Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

RISK-SENSITIVE ORDINAL OPTIMIZATION

Dohyun Ahn

Department of Systems Engineering and Engineering Management
The Chinese University of Hong Kong
Shatin, New Territories, HONG KONG

Taeho Kim

Department of Industrial & Systems Engineering
Texas A&M University
3131 TAMU
College Station, TX 77843, USA

ABSTRACT

We consider the problem of *risk-sensitive ordinal optimization*, which aims to identify the "least risky" system among a finite number of stochastic systems. Each system's riskiness is assumed to be measured by the probability that the system's loss exceeds a common threshold. Since the crude Monte Carlo estimator is highly inefficient in estimating rare-event probabilities, conventional ordinal optimization approaches coupled with that estimator show significant performance degradation in this problem, particularly for sufficiently large loss thresholds. To detour this issue, assuming that the parametric form of the underlying distribution is known, we propose to use the *tail parameter*, a function of distributional parameters, as a surrogate for the loss probability in comparing and ranking systems, which is shown to work well for many well-known distributions. Building upon this observation, we find the optimal computing budget allocation scheme that maximizes the likelihood of identifying the least risky system.

1 INTRODUCTION

In some operational problems, decision makers may prioritize tail risk over expected performance when identifying the best system in a given set of stochastic systems. For example, in the case of autonomous vehicle testing, a risk-sensitive decision maker may wish to learn the likelihood of each version of a vehicle experiencing a fatal accident via sequential testing and then choose the safest version. Also, when selecting a queueing system design, one might first construct a fixed number of designs and then test each design for a certain amount of time to figure out which one is least likely to result in significant delays. Note that events of interest in the above examples (e.g., fatal accidents and large delays) are rare events. Motivated by these potential applications, we study the problem of sequentially allocating a fixed sampling budget, i.e., the number of simulation trials, to a finite number of stochastic systems in order to select the least risky system when each system's characteristics can only be learned via simulation. This problem can be viewed as a risk-sensitive version of the traditional ordinal optimization problems (Glynn and Juneja 2004), and thus, we call it *risk-sensitive ordinal optimization*.

In this paper, we specifically consider the tail probability, i.e., the probability of system losses exceeding a large threshold, as a risk measure of each system. This choice of performance criterion is fairly common in the above-mentioned applications. An intelligent physical system such as a self-driving algorithm is commonly assessed by the probability of fatal accidents represented as the tail probability (O'Kelly et al. 2018; Norden et al. 2019; Arief et al. 2022; Xu et al. 2022). Also, when evaluating queueing system designs, a typical performance criterion is the probability of delays longer than a certain tolerance time (Juneja et al. 2007; Cahen et al. 2018). Since we focus on the case where the threshold for the tail probability is large, the estimation of the probability suffers from the lack of sufficient samples belonging to the target region, and thus, without any structural information about the underlying distribution, constructing a stable estimator requires an extremely large number of simulation efforts. In this regard, distributional information, such

as a parametric family, is typically assumed in the rare-event setting. Our paper also presumes that the parametric structure of the underlying distribution is known, but the parameter values are unknown.

Under this assumption, we show that for many well-known distributions, the said rare-event issue can be circumvented by replacing the tail probability with its surrogate, which we call *tail parameter*, defined as a function of distributional parameters. In other words, we will reformulate our main problem into the ordinal optimization of the tail parameters. To the best of our knowledge, this formulation is the first attempt to consider tail probabilities in the ordinal optimization framework. Based on that, we provide an inference method of the tail parameter and present the tractable dynamic sampling policy that is asymptotically optimal as the sampling budget grows large.

In terms of allocating a fixed sampling budget to reduce the probability of falsely selecting suboptimal systems (PFS), our work is closely related to the optimal computing budget allocation (OCBA) scheme (Chen et al. 2000). The main challenge in this scheme arises in that the PFS does not have an analytical expression in general, so it is routinely replaced with its proper approximation. Among several approximation methods of the PFS, we take the large-deviations-based approach proposed by Glynn and Juneja (2004). This approach provides a closed-form expression of the exponential decay rate of PFS that can be easily optimized by solving a convex optimization problem, and thus, it has been leveraged to develop sampling procedures using plug-in estimators inferred from the observed data. Due to its tractability, a number of follow-up studies have appeared, focusing on constrained ranking and selection (Pasupathy et al. 2014), near-optimality via two-moment approximation (Shin et al. 2018), contextual ranking and selection (Gao et al. 2019), top-two sampling (Russo 2020), feature-based selection (Ahn and Shin 2020; Ahn et al. 2023), and selection of the most probable best under input uncertainty (Kim et al. 2022). Recently, Chen and Ryzhov (2022) propose the *balancing optimal large deviations* (BOLD) algorithm and demonstrate that the algorithm asymptotically achieves the optimality condition for a general sampling distribution.

Having said that, all the aforementioned studies focus on the mean-based performance measure, and the tail-based measure has received relatively little attention in this literature despite its practical relevance. Only a few papers investigate the problem of quantile-based ordinal optimization (Bekki et al. 2007; Batur and Choobineh 2010; Pasupathy et al. 2010; Batur and Choobineh 2021; Peng et al. 2021; Shin et al. 2022). Although the methods in these papers perform well for moderate quantiles, they can hardly be used in the rare-event-focused applications (i.e., extreme quantiles) due to the said issue of insufficient samples. In contrast, by focusing on tail probabilities rather than quantiles, we develop a novel approach to address this issue using tail probability asymptotics, which has not been fully explored in the literature.

The rest of the paper is organized as follows. Section 2 formulates our risk-sensitive ordinal optimization problem using the theory of large deviations and rigorously discusses why conventional methods fail to guarantee good performance in this problem. In Section 3, we find the tail parameter for various distributions that are commonly assumed for system losses and use it to characterize asymptotically optimal allocation. Numerical results are provided to validate our analysis in Section 4. Section 5 concludes this paper. All proofs can be found in Appendix A.

2 PROBLEM FORMULATION

Assume that there exist k competing systems indexed by $i = 1, \ldots, k$. For each $i \in \{1, \ldots, k\}$, we denote by X_i a continuous random variable representing system i's loss whose distribution is a priori unknown. We define the risk measure of system i as the probability of its loss exceeding a large threshold v, i.e., $p_i(v) := \mathrm{P}(X_i > v)$, where the threshold v is commonly applied to all systems. One may understand the threshold v as a critical point triggering a catastrophic event or a maximum acceptable loss. Throughout this paper, we implicitly assume that the parametric structure of $X_1, \ldots, X_k$ is known but the corresponding parameter values are unknown in advance. We denote by b_v the index of the optimal system b_v such that $p_{b_v}(v) < \min_{i \neq b_v} p_i(v)$, i.e., the system with the smallest risk among the k systems. Our main premise is that $p_i(v)$ can be learned from independent and identically distributed (i.i.d) observations drawn from system i. Thus, we define a sampling policy π as a vector $(\pi_1, \ldots, \pi_T)$ whose t-th element $\pi_t \in \{1, 2, \ldots, k\}$

represents the index of the system sampled at time t for each $t \in \{1,\dots,T\}$. For each $i \in \{1,2,\dots,k\}$, we denote by $X_i(t)$ as an i.i.d simulation output from system i at time t. Let $\alpha_{i,t}(\pi) := \sum_{s=1}^{t} \mathbb{1}\{\pi_s = i\}/t$ be the sampling ratio of system i up to time t, where $\mathbb{1}\{C\}$ yields 1 if C is true and 0 otherwise. Based on this setup, our goal is to construct a sampling policy that maximizes the likelihood of identifying the optimal system b_ν by judiciously allocating a fixed sampling budget T to each system and sequentially learning the tail behaviors of $X_1,\dots,X_k$.

