq \ell < k$. We suppress the dependence of ℓ on δ in the notation. The value of ℓ is unknown, but can be learned over time alongside the unknown means $\mu_1, \mu_2, \dots, \mu_k$.

Let $\mathcal{D}_n = \arg \min_{1 \leq i \leq k} \mu_{i,n}$ be the system having the best sample mean after taking a total of n simulation replications. Notice that \mathcal{D}_n is random, as it depends on the simulation outputs. The PGS after taking n simulation replications, PGS_n , is then defined as

$$\text{PGS}_n = \mathbb{P}(\mathcal{D}_n \in \{1, 2, \dots, \ell\}).$$

Observe that

$$\max_{\ell+1 \leq j \leq k} \mathbb{P}(\mathcal{D}_n = j) \leq 1 - \text{PGS}_n = \mathbb{P}(\mathcal{D}_n \in \{\ell+1, \dots, k\}) \leq (k - \ell) \max_{\ell+1 \leq j \leq k} \mathbb{P}(\mathcal{D}_n = j).$$

A straightforward calculation yields that

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log(1 - \text{PGS}_n) = \min_{\ell+1 \leq j \leq k} \lim_{n \rightarrow \infty} -\frac{1}{n} \log \mathbb{P}(\mathcal{D}_n = j) = \min_{\ell+1 \leq j \leq k} \tilde{G}_j(\boldsymbol{\alpha}),$$

where $\tilde{G}_j(\boldsymbol{\alpha}) := \lim_{n \rightarrow \infty} -\frac{1}{n} \log \mathbb{P}(\mathcal{D}_n = j)$ is the LDR of the event that System j is falsely identified as the best. Observe that the relation $\{\mathcal{D}_n = j\} \subseteq \{\mu_{j,n} \leq \mu_{1,n}\}$ implies that $\tilde{G}_j(\boldsymbol{\alpha}) \geq G_j(\boldsymbol{\alpha})$. From this, we further obtain that for any $\boldsymbol{\alpha} \in \Delta^k$ and any problem instance satisfying $\mu_1 < \mu_2 \leq \dots \leq \mu_k$,

$$\min_{i \neq 1} G_i(\boldsymbol{\alpha}) \leq \min_{\ell+1 \leq i \leq k} G_i(\boldsymbol{\alpha}) \leq \min_{\ell+1 \leq i \leq k} \tilde{G}_i(\boldsymbol{\alpha}). \quad (3)$$

Equation (3) shows that the LDR of the PGS is always greater than or equal to that of the PCS. This is intuitive because all inequalities in (3) become equalities when $\ell = 1$, i.e., when the best system is uniquely good. To help formalize the optimal sampling allocation for the PGS, analogous to (2), we first present an analytical expression for $\tilde{G}_j(\boldsymbol{\alpha})$.

Theorem 2 (Theorem 1 and Lemma EC.2 in Kim et al. (2022)) Under the static allocation of $\boldsymbol{\alpha}$, $\tilde{G}_j(\boldsymbol{\alpha})$ is concave in $\boldsymbol{\alpha}$. Furthermore,

$$\tilde{G}_j(\boldsymbol{\alpha}) = \min_{x \in [\mu_1, \mu_j]} \left\{ \frac{\alpha_j}{2\sigma_j^2} (x - \mu_j)^2 + \sum_{1 \leq i < j} \frac{\alpha_i}{2\sigma_i^2} [(x - \mu_i)^+]^2 \right\} \text{ for all } j = 2, 3, \dots, k, \quad (4)$$

where $x^+ = \max(x, 0)$. Define the minimizer of the right-hand side of (4) by $x_j(\boldsymbol{\alpha})$. The partial derivatives of $\tilde{G}_j(\boldsymbol{\alpha})$ are given by

$$\frac{\partial \tilde{G}_j(\boldsymbol{\alpha})}{\partial \alpha_i} = \begin{cases} \frac{[(x_j(\boldsymbol{\alpha}) - \mu_i)^+]^2}{2\sigma_i^2}, & \text{if } i < j, \\ \frac{(x_j(\boldsymbol{\alpha}) - \mu_i)^2}{2\sigma_i^2}, & \text{if } i = j, \\ 0, & \text{if } i > j. \end{cases} \quad (5)$$

The optimal allocation for the PGS can be formulated as the solution to the optimization problem

$$\max_{\boldsymbol{\alpha} \in \Delta^k} \text{LDR}^{\text{PGS}}(\boldsymbol{\alpha}), \text{ where } \text{LDR}^{\text{PGS}}(\boldsymbol{\alpha}) = \min_{\ell+1 \leq j \leq k} \tilde{G}_j(\boldsymbol{\alpha}). \quad (6)$$

From Theorem 2, one can easily confirm that (6) is a convex program. Theorem 3 derives necessary and sufficient conditions for the optimal solution of (6) based on the KKT conditions, analogous to Theorem 1. We omit the proof of Theorem 3 and those of all other technical results due to the space limit.

