

GUARANTEES ON THE PROBABILITY OF GOOD SELECTION

David J. Eckman
Shane G. Henderson

School of Operations Research and Information Engineering
Cornell University
206 Rhodes Hall
Ithaca, NY 14853, USA

ABSTRACT

This tutorial provides an overview of guarantees on the probability of good selection (PGS), i.e., statistical guarantees on selecting – with high probability – an alternative whose expected performance is within a given tolerance of the best. We discuss why PGS guarantees are superior to more popular, related guarantees on the probability of correct selection (PCS) under the indifference-zone formulation. We review existing procedures that deliver PGS guarantees and assess several direct and indirect methods of proof. We compare the frequentist and Bayesian interpretations of PGS and highlight the differences in how procedures are designed to deliver PGS guarantees under the two frameworks.

1 INTRODUCTION

We consider the problem of selecting from among a finite number of alternatives, the performances of which cannot be evaluated exactly, but instead must be estimated using stochastic simulation. The number of alternatives under consideration is assumed to be sufficiently small, so that – subject to the available computational resources – each alternative can be simulated. This class of problems, referred to as ranking and selection (R&S), arises when the performance of an alternative involves an element of uncertainty. Much research has dealt with the design of procedures that allocate simulation effort among alternatives and ultimately select an alternative that is believed to be the best; see Kim and Nelson (2006) and Chen et al. (2015) and references therein.

A prototypical example of a R&S problem is that of allocating buffer space and service rates in a multi-station flow line with the objective of maximizing the expected steady-state throughput of the system (Pichitlamken et al. 2006). For a fixed budget of buffer space and service rates, the number of feasible, integer-valued allocations (alternatives) is finite. Furthermore, the performance of a given allocation can be estimated by running a steady-state simulation of the system and recording the throughput. R&S problems arise in virtually all simulation applications, e.g., inventory management (Koenig and Law 1985) and network reliability design (Kiekhäfer 2011).

A fundamental consideration for R&S problems is how the decision-maker assesses the quality of a selection decision. Does he or she insist on selecting the best alternative, or would he or she settle for an alternative that is, in some sense, close to the best? Requiring that the best alternative is selected with high probability could entail detecting infinitesimal differences in the performances of alternatives, a computationally intensive task. For this reason, procedures with guarantees on the probability of selecting the best alternative *for any problem instance* are uncommon, a notable exception being the indifference-zone-free procedure of Fan et al. (2016). We instead focus on the relaxed goal of selecting an alternative whose performance is deemed “good.”

A natural characterization of a good alternative is that its performance is within a given tolerance, $\delta > 0$, of that of the best alternative. Another approach, known as ordinal optimization, measures an alternative’s

goodness based on the ordering of the alternative’s performance in relation to the best (Lau and Ho 1997; Ho et al. 2000). For example, an alternative might be regarded as good if its performance puts it within the top $m\%$ or top m alternatives. From the perspective of the decision-maker, these cardinal and ordinal criteria for goodness could differ significantly. We advocate for the cardinal criterion since the optimality gap has a clear, problem-specific meaning to the decision-maker in terms of the units of the objective function. Furthermore, this criterion offers better control on the performance of a good alternative.

To formalize these ideas, let k be the number of alternatives under consideration and let θ_i be the scalar performance of alternative i , $i = 1, \dots, k$. We refer to the vector $\theta = (\theta_1, \theta_2, \dots, \theta_k)$ as the *configuration* or problem instance and use a subscript $[\cdot]$ when referring to the ordered performances $\theta_{[1]} \leq \theta_{[2]} \leq \dots \leq \theta_{[k]}$. Without loss of generality, we will assume that larger performance values are better, i.e., alternative $[k]$ is (one of) the best. Let \mathcal{D} be the (random) index of the alternative selected by a procedure.