2.1 Inefficiency of a Nonparametric Approach

Let us temporarily consider a situation where the parametric structure of the underlying distributions is ignored. Then, one natural approach to our problem would be to rely on the sample-mean estimator $p_{i,T}(\nu;\pi)$ for the tail probability $p_i(\nu)$, where

$$p_{i,T}(\nu;\pi) := \frac{\sum_{t=1}^{T} \mathbb{1}\{X_i(t) > \nu\}\mathbb{1}\{\pi_t = i\}}{\sum_{t=1}^{T} \mathbb{1}\{\pi_t = i\}}.$$

In this case, the PFS, a commonly used objective for ordinal optimization, can be defined as

$$\mathrm{P}\left(p_{b_\nu,T}(\nu;\pi) > \min_{i \neq b_\nu} p_{i,T}(\nu;\pi) \right). \tag{1}$$

Since $\mathbb{1}\{X_i > \nu\}$ is a Bernoulli random variable with success probability $p_i(\nu)$, minimizing (1) corresponds to solving a mean-based ordinal optimization problem with Bernoulli observations.

Due to the lack of the closed-form expression for the PFS, the large-deviations-based method of Glynn and Juneja (2004) has been widely adopted to address such a mean-based problem in the literature. They consider a static sampling policy $\pi(\alpha)$ satisfying $\lim_{t\to\infty} \alpha_{i,t}(\pi(\alpha)) = \alpha_i$ for all $i \in \{1,\dots,k\}$, where $\alpha = (\alpha_1,\dots,\alpha_k) \in \Delta$ denotes an allocation vector and $\Delta := \{\alpha \mid \sum_{i=1}^{k} \alpha_i = 1, \alpha_i \geq 0 \text{ for all } i\}$ is a probability simplex, and based on the large deviations principle for the sample-mean estimator, they find the asymptotically optimal allocation vector that maximizes the convergence rate of the PFS as T grows large. According to Example 2 of Glynn and Juneja (2004), the convergence rate of (1) with $\pi = \pi(\alpha)$ can be characterized as follows:

$$\lim_{T\to\infty} \frac{1}{T} \log \mathrm{P}\left(p_{b_\nu,T}(\nu;\pi) > \min_{i \neq b_\nu} p_{i,T}(\nu;\pi) \right) = -\rho^{\mathtt{NP}}(\alpha;\nu), \tag{2}$$

where $\rho^{\mathtt{NP}}(\alpha;\nu) := \min_{i\neq b_\nu} \rho_i^{\mathtt{NP}}(\alpha;\nu)$, and for each $i \neq b_\nu$, $\rho_i^{\mathtt{NP}}(\alpha;\nu)$ is given by

$$\rho_i^{\mathtt{NP}}(\alpha;\nu) = -(\alpha_{b_\nu} + \alpha_i) \log\left((1-p_{b_\nu}(\nu))^{\frac{\alpha_{b_\nu}}{\alpha_{b_\nu}+\alpha_i}} (1-p_i(\nu))^{\frac{\alpha_i}{\alpha_{b_\nu}+\alpha_i}} + p_{b_\nu}(\nu)^{\frac{\alpha_{b_\nu}}{\alpha_{b_\nu}+\alpha_i}} p_i(\nu)^{\frac{\alpha_i}{\alpha_{b_\nu}+\alpha_i}} \right). \tag{3}$$

Thus, the (nonparametric) rate-optimal allocation $\alpha^{\mathtt{NP}}(\nu)$ can be obtained by solving the following optimization problem:

$$\max_{\alpha\in\Delta} \rho^{\mathtt{NP}}(\alpha;\nu), \tag{4}$$

One may expect that the optimal convergence rate in (4) is decreasing in ν since it becomes more difficult to find the optimal system as ν increases due to insufficient samples in the tail region $[\nu,\infty)$. This intuition is formalized in the following theorem.

Theorem 1 Assume that $p_{b_\nu}(\nu)/\min_{i\neq b_\nu} p_i(\nu) \to 0$ as $\nu \to \infty$. Then, $\lim_{\nu\to\infty} \max_{\alpha\in\Delta} \rho^{\mathtt{NP}}(\alpha;\nu) = 0$.

Theorem 1 justifies that in most cases, a large threshold value deteriorates the optimal convergence rate of the PFS. The condition of the theorem is generally satisfied unless the tail probabilities of the best and

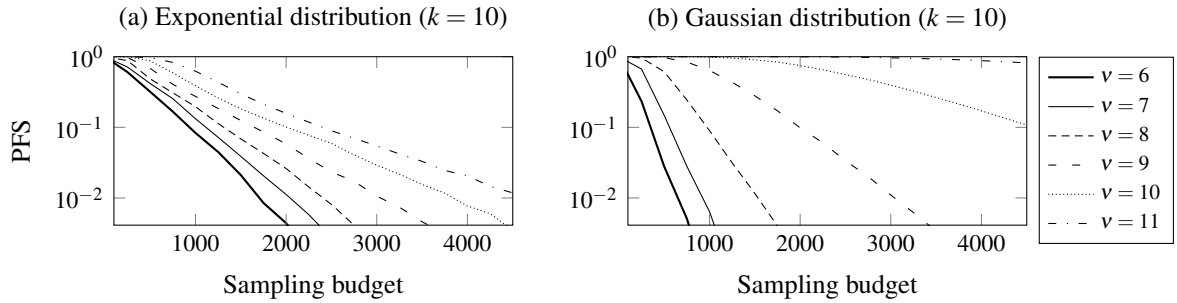

Figure 1: PFS under static allocation rules based on $\alpha^{\mathtt{NP}}(\nu)$ with $\nu = 6,7,\ldots,11$, plotted as a function of the sampling budget in a log-linear scale. For panel (a), $X_1 \sim \text{Exp}(1/2)$ and $X_i \sim \text{Exp}(1/3)$ for all $i \neq 1$. For panel (b), $X_1 \sim N(2,2^2)$ and $X_i \sim N(3,3^2)$ if $i \neq 1$. In both cases, $b_\nu = 1$ for all $\nu \geq 6$.

second-best systems decay at the exactly same rate. Thus, a dynamic sampling policy based on $\rho^{\mathtt{NP}}(\alpha;\nu)$ is expected to perform poorly when ν is large.