Theorem 3 (KKT conditions of (6)) The sampling allocation $\boldsymbol{\alpha}^{\text{PGS}}$ is the optimal solution to (6) if and only if $\boldsymbol{\alpha}^{\text{PGS}}$ satisfies the following conditions:

- (i) $\tilde{G}_{\ell+1}(\boldsymbol{\alpha}^{\text{pgs}}) = \tilde{G}_{\ell+2}(\boldsymbol{\alpha}^{\text{pgs}}) = \dots = \tilde{G}_k(\boldsymbol{\alpha}^{\text{pgs}})$; and
- (ii) for each $i \in \{1, 2, \dots, \ell\}$,

$$\sum_{j=\ell+1}^k \frac{\partial \tilde{G}_j(\boldsymbol{\alpha}) / \partial \alpha_i}{\partial \tilde{G}_j(\boldsymbol{\alpha}) / \partial \alpha_j} \Big|_{\boldsymbol{\alpha}=\boldsymbol{\alpha}^{\text{pgs}}} \begin{cases} = 1, & \text{if } \alpha_i^{\text{pgs}} > 0, \\ < 1, & \text{if } \alpha_i^{\text{pgs}} = 0. \end{cases} \quad (7)$$

Corollary 4 The optimal solution to (6), $\boldsymbol{\alpha}^{\text{pgs}}$, satisfies $\alpha_j^{\text{pgs}} > 0$ for all $j \geq \ell + 1$.

Similar to Theorem 1, Condition (i) of Theorem 3 implies that all $\tilde{G}_j(\boldsymbol{\alpha})$ s should be identical for $j = \ell + 1, \ell + 2, \dots, k$. Equation (7) is the counterpart of Condition (ii) of Theorem 1, with the exception that the gradient of the LDR for System i is balanced only if System i receives positive sampling under $\boldsymbol{\alpha}^{\text{pgs}}$. Corollary 4 is a byproduct of Theorem 3 and guarantees that all bad systems receive strictly positive sampling ratios under $\boldsymbol{\alpha}^{\text{pgs}}$. On the other hand, we will soon see that some good systems may receive a sampling ratio of zero under $\boldsymbol{\alpha}^{\text{pgs}}$.

Remark 1 We clarify that a zero sampling ratio for System i does not mean we do not allocate any simulation budget or only a finite number of replications to that system. Recall that the large deviations theory supposes that we allocate infinitely many replications to each system. Therefore, a zero sampling ratio implies that although we assign infinitely many replications to System i , i.e., $\lim_{n \rightarrow \infty} N_i^n = \infty$, this number increases sublinearly in n , i.e., $\lim_{n \rightarrow \infty} N_i^n / n = 0$.

Remark 2 Although the LDRs in (4) do not depend on δ , Program (6) does since ℓ is implicitly a function of δ .

2.3 An Empirical Study

We conduct a small empirical study to illustrate properties of the $\boldsymbol{\alpha}^{\text{pgs}}$ allocation and contrast it with the $\boldsymbol{\alpha}^{\text{pcs}}$ allocation. The clearest difference between the two allocations is that under $\boldsymbol{\alpha}^{\text{pcs}}$, all systems receive positive sampling ratios, whereas under $\boldsymbol{\alpha}^{\text{pgs}}$, some good systems may receive zero sampling ratios. At first glance, this consequence may seem counter-intuitive since it roughly means that one should not spend simulation effort on certain good systems that could improve the PGS if they were correctly identified as good systems. On the other hand, unlike in the PCS case, identifying the best system is less critical in the PGS case, so one may expect that, compared to $\boldsymbol{\alpha}^{\text{pcs}}$, $\boldsymbol{\alpha}^{\text{pgs}}$ should allocate less of the simulation budget to good systems and instead reallocate this effort to the bad systems. Loosely speaking, if one good system is especially easy to identify as good, then we need not allocate simulation effort to the other good systems since correctly identifying these other good systems is unimportant for the purposes of making a good selection. Instead, the budget is better allocated to those good systems that are easiest to identify as good, hence the zero-sampling-ratio phenomenon.

Table 1 summarizes two configurations of ten systems' means and variances that we test. The vector of means is the same in both configurations, but in Configuration 1, the bad systems have the highest variances, whereas in Configuration 2, the best system has the highest variance. Figure 1 displays the optimal LDR of PGS at $\boldsymbol{\alpha}^{\text{pgs}}$ and $(\alpha_1^{\text{pgs}}, \alpha_2^{\text{pgs}})$ as functions of μ_2 when the means and variances for all other systems are fixed, as described in Configuration 2. We set $\delta = 1$ in this example.

Table 1: Two mean-variance configurations with $k = 10$ and $\delta = 1$.

Configuration 1	$(\mu_1, \sigma_1^2) = (1, 3^2), (\mu_2, \sigma_2^2) = (1.5, 2^2),$ and $(\mu_i, \sigma_i^2) = (2.5, 6^2)$ for $i \geq 3$.
Configuration 2	$(\mu_1, \sigma_1^2) = (1, 8^2), (\mu_2, \sigma_2^2) = (1.5, 4^2),$ and $(\mu_i, \sigma_i^2) = (2.5, 3^2)$ for $i \geq 3$.

