Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

FINITE BUDGET ALLOCATION IMPROVEMENT IN RANKING AND SELECTION

Xinbo Shi[1,2], Yijie Peng[1,2], and Bruno Tuffin[3]

[1]Wuhan Inst. for Artificial Intelligence & Guanghua School of Mgmt., Peking University, Beijing, CHINA
[2]Xiangjiang Laboratory, Changsha, CHINA
[3]Inria, Univ. Rennes, CNRS, IRISA, Rennes, FRANCE

ABSTRACT

This paper introduces a new perspective on the problem of finite sample Ranking and Selection. An asymptotically equivalent approximation to the probability of correct selection in terms of power series is derived beyond the classic large deviations principle that has been widely adopted for the design and measurement of allocation policies. The novel approximation method provides more information on the finite sample performance of allocation policies, based on which a new finite computing budget allocation policy is proposed. The asymptotically equivalent approximation may also serve as an estimate of policy performance after allocating the samples. We develop a simple finite computing budget allocation policy based on our approximation and carry out experiments in various settings to show its superiority.

1 INTRODUCTION

This paper considers the problem of Ranking and Selection (R&S) that has been actively studied for decades (Bechhofer et al. 1995; Goldsman et al. 1998; Chen and Lee 2011). The classic R&S problem involves selecting the best simulation design with the largest mean output among a finite number of alternatives via Monte Carlo simulation. The primary goal of R&S is more about correct identification rather than achieving sufficiently accurate estimates for mean simulation outputs, which is why R&S is also regarded as an Ordinal Optimization problem (Ho et al. 1992). A central issue therein arises as balancing the cost of computationally expensive simulation and reducing uncertainty in comparing alternatives. Throughout this paper, we will assume a fixed-budget setting where the total number of available simulation samples for all alternatives is predetermined, and we aim to intelligently allocate the limited budget in order for correct identification. Moreover, the performance of an allocation rule will be evaluated using the well-known probability of correct selection (PCS) when the budget is exhausted.

In recent years, there has been growing research interest in the special case of a limited budget constraint. This interest stems from the increasing demand for practical applications, such as simulating complex systems like socio-cyber-physical systems (Szabo et al. 2023) and digital twin systems (Boschert and Rosen 2016), which has led to a rise in computational intensity. The budgets for these applications tend to be relatively small, driven by the practical need for quick responses. As simulation systems become more complex, the number of decision variables within these systems also scales up, giving rise to large-scale problems (Hong et al. 2021). Consequently, the cost of conducting simulations has further increased, and simulation resources are inadequate. Moreover, the behavior of allocation rules may exhibit essentially different properties with a small budget size. For instance, Peng et al. (2015) describes a low-confidence scenario where the total budget is absolutely small or relatively small, due to large variances and/or small mean gaps between alternatives, and exemplifies a counter-intuitive phenomenon: PCS may not be monotonically increasing with respect to the amount of simulation budget of some alternatives, which does not often appear when the simulation budget is sufficiently large. Therefore, the allocation policies designed for a sufficient budget should be improved for finite budget cases.

An in-depth examination of PCS evaluation is necessary to improve the finite sample performance of R&S procedures. The optimal computing budget allocation (OCBA), initialized by Chen (1995) and Chen et al. (2000), is among the most widely adopted allocation rules in fixed-budget settings. It relies on a lower bound for the PCS derived from the Bonferroni inequality. Chen et al. (2000) derives an approximately optimal sampling ratio, i.e., the fraction of the simulation budget allocated to each alternative, maximizing the lower bound at the point when no simulation samples have been collected. The approximation heavily relies on the assumption that the best alternative should receive a significantly larger budget than others asymptotically. Glynn and Juneja (2004) introduce a large deviations principle (LDP) for PCS, which asserts that PCS converges to 1 exponentially and provides an explicit expression for the convergence rate. They accordingly propose a new rate-optimal allocation (ROA) of OCBA type by maximizing the estimated exponential rate. LDP serves as a theoretical foundation for policy development for various R&S problems (Branke et al. 2007; Fu et al. 2007). The LDP rate has been extensively utilized to measure the asymptotic efficiency of R&S procedures, no matter whether they are directly derived from the LDP rate (Li and Gao 2023) or not (Peng et al. 2016; Chen and Ryzhov 2019; Russo 2020).

In this paper, we introduce a novel approximation method for PCS. Given a sampling ratio, we present an asymptotically equivalent characterization of PCS in the form of series expansions. Our approximation extends the expansion in Bahadur and Rao (1960) for the tail probability of sample means to PCS, which involves simultaneous multiple comparisons among alternatives. We demonstrate that our method provides more information about PCS compared to the LDP rate when the sample size is limited. This holds true even if only the first term in the series expansion is retained, as our approximation reflects the impact of the sampling budget and the number of alternatives on PCS, whereas the LDP rate remains invariant to these variations. We outline an optimal allocation rule and propose a new Finite Computing Budget Allocation (FCBA) policy, which outperforms OCBA and other state-of-the-art benchmarks in empirical comparisons. Our approximation method is extensible and has the potential to enhance the exploration policy for many other problems under the umbrella of R&S, including fixed-precision R&S (Cheng et al. 2023), contextual R&S (Cakmak et al. 2023; Shi et al. 2023), R&S under low-confidence scenarios (Peng et al. 2018), and large-scale R&S problems (Hong et al. 2021; Li et al. 2023). We demonstrate in numerical experiments the capability of our methodology for estimating the PCS after sample allocation.

The organization of this paper is as follows. Section 2 presents the main result compared to LDP. Section 3 sketches the proof and key ideas. Section 4 designs a simple budget allocation policy for Gaussian distributions. Section 5 demonstrates the properties of our new method and Section 6 concludes.

2 SERIES EXPANSION OF PROBABILITY OF CORRECT SELECTION

Consider the classic problem of identifying the best design with the largest mean among k alternatives. Let X_i be the population of simulation outputs from the i-th alternative for $1 \le i \le k$. We assume $m_1 > m_2 \geq \cdots \geq m_k$, where $m_i = \mathbb{E}[X_i]$ represents the mean output. Therefore, the first alternative is the unique best design.

Suppose the total number of simulation budget is equal to $T \geq 1$ and fixed. Let T_i denote the number of simulation samples allocated to the i-th alternative and $p_i := T_i/T$ indicate the sampling ratio for $1 \le i \le k$. We will ignore the minor issue that $T_i = p_i T$ might not be integer valued for an allocation rule with $\boldsymbol{p} = (p_i)_{i \le k}$ fixed as in the literature (Chen et al. 2000; Glynn and Juneja 2004). Additionally, let $X_i^{(t)}$ be the t-th simulation sample drawn from the i-th alternative, and $\bar{X}_i(\tau) := \tau^{-1} \sum_{t=1}^{\tau} X_i^{(t)}$ be the sample mean for the i-th alternative with τ simulated samples. When the budget is exhausted, the alternative with the largest sample mean $i^* := \arg\max_{1 \le i \le k} \bar{X}_i(T_i)$ is identified as the best design. Thus the probability of correct selection is defined as

$$P\{CS\} := P\left(\bar{X}_1(T_1) > \max_{2 \le j \le k} \bar{X}_j(T_j)\right) = 1 - P\left(\bar{X}_1(T_1) \le \max_{2 \le j \le k} \bar{X}_j(T_j)\right).$$

Let $\mathrm{int}(\cdot)$ denote the interior of a set, and $\mathscr{D}(\cdot)$ and $\mathscr{R}(\cdot)$ denote the domain and range of a function, respectively. Define $\Lambda_i(\lambda) := \log \mathbb{E} \exp\{\lambda X_i\}$ as the cumulant generating function (CGF) of population X_i. We reiterate the same assumption in Glynn and Juneja (2004) as follows.