To confirm this, we run simple numerical tests by implementing static sampling policies based on $\alpha^{\mathtt{NP}}(\nu)$, the optimal solution to (4). In Figure 1, we compare the PFS with respect to the threshold ν and the sampling budget under two different distribution scenarios: when $X_1,\ldots,X_k$ follow exponential distributions and when they are normally distributed. Note that, for each i, system i's mean and variance in the Gaussian case are set to match those in the exponential case. In both panels in the figure, one can clearly see that the PFS decays at a slower rate as ν increases.

The Gaussian assumption has often been employed in the related literature because it facilitates the construction of tractable and effective sampling policies for mean-based ranking and selection, supported by the central limit theorem and the use of batching (Kim and Nelson 2006). However, as demonstrated by the above results, this approach is not feasible in our setup (i.e., tail-based ranking and selection) due to the light-tailed nature of the Gaussian distribution.

2.2 Problem Reformulation with the Parametric Information

As a response to the performance degradation of the nonparametric approach observed in Section 2.1, we aim to enhance the efficiency of identifying the optimal system by leveraging the parametric information of the underlying distributions in this paper. To that end, we impose a mild condition as follows.

Assumption 1 There exists $\nu_0 > 0$ such that $b = b_\nu$ holds for all $\nu \geq \nu_0$.

Assumption 1 states that the optimal system index remains unchanged for all sufficiently large ν, which is often the case in practice. This implies that for a large threshold ν, we can characterize the optimal system b by comparing a certain distributional parameter of each system that governs the tail behavior of the underlying distribution. For each system i, we denote such a parameter by β_i and call it the *tail parameter* of system i. The construction of β_i depends on the underlying parametric distribution, and the explicit expression of β_i for various parametric distributions will be discussed in the next section.

Accordingly, our main problem is to characterize and estimate the tail parameter β_i for each system and to construct a sequential sampling scheme that ultimately identifies the optimal system

$$b = \underset{1\leq i\leq k}{\arg\min}\, \beta_i \tag{5}$$

based on that estimation. These issues will be investigated in the next two sections. While this paper focuses on ordinal optimization for tail probabilities $\{p_i(\nu)\}_{i=1}^k$, our main analysis based on (5) is not limited to the comparison of these probabilities but can be applied to that of more general tail-based risk metrics, which improves its practicality. However, the related discussion is omitted due to space constraints.

Remark 1 Since $p_i(\nu)$ is a rare-event probability, one could attempt to improve the performance of the method in Section 2.1 by estimating $p_i(\nu)$ via well-known variance reduction techniques, such as importance sampling and control variate methods, based on the known parametric structure. However, these techniques typically require tuning procedures (e.g., selecting the importance sampler and specifying the control variate) that are highly problem-dependent and/or often impractical in our context. Thus, while they may be worth exploring for some problem instances, we do not consider them in this paper.

3 ORDINAL OPTIMIZATION BASED ON TAIL PARAMETERS

In this section, based on tail parameters, we aim to find a new rate function of the PFS, different from (2), for various parametric distributions commonly encountered in the literature. This will enable us to tackle the challenge of insufficient samples in the tail region and to construct the desired sequential sampling policy in Section 4. To accomplish this, we first provide an overview of the theory of large deviations for a generic maximum likelihood estimator (MLE) and the PFS. We then make a significant observation that ordinal optimization under many parametric distributions can be reduced to that under gamma distributions. This allows us to focus on characterizing the tail parameter and the rate function for gamma distributions and extend it to other distributions.

3.1 Large Deviations Preliminaries

Consider a continuous random variable X that has a density function $f(\cdot;\vartheta)$ with a parameter ϑ. We denote by ϑ_T the MLE of ϑ calibrated from T samples independently drawn from the distribution of X. Joutard (2004) shows that under some mild regularity conditions, ϑ_T fulfills the large deviation principle in the following form: for any measurable set $G \subseteq \mathbb{R}$,

$$-\inf_{x\in G^\circ} I_\vartheta(x) \le \liminf_{T\to\infty} \frac{1}{T}\log \mathrm{P}\left(\vartheta_T \in G\right) \le \limsup_{T\to\infty} \frac{1}{T}\log \mathrm{P}\left(\vartheta_T \in G\right) \le -\inf_{x\in \bar{G}} I_\vartheta(x), \tag{6}$$

where G° and $\bar{G}$ denotes the interior and closure of G, respectively, and $I_\vartheta(x)$ is defined as

$$I_\vartheta(x) = -\inf_{u\in\mathbb{R}} \log \mathrm{E}\left[\exp\left(u\frac{\partial}{\partial x}\log f(X;x)\right)\right]. \tag{7}$$

Below we leverage the large deviation principle of a generic MLE in (6) and (7) to obtain the rate function of the PFS. Consider the static sampling policy $\pi(\alpha)$ in Section 2.1 and denote by $\beta_{i,T}$ the MLE of the tail parameter β_i based on sample observations of system i up to time t under the static policy $\pi(\alpha)$. Then, according to the reformulated problem (5) in Section 2.2, the PFS (1) can be rewritten as

$$\mathrm{P}\left(\beta_{b,T} > \min_{i\neq b} \beta_{i,T}\right) \tag{8}$$

for all thresholds $\nu \ge \nu_0$. Since this new PFS satisfies

$$\max_{i\neq b} \mathrm{P}(\beta_{b,T} > \beta_{i,T}) \le \mathrm{P}\left(\beta_{b,T} > \min_{i\neq b} \beta_{i,T}\right) \le (k-1)\max_{i\neq b} \mathrm{P}(\beta_{b,T} > \beta_{i,T}),$$

following the analysis of Glynn and Juneja (2004), we obtain the rate function of (8) as follows:

$$\lim_{T\to\infty} \frac{1}{T}\log \mathrm{P}\left(\beta_{b,T} > \min_{i\neq b} \beta_{i,T}\right) = -\rho(\alpha), \tag{9}$$

where for each $\alpha \in \Delta$, $\rho(\alpha) := \min_{i\neq b} \rho_i(\alpha)$ and

$$\rho_i(\alpha) := \inf_{x\ge\tilde{x}} \{\alpha_b I_{\beta_b}(x) + \alpha_i I_{\beta_i}(\tilde{x})\}. \tag{10}$$

It is worth noting that this newly proposed rate function $\rho(\alpha)$ is independent of the threshold ν, and therefore, it does not degenerate even if ν grows large. As we shall observe later, this property ensures that a sequential sampling policy, which ultimately attains the (parametric) rate-optimal allocation $\alpha^* \in \max_{\alpha\in\Delta}\rho(\alpha)$, exhibits a robust performance against changes in the threshold ν.

3.2 Tail Parameters and Rate Functions for Various Parametric distributions

Building upon the preliminary analysis presented in the previous subsection, the focus now shifts towards characterizing the tail parameter β_i and the corresponding rate function $\rho(\cdot)$ in (9) for a given parametric distribution. Somewhat surprisingly, this subsection reveals that for the majority of well-known distributions, the tail parameter and the rate function can be represented in either of two forms, both of which are derived from gamma distributions. As such, we begin with the analysis with respect to gamma distributions below.