Selection decisions can be described by loss functions in terms of the selected alternative and the best (Branke et al. 2005). For instance, correct selection and good selection can be modeled by the loss functions $\mathcal{L}_{CS}(\mathcal{D}, \theta) = \mathbb{1}\{\theta_{\mathcal{D}} \neq \theta_{[k]}\}$ and $\mathcal{L}_{GS, \delta}(\mathcal{D}, \theta) = \mathbb{1}\{\theta_{\mathcal{D}} \leq \theta_{[k]} - \delta\}$, respectively, where $\mathbb{1}\{\cdot\}$ denotes an indicator function taking values of 0 or 1. The probabilities of correct selection (PCS) and good selection (PGS) are then one minus the expectations of the loss functions \mathcal{L}_{CS} and $\mathcal{L}_{GS, \delta}$, respectively. Here, the expectations are with respect to repeated runs of the procedure on a fixed problem instance, i.e., the frequentist interpretation.

Another popular loss function, opportunity cost (also known as linear loss), is defined as the difference between the performances of the selected alternative and the best, i.e., $\mathcal{L}_{OC}(\mathcal{D}, \theta) = \theta_{[k]} - \theta_{\mathcal{D}}$ (Chick and Inoue 2001b). The opportunity cost loss function is analogous to simple regret in the pure exploration problem for multi-armed bandits (Bubeck et al. 2011) and is weakly related to good selection. Expected opportunity cost can be translated into a bound on PGS via Markov’s inequality (Chick and Wu 2005):

$$\text{PGS} = P_{\theta}(\theta_{\mathcal{D}} > \theta_{[k]} - \delta) = 1 - P_{\theta}(\theta_{[k]} - \theta_{\mathcal{D}} \geq \delta) \geq 1 - \mathbb{E}_{\theta}[\theta_{[k]} - \theta_{\mathcal{D}}]/\delta = 1 - \mathbb{E}_{\theta}[\mathcal{L}_{OC}]/\delta,$$

where the subscript θ indicates that the probability (or expectation) is for the fixed configuration θ . This bound, however, is quite loose. To imply $\text{PGS} \geq 1 - \alpha$ for some $0 < \alpha \leq 1 - 1/k$, the expected opportunity cost must be less than or equal to $\delta \times \alpha$. This is tantamount to selecting the best alternative with probability $1 - \alpha$ and selecting an alternative with knife-edge bad performance $\theta_{[k]} - \delta$ with probability α . While expected opportunity cost offers more control than PGS over the performance of the selected alternative, it is more commonly studied under the Bayesian interpretation of the R&S problem (see Section 4.2), and so we do not discuss it further.

Although PGS guarantees have primarily been studied within the R&S community, they have also been the subject of research on multi-armed bandits in the pure exploration setting, where they are called probably approximately correct (PAC) selection guarantees (Even-Dar et al. 2002). The term “probably” refers to the fixed confidence level $1 - \alpha$, while the term “approximately correct selection” refers to the event of selecting an alternative (arm) whose performance is within δ of the best. Much of the recent research on PAC-selection guarantees for multi-armed bandit algorithms focuses on the necessary sampling complexity and how it scales with respect to δ , $1 - \alpha$, and k (Mannor and Tsitsiklis 2004; Karnin et al. 2013; Zhou et al. 2014).

Both the R&S and multi-armed bandit communities make assumptions on the distributions of the observations taken from the alternatives. Indeed, without some regularity conditions, designing a selection procedure that delivers a fixed-confidence guarantee appears hopeless. For instance, Bubeck et al. (2011) point out that binary loss functions such as those for PCS and PGS are not suited for distribution-free analysis. In the multi-armed bandit community, two common assumptions are that the distributions of observations (rewards) either have bounded support or are sub-Gaussian. Mannor and Tsitsiklis (2004) conjecture that the complexity results for PAC-selection guarantees can hold for other special cases of reward distributions, provided the Kullback-Liebler divergence between the distributions of any two arms is finite. Meanwhile, in the R&S community, the prevalent assumption is that the observations are normally distributed. While

unrealistic, this assumption can often be approximately satisfied by batching observations. Other R&S research has avoided this assumption by studying the problem from a large-deviations perspective (Glynn and Juneja 2004; Hunter and Pasupathy 2010). In particular, Glynn and Juneja (2015) show that some regularity conditions are needed to design selection procedures that deliver a PCS guarantee of $1 - \alpha$ with complexity $O(\log(1/\alpha))$. Broadie et al. (2007) and Blanchet et al. (2008) extend the large-deviations analysis to design asymptotically optimal allocation policies for heavy-tailed distributions.