The left panel of Figure 1 shows that the optimal LDR of $\boldsymbol{\alpha}^{\text{pgs}}$ is higher than that of $\boldsymbol{\alpha}^{\text{pcs}}$ when $\mu_2 < \mu_1 + \delta = 2$, because there are two good systems, whereas the LDRs for PGS and PCS coincide when $\mu_2 > 2$, because System 1 is then the only good system. In particular, we observe the flat behavior of $\text{LDR}^{\text{pgs}}(\boldsymbol{\alpha}^{\text{pgs}})$ when μ_2 is sufficiently larger than μ_1 ; in this case, it is better to identify System 1 as a good

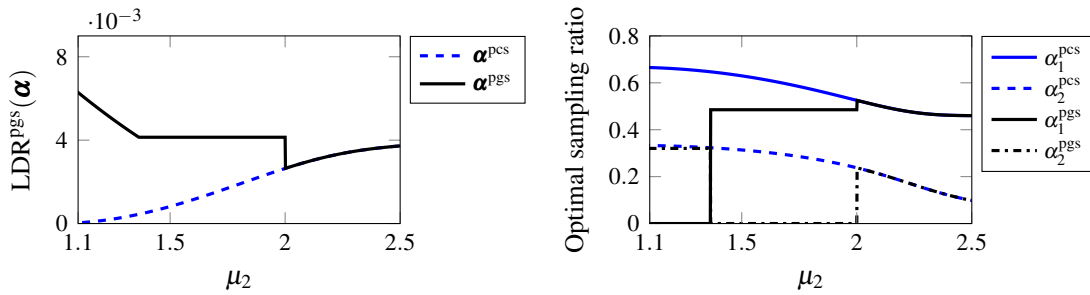


Figure 1: Comparison between α^{PCS} and α^{PGS} in terms of $\text{LDR}^{\text{PGS}}(\alpha)$ and sampling ratios allocated to Systems 1 and 2 (α_1^{PGS} and α_2^{PGS}) by shifting μ_2 in Configuration 2 when $\delta = 1$.

system instead of System 2. This makes $\alpha_2^{\text{PGS}} = 0$, and the LDR of PGS does not depend on the change of μ_2 . Furthermore, there is a jump at $\mu_2 = 2$ where System 2 becomes bad, at which point ℓ decreases from 2 to 1, and $\text{LDR}^{\text{PGS}}(\alpha^{\text{PGS}})$ changes from being the minimum of $k - 2$ LDRs to being the minimum of $k - 1$ LDRs. Thus, the optimizer, α^{PGS} , and optimal value, $\text{LDR}^{\text{PGS}}(\alpha^{\text{PGS}})$, of Problem (6) both change discontinuously. Meanwhile, the right panel of Figure 1 shows α^{PGS} at Systems 1 and 2 as a function of μ_2 . When there are multiple good systems, i.e., when $\mu_2 \leq 2$, the PGS-optimal sampling allocation can change dramatically depending on the value of μ_2 . When μ_2 is sufficiently close to μ_1 , α_1^{PGS} becomes zero, even though System 1 is the best.

We next attempt to characterize situations in which a good system receives a zero sampling ratio.

3 WHEN DOES A GOOD SYSTEM RECEIVE A ZERO SAMPLING RATIO?

We seek to understand if there are necessary and sufficient conditions dictating when a given good system receives a zero sampling ratio. If there are, then one could incorporate these conditions into a sequential sampling rule to identify such systems more easily. Even sufficient conditions alone could help to identify certain systems to not sample. We present necessary and sufficient conditions for the $k = 3$ case and sufficient conditions for the general case.

The first sufficient condition, given in Proposition 5, involves the means and variances of the good systems and the best bad system, System $\ell + 1$.

Proposition 5 For any System j , $j \leq \ell$, if $\frac{\mu_{\ell+1} - \mu_i}{\sigma_i} > \frac{\mu_{\ell+1} - \mu_j}{\sigma_j}$ and $\mu_i \leq \mu_j$ for some $i \neq j$, then $\alpha_j^{\text{PGS}} = 0$.

The fraction $\frac{\mu_{\ell+1} - \mu_j}{\sigma_j}$ can be viewed as a measure of how easy it is to distinguish System j from the best bad system. If there exists another good system, i , that is easier to distinguish from System $\ell + 1$, then the simulation budget is better allocated to System i instead of System j .

Corollary 6 states that a good system whose mean and variance are both larger than those of some other good system will receive a zero sampling ratio under α^{PGS} . This is intuitive since it would be advantageous to allocate simulation effort to the good system having the smaller variance and smaller mean.

Corollary 6 For any System j , $j \leq \ell$, if $\mu_i \leq \mu_j$ and $\sigma_i^2 < \sigma_j^2$ for any $i \neq j$, then $\alpha_j^{\text{PGS}} = 0$.

Characterizing necessary and sufficient conditions for a zero sampling ratio remains an open problem. We will, however, discuss one such set of conditions for when $k = 3$. For the $k = 3$ system case, there can be either one, two, or three good systems. If there is only one good system, then the LDR for the PGS coincides with that for the PCS, as seen in the left panel of Figure 1, thus all three systems receive positive sampling ratios. When all systems are good, the LDR becomes infinity since we will always make a good selection. We therefore focus on the case where there are two good systems and one bad system, i.e., $\mu_1 < \mu_2 \leq \mu_1 + \delta < \mu_3$.