Gamma distributions. Suppose that for $i = 1,\dots,k$, X_i follows a gamma distribution with shape parameter κ_i and scale parameter θ_i, that is, $X_i \sim \text{Gamma}(\kappa_i,\theta_i)$. Then, for each i, the density function $f(\cdot;\kappa_i,\theta_i)$ of X_i is given by $f(y;\kappa_i,\theta_i) = \Gamma(\kappa_i)^{-1}\theta_i^{-\kappa_i}y^{\kappa_i-1}\exp(-y/\theta_i)$, where $\Gamma(\kappa) \coloneqq \int_0^\infty t^{\kappa-1}\exp(-t)dt$ is the gamma function, and the tail probability $p_i(\nu)$ is written as

$$p_i(\nu) = \frac{\Gamma(\kappa_i,\nu/\theta_i)}{\Gamma(\kappa_i)}, \tag{11}$$

where $\Gamma(\kappa,x) \coloneqq \int_x^\infty t^{\kappa-1}\exp(-t)dt$ is the upper incomplete gamma function. According to Chapter 2 of Chaudhry and Zubair (2001), we have the following asymptotic properties of $\Gamma(\kappa,x)$:

$$\lim_{x\to\infty}\frac{\Gamma(\kappa,x)}{x^{\kappa-1}\exp(-x)} = 1 \text{ and } \lim_{x\to 0}\frac{\Gamma(\kappa)-\Gamma(\kappa,x)}{x^\kappa} = \frac{1}{\kappa}. \tag{12}$$

From (11) and the first equation in (12), we obtain that

$$\lim_{\nu\to\infty}\frac{1}{\nu}\log p_i(\nu) = -\frac{1}{\theta_i}, \quad i = 1,\dots,k.$$

This result implies that for each i, the scale parameter θ_i governs the tail behavior of X_i, and hence, we set the tail parameter for X_i as $\beta_i = \theta_i$. Further, we can easily see that Assumption 1 holds in this case. This observation allows us to derive the closed form of the rate function (9) for gamma-distributed systems, which is presented in Theorem 2.

Theorem 2 (Rate function with scale parameters) Suppose that X_i follows $\text{Gamma}(\kappa_i,\theta_i)$ for $i = 1,\dots,k$. Given $\kappa_1,\dots,\kappa_k$, let $\rho^\theta(\alpha)$ denote the rate function (9) of the probability of false selection (8) when $\beta_i = \theta_i$ for each i. Then, the rate function $\rho^\theta(\alpha)$ is given by

$$\rho^\theta(\alpha) = \min_{i\neq b}\left\{\alpha_b I^\theta_{\beta_b}\left(\frac{\alpha_b\kappa_b+\alpha_i\kappa_i}{\alpha_b\kappa_b/\beta_b+\alpha_i\kappa_i/\beta_i};\kappa_b\right) + \alpha_i I^\theta_{\beta_i}\left(\frac{\alpha_b\kappa_b+\alpha_i\kappa_i}{\alpha_b\kappa_b/\beta_b+\alpha_i\kappa_i/\beta_i};\kappa_i\right)\right\}, \tag{13}$$

where $I^\theta_\beta(x;\kappa) = \kappa(x/\beta - \log(x/\beta) - 1)$ for all $x,\kappa,\beta > 0$.

Since $\text{Gamma}(1,\theta)$ is the same with $\text{Exp}(1/\theta)$, identifying the gamma-distributed system with the smallest tail parameter when $\kappa_1 = \cdots = \kappa_k = 1$ is equivalent to determining the exponentially distributed system with the smallest mean. Accordingly, our rate function (13) includes the rate function in Gao and Gao (2016), which is designed for mean-based ordinal optimization with exponentially distributed systems, as a special case.

Extensions to other distributions. The above analysis for gamma distributions is useful in characterizing the tail parameter and the rate function for other distributions. Suppose, for instance, that there exists a function $g:\mathbb{R}\to\mathbb{R}^+$ such that $Y_i = g(X_i) \sim \text{Gamma}(\kappa_i,\theta_i)$ for each i. If $g(x)$ is increasing in x for all

sufficiently large x and $\lim_{x\to\infty} g(x) = \infty$, then $p_i(v) = \mathrm{P}(X_i > v) = \mathrm{P}(g(X_i) > g(v)) = \mathrm{P}(Y_i > g(v))$ for all v large enough. Thus, using the transformation g, one can set the tail parameter $\beta_i = \theta_i$ for each i and use the rate function for gamma-distributed systems, i.e., $\rho^{\theta}(\cdot)$ in (13), to find the rate-optimal allocation. The well-known parametric distributions corresponding to this case include Gaussian, lognormal, Laplace, Pareto, Weibull, and inverse Gaussian distributions.

On the other hand, if $g(x)$ is decreasing in x for all sufficiently large x and $\lim_{x\to\infty} g(x) = 0$, we have

$$p_i(v) = \mathrm{P}(Y_i < g(v)) = \frac{\Gamma(\kappa_i) - \Gamma(\kappa_i, g(v)/\theta_i)}{\Gamma(\kappa_i)}.$$

In this case, however, $\lim_{v\to\infty} g(v) = 0$, and thus, the decay rate of $p_i(v)$ is determined by the asymptotic behavior of gamma distributions near the origin. From the second equation in (12), we observe that

$$\lim_{v\to\infty} \frac{1}{\tilde{g}(v)} \log p_i(v) = -\kappa_i,$$

where $\tilde{g}(v) := -\log g(v) \to \infty$ as $v \to \infty$. This in turn suggests that using the transformation g, we can set the tail parameter β_i as $\beta_i = 1/\kappa_i$ for each i. Moreover, it can be seen that Assumption 1 is satisfied in this setting. Inverse gamma, log gamma, and scaled inverse chi-squared distributions are some of the commonly known distributions that fall under this second category. Motivated by this, the following theorem derives the rate functions for gamma-distributed systems with $\beta_i = 1/\kappa_i, i = 1,\dots,k$.

Theorem 3 (Rate function with shape parameters) Suppose that X_i follows Gamma(κ_i, θ_i) for $i = 1,\dots,k$. Let $\rho^{\kappa}(\alpha)$ denote the rate function (9) of the probability of false selection (8) when $\beta_i = 1/\kappa_i$ for each i. Then, the rate function $\rho^{\kappa}(\alpha)$ is given by

$$\rho^{\kappa}(\alpha) = \min_{i\neq b} \left\{ \alpha_b I^{\kappa}_{\beta_b}\left(\frac{\alpha_b+\alpha_i}{\alpha_b/\beta_b + \alpha_i/\beta_i}\right) + \alpha_i I^{\kappa}_{\beta_i}\left(\frac{\alpha_b+\alpha_i}{\alpha_b/\beta_b + \alpha_i/\beta_i}\right) \right\}, \tag{14}$$

where $I^{\kappa}_{\beta}(x) = \log(\Gamma(1/\beta) - \log\Gamma(1/x) + (1/x - 1/\beta)\psi(1/x)$, and $\psi(x) = (d/dx)\log\Gamma(x)$ is the digamma function.

