Proceedings of the 2007 Winter Simulation Conference
S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, eds.

IMPLICATIONS OF HEAVY TAILS ON SIMULATION-BASED ORDINAL OPTIMIZATION

Mark Broadie
Graduate School of Business
3022 Broadway, 415 Uris Hall
Columbia University
New York, NY 10027, U.S.A.

Minsup Han
Industrial Engineering and Operations Research
W. Mudd Building, 500 West 120th Street
Columbia University
New York, NY 10027, U.S.A.

Assaf Zeevi
Graduate School of Business
3022 Broadway, 406 Uris Hall
Columbia University
New York, NY 10027, U.S.A.

ABSTRACT

We consider the problem of selecting the best system using simulation-based ordinal optimization. This problem has been studied mostly in the context of light-tailed distributions, where both Gaussian-based heuristics and asymptotically optimal procedures have been proposed. The latter rely on detailed knowledge of the underlying distributions and give rise to an exponential decay of the probability of selecting the incorrect system. However, their implementation tends to be computationally intensive. In contrast, in the presence of heavy tails the probability of selecting the incorrect system only decays polynomially, but this is achieved using simple allocation schemes that rely on little information of the underlying distributions. These observations are illustrated via several numerical experiments and are seen to be consistent with asymptotic theory.

1 INTRODUCTION

Suppose there are k systems, and the performance of system i is characterized by an unknown distribution function $F^{(i)}$ with mean μ_i, $i = 1,2,\ldots,k$. For simplicity assume $\mu_1 < \mu_2 \leq \mu_3 \leq \cdots \leq \mu_k$. The user is given a simulation budget n, meaning that n independent samples can be drawn from the distributions $F^{(i)}, i = 1,2,\ldots,k$, with the goal of identifying system 1 as the best. Since the outcome of any sampling procedure is random, we wish to minimize the probability of selecting a system which does not correspond to the smallest mean μ_1. In the simplest case, we restrict our attention to static allocation policies denoted by the vector $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)^T$ with nonnegative entries that sum to one. In this policy, $\alpha_i n$ samples $(X_1^{(i)}, X_2^{(i)}, \ldots, X_{\alpha_i n}^{(i)})$ are drawn, i.e., simulated, from system i, $i = 1,2,\ldots,k$. We ignore integrality constraints for simplicity. Given that one is interested in discovering the smallest mean, an obvious statistic for doing that is the sample average. For a fixed vector α, let $\bar{X}_n^{(i)} = \sum_{j=1}^{\alpha_i n} X_j^{(i)} / \alpha_i n$. Let

$$FS_n(\alpha) \equiv \bigcup_{i=2}^{k} \{\bar{X}_n^{(1)} > \bar{X}_n^{(i)}\}$$

denote the event of selecting the wrong system, i.e., false selection, for a simulation budget n and allocation policy α. We are interested in finding an allocation policy that minimizes the probability of this event, i.e., that solves

$$\begin{aligned} \min \quad & \mathbb{P}\{FS_n(\alpha)\} \\ \text{s.t.} \quad & \sum_{i=1}^{k} \alpha_i = 1 \\ & \alpha_i \geq 0, \qquad i = 1,2,\ldots,k \end{aligned}$$

This formulation is a relaxation of the combinatorial optimization problem that includes integrality constraints on the number of samples to be simulated from each system. While we would ideally like to characterize the optimal allocation policy for each budget, this problem is hardly tractable, and the bulk of the literature pursues more relaxed notions, such as asymptotically optimal solutions, or various heuristics.

So far, there have been two main approaches to the problem described above. The first uses ranking and selection, and considers an indifference zone formulation to control the probability of false selection up to pre-specified levels. This approach relies critically on the assumption that each $X^{(i)}$ has a Gaussian distribution; see, e.g., Goldsman and Nelson (2001), Kim and Nelson (2003) for an overview. The second approach, which this paper builds on, uses ordinal optimization to determine the number of samples taken from each system to minimize the probability of false selection. Under suitable conditions, the probability of false selection converges to zero at an exponential rate; see Dai (1996). Based on Dai (1996), Chen et al. (2000) suggest Gaussian-based heuristics to minimize the probability of false selection. Glynn and Juneja (2004) use

large deviations theory to develop a mathematically rigorous framework for more general non-Gaussian settings.

Glynn and Juneja (2004) propose allocation policies that are asymptotically optimal, a relaxed version of optimality that pertains to large simulation budgets. While these policies do not rely on Gaussian assumptions (as in Chen et al. (2000)), they suffer from two important shortcomings: they rely on finer structure of the underlying distribution functions, as opposed to just first and second moment information that characterizes Gaussian heuristics; and, they are computationally intensive and hence difficult to implement.

It is important to note that the approaches outlined above are only rigorously justified (and in the case of Glynn and Juneja (2004), feasible), if the underlying distributions are light-tailed. (We say that a distribution is light-tailed if its moment generating function exists in a neighborhood of zero.) This opens up the question, which is the main topic of this paper, of characterizing achievable performance in the best system selection problem with heavy-tailed distributions, and the derivation of policies that perform well in that environment.