Assumption 1 (Light-tailed distribution) The domain of $\Lambda_i(\cdot)$ has a non-empty interior containing the origin, i.e., $0 \in \mathrm{int}(\mathscr{D}(\Lambda_i))$. Moreover, assume $[m_k, m_1] \subseteq \mathrm{int}(\mathscr{R}(\Lambda_i'))$, where Λ_i' denotes the derivative.

The first statement in Assumption 1 ensures that the underlying distributions of simulation outputs are light-tailed with finite CGFs. As for the second statement, notice that the derivative functions of CGFs, if existing, range in the closure of the support of the underlying distribution. The second statement, informally speaking, rules out the trivial case where certain alternatives are deterministically inferior, i.e., their distributions' supports are strictly bounded above by m_1, or vice versa.

Assumption 1 guarantees the existence of a rate function. We reiterate the result in Glynn and Juneja (2004) to facilitate comparison.

Theorem 1 (LDP, Glynn and Juneja 2004) Under Assumption 1, there exist bivariate functions $G_j : \mathbb{R}_+^2 \to \mathbb{R}_+$ for $2 \leq j \leq k$, satisfying

$$\lim_{T\to\infty} -\frac{1}{T} \log(1 - P\{CS\}) = \min_{2\leq j\leq k} G_j(p_1, p_j).$$

The functions $G_j(\cdot,\cdot)$ are known as rate functions.

The rate functions can be explicitly expressed, in the form of the infimum of Fenchel-Legendre transformations of $\bar{X}_1(T_1) - \bar{X}_j(T_j)$, using the Gartner-Ellis Theorem (Dembo 2009). The closed form can be derived for common distribution families or estimated in a data-driven approach (Chen 2023). However, the validity of the LDP rate does not imply that this seemingly correct asymptotically equivalent characterization of PCS, i.e., $1 - P\{CS\} = \exp\{-T \cdot \min_{2\leq j\leq k} G_j(p_1, p_j)\} \cdot (1 + o(1))$ holds true. To be specific, assume the following functional form of PCS

$$1 - P\{CS\} = a(T) \cdot \exp\{-T \cdot \min_{2\leq j\leq k} G_j(p_1, p_j)\} \cdot (1 + o(1)). \tag{1}$$

Regardless of the form of $a(T)$, as long as it is a function that decays slower than exponential functions asymptotically, i.e., $-\frac{1}{T}\log a(T) \to 0$, $a(T)$ will have no impact on the LDP rate. In other words, the large deviations principle may adopt any functional form of PCS as delineated in (1) with a sub-exponential function $a(T)$.

Motivated by equation (1), we aim to find an asymptotical equivalent expression for PCS. To this end, we make the following assumption.

Assumption 2 The density functions of X_i for $i = 1, \ldots, k$ exist, and have bounded total variation (BTV).

The first statement excludes distributions with discrete or singular components. The BTV assumption ensures that the error introduced by approximating the distribution function of X_i's using the Edgeworth expansion is controllable (see Cramer 1970). Assumption 2 is mild and satisfied by a majority of common statistical models. Our main result then follows.

Theorem 2 Under Assumptions 1 and 2, for any $\ell \in \mathbb{N}$ and $\boldsymbol{p} > 0$ with $\sum_{i=1}^k p_i = 1$, the following expansion is valid with $G_j(\cdot,\cdot)$'s as in Theorem 1:

$$1 - P\{CS\} = \sum_{j=2}^{k} \exp\left\{-T \cdot G_j(p_1, p_j)\right\} \cdot \frac{1}{\sqrt{2\pi} \cdot \lambda_j^* p_j \tilde{\sigma}_{1,j} \sqrt{T}} \cdot \left(1 + \sum_{l=1}^{\ell} \frac{c_{j,l}}{T^l} + \mathscr{O}(T^{-(\ell+1)})\right). \tag{2}$$

The symbols λ_j^*, $\tilde{\sigma}_{1,j}$ and $c_{j,l}$ in Theorem 2, which will be introduced in Section 3, stand for constants that may vary depending on $\boldsymbol{p}$ and the specific problem instance, but remain independent of budget T. Theorem 2 extends Theorem 1 by offering an asymptotically equivalent approximation to PCS. In other words, the ratio of $1 - P\{CS\}$ to the right-hand side of (2) without error terms $\mathscr{O}(1/T^{\ell+1})$ approaches 1 as T tends to infinity. Unlike the LDP rate, the novel approximation to the PCS is a nonlinear function of T. Consequently, the optimal sampling ratio derived from our approximation depends on T. Besides, each summand in Equation (2) reflects a probability of correct binary comparison between the means of the optimal alternative and the corresponding sub-optimal alternative, which jointly affects PCS. For instance, let a be an integer and $0 < s \ll S$ be two reals. Consider a sampling ratio $\boldsymbol{p}^a$ indexed by $2 \leq a < k$ such that $G_j(p_1^a, p_j^a) = s$ for $2 \leq j \leq a$, $G_j(p_1^a, p_j^a) = S$ for $a+1 \leq j \leq k$. Given $\boldsymbol{p}^a$, the ordinal relationships between the best alternative and alternatives 2 to a are harder to tell, while the differences in sample means between the best alternative and alternatives $a+1$ to k are relatively large with a high probability. Loosely speaking, the value of a can be interpreted as the number of competitive alternatives, wherein a *competitive alternative* denotes a design with a comparatively higher likelihood of being optimal. The LDP rate does not depend on a directly but depends on the value of s. In contrast, even if the value of s stays unchanged, an increase in the number of competitive alternatives would result in a decrease in our approximation, reflecting the influence of competitive alternatives on PCS.