Based on the above theorem, we can employ the rate function $\rho^{\kappa}(\alpha)$ in (14) to determine the rate-optimal allocation for the distributions classified under the second category. Interestingly, $\rho^{\kappa}(\alpha)$ is independent of the scale parameter θ_i, which differs from $\rho^{\theta}(\alpha)$ that depends on both scale and shape parameters.

Table 1 presents a list of the distributions included in the aforementioned two classes, along with the associated transformation g, the scale and shape parameters induced by g, and the corresponding rate function (either $\rho^{\theta}(\cdot)$ or $\rho^{\kappa}(\cdot)$). It is important to note that the composition or monotone transformation of the distributions listed in Table 1 may also belong to either of the two categories.

4 NUMERICAL EXPERIMENTS

4.1 Tail Parameter Estimation and Dynamic Sampling Policy

The rate functions developed in Section 3 involve unknown distributional parameters and tail parameters that should be sequentially estimated via simulation. Alternatively, we use pseudo rate functions $\rho^{\theta}_t(\cdot)$ and $\rho^{\kappa}_t(\cdot)$ constructed by replacing the unknown parameters in (13) and (14), respectively, with their estimates in each stage t. In particular, Table 1 shows that β_i can be expressed as a function of the distributional parameters, which allows us to find the MLE of β_i using the MLEs of distributional parameters owing to the MLE's invariance property (Casella and Berger 2002). As an example, for inverse Gaussian systems, if we have the stage-t MLEs $\mu_{i,t}$ and $\lambda_{i,t}$ of the distributional parameters μ_i and λ_i respectively, we can directly compute the stage-t MLE of β_i as $\beta_{i,t} = 2\mu^2_{i,t}/\lambda_{i,t}$, and hence, the corresponding pseudo rate function $\rho^{\theta}_t(\cdot)$ in stage t can be found by setting $\beta_i = \beta_{i,t}$, $\kappa_i = 1/2$, and $b = \arg\min_i \beta_{i,t}$ in (13) for each i. Similarly,

Table 1: Tables of distributions that can be suitably transformed into gamma distributions.

Distribution	Density function	$g(x)$	θ_i	κ_i	$\rho(\cdot)$
Gamma	$\frac{1}{\Gamma(\kappa)\theta^\kappa} y^{\kappa-1}\exp\left(-\frac{y}{\theta}\right),\ x\in\mathbb{R}$	x	θ	κ	$\rho^\theta(\cdot)$ $(\beta_i=\theta_i)$
Gaussian	$\frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right),\ x\in\mathbb{R}$	$\frac{(x-\mu)^2}{2}$	σ^2	$\frac{1}{2}$	
Lognormal	$\frac{1}{x\sqrt{2\pi}\sigma}\exp\left(-\frac{(\log x-\mu)^2}{2\sigma^2}\right),\ x\in\mathbb{R}$	$\frac{(\log x-\mu)^2}{2}$	σ^2	$\frac{1}{2}$	
Laplace	$\frac{1}{2\theta}\exp\left(-\frac{\lvert x-\mu\rvert}{\theta}\right),\ x\in\mathbb{R}$	$\lvert x-\mu\rvert$	θ	1	
Pareto	$\frac{\gamma x_m^\gamma}{x^{\gamma+1}},\ x\geq x_m>0$	$\log\left(\frac{x}{x_m}\right)$	γ^{-1}	1	
Weibull	$\frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1}\exp\left(-\left(\frac{x}{\lambda}\right)^k\right),\ x\geq 0$	x^k	λ^k	1	
Inverse Gaussian	$\left(\frac{\lambda}{2\pi x^3}\right)^{1/2}\exp\left(-\frac{\lambda(x-\mu)^2}{2\mu^2 x}\right),\ x\geq 0$	$\frac{(x-\mu)^2}{x}$	$\frac{2\mu^2}{\lambda}$	$\frac{1}{2}$	
Inverse gamma	$\frac{\theta^\gamma}{\Gamma(\gamma)}x^{-\gamma-1}\exp\left(-\frac{\theta}{x}\right),\ x>0$	$1/x$	θ	γ	$\rho^\kappa(\cdot)$ $(\beta_i=1/\kappa_i)$
Log gamma	$\frac{1}{\Gamma(\gamma)\theta^\gamma}\exp\left(-\frac{\exp(-x)}{\theta}\right)\exp(-\gamma x),\ x\in\mathbb{R}$	$\exp(-x)$	θ	γ	
Scaled inverse χ^2	$\frac{(\tau^2\nu/2)^{\nu/2}}{\Gamma(\nu/2)}x^{-\frac{\nu}{2}-1}\exp\left(-\frac{\nu\tau^2}{2x}\right),\ x>0$	$1/x$	$\nu\tau^2/2$	$\nu/2$	

Algorithm 1 Dynamic sampling policy $\pi^\beta(n_0,m)$

Generate n_0 i.i.d. replications and estimate distributional parameters for each system, and set $t=kn_0$
while $t<T$ **do**
  if the underlying systems satisfy $\beta_i=\theta_i$ for each i **then**
    Compute $\beta_{i,t}=\theta_{i,t}$ using the estimated distributional parameters for each i
    Find $\hat{\alpha}=\arg\max_{\alpha\in\Delta}\rho_t^\theta(\alpha)$, where $\rho_t^\theta(\alpha)$ is the stage-t version of $\rho^\theta(\alpha)$
  else if the underlying systems satisfy $\beta_i=\kappa_i^{-1}$ for each i **then**
    Compute $\beta_{i,t}=\kappa_{i,t}^{-1}$ using the estimated distributional parameter for each i
    Find $\hat{\alpha}=\arg\max_{\alpha\in\Delta}\rho_t^\kappa(\alpha)$, where $\rho_t^\kappa(\alpha)$ is the stage-t version of $\rho^\kappa(\alpha)$
  end if
  Generate $[m\hat{\alpha}_i]$ replications for each system i, and set $t=t+m$
end while

for inverse gamma systems, the stage-t MLE is given by $\beta_{i,t}=\gamma_{i,t}^{-1}$, where $\gamma_{i,t}$ is the stage-t MLE of the distributional parameter γ_i, and the associated pseudo rate function $\rho_t^\kappa(\cdot)$ can be obtained by plugging $\beta_i=\beta_{i,t}$ and $b=\arg\min_i\beta_{i,t}$ into (14) for each i.

Given these pseudo rate functions, one can construct sampling policies that sequentially allocate samples based on the optimizers of the pseudo rate functions and thus eventually achieve rate optimality. For our numerical experiments, we use a batch-based allocation rule among those asymptotically optimal policies, which is described in Algorithm 1. While some fully sequential algorithms, such as Balancing Optimal Large Deviations proposed by Chen and Ryzhov (2022), can also be used with some modifications based on our rate functions, in this section, we restrict our focus on numerically validating the superiority of our tail-parameter-based approach over the nonparametric method in the case of large thresholds, rather than comparing the performance of different rate-optimal policies based on tail parameters.