2 REVIEW OF RESULTS IN LIGHT-TAILED ENVIRONMENTS

2.1 A LARGE DEVIATION-BASED APPROACH

Under technical regularity conditions, Glynn and Juneja (2004) proved that for almost all static allocation policies $\mathbb{P}\{FS_n(\alpha)\}$ converges to zero at an exponential rate as n grows large. In particular, they showed that

$$\log \mathbb{P}\{FS_n(\alpha)\} \sim -G(\alpha)n \quad \text{as} \quad n \to \infty$$

where $G(\alpha)$ can be characterized using tools from large deviations theory. Here sequences of real numbers $\{a_n\}$, $\{b_n\}$ are said to be asymptotically equivalent, denoted $a_n \sim b_n$, if $a_n/b_n \to 1$ as $n \to \infty$. Glynn and Juneja (2004) then formulated the problem of selecting the best system as:

$$\begin{aligned} \min \quad & -G(\alpha) \\ \text{s.t.} \quad & \sum_{i=1}^{k} \alpha_i = 1 \\ & \alpha_i \geq 0, \qquad i = 1,2,\ldots,k \end{aligned}$$

Hence, the goal is to find a vector α^* which maximizes $G(\alpha)$, the asymptotic decay rate of the probability of false selection. Glynn and Juneja (2004) establish that, in the light-tailed environment, there exist static allocation policies that are asymptotically optimal, in the sense that they asymptotically minimize the probability of false selection. The caveat is that computing this policy is quite complicated. In practice, $G(\alpha)$ relies on knowledge of the moment generating function of each underlying distribution. If these are to be estimated, the computation of $G(\alpha)$ can be prohibitive. In particular, setting $I_i(x) = \sup_\theta\{\theta X - \log E[\exp(\theta X^{(i)})]\}$, the rate function $G(\alpha)$ is constructed as follows:

$$\begin{aligned} G(\alpha) &= \min_{2\leq i\leq k} G_i(\alpha_1, \alpha_i) \\ &= \min_{2\leq i\leq k} \inf_x (\alpha_1 I_1(x) + \alpha_i I_i(x)) \,. \end{aligned}$$

For purposes of implementation in our paper, we take as primitive knowledge of the rate function $I_i(x), i = 1,2,\ldots,k$. Of course, knowledge of the rate function is an unrealistic assumption in practice, but it allows us to explore the best possible results obtainable with this approach. We will refer to this as the modified Glynn-Juneja algorithm.

The Modified GJ Algorithm

Step 0. Distribute the initial computing budget n_0 to each system equally, i.e., set $n_i = n_0/k$ for $i = 1,2,\ldots,k$.

Step 1. If $\sum_{i=1}^{k} n_i \geq n$, then stop.

Step 2. Compute $\bar{X}_{n_1}^{(1)}, \ldots, \bar{X}_{n_k}^{(k)}$ and find $b = \text{argmin}_{i=1,2,\ldots,k}\{\bar{X}_{n_i}^{(i)}\}$.

Step 3. Calculate $\alpha_1, \ldots, \alpha_k$ by solving the equations:

$$\begin{aligned} \min \quad & -\min_{i\neq b} G_i(\alpha_b, \alpha_i) \\ \text{s.t.} \quad & \sum_{i=1}^{k} \alpha_i = 1 \\ & \alpha_i \geq 0, \qquad i = 1,2,\ldots,k \end{aligned}$$

where

$$G_i(\alpha_b, \alpha_i) = \inf_x (\alpha_b I_b(x) + \alpha_i I_i(x)) \,. \tag{1}$$

Step 4. Increase the simulation budget by Δ and set $n_i \leftarrow n_i + \Delta\alpha_i$ for $i = 1,2,\ldots,k$. Go to Step 1.

In our numerical examples that follow in section 4, we test this algorithm for two cases: normal and exponential distributions. In the normal distribution case, (μ_i, σ_i^2) are, respectively, the mean and variance of the ith system. In the exponential distribution case, we start with a standard, unit rate, exponential random variable Z, and then $X^{(i)} \stackrel{d}{=} \sigma_i(Z-1) + \mu_i$. In this manner, the mean of the ith system is μ_i, and its variance is σ_i^2. The respective rate functions

that we take as primitives are easily seen to be

$$I_i(x) = \begin{cases} \frac{(x-\mu_i)^2}{2\sigma_i^2} & \text{Normal distribution} \\ \frac{x-\mu_i}{\sigma_i} + \log\frac{\sigma_i}{x-\mu_i+\sigma_i} & \text{Exponential distribution} \end{cases}$$

2.2 A GAUSSIAN-BASED APPROACH

Chen et al. (2000) assumed each $\bar{X}^{(i)}$ follows a normal distribution with the mean μ_i and variance $\sigma_i^2/(\alpha_i n)$. They minimize an upper bound on $\mathbb{P}\{FS_n(\alpha)\}$ and formulate the problem of selecting the best system as:

$$\begin{aligned} \min \quad & \sum_{i=2}^{k} \bar{\Phi}\left(\frac{\mu_i - \mu_1}{\sqrt{\frac{\sigma_1^2}{\alpha_1 n} + \frac{\sigma_i^2}{\alpha_i n}}}\right) \\ \text{s.t.} \quad & \sum_{i=1}^{k} \alpha_i = 1 \\ & \alpha_i \geq 0, \qquad i = 1,2,\ldots,k \end{aligned}$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution function and $\bar{\Phi}(x) = 1 - \Phi(x)$.

Using the first order conditions, they suggest approximate solutions and the following sequential method for finding an approximately optimal computing budget allocation (OCBA).

A Sequential Algorithm for Optimal Computing Budget Allocation (OCBA)

Step 0. Distribute the initial computing budget n_0 to each system equally, i.e., set $n_i = n_0/k$ for $i = 1,2,\ldots,k$.

Step 1. If $\sum_{i=1}^{k} n_i \geq n$, then stop.

Step 2. Compute $\bar{X}_{n_1}^{(1)},\ldots,\bar{X}_{n_k}^{(k)}$ and find $b = \text{argmin}_{i=1,2,\ldots,k}\{\bar{X}_{n_i}^{(i)}\}$.

Step 3. Calculate $\alpha_1,\ldots,\alpha_k$ by solving the equations:

$$\frac{\alpha_i}{\alpha_j} = \left(\frac{\sigma_i/(\bar{X}^{(b)} - \bar{X}^{(i)})}{\sigma_j/(\bar{X}^{(b)} - \bar{X}^{(j)})}\right)^2 \tag{2}$$

for $i,j \in 1,2,\ldots,k$ and $i \neq j \neq b$, and

$$\alpha_b = \sigma_b \sqrt{\sum_{i \neq b} \frac{\alpha_i^2}{\sigma_i^2}} \tag{3}$$

Step 4. Increase the simulation budget by Δ and set $n_i \leftarrow n_i + \Delta\alpha_i$ for $i = 1,2,\ldots,k$. Go to Step 1.