Our approximation better captures the behavior of PCS in the finite sample case compared to LDP. Notice that LDP characterizes the exponential rate at which PCS converges. Although the exponential terms in Equation (2) predominate the decrease of $1 - P\{CS\}$ in the asymptotic case, it is the polynomial term involving $T^{-1/2}$ that decays more rapidly in the finite sample case, particularly when $T \cdot G_j(\cdot,\cdot)$ is close to zero. However, this distinction should never be interpreted as a contradiction between our result and Theorem 1, which is in fact a direct corollary from Theorem 2 if the assumptions in the latter are valid. Given $\boldsymbol{p}$, and for any positive number $\varepsilon > 0$, the inequality

$$\exp\{-\varepsilon T\} \leq \frac{1}{\sqrt{2\pi}\cdot\lambda_j^* p_j \tilde{\sigma}_{1,j}\sqrt{T}} \cdot \left(1 + \sum_{l=1}^{\ell} \frac{c_{j,l}}{T^l} + \mathscr{O}(T^{-(\ell+1)})\right) \leq 1.$$

holds true for T sufficiently large because the intermediate term is roughly a polynomial in $T^{-1/2}$, which is a lower-order infinitesimal of $\exp\{-\varepsilon T\}$. Therefore, for T sufficiently large, the right hand side of Equation (2) can be lower bounded by

$$\max_{2\leq j\leq k} \exp\left\{-T\cdot G_j(p_1,p_j)\right\}\exp\{-\varepsilon T\} = \exp\left\{-T\cdot \min_{2\leq j\leq k} G_j(p_1,p_j)\right\}\exp\{-\varepsilon T\},$$

and upper bounded by

$$(k-1)\max_{2\leq j\leq k} \exp\left\{-T\cdot G_j(p_1,p_j)\right\} = (k-1)\exp\left\{-T\cdot \min_{2\leq j\leq k} G_j(p_1,p_j)\right\}.$$

By taking logarithms on these bounds followed by dividing them with $-T$, we see that

$$\min_{2\leq j\leq k} G_j(p_1,p_j) \leq \lim_{T\to\infty} -\frac{1}{T}\log(1-P\{CS\}) \leq \min_{2\leq j\leq k} G_j(p_1,p_j) + \varepsilon.$$

Then Theorem 1 follows from the arbitrariness of ε.

The following example illustrates our approximation to PCS in the widely used Gaussian setting.

Example 1 (Gaussian) Suppose $X_i \sim N(m_i, \sigma_i^2)$ with mean m_i and variance σ_i^2, $i = 1,2,\ldots,k$. Since $\Lambda_i(\lambda) = m_i\lambda + \frac{1}{2}\sigma_i^2\lambda^2$, Assumption 1 is valid. Assumption 2 is also clearly valid.

Define $R_j(p_1,p_j) = (m_1 - m_j)^2/(\sigma_1^2/p_1 + \sigma_j^2/p_j)$ as the signal-to-noise ratio, $j = 2,\ldots,k$. Then we have $G_j(p_1,p_j) = \frac{1}{2}R_j(p_1,p_j)$ and $\lambda_j^* p_j \tilde{\sigma}_{1,j} = \sqrt{R_j(p_1,p_j)}$ for $j = 2,\ldots,k$, and $c_{j,l} = \frac{(-1)^l(2l-1)!!}{R_j^l(p_1,p_j)}$ for $j = 2,\ldots,k$ and $l \geq 1$. Theorem 2 implies that for any $\ell \geq 0$,

$$1 - P\{CS\} = \sum_{j=2}^{k} \exp\left\{-\frac{1}{2}T\cdot R_j(p_1,p_j)\right\}\cdot\frac{1}{\sqrt{2\pi T}\sqrt{R_j(p_1,p_j)}}\cdot\left(1 + \sum_{l=1}^{\ell}\frac{(-1)^l(2l-1)!!}{R_j^l(p_1,p_j)}\frac{1}{T^l} + \mathscr{O}(T^{-(\ell+1)})\right).$$

3 SKETCH OF PROOF OF THEOREM 2

In this section, we will outline the proof of the main result. A rigorous proof will be deferred to an extended version of this work. We will commence by introducing the necessary notations in the large deviations theory. Then, we will delve into the PCS for binary comparisons and extend the result to encompass generic multiple comparisons.

3.1 Large Deviations of Sample Means

We here introduce notations similar to those in Glynn and Juneja (2004). The PCS can be easily expressed by its complementary probability

$$1-P\{CS\}=P\left(\bigcup_{j=2}^{k}\{\bar{X}_1(T_1)\leq\bar{X}_j(T_j)\}\right).$$

It follows from the sub-additivity of the probability measure that

$$\max_{2\leq j\leq k}P\left(\bar{X}_1(T_1)\leq\bar{X}_j(T_j)\right)\leq 1-P\{CS\}\leq\sum_{j=2}^{k}P\left(\bar{X}_1(T_1)\leq\bar{X}_j(T_j)\right)\leq(k-1)\max_{2\leq j\leq k}P\left(\bar{X}_1(T_1)\leq\bar{X}_j(T_j)\right). \quad (3)$$

For $2\leq j\leq k$ fixed, the exponential rate of $P(\bar{X}_1(T_1)\leq\bar{X}_j(T_j))$ can be characterized utilizing the CGF of the vector $(\bar{X}_1(T_1),\bar{X}_j(T_j))$, i.e., $\Lambda^{(T)}(\lambda_1,\lambda_j):=\log\mathbb{E}\exp\{\lambda_1\bar{X}_1(T_1)+\lambda_j\bar{X}_j(T_j)\}$. Moreover, define

$$I_{1,j}(x_1,x_j):=\sup_{\lambda_1,\lambda_j}(\lambda_1x_1+\lambda_jx_j-\lim_{T\to\infty}\frac{1}{T}\Lambda^{(T)}(T\lambda_1,T\lambda_j))=p_1I_1(x_1)+p_jI_j(x_j),$$

where $I_i(x_i):=\sup_\lambda \lambda x_i-\Lambda_i(\lambda)$, $i=1,2,\ldots,k$, denotes the Legendre-Fenchel transformation of Λ_i. Formally, $I_{1,j}(x_1,x_j)$ can be interpreted as the exponential rate of $P(\{\bar{X}_1(T_1)\leq x_1\}\cup\{x_j\leq\bar{X}_j(T_j)\})$, which serves as a lower bound for the focal probability if $x_1\leq x_j$. Following Glynn and Juneja (2004), the LDP rate for comparing $\bar{X}_1(T_1)$ and $\bar{X}_j(T_j)$ is $G_j(p_1,p_j)=I_{1,j}(\mu_j,\mu_j)=\inf_{x_1\leq x_j}I_{1,j}(x_1,x_j)$ with $\mu_j\in[m_j,m_1]$. It follows from the two-sided bound in (3) that $1-P\{CS\}$ decays exponentially with rate $\min_{2\leq j\leq k}G_j(p_1,p_j)$.

3.2 Binary Comparisons

Focus now on the binary comparison $P(\bar{X}_1(T_1)\leq\bar{X}_j(T_j))$. We take inspiration from Bahadur and Rao (1960) who evaluate the tail probability $P(\bar{X}(T)\geq x)$ for sample means utilizing a distribution shifting technique. Suppose X and Z are two random variables with probability measure ν_X and ν_Z, respectively, and are absolutely continuous with respect to each other. Then for any measurable function f and g, the following equations are valid as long as the expectations on at least one side are well-defined (Yan 1998):

$$\mathbb{E}[f(X)]=\mathbb{E}\left[f(Z)\frac{d\nu_X}{d\nu_Z}(Z)\right],\quad\text{and}\quad\mathbb{E}[g(Z)]=\mathbb{E}\left[g(X)\frac{d\nu_Z}{d\nu_X}(X)\right], \quad (4)$$

where $\frac{d\nu_X}{d\nu_Z}$ and $\frac{d\nu_Z}{d\nu_X}$ represent the Radon-Nikodym derivatives. Specially, we will take advantage of the exponential tilting technique which defines Z by $\frac{d\nu_Z}{d\nu_X}(x)=C^{-1}e^{\lambda x}$ for some $\lambda\in\mathbb{R}$ and normalizing constant C. Since μ_Z is a probability measure, it follows by (4) that $1\equiv\mathbb{E}_Z[1]=C^{-1}\mathbb{E}[e^{\lambda X}]$, and thus $C=\mathbb{E}[e^{\lambda X}]$.