Figure 2: α^{pgs} as a function of (μ_2, σ_2) when $\delta = 1$, $(\mu_1, \sigma_1) = (1, 2)$ and $(\mu_3, \sigma_3) = (2.5, 3)$. The equations of solid and dashed lines are given as $\frac{\mu_3 - \mu_2}{\sigma_3 + \sigma_2} = \frac{\mu_3 - \mu_1}{\sigma_3 + \sigma_1}$ and $\frac{\mu_3 - \mu_2}{\sigma_2} = \frac{\mu_3 - \mu_1}{\sigma_1}$, respectively.

In this case, Program (6) can be reduced to the maximization of $\tilde{G}_3(\alpha)$ with respect to α . Since System 3 is a bad system, $\alpha_3^{\text{pgs}} > 0$. Proposition 7 provides necessary and sufficient conditions – in terms of $\{(\mu_i, \sigma_i^2)\}_{1 \leq i \leq 3}$ – for which α_1^{pgs} or α_2^{pgs} are zero.

Proposition 7 Assume $\mu_1 < \mu_2 \leq \mu_1 + \delta < \mu_3$ and let $\alpha^{\text{pgs}} = \arg \max_{\alpha} \tilde{G}_3(\alpha)$.

- $\alpha_1^{\text{pgs}} > 0$ and $\alpha_2^{\text{pgs}} > 0$ if and only if $\frac{\mu_3 - \mu_1}{\sigma_3 + \sigma_1} = \frac{\mu_3 - \mu_2}{\sigma_3 + \sigma_2}$;
- $\alpha_1^{\text{pgs}} > 0$ and $\alpha_2^{\text{pgs}} = 0$ if and only if $\frac{\mu_3 - \mu_1}{\sigma_3 + \sigma_1} > \frac{\mu_3 - \mu_2}{\sigma_3 + \sigma_2}$; and
- $\alpha_1^{\text{pgs}} = 0$ and $\alpha_2^{\text{pgs}} > 0$ if and only if $\frac{\mu_3 - \mu_1}{\sigma_3 + \sigma_1} < \frac{\mu_3 - \mu_2}{\sigma_3 + \sigma_2}$.

Proposition 7 involves differences in means divided by the sum of standard deviations. This differs from Proposition 5 in that the denominators also involve the standard deviation of the bad system. This means that we need to account for the variability of bad systems to understand whether a good system should receive a zero sampling ratio. The proof of Proposition 7 exploits the necessary and sufficient conditions of Theorem 3.

To illustrate the various conditions, we perform a small numerical experiment with $k = 3$ systems wherein we vary the mean and variance of the second good system, System 2. For each value of μ_2 and σ_2 , we compute $\alpha^{\text{pgs}} = \arg \max_{\alpha \in \Delta^k} \tilde{G}_3(\alpha)$. The three components of α^{pgs} sum to 1 and are shown separately in the three panels of Figure 2. The solid line in each panel represents the boundary described by Proposition 7 and reaffirms the necessity and sufficiency of this condition. On the other hand, the dashed line indicates the sufficient condition provided by Proposition 5, which only implies that $\alpha_2^{\text{pgs}} = 0$ if (μ_2, σ_2) lies above the dashed line. The gap between the two lines is indicative of the looseness the conditions of Proposition 5 might have in the general case.

4 CHALLENGES IN DESIGNING ADAPTIVE SAMPLING PROCEDURES

Building upon the large-deviations analysis of the PGS, this section investigates adaptive procedures designed to identify good systems. When testing adaptive procedures, we will assume that the variances σ_i^2 , $i = 1, 2, \dots, k$, are unknown, as would be the case in practice; however, we notice similar empirical results when using the true variances, suggesting that the performance issues that we observe are not a result of needing to estimate the variances. We focus on a straightforward procedure that iteratively solves a plug-in version of Program (6) and allocates a batch of replications according to the estimated optimal sampling allocation. We refer to this procedure as $\pi^{\text{B}}(\delta)$ (B for batch) and provide details in Algorithm 1.

More specifically, Algorithm 1 repeatedly sets up and solves (8), a plug-in version of (6) in which the unknown μ_i and σ_i^2 are replaced with their respective estimates $\mu_{i,n}$ and $\sigma_{i,n}^2$. Program (8) is a convex program in α and can be solved efficiently via a first-order method, e.g., subgradient descent, using the

Algorithm 1 Plug-in Algorithm with Batch Allocation ($\pi^B(\delta)$)

- 1: **Initialization:** initial allocations n_0 , total budget N , batch size B , and budget counter $n = 0$.
- 2: For each System i , $i = 1, 2, \dots, k$, run n_0 replications and update $\mu_{i,n}$ and $\sigma_{i,n}^2$. Set $N_i^n = n_0$ for all $i = 1, 2, \dots, k$ where $n = kn_0$.
- 3: **while** $n \leq N$ **do**
- 4: Define $\mathcal{J}_n := \{i : \mu_{i,n} \leq \min_{1 \leq j \leq k} \mu_{j,n} + \delta\}$, the collection of all systems that look good based on their sample means.
- 5: **if** $\mathcal{J}_n = \{1, 2, \dots, k\}$, i.e., all systems look good, **then**
- 6: Set $\alpha_n = \frac{1}{k} \mathbf{1}_k$.
- 7: **else**
- 8: Solve the following optimization problem to obtain the estimated optimal sampling allocation:

$$\alpha_n = \arg \max_{\alpha \in \Delta^k} \min_{j \notin \mathcal{J}_n} \tilde{G}_{j,n}(\alpha), \tag{8}$$

where $\tilde{G}_{j,n}(\alpha)$ is a plug-in version of $\tilde{G}_j(\alpha)$ in which μ_i and σ_i^2 are replaced with $\mu_{i,n}$ and $\sigma_{i,n}^2$.