The remainder of this tutorial is organized as follows. In Section 2, we draw connections between PGS guarantees and more popular PCS guarantees under the indifference-zone formulation and argue that PGS guarantees are preferable. In Section 3, we review past research about the design of procedures that deliver PGS guarantees and the various proof techniques that have been used. In Section 4, we discuss PGS guarantees that result from bootstrapping and those that arise under a Bayesian framework. In Section 5, we outline unresolved research questions and possible extensions of PGS guarantees to other stochastic optimization problems. We direct the reader to the working paper of Eckman and Henderson (2018a) for detailed proofs and additional discussion.

2 CONNECTIONS TO THE INDIFFERENCE-ZONE FORMULATION

Ever since the earliest research on R&S problems, the emphasis has been on designing selection procedures that deliver guarantees on PCS. Nevertheless, this insistence on selecting the best alternative with high probability runs counter to the fact that a simulation model is, unavoidably, an imprecise representation of reality. Consequently, the best alternative according to a simulation model is not necessarily the best alternative in the real world. In addition, the decision-maker may have secondary criteria for evaluating an alternative and some differences in performance may simply be too small to care about. Some leniency in selecting a high-quality alternative is, therefore, justified.

As previously mentioned, designing a procedure that guarantees to select the best alternative with high probability is difficult, because detecting a small difference in performance is computationally expensive. An early approach that circumvents this issue is the indifference-zone (IZ) formulation of Bechhofer (1954). The IZ formulation uses a scalar parameter $\delta > 0$ to partition the space of configurations into two regions: the preference zone and the indifference zone. For configurations in the preference zone, the performance of the best alternative is at least δ better than those of all other alternatives, i.e., $\theta_{[k]} - \theta_{[k-1]} \geq \delta$. For these problem instances, the decision-maker clearly prefers the best alternative to the others. For configurations in the indifference zone, there is at least one other alternative whose performance is within δ of the best, i.e., $\theta_{[k]} - \theta_{[k-1]} < \delta$. For these problem instances, the decision-maker is believed to be indifferent to selecting any alternative whose performance is within δ of the best.

Under the indifference-zone formulation, selection procedures are typically designed to deliver PCS guarantees for configurations in the preference zone, i.e., $P_{\theta}(\theta_{\mathcal{D}} = \theta_{[k]}) \geq 1 - \alpha$ for all θ such that $\theta_{[k]} - \theta_{[k-1]} \geq \delta$. In other words, when there is a clear-cut best alternative, it will be selected with high probability. This IZ-inspired PCS guarantee has remained the primary fixed-confidence guarantee sought for R&S procedures. Although IZ-inspired PCS guarantees are at times difficult to prove, the IZ formulation often makes the analysis more tractable. In particular, the lower bound on the difference in performances between the best and second-best alternatives can frequently be used to develop upper bounds on the sample sizes needed to detect such differences.

Despite their prevalence in the R&S literature, IZ-inspired PCS guarantees are unsatisfactory. The most glaring omission of IZ-inspired PCS guarantees is that they say nothing about the performance of a selection procedure when the configuration is not in the preference zone. Does the procedure guarantee to select a δ -good alternative with high probability? Does the procedure even guarantee to terminate in finite time? This deficiency is only exacerbated by the fact that real-world problem instances are unlikely to belong to the preference zone, especially when the number of alternatives is large. Moreover, when a simulation-optimization search is used to identify alternatives, one would expect the visited alternatives to have increasingly similar performance as the search progresses (Eckman and Henderson 2018b).