4.2 Numerical Results

In this subsection, as alluded to earlier, our goal is to validate the numerical performance of Algorithm 1. We use two alternative policies for comparison: The first method is a batch-based adaptive version of the

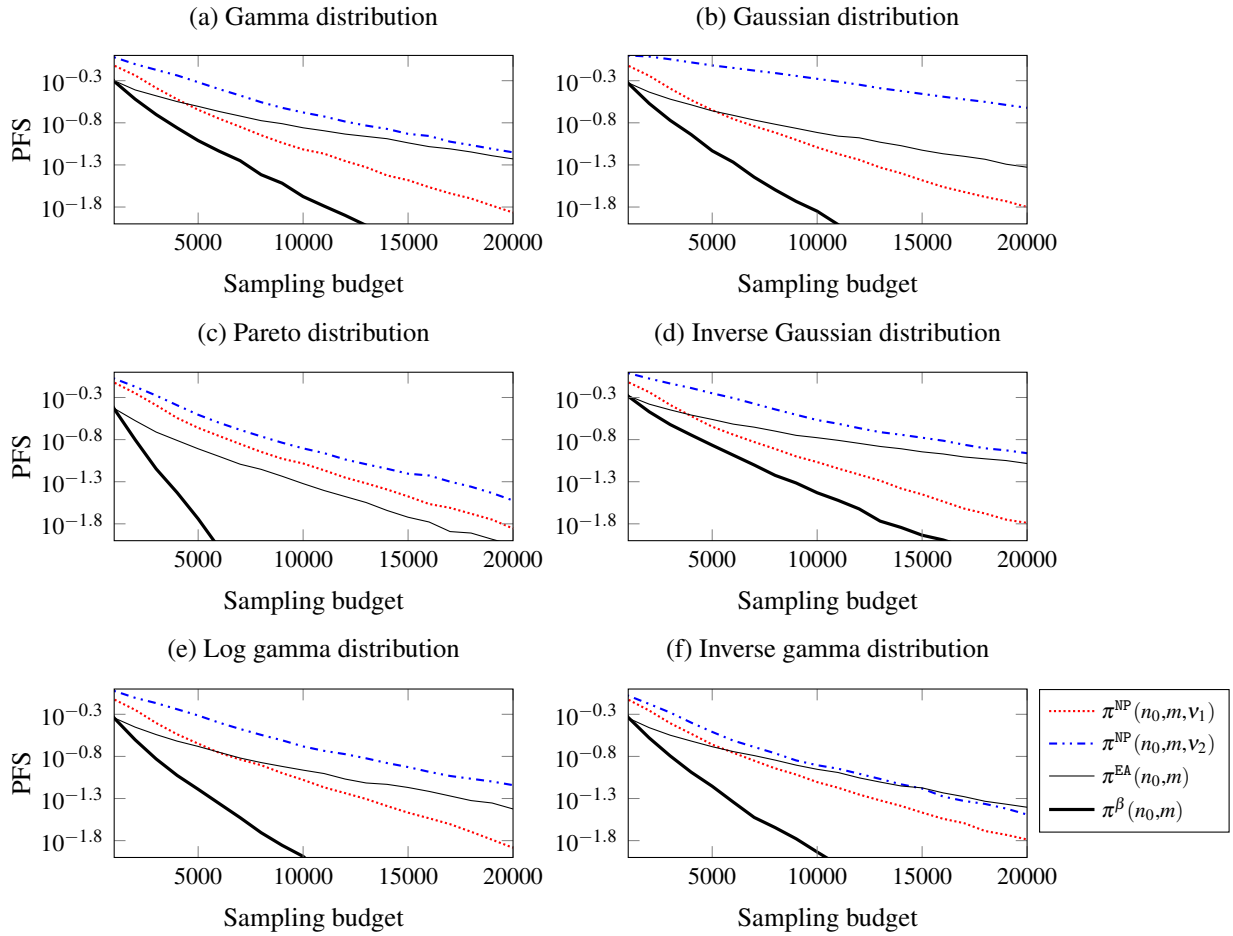

Figure 2: Probability of false selection is plotted as a function of the sampling budget in a log-linear scale for all six testing examples.

nonparametric approach in Section 2.1 (denoted by $\pi^{\mathrm{NP}}(n_0,m,\nu)$) that is constructed in a similar manner to Algorithm 1: we take n_0 initial samples from each system and allocate the sampling budget using the rate function (4) for each batch of size m. The second alternative policy is a batch-based equal allocation rule (denoted by $\pi^{\mathrm{EA}}(n_0,m)$) that takes n_0 initial samples from each system and distributes the sampling budget evenly for each batch of size m.

We fix $k=10$ and consider six different distributions for $\{X_i\}_{1\le i\le k}$: gamma, Gaussian, Pareto, inverse Gaussian, log gamma, and inverse gamma. The first four distributions correspond to the case where $\beta_i=\theta_i$ and $\rho(\alpha)=\rho^{\theta}(\alpha)$, whereas the last two distributions are associated with $\beta_i=1/\kappa_i$ and $\rho(\alpha)=\rho^{\kappa}(\alpha)$. For all sampling policies, we set $n_0=100$ and $m=1{,}000$. For the nonparametric policy $\pi^{\mathrm{NP}}(n_0,m,\nu)$, we consider two scenarios of ν: $\nu_1=2$ and $\nu_2=2.5$. We fix $\beta_i=1/(3-0.2i)$ for $i=1,\ldots,k$ and use the distributional parameters satisfying $p_i(\nu_1)=0.01+0.005i$ for $i=1,\ldots,k$. We use the sample-mean-based PFS in (1) as a performance criterion for the nonparametric policy π^{NP}, and the tail-parameter-based PFS in (8) for our policy π^{β} and the equal allocation rule π^{EA}. We estimate both types of the PFS through Monte Carlo simulation with 10^4 simulation trials.

Figure 2 visualizes the numerical evaluation of four dynamic sampling policies for the six distributional cases. In the figure, it is obvious that our policy π^{β} outperforms the other policies in all cases. In addition, by comparing π^{β} and π^{EA}, we first find that the rate-optimal allocation is more effective than the equal allocation in identifying the system with the smallest tail parameter. Second, when the threshold ν is

large enough, the sample-mean-based PFS is not a good criterion for selecting the optimal system. This is demonstrated by the observation that the performance of $\pi^{\mathtt{NP}}$ is significantly weaker than π^{β} and often comparable to that of $\pi^{\mathtt{EA}}$. However, if ν is not large, $\pi^{\mathtt{NP}}$ may perform better than π^{β} since the performance of $\pi^{\mathtt{NP}}$ improves as ν decreases, while that of π^{β} remains independent of ν. Thus, it would be interesting to characterize the critical points ν_c and $\bar{\nu}_c$ with $\nu_c < \bar{\nu}_c$ such that π^{β} outperforms $\pi^{\mathtt{NP}}$ if $\nu > \bar{\nu}_c$ and the opposite holds if $\nu < \nu_c$, which is left for further research. Third, the performance of π^{β} varies across different distributions, which can be attributed to differences in the rate functions and estimation errors for the MLEs of the distributional parameters.