The above algorithm is, essentially, based on the assumption that the distributions $F^{(i)}$ are Gaussian. Hence, the variances are the only information necessary to implement the algorithm.

3 SELECTING THE BEST SYSTEM IN THE PRESENCE OF HEAVY-TAILED DISTRIBUTIONS

One can find in the literature many classes and definitions of heavy-tailed distributions, see, e.g., Embrechts et al. (1997). For our purposes, it will be convenient to focus on a class of distributions which is characterized by a polynomial-like tail. Although this is not the most general notion possible, the salient features of our problem can be easily illustrated within the confines of this family. To formalize the notion of a polynomial-like tail, we use the following definition.

Definition 1. *A positive Lebesque measurable function h on (0,∞) is* ***regularly varying*** *at ∞ with index $\nu \in \mathscr{R}$, written $h \in \mathscr{R}_\nu$, if*

$$\lim_{x\to\infty} \frac{h(tx)}{h(x)} = t^\nu, \quad t > 0.$$

3.1 MAIN THEORETICAL RESULT

The main result of this paper is stated next.

Theorem 1. *Assume $\bar{F}^{(i)}(x) = \mathbb{P}\{X^{(i)} > x\} \in \mathscr{R}_{-\nu_i}$ and $F^{(i)}(-x) = \mathbb{P}\{-X^{(i)} > x\} \in \mathscr{R}_{-\nu_i}$ for all $i = 1,2,\ldots,k$. Then, for any allocation vector $\alpha = (\alpha_1,\ldots,\alpha_k)^T$ with strictly positive elements*

$$\log \mathbb{P}\{FS_n(\alpha)\} \sim -\min_{1\leq i\leq k}\{\nu_i - 1\}\log n \quad \textit{as} \quad n \to \infty\,.$$

Proof. See Appendix. □

Theorem 1 states that the probability of false selection, $\mathbb{P}\{FS_n(\alpha)\}$, converges to zero at a polynomial rate, namely, $\mathbb{P}\{FS_n(\alpha)\} \approx n^{-r}$. Moreover, the rate r does not depend on α so long as all elements of α are strictly positive. In that case, it is determined by the minimal tail index $\min_{1\leq i\leq k}\{\nu_i\}$, i.e., the heaviest tail. This result is quite different from the type of results that one sees in light-tailed environments, where it is possible to control the decay rate using static allocation policies, and these policies are in fact asymptotically optimal (see Glynn and Juneja (2004)). On the other hand, to optimize the rate at which the probability of false selection decays to zero in the light-tailed setting, one requires detailed information characterizing the underlying distributions (in the form of the moment generating func-

tions). In contrast, in the heavy-tailed setting one can obtain the "best" asymptotic performance with static policies that are oblivious to the structure of the underlying distributions. While Theorem 1 asserts that it is impossible to control the polynomial convergence rate on logarithmic scale, improved convergence may be obtained when more refined tail information is available. In particular, if one can approximate $\log \mathbb{P}\{FS_n(\alpha)\}$ by $\log(c(\alpha)) - \min_{1\le i\le k}\{\nu_i - 1\}\log n$, where $c(\alpha)$ is a function of the policy α, then it may be possible to minimize $c(\alpha)$ over feasible policies α. An illustration of this observation using the t-distribution is given next.

Example: The t-distribution. Suppose that $X^{(i)}$ follows a t-distribution with mean μ_i and degree of freedom (df) parameter ν_i, for $i = 1, 2, \ldots, k$. The density function of $X^{(i)}$ is:

$$f(x) = c_i \left(1 + \frac{(x-\mu_i)^2}{\nu_i}\right)^{-\frac{\nu_i+1}{2}}$$

where

$$c_i = \frac{\Gamma(\frac{\nu_i+1}{2})}{\sqrt{\pi\nu_i}\,\Gamma(\frac{\nu_i}{2})}.$$

Clearly $F^{(i)}(n) \in \mathscr{R}_{-\nu_i}$ for all $i = 1, 2, \ldots, k$. Tedious but straightforward algebra gives the following bound:

$$\begin{aligned}\mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\} \quad &\le \quad c_1 \nu_1^{\frac{\nu_1+1}{2}} ((\mu_i - \mu_1)\alpha_i/2)^{-\nu_1} n^{-\nu_1+1} \\ &\quad + c_i \nu_i^{\frac{\nu_i+1}{2}} ((\mu_i - \mu_1)\alpha_1/2)^{-\nu_i} n^{-\nu_i+1}.\end{aligned}$$

Hence,

$$\begin{aligned}\mathbb{P}\{FS_n(\alpha)\} \quad &\le \quad (k-1) \max_{2\le i\le k} \Big(c_1 \nu_1^{\frac{\nu_1+1}{2}} ((\mu_i - \mu_1)\alpha_i/2)^{-\nu_1} \\ &\quad n^{-\nu_1+1},\ c_i \nu_i^{\frac{\nu_i+1}{2}} ((\mu_i - \mu_1)\alpha_1/2)^{-\nu_i} n^{-\nu_i+1} \Big).\end{aligned}$$

Lower bounds can be derived in a similar manner. Here there is a potential to improve the probability of false selection $\mathbb{P}\{FS_n(\alpha)\}$ by controlling the pre-multiplier $c_i \nu_i^{(\nu_i+1)/2}((\mu_i - \mu_1)\alpha_1/2)^{-\nu_i}$ for each $i = 1, 2, \ldots, k$. Of course this requires significant knowledge of the underlying distribution functions.