IZ-inspired PCS guarantees are further undermined by the inherent uncertainty about the configuration. Merely verifying that the IZ assumption is satisfied is impractical (Parnes and Srinivasan 1986). If the decision-maker had a prior belief that the configuration was in the preference zone, he or she might be better served by using this information within a Bayesian R&S procedure, thereby avoiding the IZ assumption altogether.

Even if it were possible to assert a lower bound on the difference in the performances between the best and second-best alternatives, that quantity may be smaller than the decision-maker's tolerance towards making a suboptimal decision. To be clear, although the IZ parameter δ is often regarded as the minimum difference in performance worth detecting, it does *not* play that role for selection procedures with IZ-inspired PCS guarantees. Instead, δ demarcates the problem instances on which a selection procedure can be relied upon to deliver a statistical guarantee. This disconnection between the IZ parameter and the decision-maker's tolerance can lead to undesirable consequences. For example, the decision-maker may choose an IZ parameter much smaller than his or her tolerance in order to make an IZ-inspired PCS guarantee more reliable, with the cost of this conservativeness being excessive simulation effort.

Because of these shortcomings of IZ-inspired PCS guarantees, PGS guarantees are a superior goal for which selection procedures should instead be designed. For all their deficiencies as a stand-alone goal, IZ-inspired PCS guarantees are not without purpose. As we will show in Section 3.2, under certain sufficient conditions, IZ-inspired PCS guarantees can be lifted to PGS guarantees.

3 METHODS FOR PROVING PGS GUARANTEES

In this section, we review several techniques that have been used to prove PGS guarantees for selection procedures. We first introduce a notation to describe the problem formulation and selection procedures.

For each alternative $i = 1, \dots, k$, let X_{ij} denote the j th observation. We assume that the vectors of the j th observations from all alternatives, $(X_{1j}, X_{2j}, \dots, X_{kj})$, are independent and identically distributed (i.i.d.) for $j \geq 1$, with each being drawn from a joint distribution F having marginal distributions F_i , $i = 1, \dots, k$. We make no assumption that the observations $X_{1j}, X_{2j}, \dots, X_{kj}$ are independent, and so allow for the use of common random numbers across alternatives.

A key assumption we make is that the joint distribution F is fully specified by the configuration θ . In other words, given a vector of performances θ , one can generate the observations $(X_{1j}, X_{2j}, \dots, X_{kj})$ for $j \geq 1$. This assumption enables us to make well-defined statements relating the probabilities of selection decisions under different configurations. It also aligns with the indifference-zone formulation, in which problem instances are characterized only by the vector of performances θ , and not by any other distributional parameters of F . For example, when observations are normally distributed with mean vector θ and covariance matrix Σ , Σ is assumed to be the same under any two configurations θ and $\tilde{\theta}$.

Given this setup, we define a selection procedure as follows. For each alternative $i = 1, \dots, k$, a selection procedure takes observations X_{i1}, \dots, X_{in_i} and calculates an estimator Y_i for θ_i based on these observations. The sample size of alternative i , n_i , may depend on the observations of other alternatives and can be random. We assume that a selection procedure then selects an alternative $\mathcal{D} \in \arg \max_i Y_i$ as the best, though there exist procedures that use other selection criteria (Peng et al. 2016). In the event of a tie among estimators, a selection procedure uses a tie-breaking rule to make the final selection. Throughout our discussion, we assume that the estimators (Y_1, \dots, Y_n) have a joint probability density function, so that ties occur with probability zero. This assumption is made for convenience; all of the results should extend to the case when the estimators are discrete by using a randomized tie-breaking rule and properly accounting for the events of ties.

When necessary, we distinguish between selection procedures that use some form of screening – to eliminate inferior alternatives before the end of the procedure – and those that do not. Many early selection procedures do not use screening, e.g., those of Bechhofer (1954), Dudewicz and Dalal (1975), and Rinott (1978), while most contemporary selection procedures use screening to achieve greater sampling efficiency,

e.g., those of Nelson et al. (2001), Kim and Nelson (2001), Kim and Dieker (2011), Frazier (2014), and Zhong and Hong (2017).