5 CONCLUDING REMARKS

We formulate the problem of risk-sensitive ordinal optimization based on tail probabilities and investigate the associated sequential sampling rule when the distribution of each system is known except its parameters. Our analysis can be extended in several directions. For example, it would be interesting to explore situations where the distribution is fully unknown or only partially known. Also, employing more general risk measures than tail probabilities for risk-sensitive ordinal optimization could be done via a variant of our approach.

ACKNOWLEDGMENTS

The first author acknowledges the support of the Early Career Scheme from the Research Grants Council of Hong Kong (Project No. 24210420) and CUHK Direct Grant for Research (Project No. 4055206).

A PROOFS OF THEORETICAL RESULTS

Lemma 1 If $q_1, \ldots, q_k$ are positive constants satisfying $q_1 < \min_{i \neq 1} q_i$, then

$$\max_{\alpha \in \Delta} \min_{i \neq 1} \{\alpha_1 q_1 + \alpha_i q_i\} = \max \left\{ q_1, \left(\sum_{i \neq 1} \frac{1}{q_i} \right)^{-1} \right\}. \tag{15}$$

Proof of Lemma 1. Let $q^* := (\sum_{i \neq 1} q_i^{-1})^{-1}$. We note that $\tilde{\alpha} := (1, 0, \ldots, 0)$ and $\hat{\alpha} := (0, q_2^{-1} q^*, \ldots, q_k^{-1} q^*)$ are feasible solutions to the left-hand side of (15) since both vectors belong to the set Δ. Furthermore, it is easy to see that $\min_{i \neq 1} \{\tilde{\alpha}_1 q_1 + \tilde{\alpha}_i q_i\} = q_1$ and $\min_{i \neq 1} \{\hat{\alpha}_1 q_1 + \hat{\alpha}_i q_i\} = q^*$. Therefore, we have

$$\max_{\alpha \in \Delta} \min_{i \neq 1} \{\alpha_1 q_1 + \alpha_i q_i\} \geq \max \{q_1, q^*\}. \tag{16}$$

On the other hand, the left-hand side in (15) is equal to $\max \{z \,|\, \alpha \in \Delta, z \leq \alpha_1 q_1 + \alpha_i q_i \; \forall i = 2, \ldots, k\}$. Since for all $i \neq 1$, the condition $z \leq \alpha_1 q_1 + \alpha_i q_i$ is equivalent to $(z - \alpha_1 q_1)/q_i \leq \alpha_i$, taking the summation of both sides of this inequality over $i \neq 1$ and rearranging the terms yield $z \leq q^* + (q_1 - q^*)\alpha_1$, we obtain

$$\max_{\alpha \in \Delta} \min_{i \neq 1} \{\alpha_1 q_1 + \alpha_i q_i\} \leq \max \{z \,|\, \alpha \in \Delta, z \leq q^* + (q_1 - q^*)\alpha_1\} = \max_{\alpha \in \Delta} \{q^* + (q_1 - q^*)\alpha_1\} = \max\{q_1, q^*\}.$$

Combining this result with (16) completes the proof. □

Proof of Theorem 1. From (3), we observe that for each $i \neq b_\nu$,

$$\rho_i^{\mathtt{NP}}(\alpha; \nu) \leq -(\alpha_{b_\nu} + \alpha_i) \log \left((1 - p_{b_\nu}(\nu))^{\frac{\alpha_{b_\nu}}{\alpha_{b_\nu} + \alpha_i}} (1 - p_i(\nu))^{\frac{\alpha_i}{\alpha_{b_\nu} + \alpha_i}} \right) = \alpha_{b_\nu} q_{b_\nu}(\nu) + \alpha_i q_i(\nu),$$

where $q_j(\nu) = -\log(1 - p_j(\nu))$ for $j = 1, \ldots, k$. We note that for each $j = 1, \ldots, k$,

$$\lim_{\nu \to \infty} \frac{q_j(\nu)}{p_j(\nu)} = 1. \tag{17}$$

Since $0 < q_{b_\nu}(\nu) < \min_{i \neq b_\nu} q_i(\nu)$, by Lemma 1, $\max_{\alpha \in \Delta} \rho^{\mathrm{NP}}(\alpha;\nu) \leq \max\{q_{b_\nu}(\nu), (\sum_{i \neq b_\nu} q_i(\nu)^{-1})^{-1}\}$. Furthermore, according to (17) and by assumption, $q_{b_\nu}(\nu)/\min_{i \neq b_\nu} q_i(\nu) \to 0$ as $\nu \to \infty$. Consequently, we have $(\min_{i \neq b_\nu} q_i(\nu))^{-1} \leq \sum_{i \neq b_\nu} q_i(\nu)^{-1} < q_{b_\nu}(\nu)^{-1}$ for all sufficiently large ν, and thus, we obtain that $\max_{\alpha \in \Delta} \rho^{\mathrm{NP}}(\alpha;\nu) \leq \min_{i \neq b_\nu} q_i(\nu)$. for all sufficiently large ν. Then, by (17), the desired result follows. □

Proof of Theorem 2. Observe that $(\partial/\partial x) \log f(y;\kappa_i,x) = y/x^2 - \kappa_i/x$. Then, we have

$$\log \mathrm{E}\left[\exp\left(u\frac{\partial}{\partial x}\log f(X_i;\kappa_i,x)\right)\right] = \log \mathrm{E}\left[\exp\left(\frac{uX_i}{x^2} - \frac{u\kappa_i}{x}\right)\right] = -\frac{\kappa_i u}{x} - \kappa_i \log\left(1 - \frac{\beta_i u}{x^2}\right), \tag{18}$$

where the second equality exploits that the moment generating function of X_i, $\mathrm{E}[\exp(uX_i)]$, is given by $(1-\beta_i u)^{-\kappa_i}$. From the first-order condition, one can easily see that the infimum of (18) over $u \in \mathbb{R}$ is attained at $u = (x^2 - \beta_i x)/\beta_i$ and thus is equal to $I^{\theta}_{\beta_i}(x;\kappa_i) = \kappa_i(x/\beta_i - \log(x/\beta_i) - 1)$.