4 NUMERICAL RESULTS

The numerical experiments include tests of several algorithms, both light-tailed and heavy-tailed. For each budget n, we estimate $\mathbb{P}\{FS_n(\alpha)\}$ by counting the number of times a method identifies the true smallest mean system (by selecting the smallest sample average) out of the total number of replications. The number of replications, m, is chosen so that:

$$\sqrt{\frac{P_n(1-P_n)}{m}} \le \frac{P_n}{10}$$

where P_n is the order of magnitude of the probability of false selection for the given budget n. This implies that the standard error for each estimate of $\mathbb{P}\{FS_n(\alpha)\}$ is at least ten times smaller than the value of $\mathbb{P}\{FS_n(\alpha)\}$. For example, in Figure 1 the minimum value of P_n is 10^{-4}, so we choose m to be at least 10^6 by this argument. Consequently, we have high confidence that the results are not driven by simulation error for all examples with $\mathbb{P}\{FS_n(\alpha)\}$ larger than 10^{-4}.

4.1 PERFORMANCE OF LIGHT-TAILED ALGORITHMS

Suppose there are five systems and the random variable $X^{(i)}$ representing the performance of the ith system follows a normal distribution with variance and mean values given in Table 1. Three different algorithms are compared. The first equally distributes samples across systems, i.e., $\alpha = (1/5, \ldots, 1/5)^T$. We refer to this as the equal allocation (EA) algorithm. The second is the Gaussian-based method suggested by Chen et al. (2000). The third is the modified GJ algorithm based on Glynn and Juneja (2004) which was discussed in section 2.1.

Table 1: Mean and variance of each system.

System (i)	1	2	3	4	5
μ_i	0	0.5	0.6	0.8	1.0
σ_i	1.0	1.0	1.0	1.0	1.0

The graph (a) in Figure 1 illustrates that $\mathbb{P}\{FS_n(\alpha)\}$ converges to zero exponentially in a strictly idealized case where the true mean values (instead of $\bar{X}_i$) are used to allocate samples in (2) in OCBA, and (1) in modified GJ algorithm. The exponents for OCBA and the modified GJ algorithm are almost identical, while the EA algorithm exhibits a slower decay (the exponents are the slopes of the graphs in Figure 1). This illustrates that one can indeed control the exponential convergence rate by the choice of the static allocation policy.

Graphs (b), (c) and (d) in Figure 1 deal with cases where preliminary estimates of μ are first obtained using different initial resource allocations ($n_0 = 50, 100, 150$). These values are then used to initialize the algorithms. When the initial resource allocation n_0 is 50, both the OCBA and modified GJ algorithms perform poorly compared to the naive EA algorithm as n grows. For n_0 equal to 150 in graph (d), both algorithms exhibit good performance. The slope of the modified GJ algorithm in Figure 1(d) is almost identical to that in Figure 1(a), that is, with sufficient samples allocated to

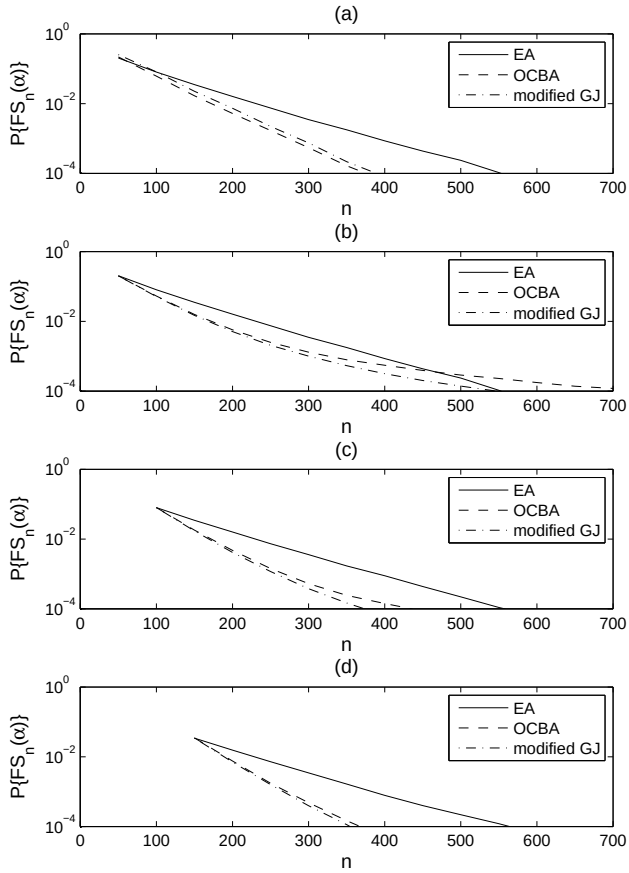

Figure 1: Probability of false selection as a function of the simulation budget n plotted in log-linear scale. The three allocation algorithms are: equal allocation (EA); Gaussian-based approach (OCBA); large deviations-based approach (modified GJ). The four graphs are: (a) μ assumed to be known; (b) μ unknown and the initial resource allocation is $n_0 = 50$; (c) μ unknown and $n_0 = 100$; and (d) μ unknown and $n_0 = 150$.