- 9: **end if**
 - 10: Draw a sample of size B from a multinomial distribution with probability α_n , i.e., $(m_i)_{1 \leq i \leq k} \sim \text{MN}(B, \alpha_n)$.
 - 11: For each System i , $i = 1, 2, \dots, k$, take m_i replications and update $N_i^n \leftarrow N_i^n + m_i$, $n \leftarrow n + B$, $\mu_{i,n}$, and $\sigma_{i,n}^2$ accordingly.
 - 12: **end while**
 - 13: **return** $\arg \min_{1 \leq i \leq k} \mu_{i,N}$.
-

closed-form expression of the gradient of $\tilde{G}_j(\alpha)$ given in (5). The algorithm then randomly apports a batch of B replications according to the optimal allocation from the plug-in problem. The multinomial distribution is used to address the issue of integrality. These steps are repeated using the updated statistics until the budget is exhausted, at which time the system having the best sample mean is selected. Setting $B = 1$ makes Algorithm 1 a fully sequential procedure. However, this may be computationally expensive due to the need to repeatedly solve convex optimization problems.

Remark 3 If $\delta = 0$, Algorithm 1 becomes a sequential sampling procedure for the PCS, using plug-in estimates and batch allocations.

Remark 4 We also introduce a “static” allocation with a given α by modifying Lines 5–9 of Algorithm 1: instead of solving Program (8), we simply let $\alpha_n = \alpha$. We denote this procedure by $\pi^S(\alpha)$. This algorithm will be employed to confirm some of the theoretical results presented in Section 2.

4.1 Inconsistency of $\pi^B(\delta)$

In this section, we discuss how the aforementioned zero-sampling-ratio property can degrade the PGS of Algorithm 1. For the vast majority of existing adaptive procedures that do not eliminate systems from contention, every system is sampled infinitely often as the simulation budget goes to infinity. Combined with the law of large numbers, this implies that the procedure is consistent, i.e., will asymptotically learn the expected performances of all systems. Observe that Algorithm 1 would assign a simulation replication to an estimated good system, say System i , only if $\alpha_{i,n} > 0$. This may result in System i receiving only a finite number of samples asymptotically if (i) \mathcal{J}_n is not a singleton, (ii) System i looks good, and (iii) $\alpha_{i,n} = 0$. This kind of behavior can deteriorate the convergence rate of the empirical PGS of Algorithm 1.

In Figure 3, we illustrate two representative scenarios in which $\pi^B(\delta)$ struggles to identify a good system due to the zero-sampling-ratio property of the optimal solution to Program (8). For simplicity, we illustrate the case where Systems 1 and 2 are good, System 3 is bad, and omit the other systems. In Bad




Figure 3: Two bad scenarios under which the empirical performance of $\pi^B(\delta)$ could be poor. The bell curves indicate the probability density functions of $\{N(\mu_i, \sigma_i^2)\}_{1 \leq i \leq 3}$ under each scenario.