3.1 Concentration Inequalities

As mentioned earlier, in the multi-armed bandit literature, it is commonly assumed that the marginal distributions F_i have bounded support or are sub-Gaussian with known scale (bound on the variance). These regularity conditions allow the use of concentration inequalities to prove PGS guarantees for multi-armed bandit algorithms. The general approach is two-staged: First, bound the probability of a bad alternative having a higher estimator than the best alternative using a concentration inequality. Second, use Bonferroni's inequality to sum over all pairwise comparisons with the best alternative. As an illustration, suppose that each marginal distribution has a support of $[0, 1]$ and that $\theta_i = \mathbb{E}[X_{i1}]$. Consider the "Naive" algorithm of Even-Dar et al. (2006) that takes $n = (2/\delta^2) \ln(2k/\alpha)$ observations independently from each alternative and selects the alternative with the highest sample mean, $Y_i = n^{-1} \sum_{j=1}^n X_{ij}$.

If a given bad alternative is selected, then its estimator may have exceeded its performance by $\delta/2$ or the estimator of the best alternative may have been less than $\delta/2$ below its performance. If neither of these events happened, then the bad alternative would not have been selected since the best alternative's estimator would have been higher. Thus, for a fixed bad alternative $i \neq [k]$,

$$\begin{aligned} P_\theta(Y_i > Y_{[k]}) &\leq P_\theta(Y_i > \theta_i + \delta/2 \text{ or } Y_{[k]} < \theta_{[k]} - \delta/2) \\ &\leq P_\theta(Y_i > \theta_i + \delta/2) + P_\theta(Y_{[k]} < \theta_{[k]} - \delta/2) \\ &\leq 2 \exp(-2(\delta/2)^2 n), \end{aligned}$$

where the last inequality follows from applying Hoeffding's inequality twice. From the choice of n , it follows that $P_\theta(Y_i > Y_{[k]}) \leq \alpha/k$.

If a bad alternative is selected, then it must be the case that the estimator for that alternative is greater than that of Alternative $[k]$. As a worst case, assume that all alternatives other than Alternative $[k]$ are bad. By Bonferroni's inequality and the above inequality,

$$P_\theta(\text{Bad Selection}) \leq P_\theta(\cup_{i \neq [k]} \{Y_i > Y_{[k]}\}) \leq \sum_{i \neq [k]} P_\theta(Y_i > Y_{[k]}) \leq (k-1)(\alpha/k) < \alpha.$$

Therefore, the Naive algorithm has a PGS guarantee.

In summary, the key step is using Hoeffding's inequality to bound the probability that an alternative's estimator differs from the alternative's performance by at least $\delta/2$ in a particular direction. If the marginal distributions do not have bounded supports, but are instead sub-Gaussian with known scale, then other concentration inequalities such as Chernoff's bound can be used for the same purpose. For some multi-armed bandit algorithms that eliminate alternatives in stages, e.g., the Successive Elimination and Median Elimination algorithms of Even-Dar et al. (2002) and Even-Dar et al. (2006), concentration inequalities have been used similarly to prove PGS guarantees.

In contrast, the R&S literature typically assumes that observations are normally distributed, i.e., $X_{i1} \sim \mathcal{N}(\theta_i, \sigma_i^2)$. While the aforementioned concentration inequalities are convenient for proving PGS guarantees under the standard multi-armed bandit assumptions, they cannot be applied in this setting without knowing the variances, or at least an upper bound on the variances. Instead, the idea of constructing confidence bands around each alternative's estimator has been used to prove a PGS guarantee for the fully sequential envelope procedure (EP) of Ma and Henderson (2017). Unlike procedures that use screening, EP keeps all alternatives in contention until the procedure terminates, but still offers flexibility in how observations are allocated across alternatives. The procedure constructs lower and upper confidence limits for the performances of each alternative and updates these limits as more observations are taken. The confidence limits are designed to guarantee that with a probability of at least $1 - \alpha$, the upper confidence

limit of the best alternative stays above its true performance while the lower confidence limits of the other alternatives stay below their true performances *throughout the entire procedure*. When these bounds hold, a good selection is implied by terminating when the lower confidence limit of the estimated best alternative exceeds the highest upper confidence limit of the other alternatives minus δ (Ma and Henderson 2017).