For each $1\leq t\leq T_1$, with X substituted by $X_1^{(t)}$ and $\lambda=\lambda_{1,j}^*:=\arg\max_\lambda \lambda\mu_j-\Lambda_1(\lambda)$, we have $C=e^{\Lambda_1(\lambda_{1,j}^*)}$ by the definition of Λ_1. Let $Z_1^{(t)}$ denote the corresponding Z variable constructed by the exponential tilting for each t. Then, the first equality in Equation (4) reduces to

$$\begin{aligned}\mathbb{E}[f(X_1^{(t)})]=e^{\Lambda_1(\lambda_{1,j}^*)}\mathbb{E}\left[f(Z_1^{(t)})e^{-\lambda_{1,j}^*Z_1^{(t)}}\right]&=e^{\Lambda_1(\lambda_{1,j}^*)-\lambda_{1,j}^*\mu_j}\mathbb{E}\left[f(Z_1^{(t)})e^{-\lambda_{1,j}^*(Z_1^{(t)}-\mu_j)}\right]\\&=e^{-I_1(\mu_j)}\mathbb{E}\left[f(Z_1^{(t)})e^{-\lambda_{1,j}^*(Z_1^{(t)}-\mu_j)}\right],\end{aligned} \quad (5)$$

where the last equality follows from the definition of $I_1(\cdot)$.

Now, define $\Omega = \{T_1^{-1}(x_1^{(1)} + \ldots x_1^{(T_1)}) \leq T_j^{-1}(x_j^{(1)} + \ldots x_j^{(T_j)})\} \subseteq \mathbb{R}^{T_1+T_j}$ be a set leading to incorrect binary comparison and $\mathbf{1}_\Omega(X_1^{(1)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})$ denote the associated indicator function. Then the focal probability equals to

$$
\begin{aligned}
&P(\bar{X}_1(T_1) \leq \bar{X}_j(T_j)) \\
=&\mathbb{E}[\mathbf{1}_\Omega(X_1^{(1)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})] \\
=&\mathbb{E}[\mathbb{E}[\mathbf{1}_\Omega(X_1^{(1)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})|X_1^{(2)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)}]] \\
=&\mathbb{E}[e^{-I_1(\mu_j)}\mathbb{E}[\mathbf{1}_\Omega(Z_1^{(1)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})e^{\lambda_{1,j}^*(Z_1^{(1)}-\mu_j)}|X_1^{(2)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)}]] \\
=&e^{-I_1(\mu_j)}\mathbb{E}[\mathbf{1}_\Omega(Z_1^{(1)}, \ldots, X_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})e^{-\lambda_{1,j}^*(Z_1^{(1)}-\mu_j)}] \\
=&e^{-T_1\cdot I_1(\mu_j)}\mathbb{E}\left[\mathbf{1}_\Omega(Z_1^{(1)}, \ldots, Z_1^{(T_1)}, X_j^{(1)}, \ldots, X_j^{(T_j)})e^{-\lambda_{1,j}^*\sum_{t=1}^{T_1}(Z_1^{(t)}-\mu_j)}\right].
\end{aligned}
$$

The first equality holds by definition. The second and the fourth equality follow from the tower law of iterated conditional probability. The third equality follows from (5) with $t = 1$, and the last follows by repeating the process for $t = 2, \ldots, T_1$.

Similarly, we define $Z_j^{(t)}$ by exponential tilting with X replaced by $X_j^{(t)}$, $\lambda = \lambda_j^* := \arg\max_\lambda \lambda\mu_j - \Lambda_j(\lambda)$. It turns out that

$$
\begin{aligned}
&P(\bar{X}_1(T_1) \leq \bar{X}_j(T_j)) \\
=&e^{-T_1\cdot I_1(\mu_j)-T_j\cdot I_j(\mu_j)}\mathbb{E}\left[\mathbf{1}_\Omega(Z_1^{(1)}, \ldots, Z_1^{(T_1)}, Z_j^{(1)}, \ldots, Z_j^{(T_j)})e^{-\lambda_{1,j}^*\sum_{t=1}^{T_1}(Z_1^{(t)}-\mu_j)-\lambda_j^*\sum_{t=1}^{T_j}(Z_j^{(t)}-\mu_j)}\right].
\end{aligned} \tag{6}
$$

Recall that $T_i = p_iT$, and thus $-T_1 \cdot I_1(\mu_j) - T_j \cdot I_j(\mu_j) = -T(p_1I_1(\mu_j) + p_jI_j(\mu_j)) = -T \cdot G_j(p_1, p_j)$. It can also be checked that $p_1\lambda_{1,j}^* + p_j\lambda_j^* = 0$ by the definition of μ_j, and consequently,

$$
-\lambda_{1,j}^*\sum_{t=1}^{T_1}(Z_1^{(t)} - \mu_j) - \lambda_j^*\sum_{t=1}^{T_j}(Z_j^{(t)} - \mu_j) = -\lambda_j^*p_jT\left(T_j^{-1}\sum_{t=1}^{T_j}(Z_j^{(t)} - \mu_j)) - T_1^{-1}\sum_{t=1}^{T_1}(Z_1^{(t)} - \mu_j)\right).
$$

Let H_j denote the term in the parenthesis, then we see that

$$
P(\bar{X}_1(T_1) \leq \bar{X}_j(T_j)) = e^{-T\cdot G_j(p_1,p_j)}\mathbb{E}\left[\mathbf{1}\{H_j \geq 0\}e^{-\lambda_j^*p_j\sqrt{T}\cdot\sqrt{T}H_j}\right].
$$

Now, it remains to characterize the expectation. Let $f_{\tilde{H}_j}(x)$ denote the probability density function of $\tilde{H}_j = \sqrt{T}H_j$, and $\mathscr{F}$ denote the Fourier transformation operator. Immediately, by the definition of H_j,

$$
\mathscr{F}f_{\tilde{H}_j}(\tau) := \int e^{i\tau x}f_{\tilde{H}_j}(x)dx = \mathbb{E}\left[e^{i\tau\tilde{H}_j}\right] = \mathbb{E}\left[\exp\left\{i\tau\sqrt{T}(T_j^{-1}\sum_{t=1}^{T_j}Z_j^{(t)} - T_1^{-1}\sum_{t=1}^{T_1}Z_1^{(t)})\right\}\right].
$$

It follows again by (4) that

$$
\mathscr{F}f_{\tilde{H}_j}(t) = \exp\{T_j\Lambda_j(\lambda_j^* + i\tau p_j^{-1}T^{-1/2}) + T_1\Lambda_1(\lambda_1^* - i\tau p_1^{-1}T^{-1/2})\}.
$$