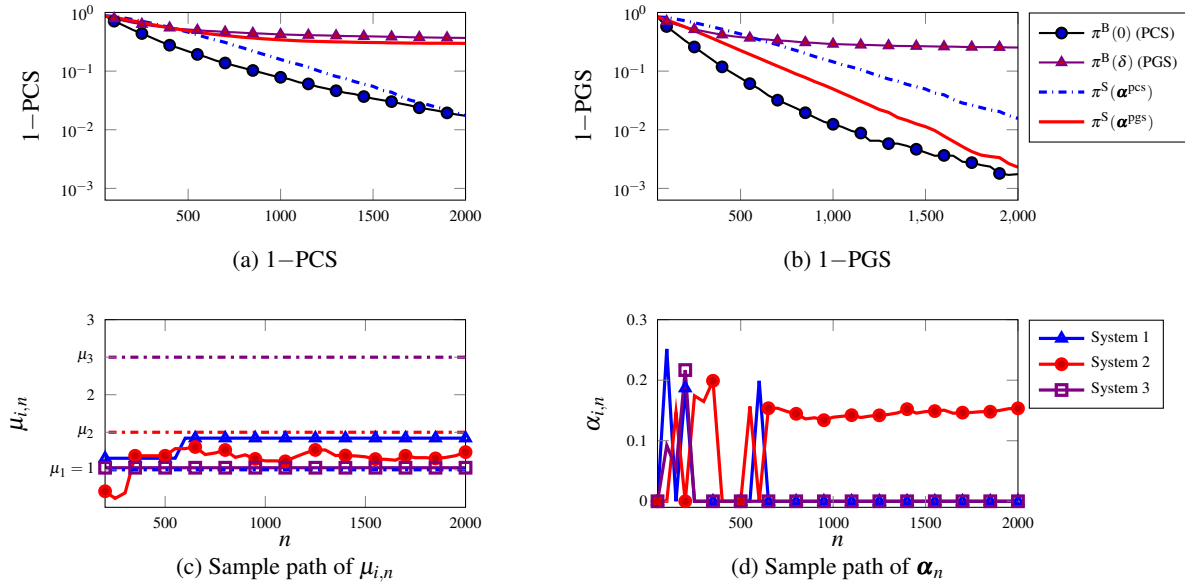


Figure 4: Numerical results under Configuration 1. (First row) PCS and PGS estimated from 10,000 macroreplications. (Second row) Sample paths of $\mu_{i,n}$ and α_n when in Bad Scenario 1. The dashed lines in panel (c) indicate $\{\mu_i\}_{1 \leq i \leq 3}$, so this shows that the sample path falls into Bad Scenario 1.

Scenario 1, System 3 has a high variance and has been significantly underestimated and falsely identified as the best. Depending on the sample means and variances, $\pi^B(\delta)$ may choose to not sample System 3, especially if the estimated variance of System 3 is large. In Bad Scenario 2, both Systems 1 and 2 have high variances and their means have been significantly overestimated. In this scenario, $\pi^B(\delta)$ may choose to not samples Systems 1 and 2 because all systems look good and System 3 is then the system having the smallest mean and variance. Bad Scenarios 1 and 2 are most likely to arise from Configurations 1 and 2, respectively, from Table 1.

Figure 4 shows numerical results under Configuration 1 for four algorithms: two static sampling algorithms, $\pi^S(\alpha^{PCS})$ and $\pi^S(\alpha^{PGS})$, which serve as benchmarks, and two adaptive sampling algorithms, $\pi^B(0)$ and $\pi^B(\delta)$. We set $B = 50$ and repeat each algorithm 10,000 times to estimate the PCS and the PGS. In the first row, we plot $1 - \text{PCS}$ and $1 - \text{PGS}$ for each algorithm. In terms of the empirical PCS, $\pi^B(0)$ and $\pi^S(\alpha^{PCS})$ clearly outperform $\pi^B(\delta)$ and $\pi^S(\alpha^{PGS})$. In terms of the empirical PGS, $\pi^S(\alpha^{PGS})$ outperforms $\pi^S(\alpha^{PCS})$, as expected. Although $\pi^B(\delta)$ performs worse than $\pi^S(\alpha^{PGS})$, indicating some performance

loss due to the use of plug-in estimates, $\pi^B(0)$ surprisingly outperforms the other three algorithms, even though it is designed for the PCS goal. We note that we have observed $\pi^B(0)$ performing poorly in other configurations, which will be discussed later in this paper.

We further investigate the poor performance of $\pi^B(\delta)$ via a sample-path analysis. Figures 4(c) and (d) show sample paths of $\{\mu_{i,n}\}_{1 \leq i \leq 3}$ and $\{\alpha_{i,n}\}_{1 \leq i \leq 3}$ from one macroreplication corresponding to Bad Scenario 1. Figure 4(c) shows that $\mu_{1,n}$ and $\mu_{3,n}$ are not updated after time 600. This is because all three systems consistently look good, yet $\alpha_{1,n} = \alpha_{3,n} = 0$, as shown in panel (d). Algorithm $\pi^B(\delta)$ struggles to escape this situation because there is no mechanism to compel exploration. Theorem 8 below shows that bad sample paths like those illustrated in panel (d) can be made to occur with a positive probability under Bad Scenario 1 by choosing n_0 appropriately, subject to some regularity conditions when $k = 3$.

Theorem 8 Consider a problem with $k = 3$ systems where $\mu_1 < \mu_2 < \mu_1 + \delta < \mu_3$. If $\sigma_3 > \max\{\sigma_1, \sigma_2\}$ and $\delta < \frac{\sigma_3 + \sigma_2}{\sigma_2 + \sigma_1}(\mu_2 - \mu_1)$, then there exists n_0 such that the probability of Algorithm 1 falsely identifying System 3 as the sample best for all $n \geq n_0 k$ is positive.

Typically, the convergence of the sample best system to the true best system is guaranteed as long as $\lim_{n \rightarrow \infty} N_i^n = \infty$, regardless of n_0 and $\{(\mu_i, \sigma_i)\}_{1 \leq i \leq k}$. However, Theorem 8 shows that we can find a problem instance $\{(\mu_i, \sigma_i)\}_{1 \leq i \leq k}$ and initial sample size n_0 for which Algorithm 1 fails to select a good system with positive probability, even if the total simulation budget goes to infinity. In light of this issue, we explore one potential modification to $\pi^B(\delta)$ that seeks to improve its performance by mixing the optimal allocations for the PCS and PGS plug-in problems.