3.2 Lifting IZ-Inspired PCS Guarantees

Another approach to proving PGS guarantees is to first prove that a selection procedure has an IZ-inspired PCS guarantee and then show that this implies a PGS guarantee. In terms of the two guarantees, a PGS guarantee implies an IZ-inspired PCS guarantee because for configurations in the preference zone, the only good alternative is the best. One might naturally wonder whether the converse holds: does an IZ-inspired PCS guarantee imply a PGS guarantee over the entire space of configurations? Intuitively, one might expect this to be the case since for problem instances with more than one good alternative, it should be more likely that one of them is selected. Put differently, if a procedure can select the best alternative when all other alternatives have δ -worse performance, why would it do any worse when there are more good alternatives?

Unfortunately, the answer to the posed question is negative. In terms of a counter-example, it is possible to contrive a selection procedure with an IZ-inspired PCS guarantee but without a PGS guarantee (Eckman and Henderson 2018a). Specifically, when the estimators of the two best-looking alternatives are similar in value, the procedure paradoxically selects one of the other alternatives. Thus when there is a single δ -best alternative, the PCS is high, but when there are multiple good alternatives, the PGS can suffer.

Although the converse implication does not universally hold, many R&S procedures with IZ-inspired PCS guarantees have been shown to have PGS guarantees (Matejcik and Nelson 1995), and many others may very well have (yet unproven) PGS guarantees, simply by not behaving bizarrely.

A logical approach to proving the converse is to show that for any configuration in the indifference zone, there is a corresponding configuration in the preference zone with a lower PGS, which is equivalent to the PCS. In particular, we consider pairing an arbitrary configuration θ in the indifference zone with another configuration θ^* for which all of the good – but not the best – alternatives of θ are shifted to become bad alternatives; thus θ^* is in the preference zone. To be precise, for all $i \in \{1, \dots, k-1\}$ for which $\theta_{[i]} > \theta_{[k]} - \delta$, we set $\theta_{[i]}^* = \theta_{[k]} - \delta$, and for all other i , including k , we take $\theta_{[i]}^* = \theta_{[i]}$. The critical question is then: What are sufficient conditions under which for any configuration θ in the indifference zone, the PGS in θ is higher than the PGS in its paired configuration θ^* ?

Research into sufficient conditions dates back to the work of Fabian (1962) and has led to various results, many of them involving overly restrictive assumptions on the selection procedures. For example, Chiu (1974a) and Chiu (1974b) assume that the probability that the best alternative has a higher estimator than any bad alternative is monotone with respect to the differences between performances of consecutively ordered alternatives. Similarly, Feigin and Weissman (1981) and Chen (1982) assume that the estimators of the alternatives' performances are mutually independent and come from the same family of stochastically increasing distributions having different location parameters. Furthermore, Nelson and Matejcik (1995) effectively assume that the estimators of the alternatives' performances are shift-invariant with respect to the vector of performances. These assumptions can all be generalized into the following condition:

Condition 1 (Guiard 1996) For any subset $A \subset \{1, \dots, k\}$, the joint distribution of the estimators $\{Y_i : i \in A\}$ does not depend on θ_j , for all $j \notin A$.

Condition 1 states that changing the performance of an alternative or subset of alternatives does not change the joint distribution of the estimators of the other alternatives. Condition 1 thus allows one to compare events involving the estimators of a subset of alternatives between two configurations in which the performances of alternatives in that subset are the same.