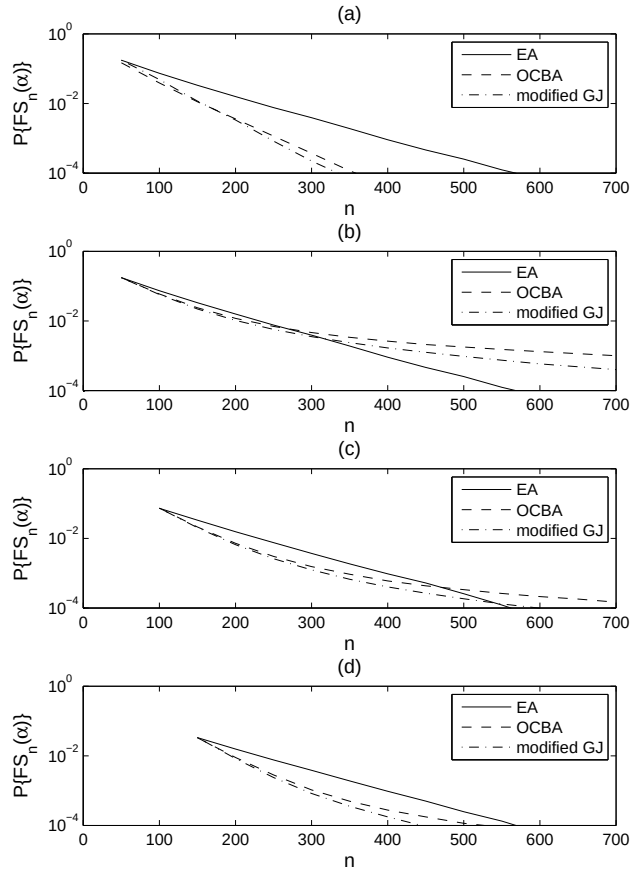

Figure 2: Probability of false selection as a function of the simulation budget n plotted in log-linear scale for five systems with an underlying exponential distribution. The three allocation algorithms are: equal allocation (EA); Gaussian-based approach (OCBA); large deviations-based approach (modified GJ). The four graphs are: (a) μ assumed to be known; (b) μ unknown and the initial resource allocation is $n_0 = 50$; (c) μ unknown and $n_0 = 100$; and (d) μ unknown and $n_0 = 150$.

the initial phase, the OCBA and the modified GJ algorithms perform almost as if they had knowledge of the system means a priori. (Chen et al. (2000) recommend allocating about 10% of the simulation budget to the initial phase.)

Since the OCBA algorithm of Chen et al. (2000) is based on the premise of underlying normal distributions, its application in environments where the distribution is light-tailed but non-Gaussian can be viewed as a heuristic that is based on the central limit theorem. In contrast, the large deviations-based approach of Glynn and Juneja (2004) has rigorous justification in the non-Gaussian case. Figure 2 compares the performance of the three algorithms studied previously, when the underlying distributions are exponential with means given in Table 1. Again, we observe that when the allocation budget is constructed based on known mean values (Figure 2(a)), both OCBA and modified GJ algorithms outperform the naive equal allocation (EA) scheme. However, when the algorithms do not assume knowledge of the true means, and the initial allocation is "small" (e.g., $n_0 = 50$ in Figure 2(b)), then both the OCBA and modified GJ algorithms perform poorly compared to EA. Only when the initial allocation is increased significantly (Figure 2(c) and, in particular, Figure 2(d)) do the two

light-tailed algorithms start to outperform EA. In graphs (a) and (d) in Figure 2 we observe that the modified GJ algorithm performs slightly better than the Gaussian-based OCBA method, consistent with the fact that the underlying distributions are non-Gaussian.

From Figures 1 and 2 we conclude that the light-tailed algorithms perform well either (i) when they are constructed based on the known μ values, or (ii) when the initial allocation (n_0) is fairly large. The former is of course not a situation that arises in practice, and the latter is a significant restriction on the application of these algorithms. Moreover, it is worth emphasizing again that the algorithm of Glynn and Juneja (2004) that is based on a large deviations principle is quite complicated to implement. The version tested here is a simplified and idealized modification that assumes the structural form of the rate function is known, and thus the only issue is to estimate its parameters. In reality, since such information is not available, the algorithm would need to estimate the entire rate function non-parametrically, and this raises significant implementation issues. As seen in Table 2, the modified GJ algorithm takes 80 times more CPU time than the EA algorithm and 25 times more CPU time compared to the OCBA algorithm. Even though the implementation of each algorithm was not optimized for performance, the relative CPU times illustrate the dramatic differences in complexity between the methods.

Table 2: CPU time for each algorithm. (Operating system: Windows XP; processor: Intel core duo 1.8GHz; memory: 1GB RAM; language: C using the GNU scientific library.)

Algorithm	EA	OCBA	modified GJ
CPU time (minutes)	4	13	313

4.2 PERFORMANCE OF HEAVY-TAILED ALGORITHMS

4.2.1 TWO SYSTEMS

We consider examples with two systems where the performance of each system follows a t-distribution. In graphs (a) and (b) in Figure 3, the first system follows a t-distribution with $\mu_1 = 0, \nu_1 = 3$. In Figure 3(a), the second system follows a t-distribution with $\mu_2 = 3, \nu_2 = 3$ and in Figure 3(b), the second system follows a t-distribution with $\mu_2 = 3, \nu_2 = 5$. First, note that $\mathbb{P}\{FS_n(\alpha)\}$ in Figure 3 asymptotically converges to zero at a polynomial rate for all static allocation policies considered. This observation is consistent with the result given in Theorem 1. The polynomial convergence rates are given by the limiting slopes in Figure 3. The slopes are all -2, which is equal to $-\min(\nu_1, \nu_2) + 1$, and indeed the slopes depend only on the system with the heaviest tail and not on the allocation policy.

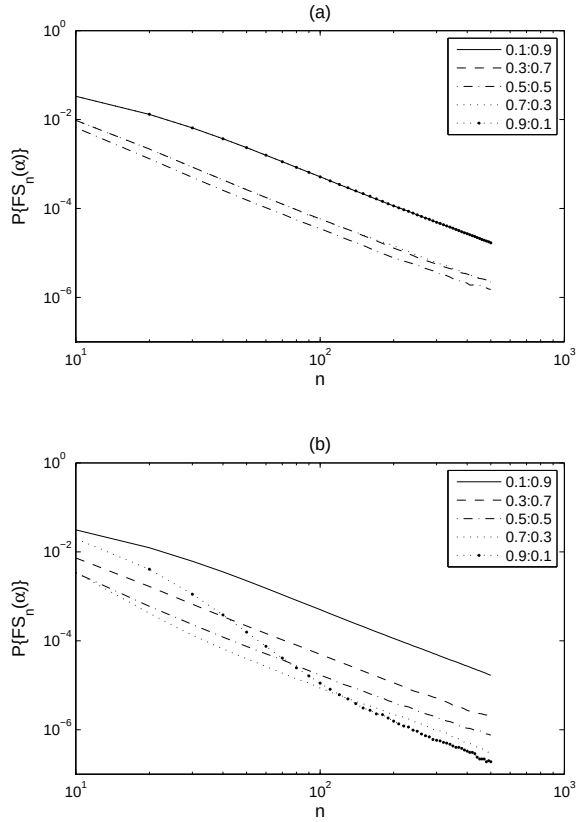

Figure 3: Probability of false selection as a function of the simulation budget n plotted in log-log scale for two systems following t-distributions. The five asymptotic allocation vectors that are used for each graph are: α = (0.1, 0.9), (0.3, 0.7), (0.5, 0.5), (0.7, 0.3) and (0.9, 0.1). The parameters for the two graphs are: (a) $\mu_1 = 0, \mu_2 = 3$ and $\nu_1 = 3, \nu_2 = 3$; and (b) $\mu_1 = 0, \mu_2 = 3$ and $\nu_1 = 3, \nu_2 = 5$.