Let $g(x) = \mathbf{1}\{x \geq 0\}e^{-\lambda_j^*p_j\sqrt{T}\cdot x}$. Then, it follows from the Parseval–Plancherel identity that

$$
\mathbb{E}\left[\mathbf{1}\{H_j \geq 0\}e^{-\lambda_j^*p_j\sqrt{T}\cdot\sqrt{T}H_j}\right] = \int g(x)\cdot f_{\tilde{H}_j}(x)dx = \frac{1}{2\pi}\int \text{conj}(\mathscr{F}g(\tau))\cdot\mathscr{F}f_{\tilde{H}_j}(\tau)d\tau.
$$

Since the power series expansions of $\mathscr{F}g(\tau)$, and $\Lambda_1(\lambda_1^*+\cdot)$ and $\Lambda_j(\lambda_j^*+\cdot)$ in the expression of $\mathscr{F}f_{\tilde{H}_j}(\tau)$, are easily available, we can easily arrive at a power series expansion for the expectation in (6) by putting pieces together. The calculations and the truncation of the series will be deferred to the extended version of this work. The following proposition concludes the result in Section 3.2.

Proposition 1 (Binary Comparisons) Under Assumptions 1 and 2, the probability of incorrect binary comparison decays exponentially. Specially, for $2 \le j \le k$ and $\ell \ge 1$, there exists constants $c_1, \dots, c_\ell$, satisfying

$$P(\bar{X}_1(T_1) \le \bar{X}_j(T_j)) = \exp\{-T \cdot G_j(p_1, p_j)\} \cdot \frac{1}{\sqrt{2\pi} \cdot \lambda_j^* p_j \tilde{\sigma}_{1,j} \sqrt{T}} \cdot \left(1 + \frac{c_1}{T} + \cdots + \frac{c_\ell}{T^\ell} + \mathscr{O}(T^{-(\ell+1)})\right),$$

where $\tilde{\sigma}_{1,j} = \sqrt{\mathrm{var}(X_1)/p_1 + \mathrm{var}(X_j)/p_j}$.

3.3 Multiple Comparisons

Inequality (3) together with Proposition 1 does not yield Theorem 2 directly. We can make use of the inclusion-exclusion principle:

$$\sum_{j=2}^{k} P(\bar{X}_1(T_1) \le \bar{X}_j(T_j)) - \sum_{2 \le j_1 < j_2 \le k} P(\bar{X}_1(T_1) \le \bar{X}_{j_1}(T_{j_1}) \wedge \bar{X}_{j_2}(T_{j_2})) \le 1 - P\{CS\} \le \sum_{j=2}^{k} P(\bar{X}_1(T_1) \le \bar{X}_j(T_j)). \tag{7}$$

In other words, the sum of the probabilities of incorrect binary comparisons overestimates $1-P\{CS\}$ by at most the probability of simultaneous incorrect binary comparisons. However, we can show that the overestimation is negligible.

Proposition 2 (Simultaneous Comparisons) Under Assumptions 1 and 2, the probability of simultaneous incorrect binary comparisons decays exponentially. For any $2 \le j_1 < j_2 \le k$, there exists a rate function $G_{j_1,j_2}(p_1, p_{j_1}, p_{j_2})$ such that

$$P(\bar{X}_1(T_1) \le \bar{X}_{j_1}(T_{j_1}) \wedge \bar{X}_{j_2}(T_{j_2})) = \exp\{-T \cdot G_{j_1,j_2}(p_1, p_{j_1}, p_{j_2})\} \cdot \mathscr{O}(T^{-1/2}).$$

Moreover, the rate functions satisfy

$$G_{j_1,j_2}(p_1, p_{j_1}, p_{j_2}) > G_{j_1}(p_1, p_{j_1}) \wedge G_{j_2}(p_1, p_{j_2}).$$

The argument is similar to that for binary comparisons and is skipped. Proposition 2 implies that $P(\bar{X}_1(T_1) \le \bar{X}_{j_1}(T_{j_1}) \wedge \bar{X}_{j_2}(T_{j_2}))$ is exponentially negligible to as least one of $P(\bar{X}_1(T_1) \le \bar{X}_{j_1}(T_{j_1}))$ and $P(\bar{X}_1(T_1) \le \bar{X}_{j_2}(T_{j_2}))$. Putting Propositions 1 and 2 and (7) together, we see that Theorem 2 is valid.

4 FINITE SAMPLE BUDGET ALLOCATION

It takes the explicit form of $G_j(\cdot,\cdot)$, $\lambda_j^* p_j \tilde{\sigma}_{1,j}$ and $c_{j,l}$'s in Theorem 2 to develop allocation policies. Henceforth, we will stick to the Gaussian case in Example 1 for policy development.

For $\ell \ge 0$, define $U_\ell : \mathbb{R}_+^2 \to \mathbb{R}_+$ by $U_\ell(x) = \exp\{-\frac{1}{2}Tx - \frac{1}{2}\ln x\} \cdot (1 + \sum_{l=1}^{\ell} \frac{(-1)^l (2l-1)!!}{x^l} \frac{1}{T^l})$. We adhere to the convention that $\sum_{l=1}^{0} \cdot = 0$. Then $V_\ell(\boldsymbol{p}) := \sum_{j=2}^{k} U_\ell(R_j(p_1, p_j))$ is an approximation to $1-P\{CS\}$ of order ℓ by our main result. It is worth mentioning that the constants $c_{j,l}$ are typically computationally intractable except for the Gaussian case. However, we will empirically show that the approximation of order 0 is good enough for static allocation rules. In the following, we prove the convexity property of $V_\ell(\boldsymbol{p})$ by utilizing the concavity of $R_j(\cdot,\cdot)$.

Proposition 3 (Convexity) For any even $\ell \ge 0$, $V_\ell(\boldsymbol{p})$ is a convex function of $\boldsymbol{p}$. For any odd $\ell \ge 1$, $V_\ell(\boldsymbol{p})$ is asymptotically almost convex of $\boldsymbol{p}$, in a sense that there exists a sequence of sets $(E_T)_{T\ge 1}$, such that the complements of E_T in $\{\boldsymbol{p} \ge 0 : p_1 + \cdots + p_k = 1\}$ converge to the empty set as $T \to \infty$, and for any $T \ge 1$, $V_\ell(\boldsymbol{p})$ is convex in E_T.