Next, since $I^{\theta}_{\beta_i}(x;\kappa_i)$ is increasing when $x > \beta_i$ and is decreasing when $x < \beta_i$, one can confirm that for $I_{\beta_i}(\cdot) = I^{\theta}_{\beta_i}(\cdot)$, the optimal solution in (10) should satisfy $x = \tilde{x}$. Thus, (10) can be recast as

$$\rho^{\theta}_i(\alpha) = \inf_x \left\{ \alpha_b I^{\theta}_{\beta_b}(x;\kappa_b) + \alpha_i I^{\theta}_{\beta_i}(x;\kappa_i) \right\}. \tag{19}$$

By the first-order condition for (19), the optimum is achieved at $x = (\alpha_1\kappa_1 + \alpha_i\kappa_i)/(\alpha_1\kappa_1/\beta_b + \alpha_i\kappa_i/\beta_i)$. □

Proof of Theorem 3. Recall that $\beta_i = 1/\kappa_i$. We first observe that the derivative of the log-likelihood with respect to β_i is given by $(\partial/\partial x)\log f(y;1/x,\theta_i) = -(\log y)/x^2 + (\log\theta_i)/x^2 + x^{-2}\psi(1/x)$. Thus,

$$\begin{aligned}\log \mathrm{E}\left[\exp\left(u\frac{\partial}{\partial x}\log f(X_i;1/x,\theta_i)\right)\right] &= \log \mathrm{E}\left[\exp\left(-u\frac{\log X_i}{x^2} + \frac{u}{x^2}\log\theta_i + \frac{u}{x^2}\psi\left(\frac{1}{x}\right)\right)\right] \\ &= \begin{cases} \log\Gamma\left(-\frac{u}{x^2} + \frac{1}{\beta_i}\right) - \log\Gamma\left(\frac{1}{\beta_i}\right) + \frac{u}{x^2}\psi\left(\frac{1}{x}\right) & \text{if } u < x^2/\beta_i; \\ \infty & \text{otherwise,} \end{cases}\end{aligned} \tag{20}$$

where (20) holds since it is well known that $\mathrm{E}[X_i^{\eta}] = \theta_i^{\eta}\Gamma(\eta + 1/\beta_i)/\Gamma(1/\beta_i)$ if $\eta + 1/\beta_i > 0$ and ∞ otherwise. Clearly, the infimum of the minimization of (20) over $u \in \mathbb{R}$ is attained at $u = -x + x^2/\beta_i < x^2/\beta_i$. Thus, a straightforward calculation yields $I^{\kappa}_{\beta_i}(x) = \log\Gamma(1/\beta_i) - \log\Gamma(1/x) + (1/x - 1/\beta_i)\psi(1/x)$. Using the same argument as in the proof of Theorem 2, we find that for $I_{\beta_i}(\cdot) = I^{\kappa}_{\beta_i}(\cdot)$, the infimum in (10) is achieved at $x = \tilde{x} = (\alpha_b + \alpha_i)/(\alpha_b/\beta_1 + \alpha_i/\beta_i)$. Plugging this into $\alpha_b I^{\kappa}_{\beta_b}(x) + \alpha_i I^{\kappa}_{\beta_i}(\tilde{x})$ completes the proof. □

REFERENCES

Ahn, D., and D. Shin. 2020. "Ordinal Optimization with Generalized Linear Model". In *Proceedings of the 2020 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 3008–3019. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Ahn, D., D. Shin, and A. Zeevi. 2023. "Feature Misspecification in Sequential Learning Problems". Available at SSRN: https://ssrn.com/abstract=3860650.

Arief, M., Z. Cen, Z. Liu, Z. Huang, B. Li, H. Lam, and D. Zhao. 2022. "Certifiable Evaluation for Autonomous Vehicle Perception Systems using Deep Importance Sampling (Deep IS)". In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 1736–1742. Institute of Electrical and Electronics Engineers, Inc.

Batur, D., and F. Choobineh. 2010. "A Quantile-Based Approach to System Selection". *European Journal of Operational Research* 202(3):764–772.

Batur, D., and F. F. Choobineh. 2021. "Selecting the Best Alternative based on Its Quantile". *INFORMS Journal on Computing* 33(2):657–671.

Bekki, J. M., J. W. Fowler, G. T. Mackulak, and B. L. Nelson. 2007. "Using Quantiles in Ranking and Selection Procedures". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1722–1728. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Cahen, E. J., M. Mandjes, and B. Zwart. 2018. "Estimating Large Delay Probabilities in Two Correlated Queues". *ACM Transactions on Modeling and Computer Simulation* 28(1):1–19.

Casella, G., and R. L. Berger. 2002. *Statistical Inference*. 2nd ed. Pacific Grove, CA: Duxbury Thomson Learning.

Chaudhry, M. A., and S. M. Zubair. 2001. *On a Class of Incomplete Gamma Functions with Applications*. London, UK: Chapman & Hall.

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* 10(3):251–270.

Chen, Y., and I. O. Ryzhov. 2022. "Balancing Optimal Large Deviations in Sequential Selection". *Management Science*. Article in advance.

Gao, F., and S. Gao. 2016. "Optimal Computing Budget Allocation with Exponential Underlying Distribution". In *2016 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 682–689. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Gao, S., J. Du, and C.-H. Chen. 2019. "Selecting the Optimal System Design under Covariates". In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, 547–552. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Glynn, P., and S. Juneja. 2004. "A Large Deviations Perspective on Ordinal Optimization". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 577–585. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Joutard, C. 2004. "Large Deviations for M-estimators". *Mathematical Methods of Statistics* 13(2):179–200.

Juneja, S., R. L. Karandikar, and P. Shahabuddin. 2007. "Asymptotics and Fast Simulation for Tail Probabilities of Maximum of Sums of Few Random Variables". *ACM Transactions on Modeling and Computer Simulation* 17(2):1–35.

Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, Volume 13 of *Handbooks in Operations Research and Management Science*, 501–534. Elsevier.

Kim, T., K.-K. Kim, and E. Song. 2022. "Selection of the Most Probable Best". arXiv preprint arXiv:2207.07533.

Norden, J., M. O'Kelly, and A. Sinha. 2019. "Efficient Black-Box Assessment of Autonomous Vehicle Safety". arXiv preprint arXiv:1912.03618.

O'Kelly, M., A. Sinha, H. Namkoong, R. Tedrake, and J. C. Duchi. 2018. "Scalable End-to-End Autonomous Vehicle Testing via Rare-event Simulation". *Advances in Neural Information Processing Systems* 31:9827–9838.

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* 25(1):1–26.

Pasupathy, R., R. Szechtman, and E. Yücesan. 2010. "Selecting Small Quantiles". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yüucesan, 2762–2770. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Peng, Y., C.-H. Chen, M. C. Fu, J.-Q. Hu, and I. O. Ryzhov. 2021. "Efficient Sampling Allocation Procedures for Optimal Quantile Selection". *INFORMS Journal on Computing* 33(1):230–245.

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.

Shin, D., M. Broadie, and A. Zeevi. 2022. "Practical Nonparametric Sampling Strategies for Quantile-Based Ordinal Optimization". *INFORMS Journal on Computing* 34(2):752–768.

Xu, M., P. Huang, F. Li, J. Zhu, X. Qi, K. Oguchi, Z. Huang, H. Lam, and D. Zhao. 2022. "Scalable Safety-Critical Policy Evaluation with Accelerated Rare Event Sampling". In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 12919–12926. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

DOHYUN AHN is an Assistant Professor in the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong. His research interests include, but are not limited to, quantitative finance, Monte Carlo simulation methodologies, decision making under uncertainty, and network analysis in finance and operations. His email address is dohyun.ahn@cuhk.edu.hk and his website is https://sites.google.com/view/dohyun.

TAEHO KIM is a postdoctoral researcher in the Department of Industrial and Systems Engineering at Texas A&M University. His research interests are simulation optimization and data-driven analysis of stochastic models. His email address is thk5594@gmail.com and his webpage can be found at https://sites.google.com/view/taehokim.