4.3 MULTIPLE SYSTEMS

We now consider five systems with t-distributions and with the means and degrees of freedom (df) parameters given in Table 3. We denote the mean and the degree of freedom of each system using vectors $\mu = (\mu_1, \ldots, \mu_k)$ and $\nu = (\nu_1, \ldots, \nu_k)$.

For this experiment, the total number of replications, m, is set to 10^7, so the estimates of $\mathbb{P}\{FS_n(\alpha)\}$ with values up to 10^{-5} are accurate with high confidence. The variances $\sigma_i^2 = \nu_i/(\nu_i - 2), i = 1, 2, \ldots, 5$, are assumed known and are used in the OCBA algorithm. In Figure 4, $\mathbb{P}\{FS_n(\alpha)\}$ with the equal allocation method converges to zero at a

Table 3: Mean and df for each system.

System (i)	1	2	3	4	5
μ_i	0.0	0.7	1.0	1.3	1.5
ν_i	3	3	3	3	3

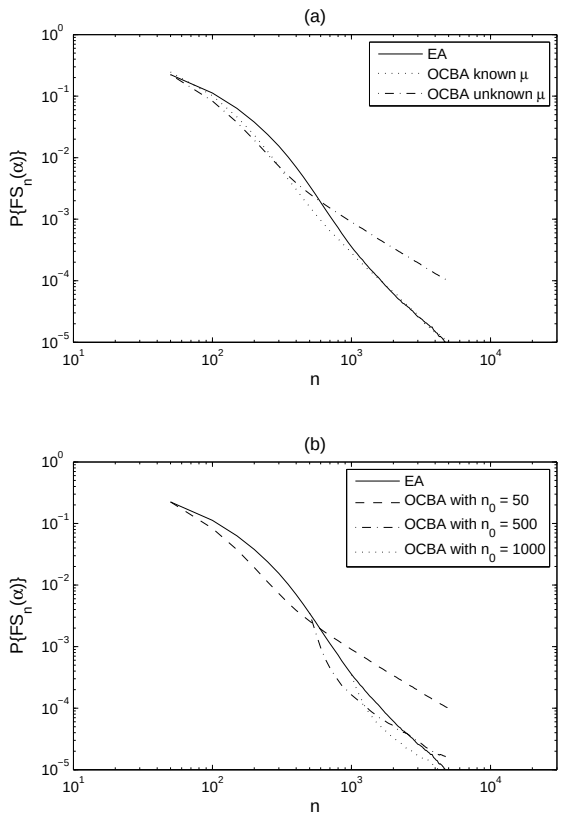

Figure 4: Probability of false selection as a function of the simulation budget n plotted in log-log scale for five t-distributions with μ = (0.0, 0.7, 1.0, 1.3, 1.5) and ν = (3,3,3,3,3). In (a), three allocation rules (EA, OCBA when μ is known, and OCBA when μ is unknown) are shown. In (b), three initial resource allocations ($n_0 = 50, 100, 150$) for the OCBA algorithm are tested and compared to equality algorithm (EA).

polynomial rate as anticipated. When the true μ values and dfs are assumed to be known, $\mathbb{P}\{FS_n(\alpha)\}$ for the OCBA algorithm also converges to zero at a polynomial rate as anticipated in Theorem 1, since the allocation vector α becomes constant. When the μ values are unknown but dfs are known, $\mathbb{P}\{FS_n(\alpha)\}$ for the OCBA algorithm converges to zero but the convergence rate is not same as

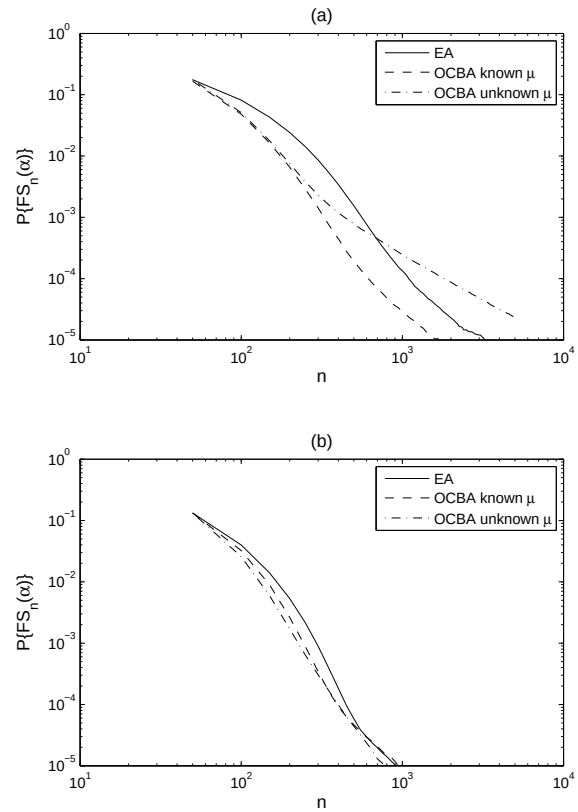

Figure 5: Probability of false selection as a function of the simulation budget n plotted in log-log scale for the five t-distributions with μ = (0.0, 0.7, 1.0, 1.3, 1.5) and (a) ν = (3,4,5,6,7); and (b) ν = (7,6,5,4,3).

the rate when μ is known. Even though the initial resource allocation increases, the slope does not change much as seen in Figure 4(b). This suggests that OCBA is not very well suited for heavy-tailed systems.