Proof. Since the conic combination of convex functions is still convex, it suffices to show that each summand in $V_\ell(\boldsymbol{p})$ is convex of $\boldsymbol{p}$. Fix any alternative j ranging from 2 to k. It is well-known that $R_j(p_1,p_j)$ is a concave function of (p_1,p_j) (Glynn and Juneja 2004). After some calculations, we see that

$$\begin{aligned}U_\ell'(x) &= \exp\left\{-\frac{1}{2}Tx-\frac{1}{2}\ln x\right\}\cdot\left(-\frac{1}{2}(T+\frac{1}{x})-\frac{1}{2}\sum_{l=0}^{\ell-1}\frac{(-1)^{l+1}(2l+1)!!}{x^{l+1}}\frac{1}{T^l}+\frac{1}{2}\sum_{l=1}^{\ell}\frac{(-1)^{l+1}(2l+1)!!}{x^{l+1}}\frac{1}{T^l}\right)\\ &= \exp\left\{-\frac{1}{2}Tx-\frac{1}{2}\ln x\right\}\cdot\left(-\frac{1}{2}T+\frac{1}{2}\frac{(-1)^{\ell+1}(2\ell+1)!!}{x^{\ell+1}}\frac{1}{T^\ell}\right).\end{aligned} \tag{8}$$

For even $\ell\geq 0$, U_ℓ' is strictly negative and increasing in $x>0$, and thus U_ℓ is strictly decreasing and convex. Therefore, it follows that $U_\ell(R_j(p_1,p_j))$ is a convex function of $\boldsymbol{p}$ from the property that the composition of a concave function and a decreasing convex function is convex (Boyd and Vandenberghe 2004). For odd $\ell\geq 1$, we have

$$U_\ell''(x)=\exp\left\{-\frac{1}{2}Tx-\frac{1}{2}\ln x\right\}\cdot\left(\frac{1}{4}\frac{T}{x}+\frac{1}{4}T^2-\frac{1}{4}\frac{(2\ell+1)!!}{x^{\ell+1}}\frac{1}{T^{\ell-1}}-\frac{1}{4}\frac{(2\ell+3)!!}{x^{\ell+2}}\frac{1}{T^\ell}\right).$$

It can be seen that $U_\ell'(x)<0$ and $U_\ell''(x)>0$ for $x\in\left[((2\ell+3)!!)^{1/(\ell+1)}/T,\infty\right)$. Again, it follows from Boyd and Vandenberghe (2004) that $U_\ell(R_j(p_1,p_j))$ is convex in the domain $\{(p_1,p_j)\in\mathbb{R}_+^2 : R_j(p_1,p_j)\geq ((2\ell+3)!!)^{1/(\ell+1)}/T\}$. This completes the proof. □

This allows us to formulate the optimal allocation problem that minimizes $V_\ell(\boldsymbol{p})$ in $\{\boldsymbol{p}\geq 0 : p_1+\cdots+p_k=1\}$ as an (asymptotic) convex program. The following proposition characterizes a necessary optimality condition.

Proposition 4 (Necessary Condition) If $\boldsymbol{p}\in\{\boldsymbol{p}\geq 0 : p_1+\cdots+p_k=1\}$ minimizes $V_\ell(\boldsymbol{p})$, then

$$\begin{cases} U_\ell'(R_i(p_1,p_i))\cdot R_i(p_1,p_i)\frac{\sigma_i^2/p_i^2}{\sigma_1^2/p_1+\sigma_i^2/p_i}=U_\ell'(R_j(p_1,p_j))\cdot R_j(p_1,p_j)\frac{\sigma_j^2/p_j^2}{\sigma_1^2/p_1+\sigma_j^2/p_j},\ \forall\, 2\leq i,j\leq k,\\ \frac{p_1^2}{\sigma_1^2}=\sum_{j=2}^k\frac{p_j^2}{\sigma_j^2}.\end{cases}$$

According to Proposition 4, we offer a series of finite computing budget allocation (FCBA) policies, which sequentially balance the two sets of necessary conditions. See Algorithm 1.

Algorithm 1 Finite Computing Budget Allocation of Order ℓ (FCBA(ℓ))

Initialize: Set budget T, number of alternatives k, initial samples.
while the budget is not exhausted **do**
  Update sample means $\hat{m}_i$ and variances $\hat{\sigma}_i^2$, and sampling ratios p_j and numbers of past samples T_j.
  Estimate the best alternative using plug-in $j^*\leftarrow\arg\max_{1\leq j\leq k}\hat{m}_j$.
  Use plug-in estimation based on sample means and sample variances to calculate
  $j'\leftarrow\arg\max_{j\neq j^*}\hat{U}_\ell'(\hat{R}_j(p_1,p_j))\cdot\hat{R}_j(p_1,p_j)\cdot\frac{\hat{\sigma}_j^2/p_j^2}{\hat{\sigma}_{j^*}^2/p_{j^*}+\hat{\sigma}_j^2/p_j}$
  if $p_{j^*}^2/\hat{\sigma}_{j^*}^2>\sum_{j\neq j^*}p_j^2/\hat{\sigma}_j^2$ **then**
    Simulate one additional replication for alternative j', **else** for alternative j^*
  end if
end while

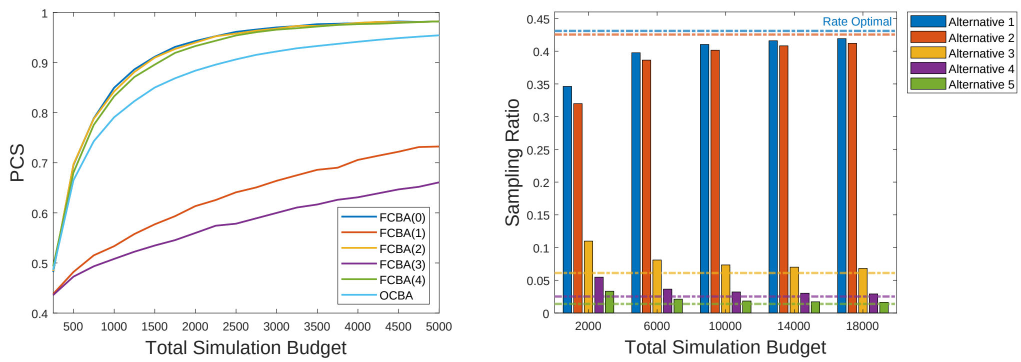

Figure 1: Left: PCS of FCBA(ℓ) versus OCBA under varying number of total simulation budget based on 5,000 macro-replications. Right: Theoretically optimal allocation based on $V_0(\boldsymbol{p})$.

5 EXPERIMENTS

In Section 5, we conduct numerical experiments on FCBA policies using testing instances where simulation outputs are assumed to follow Gaussian distributions. The means of alternatives are predetermined by either (I) the **stepping** setting: $m_i = 0.1*(k+1-i)$ for $i = 1,2,\dots,k$; or (II) the **random** setting: $m_i \overset{i.i.d.}{\sim} 0.1*k*\text{Uniform}([0,1])$. Regarding variances, each instance is configured as either (i) the **equal** setting, where all variances are fixed at 4; or (ii) the **increasing** setting: the alternatives are ranked by means and divided equally into five groups, and the variances of the top 20% of alternatives to the bottom 20% of alternatives are 2, 3, 4, 5, and 6, respectively; or (iii) the **decreasing** setting: the variances of the top 20% of alternatives to the bottom 20% of alternatives are 6, 5, 4, 3, and 2, respectively.