4.2 A Mixture-Based Approach

To avoid situations in which the zero-sampling phenomenon causes $\pi^B(\delta)$ to stop sampling certain systems that look good, we propose adjusting α_n^{pgs} to ensure that all systems that look good have a positive probability of being sampled. Motivated by the fact that every component of α^{pcs} is positive, we propose sampling according to a convex combination of α^{pgs} and α^{pcs} , specifically

$$\alpha(\varepsilon) := (1 - \varepsilon)\alpha^{\text{pgs}} + \varepsilon\alpha^{\text{pcs}}$$

for some $\varepsilon \in (0, 1)$. From the definition of α^{pgs} as the PGS-optimal allocation, $\alpha(\varepsilon)$ cannot outperform α^{pgs} in terms of the LDR. Theorem 9 bounds the loss from mixing α^{pgs} with α^{pcs} .

Theorem 9 For any $\varepsilon \in (0, 1)$,

$$\text{LDR}^{\text{pgs}}(\alpha^{\text{pcs}}) \leq \text{LDR}^{\text{pgs}}(\alpha(\varepsilon)) \leq \text{LDR}^{\text{pgs}}(\alpha^{\text{pgs}}) \text{ and } \frac{\text{LDR}^{\text{pgs}}(\alpha^{\text{pgs}}) - \text{LDR}^{\text{pgs}}(\alpha(\varepsilon))}{\text{LDR}^{\text{pgs}}(\alpha^{\text{pgs}}) - \text{LDR}^{\text{pgs}}(\alpha^{\text{pcs}})} \leq \varepsilon.$$

The first part of Theorem 9 shows that $\alpha(\varepsilon)$ always performs better than α^{pcs} as long as $\varepsilon \in (0, 1)$. The second part shows that the ratio between the optimality gaps from sampling according to $\alpha(\varepsilon)$ and α^{pcs} is bounded above by ε .

Figure 5 displays the numerical results from implementing a static allocation $\pi^S(\alpha(\varepsilon))$ for different values of ε for the two configurations in Table 1. These results match our intuition in the sense that α^{pgs} performs best, and the PGS grows more slowly as ε increases.

We also propose a modification to $\pi^B(\delta)$ in which the optimization procedure in Line 8 of Algorithm 1 is replaced with a ε -mixture of α_n^{pgs} and α_n^{pcs} , as shown in Algorithm 2, which we refer to as $\pi^{\text{mix}}(\delta, \varepsilon)$.

Algorithm 2 ε -Modification of Line 8 in Algorithm 1 ($\pi^{\text{mix}}(\delta, \varepsilon)$)

- 8: Let $i_n = \min_{1 \leq i \leq k} \mu_{i,n}$ and set $\alpha_n = (1 - \varepsilon)\alpha_n^{\text{pgs}} + \varepsilon\alpha_n^{\text{pcs}}$, where $\alpha_n^{\text{pgs}} = \arg \max_{\alpha \in \Delta^k} \min_{j \notin \mathcal{J}_n} \tilde{G}_{j,n}(\alpha)$ and $\alpha_n^{\text{pcs}} = \arg \max_{\alpha \in \Delta^k} \min_{j \neq i_n} \tilde{G}_{j,n}(\alpha)$.
-

Figure 6 shows the estimated 1-PGS for variations of the modified algorithm with different values of ε . Whereas $\pi^B(\delta)$ performed poorly under Configuration 1, Figure 6(a) shows that this issue can be overcome




Figure 5: Numerical comparison of static allocations with mixing with $n_0 = 5$.




Figure 6: Numerical comparison between Algorithms 1 and 2 for Configurations 1 and 2. All results are averaged over 10,000 macroreplications, and we set $B = 50$.

by mixing α_n^{PGS} with α_n^{PCS} . In this case, $\pi^{\text{mix}}(\delta, 0.25)$ shows a significant improvement and $\pi^{\text{mix}}(\delta, 0.75)$ even outperforms $\pi^{\text{B}}(0)$, which was previously the best. In Figure 6(b), $\pi^{\text{B}}(\delta)$ already outperforms $\pi^{\text{B}}(0)$, and $\pi^{\text{mix}}(\delta, 0.25)$ and $\pi^{\text{mix}}(\delta, 0.75)$ exhibit slightly better performance than $\pi^{\text{B}}(\delta)$. Together, these results suggest that mixing may help achieve more robust performance across problem instances.

5 CONCLUSION

This paper investigates the PGS-optimal sampling allocation and associated optimality conditions based on large deviations theory. We provide some sufficient conditions for identifying the good systems that receive a zero sampling ratio under the PGS-optimal allocation. We also study adaptive sampling algorithms that allocate a budget in batches by solving the plug-in version of the optimal allocation problem. The initial PGS-based plug-in algorithm performs poorly in some circumstances. Based on a sample-path analysis, we attribute this behavior to the zero-sampling-ratio property of the PGS-optimal allocation. We propose and test a possible correction that mixes the PGS- and PCS-optimal allocations, observing some improvements in empirical performance.

Several open questions remain. Firstly, one could devise an adaptive calibration of ϵ that eventually drives ϵ to zero as n increases. A natural follow-up question is then how best to update ϵ . Secondly, one can alternatively consider a fully sequential procedure for PGS that exploits the optimality condition in Theorem 3. For instance, one may extend the BOLD algorithm of Chen and Ryzhov (2023) to the PGS setting. Finally, our “selection rule” is to choose the system with the best sample mean, which is intuitive when the goal is the correction selection. However, other selection rules that incorporate the sample variance information and consider the set of systems that look good could be advantageous.

ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation Grant CMMI-2206972.

REFERENCES

- Ahn, D., D. Shin, and A. Zeevi. 2021. "Feature Misspecification in Sequential Learning Problems". Accepted to *Management Science*.
- Applegate, E. A., G. Feldman, S. R. Hunter, and R. Pasupathy. 2020. "Multi-Objective Ranking and Selection: Optimal Sampling Laws and Tractable Approximations via SCORE". *Journal of Simulation* 14(1):21–40.
- Avci, H., B. L. Nelson, E. Song, and A. Wächter. 2023. "Using Cache or Credit for Parallel Ranking and Selection". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 33(4):1–28.
- Bechhofer, R. E. 1954. "A Single-Sample Multiple Decision Procedure for Ranking Means of Normal Populations with Known Variances". *The Annals of Mathematical Statistics* 25(1):16–39.
- Cakmak, S., Y. Wang, S. Gao, and E. Zhou. 2023. "Contextual Ranking and Selection with Gaussian Processes and OCBA". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 34(2):1–24.
- Chen, Y. and I. O. Ryzhov. 2019. "Complete Expected Improvement Converges to an Optimal Budget Allocation". *Advances in Applied Probability* 51(1):209–235.
- Chen, Y. and I. O. Ryzhov. 2023. "Balancing Optimal Large Deviations in Sequential Selection". *Management Science* 69(6):3457–3473.
- Du, J., S. Gao, and C.-H. Chen. 2024. "A Contextual Ranking and Selection Method for Personalized Medicine". *Manufacturing & Service Operations Management* 26(1):167–181.
- Eckman, D. J. and S. G. Henderson. 2021. "Fixed-Confidence, Fixed-Tolerance Guarantees for Ranking-and-Selection Procedures". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 31(2):1–33.
- Feldman, G. and S. R. Hunter. 2018. "SCORE Allocations for Bi-Objective Ranking and Selection". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 28(1):1–28.
- Glynn, P. and S. Juneja. 2004. "A Large Deviations Perspective on Ordinal Optimization". In *2004 Winter Simulation Conference (WSC)*, 577–585 <https://doi.org/10.1109/WSC.2004.1371364>.
- Hong, L. J. 2006. "Fully Sequential Indifference-Zone Selection Procedures with Variance-Dependent Sampling". *Naval Research Logistics (NRL)* 53(5):464–476.
- Hong, L. J., W. Fan, and J. Luo. 2021. "Review on Ranking and Selection: A New Perspective". *Frontiers of Engineering Management* 8(3):321–343.
- Jourdan, M., R. Degenne, and E. Kaufmann. 2024. "An ϵ -Best-Arm Identification Algorithm for Fixed-Confidence and Beyond". In *Advances in Neural Information Processing Systems*, Volume 36, 16578–16649.
- Kim, S.-H. and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-Zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11(3):251–273.
- Kim, S.-H. and B. L. Nelson. 2006. "On the Asymptotic Validity of Fully Sequential Selection Procedures for Steady-State Simulation". *Operations Research* 54(3):475–488.
- Kim, T., K.-K. Kim, and E. Song. 2022. "Selection of the Most Probable Best". <https://arxiv.org/abs/2207.07533>.
- Li, H., H. Lam, and Y. Peng. 2022. "Efficient Learning for Clustering and Optimizing Context-Dependent Designs". *Operations Research* 72(2):617–638.
- Ma, S. and S. G. Henderson. 2017. "An Efficient Fully Sequential Selection Procedure Guaranteeing Probably Approximately Correct Selection". In *2017 Winter Simulation Conference (WSC)*, 2225–2236 <https://doi.org/10.1109/WSC.2017.8247954>.
- Pasupathy, R., S. R. Hunter, N. A. Pujowidianto, L. H. Lee and C.-H. Chen. 2014. "Stochastically Constrained Ranking and Selection via SCORE". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 25(1):1–26.
- Peng, Y., E. K. Chong, C.-H. Chen, and M. C. Fu. 2018. "Ranking and Selection as Stochastic Control". *IEEE Transactions on Automatic Control* 63(8):2359–2373.
- Russo, D. 2020. "Simple Bayesian Algorithms for Best-Arm Identification". *Operations Research* 68(6):1625–1647.
- Shin, D., M. Broadie, and A. Zeevi. 2018. "Tractable Sampling Strategies for Ordinal Optimization". *Operations Research* 66(6):1693–1712.

AUTHOR BIOGRAPHIES

TAEHO KIM is a postdoctoral researcher in the Department of Information Systems, Business Statistics, and Operations Management at the Hong Kong University of Science and Technology (HKUST) Business School. His research interests lie in simulation optimization and stochastic system analysis. His website is <https://sites.google.com/view/taehokim> and his email address is thk5594@gmail.com.

DAVID J. ECKMAN is an Assistant Professor in the Wm Michael Barnes '64 Department of Industrial and Systems Engineering at Texas A&M University. His research interests deal with optimization and output analysis for stochastic simulation models. He is a co-creator of SimOpt, a testbed of simulation optimization problems and solvers. His e-mail address is eckman@tamu.edu.