Figure 5 shows the convergence of $\mathbb{P}\{FS_n(\alpha)\}$ for the same mean parameter values but with different dfs. In Figure 5(a), the dfs are ν = (3,4,5,6,7), and in Figure 5(b), they are ν = (7,6,5,4,3). The results are quite similar to the case where the dfs were ν = (3,3,3,3,3).

Even though the polynomial convergence rate is determined by the system with the heaviest-tailed distribution, it might be possible to improve performance by allowing for non-static allocation policies. We propose the following non-static heuristic algorithm. Set $i^* = \text{argmin}_{1\le i\le k}\{\nu_i\}$. Then calculate the allocation vector α as follows:

$$\alpha_i = \begin{cases} n^{-\frac{\nu_i-\nu_{i^*}}{\nu_i-1}} & i \neq i^* \\ 1-\sum_{j\neq i^*}\alpha_j & i = i^* \end{cases}$$

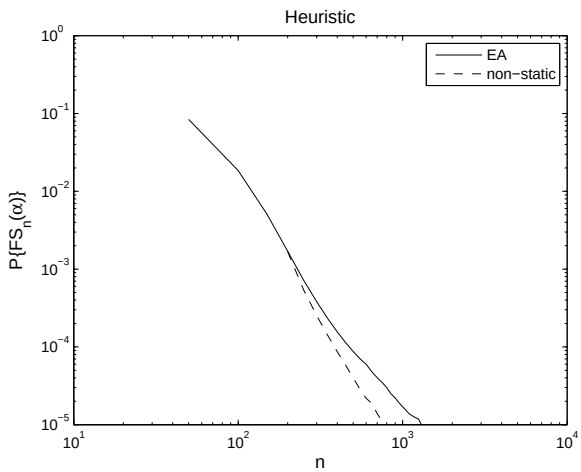

Figure 6: Probability of false selection as a function of the simulation budget n plotted in log-log scale for five t-distributions with μ = (0.0, 1.0, 1.2, 1.3, 1.5) and ν = (4,6,3,5,7). The two algorithms are: equal allocation (EA); and a non-static heuristic with initial resource allocation $n_0 = 200$.

These values are then updated at every point in time when the budget is incremented (as in Step 4 of OCBA and modified GJ algorithms). Since the system with the heaviest-tailed distribution dominates $\mathbb{P}\{FS_n(\alpha)\}$ for large n, allocating more samples to this system is natural. Note that the policy does not require any sample average information. Figure 6 compares the EA algorithm with the non-static heuristic for five systems which follow t-distributions with μ = (0.0, 1.0, 1.2, 1.3, 1.5) and ν = (4,6,3,5,7). As seen in Figure 6, the non-static heuristic works well for the initial resource allocation $n_0 = 200$. Similar to the OCBA algorithm, this non-static heuristic is also sensitive to the initial resource allocation. We have observed similar results to these shown in Figure 6 in other test cases which are not repeated here because of space considerations.

5 CONCLUSIONS

In the case of light-tailed distributions and static allocation policies, the probability of false selection, $\mathbb{P}\{FS_n(\alpha)\}$, converges to zero at an exponential rate which is controlled by the static allocating policy. In this setting, OCBA (a Gaussian-based heuristic) is easily implementable, but its convergence rate depends heavily on the initial resource allocation. In contrast, the GJ algorithm is more general and can be rigorously justified, yet is computationally intensive and difficult to implement in practice, since it requires considerable information about the underlying distributions.

In the case of heavy-tailed distributions, $\mathbb{P}\{FS_n(\alpha)\}$ converges to zero at a polynomial rate for every static policy with strictly positive components. In other words, essentially all static allocation policies eventually achieve the same polynomial convergence rate, which is determined by the system with the heaviest tail. In sharp contrast to the light-tailed case, no information on the distributions is needed for the design of such policies.

A APPENDIX

Let $\{X_j\}_{j=1}^n$ be i.i.d. samples and define $S_n = \sum_{j=1}^n X_j$. We use the following large deviations result for heavy-tailed distributions given in the next theorem in our proof of Theorem 1.

Theorem 2. *[Nagaev (1979)] Suppose that $\bar{F} \in \mathscr{R}_{-\nu}$ for some $\nu > 2$, $E|X|^{2+\delta} < \infty$ for some $\delta > 0$ and that var(X) = 1. Then*

$$\mathbb{P}\{S_n - n\mu > x\} = \bar{\Phi}(\frac{x}{\sqrt{n}})(1+o(1)) + n\bar{F}(x)(1+o(1))$$

uniformly for $x \geq \sqrt{n}$. In particular,

$$\mathbb{P}\{S_n - n\mu > x\} = \bar{\Phi}(\frac{x}{\sqrt{n}})(1+o(1))$$

for $\sqrt{n} \leq x \leq a\sqrt{n \ln n}$ and $a < \sqrt{\nu - 2}$, and

$$\mathbb{P}\{S_n - n\mu > x\} = n\bar{F}(x)(1+o(1))$$

for $x > a\sqrt{n \ln n}$ and $a > \sqrt{\nu - 2}$.