5.1 FCBA(ℓ) versus OCBA

We first test FCBA(ℓ) policies against OCBA. Our comparison involves $k = 50$ alternatives with means in the **stepping** setting and with variances in the **equal** setting set at 1 as described above. The left panel of Figure 1 shows the consistent superiority of FCBA(ℓ) with even ℓ over OCBA. However, FCBA(ℓ) policies with odd ℓ exhibit poor performance. This result is possibly attributed to the non-convex nature of $V_\ell(\boldsymbol{p})$ for odd ℓ, as FCBA(ℓ) allocates samples by following the descending direction of $V_\ell(\boldsymbol{p})$ whereas tracking the descent of a non-convex function does not guarantee reaching to its global minimum. FCBA(0) and FCBA(2) perform comparably, slightly outperforming FCBA(4) with small total budgets. These findings suggest that FCBA(0) suffices for real applications. Therefore, we will focus on the FCBA of order $\ell = 0$. Notice that FCBA simply ignores the probability of simultaneous incorrect selection in (7), according to Proposition 2. This ignored term is actually larger than the higher-order correction terms in absolute values when the budget is limited. Nevertheless, our approximation remains asymptotically valid. To obtain higher-order algorithms, it is necessary to characterize the ignored probability, which is deferred to the extended version of this work.

We proceed by examining the behavior of FCBA(0) through its corresponding theoretically optimal allocation, i.e., the optimal sampling ratio maximizing $V_0(\boldsymbol{p})$ under the presumption of known underlying distributional parameters. The right panel of Figure 1 illustrates the theoretically optimal allocation based on $V_0(\boldsymbol{p})$ for the first five alternatives. Dashed lines denote the rate-optimal allocation that maximizes the LDR rate. With a small budget of T, FCBA(0) is more conservative, inclined to allocate fewer samples to the best alternative compared to OCBA. However, as T approaches infinity, the allocation of FCBA(0) converges to that of OCBA, empirically demonstrating the asymptotic rate-optimality of FCBA(0).

5.2 Finite Sample Performances

The finite sample performance of FCBA(0) consistently outperforms all modern R&S procedures under comparison, including equal allocation (EA), OCBA, ROA, the approximately optimal allocation policy (AOAP, Peng et al. 2016), and mCEI (Chen and Ryzhov 2019), a variant of the expected improvement method tailored for rate optimality. Table 1 presents the final PCS for every combination of policies and problem instances. Each instance contains $k = 50$ alternatives, with a simulation budget $T = 1,000$.

Table 1: PCS estimated by $100,000$ macro-replications for all instances with $k = 50$ and $T = 1,000$.

Instances	**FCBA**(0)	AOAP	ROA	OCBA	mCEI	EA
Stepping + Equal	**0.5767**	0.5537	0.5487	0.5641	0.5652	0.3027
Stepping + Increasing	**0.7042**	0.6702	0.6659	0.6847	0.6669	0.3805
Stepping + Decreasing	**0.5050**	0.4902	0.4851	0.4975	0.5045	0.2641
Random + Equal	**0.9155**	0.8140	0.7899	0.8129	0.8744	0.4897
Random + Increasing	**0.9698**	0.8965	0.8838	0.8993	0.9308	0.6819
Random + Decreasing	**0.8468**	0.7502	0.7206	0.7459	0.8194	0.3971

Notice that FCBA(0) is developed to maximize the final PCS and thus depends on the total budget. The following experiment shows the PCS of FCBA(0) against compared policies if the identification of the best alternative is made before the simulation budget is exhausted. The success of FCBA(0) can also be recognized in Figure 2, with the only exception in the comparison against AOAP in the special case when the simulation is stopped prematurely. It is because AOAP targets a nearly optimal dynamic policy while FCBA merely targets the final PCS. The good performance of FCBA for cases when the budget is nearly exhausted allows for the robust extensions of FCBA policies to applications where the total budget is not completely predetermined, for example, in a fixed-precision setting. Figure 2 also demonstrates the scalability of FCBA in an instance with means in the **stepping** setting as described above and with variances in the **equal** setting set at 1 as described above. Three initial alternatives for each alternative are counted in the total budget. Even if the simulation budget is given fixed, the difference in PCS between FCBA(0) and other benchmarks increases as the size of the problem grows. To be specific, the gaps in the final PCS are equal to 0.26%, 0.62%, 1.19%, 2.10% for $k = 50, 100, 200, 400$, respectively.

Finally, we demonstrate the potential of our approximation as a rough estimate for the PCS. In the $k = 100$ instance mentioned earlier, we implement FCBA(0) across various total budgets, ranging from $T = 1,000$ to $5,000$. Upon exhausting the simulation budget, we construct a post hoc estimation for the PCS using $1 - \hat{V}_0(\boldsymbol{p})$, where $\hat{V}_0(\boldsymbol{p}) := \sum_{j' \neq j^*} \exp\{-\frac{1}{2}\hat{R}_{j'}(p_{j^*}, p_{j'})T - \frac{1}{2}\ln \hat{R}_{j'}(p_{j^*}, p_{j'})\}$. Therein, $j^* := \arg\max_{1 \leq j \leq k} \hat{m}_i$, $\hat{R}_{j'}(p_{j^*}, p_{j'}) := (m_{j^*} - m_{j'})^2 / (\hat{\sigma}^2_{j^*}/p_{j^*} + \hat{\sigma}^2_{j'}/p_{j'})$, and $\hat{m}_i$'s and $\hat{\sigma}_i^2$'s represent sample means and variances, respectively. Consistent with our findings from the initial experiment, it turns out as in Table 2 that $\hat{V}_0(\boldsymbol{p})$ tends to be conservative, underestimating the PCS. However, the accuracy of this estimation improves with the sample size growing.

Table 2: The true PCS and the statistical characteristics of the estimation $\hat{V}_0(\boldsymbol{p})$, estimated via $100,000$ macro-replications.

T	1,000	2,000	3,000	4,000	5,000
True PCS	0.8485	0.9437	0.9706	0.9789	0.9826
Mean Estimation	0.6381	0.8336	0.9205	0.9605	0.9801
$\pm$ Standard Error	(0.3586)	(0.2667)	(0.1799)	(0.1199)	(0.0789)