Proof of Theorem 1

Proof. Fix an allocation vector $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)^T$ with strictly positive elements and $n \in \mathbb{N}$.
Step 1. Since $\mathbb{P}\{FS_n(\alpha)\} = \mathbb{P}\{\cup_{i=2}^k (\bar{X}^{(1)} > \bar{X}^{(i)})\}$,

$$\mathbb{P}\{FS_n(\alpha)\} \leq (k-1) \max_{2 \leq i \leq k} \mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\}$$

also

$$\mathbb{P}\{FS_n(\alpha)\} \geq \max_{2 \leq i \leq k} \mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\}.$$

Thus, if $\mathbb{P}\{FS_n(\alpha)\} \to 0$ as $n \to \infty$, then

$$\log \mathbb{P}\{FS_n(\alpha)\} \sim \log[\max_{2 \leq i \leq k} \mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\}] \tag{4}$$

Step 2. Fix i such that $2 \le i \le k$.

$$\begin{aligned}
&\mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\} \\
&\quad = \mathbb{P}\{\alpha_i S^{(1)}_{\alpha_1 n} - \alpha_1 S^{(i)}_{\alpha_i n} > 0\} \\
&\quad = \mathbb{P}\{\alpha_i (S^{(1)}_{\alpha_1 n} - \mu_1 \alpha_1 n) - \alpha_1 (S^{(i)}_{\alpha_i n} - \mu_i \alpha_i n) > \mu \alpha_1 \alpha_i n\} \\
&\quad = \mathbb{P}\{\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} - \alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > \mu \alpha_1 \alpha_i n\} \qquad (5)
\end{aligned}$$

where

$$\tilde{S}^{(i)}_{\alpha_i n} = \sum_{j=1}^{\lfloor \alpha_i n \rfloor} X^{(i)}_j - \lfloor \alpha_i n \rfloor \mu_i$$

and $\mu \equiv \mu_i - \mu_1 > 0$. Since $\mathbb{P}\{x+y > n\} \le \mathbb{P}\{x > \frac{n}{2}\} + \mathbb{P}\{y > \frac{n}{2}\}$,

$$\begin{aligned}
&\mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} - \alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > n\mu \alpha_1 \alpha_i\} \\
&\quad \le \mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > \frac{n\mu \alpha_1 \alpha_i}{2}\} + \mathbb{P}\{-\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > \frac{n\mu \alpha_1 \alpha_i}{2}\} \qquad (6)
\end{aligned}$$

Also, since the two random variables are independent we have

$$\begin{aligned}
&\mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} - \alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > n\mu \alpha_1 \alpha_i\} \\
&\quad \ge \max\{\mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > n\mu \alpha_1 \alpha_i\}\mathbb{P}\{-\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > 0\}, \\
&\qquad \mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > 0\}\mathbb{P}\{-\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > n\mu \alpha_1 \alpha_i\}\}. \qquad (7)
\end{aligned}$$

Since all $\alpha_i > 0$, by Theorem 2, we have that

$$\mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > n\mu \alpha_1 \alpha_i\} = n\bar{F}^{(i)}(n\mu \alpha_i)(1+o(1))$$

and

$$\mathbb{P}\{-\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > n\mu \alpha_1 \alpha_i\} = n\bar{F}^{(1)}(n\mu \alpha_1)(1+o(1)) \qquad (8)$$

since $n > \sqrt{n \ln n}$ for large n. Also,

$$\begin{aligned}
\mathbb{P}\{-\alpha_i \tilde{S}^{(1)}_{\alpha_1 n} > 0\} &= \mathbb{P}\{\frac{1}{\sqrt{\alpha_1 n}} \tilde{S}^{(1)}_{\alpha_1 n} < 0\} \\
&= \mathbb{P}\{\sigma N(0,1) \le 0\}(1+o(1))
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{P}\{\alpha_1 \tilde{S}^{(i)}_{\alpha_i n} > 0\} &= \mathbb{P}\{\frac{1}{\sqrt{\alpha_i n}} \tilde{S}^{(i)}_{\alpha_i n} > 0\} \\
&= \mathbb{P}\{\sigma N(0,1) \ge 0\}(1+o(1)) \qquad (9)
\end{aligned}$$

by the CLT.

Using (5), (6), (8) and (9) we obtain

$$\mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\} \le c_1 n\bar{F}^{(1)}(\frac{n\mu \alpha_i}{2}) + c_2 n\bar{F}^{(i)}(\frac{n\mu \alpha_1}{2}) \qquad (10)$$

for some positive finite c_1 and c_2.

Using (2), (4), (5) & (6) we also obtain

$$\mathbb{P}\{\bar{X}^{(1)} > \bar{X}^{(i)}\} \ge c_3 n\bar{F}^{(1)}(n\mu \alpha_i) + c_4 n\bar{F}^{(i)}(n\mu \alpha_1) \qquad (11)$$

for some positive finite c_3 and c_4.

From (4), (10) and (11) and the assumption of $\bar{F}^{(i)} \in \mathscr{R}_{-v_i}$, we have

$$\log \mathbb{P}\{FS_n(\alpha)\} \sim -\min_{1 \le i \le k}\{v_i - 1\} \log n$$

as $n \to \infty$. □

B ACKNOWLEDGMENT

This work was partly supported by NSF grants DMS-0410234 and DMI-044765.

C REFERENCES

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

Dai, L. 1996. Convergence properties of ordinal comparison in the simulation of discrete event dynamic systems. *Journal of Optimization Theory and Applications* 91: 363-388.

Embrechts, P., C. Kluppelberg, and T. Mikosch. 2000. Modelling Extremal Events, Springer.

Goldie, C.M., N.H. Bingham, J.L. Teugels. 1987. Regular Variation, Cambridge University Press.

Goldsman, D., and B.L. Nelson. 2001. Statistical selection of best system. *Proceedings of the 2001 Winter Simulation Conference*, Eds. B.A. Peters, J.S. Smith, D.J. Medeiros, and M.W. Rohrer. 139-146. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Glynn, P., S. Juneja. 2004. A Large deviations perspective on ordinal optimization. *Proceedings of the 2004 Winter Simulation Conference*, Eds. R.G. Ingalls, M.D. Rossetti, J.S. Smith, and B.A. Peters. 577-585. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Kim, S.H., and B.L. Nelson. 2003. Selecting the best system: theory and methods. *Proceedings of the 2003 Winter Simulation Conference*, Eds. S. Chick, P.J. Sanchez, D. Ferrin, and D.J. Morrice. 101-112. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Nagaev, S.V. 1979. Large deviations of sums of independent random variables. *Ann. Probab*. 7: 745-789.