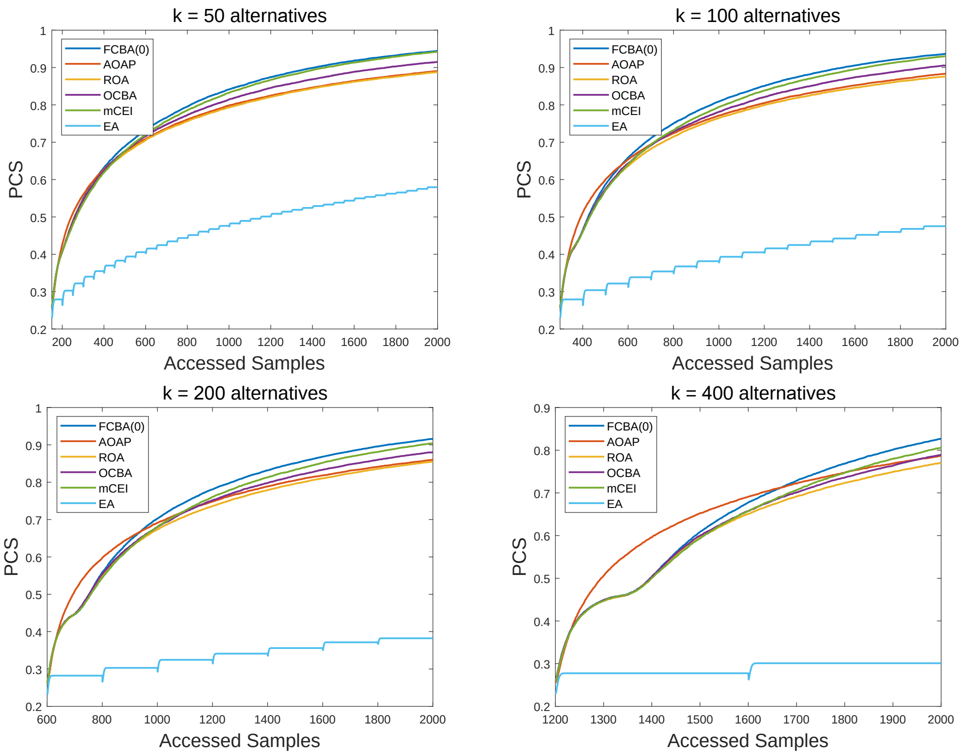

Figure 2: PCS before the budget is exhausted estimated by $100,000$ macro-replications.

6 CONCLUSION

We offer a new method to approximate the performance metric for R&S policies, i.e., the probability of correct selection. The new method, in the form of the Edgeworth expansion, extends the classic LDP rate to an asymptotical equivalent expression for PCS, allowing for finer characterization of the finite sample performance. Based on the naive inclusion-exclusion principle, a simple approximation to PCS is available. Although the approximation is conservative, the consequent allocation policy is empirically efficient.

ACKNOWLEDGMENTS

The work was supported in part by the National Natural Science Foundation of China (NSFC) under Grants 72325007, 72250065, and 72022001.

REFERENCES

Bahadur, R. R. and R. R. Rao. 1960. "On deviations of the sample mean". *The Annals of Mathematical Statistics* 31(4):1015–1027.

Bechhofer, R., T. Santner, and D. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. Wiley Series in Probability and Statistics. Wiley.

Boschert, S. and R. Rosen. 2016. "Digital twin—the simulation aspect". *Mechatronic Futures: Challenges and Solutions for Mechatronic Systems and Their Designers*:59–74.

Boyd, S. and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.

Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a selection procedure". *Management Science* 53(12):1916–1932.

Cakmak, S., Y. Wang, S. Gao, and E. Zhou. 2023. "Contextual Ranking and Selection with Gaussian Processes and OCBA". *ACM Trans. Model. Comput. Simul.*.

Chen, C.-H. 1995. "An effective approach to smartly allocate computing budget for discrete event simulation". In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, Volume 3, 2598–2603.

Chen, C.-H. and L. H. Lee. 2011. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*, Volume 1. World scientific.

Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. "Simulation budget allocation for further enhancing the efficiency of ordinal optimization". *Discrete Event Dynamic Systems: Theory and Applications* 10(3):251–270.

Chen, Y. 2023. "Data-Driven Optimal Allocation for Ranking and Selection under Unknown Sampling Distributions". In *2023 Winter Simulation Conference (WSC)*, 3382–3387 https://doi.org/10.1109/WSC60868.2023.10407177.

Chen, Y. and I. O. Ryzhov. 2019. "Complete expected improvement converges to an optimal budget allocation". *Advances in Applied Probability* 51(1):209–235.

Cheng, Z., J. Luo, and R. Wu. 2023. "On the finite-sample statistical validity of adaptive fully sequential procedures". *European Journal of Operational Research* 307(1):266–278.

Cramer, H. 1970. *Random Variables and Probability Distributions*, 70–88. Cambridge Tracts in Mathematics. Cambridge University Press.

Dembo, A. 2009. *Large Deviations Techniques and Applications*. Springer.

Li et al. 2023. "The (Surprising) Sample Optimality of Greedy Procedures for Large-Scale Ranking and Selection". *arXiv preprint arXiv: 2303.02951*.

Fu, M. C., J. Q. Hu, C. H. Chen, and X. Xiong. 2007. "Simulation allocation for determining the best design in the presence of correlated sampling". *INFORMS Journal on Computing* 19(1):101–111.

Glynn, P. and S. Juneja. 2004. "A large deviations perspective on ordinal optimization". In *Proceedings of the 2004 Winter Simulation Conference, 2004.*, Volume 1. IEEE.

Goldsman, D., B. L. Nelson, and J. Banks. 1998. "Comparing systems via simulation". *Handbook of Simulation*:273–306.

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

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.

Li, Y. and S. Gao. 2023. "Convergence rate analysis for optimal computing budget allocation algorithms". *Automatica* 153:111042.

Peng, Y., C.-H. Chen, M. C. Fu, and J.-Q. Hu. 2015. "Non-monotonicity of probability of correct selection". In *2015 Winter Simulation Conference (WSC)*, 3678–3689 https://doi.org/10.1109/WSC.2015.7408526.

Peng, Y., C.-H. Chen, M. C. Fu, and J.-Q. Hu. 2016. "Dynamic sampling allocation and design selection". *INFORMS Journal on Computing* 28(2):195–208.

Peng, Y., C.-H. Chen, M. C. Fu, and J.-Q. Hu. 2018. "Gradient-Based Myopic Allocation Policy: An Efficient Sampling Procedure in a Low-Confidence Scenario". *IEEE Transactions on Automatic Control* 63(9):3091–3097.

Russo, D. 2020. "Simple Bayesian Algorithms for Best-Arm Identification". *Operations Research* 68(6):1625–1647.

Shi, X., Y. Peng, and G. Zhang. 2023. "Top-Two Thompson Sampling for Selecting Context-Dependent Best Designs". In *2023 Winter Simulation Conference (WSC)*, 3400–3411 https://doi.org/10.1109/WSC60868.2023.10407595.

Szabo, C., R. Castro, J. Denil, and S. M. Sanchez. 2023. "Resilience and Complexity in Socio-Cyber-Physical Systems". In *2023 Winter Simulation Conference (WSC)*, 660–670 https://doi.org/10.1109/WSC60868.2023.10408660.

Yan, J.-a. 1998. *Lectures on Measure Theory*. Beijing: Science Press. 3rd edition, 2021. (in Chinese).

AUTHOR BIOGRAPHIES

XINBO SHI is a Ph.D. candidate in Guanghua School of Management at Peking University, Beijing, China. His research interest includes simulation optimization and machine learning. His email address is xshi@stu.pku.edu.cn.

YIJIE PENG is an Associate Professor in Guanghua School of Management at Peking University, Beijing, China. His research interests include stochastic modeling and analysis, simulation optimization, machine learning, data analytics, and healthcare. His email address is pengyijie@pku.edu.cn.

BRUNO TUFFIN received his PhD degree in applied mathematics from the University of Rennes 1, France, in 1997. Since then, he has been with Inria in Rennes. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems and telecommunication-related economical models. He is currently Area Editor for *INFORMS Journal on Computing* and Associate Editor for *ACM Transactions on Modeling and Computer Simulation* and *Queuing Systems*. His email: bruno.tuffin@inria.fr.