A short proof that Condition 1, when combined with an IZ-inspired PCS guarantee, implies a PGS guarantee is as follows. Let P_θ and P_{θ^*} denote the probability measures under configurations θ and θ^* , respectively, and for a given alternative, i , let Y_i and Y_i^* be the estimators of Alternative i under configurations

θ and θ^* . One way for a good selection to occur under configuration θ is for Alternative $[k]$ to have a higher estimator than every bad alternative. Condition 1 can then be applied, taking $A = \{[k]\} \cup \{i : \theta_i \leq \theta_{[k]} - \delta\}$:

$$\begin{aligned} P_\theta(\text{Good Selection}) &\geq P_\theta(Y_{[k]} > Y_i \text{ for all } i \text{ such that } \theta_i \leq \theta_{[k]} - \delta) \\ &= P_{\theta^*}(Y_{[k]}^* > Y_i^* \text{ for all } i \text{ such that } \theta_i \leq \theta_{[k]} - \delta) \\ &\geq P_{\theta^*}(Y_{[k]}^* > Y_i^* \text{ for all } i \neq [k]) \\ &= P_{\theta^*}(\text{Correct Selection}) \geq 1 - \alpha. \end{aligned}$$

An important special case that implies Condition 1 is when the estimators of the performance of each alternative are mutually independent, as is the case for many selection procedures in which alternatives are sampled independently and no screening is used, e.g., those of Bechhofer (1954), Rinott (1978) and Dudewicz and Dalal (1975). For this reason, Condition 1 can even be used for problems in which the vector of performances is not a location parameter of the joint distribution F . For example, it can be used to show that the procedure of Sobel and Huyett (1957) for selecting the best Bernoulli population delivers a PGS guarantee. Condition 1 can also be satisfied for some procedures that use common random numbers across alternatives, such as those of Clark and Yang (1986) and Nelson and Matejcek (1995).

Another special case of Condition 1 is when the joint distribution of the estimators is shift-invariant with respect to the configuration, i.e., for any configurations θ and $\tilde{\theta}$,

$$(Y_1 - \theta_1, \dots, Y_k - \theta_k) \stackrel{d}{=} (\tilde{Y}_1 - \tilde{\theta}_1, \dots, \tilde{Y}_k - \tilde{\theta}_k),$$

where $\stackrel{d}{=}$ denotes equivalence in distribution. In this case, a shift in the vector of performances θ produces an equivalent shift in the joint distribution of the estimators Y_1, \dots, Y_k . This situation is slightly stronger than the condition in Theorem 1 of Nelson and Matejcek (1995). In order for this special case to be satisfied, θ would likely need to be a location parameter of the joint distribution of the observations F . This would be the case, for example, if $X_{i1} \sim \mathcal{N}(\theta_i, \sigma_i^2)$ and alternatives were sampled independently. In this setting, θ_i could also represent some other location parameter of F_i , such as the median or q -quantile.

A major shortcoming of Condition 1 is that it is unlikely to hold for modern procedures that use screening to eliminate inferior alternatives. This is because changing the performance of a given alternative can either change the set of alternatives that survive the screening stage, or affect the sequencing of alternatives that are eliminated (Nelson et al. 2001; Kim and Nelson 2001). We now discuss another condition that lifts an IZ-inspired PCS guarantee to a PGS guarantee and seems more likely to hold for procedures that use screening.

Condition 2 (Hayter 1994): For all alternatives $i = 1, \dots, k$, $P_\theta(\text{Select Alternative } i)$ is non-increasing in θ_j for every $j \neq i$.

Condition 2 states that increasing the performance of an alternative will not improve the likelihood that any other alternative is selected. And since

$$P_\theta(\text{Select Alternative } j) = 1 - \sum_{i \neq j} P_\theta(\text{Select Alternative } i),$$

Condition 2 also implies that increasing the performance of an alternative does not decrease the probability of selecting that alternative.

The proof that an IZ-inspired PCS guarantee implies a PGS guarantee for selection procedures satisfying Condition 2 makes use of the following consequence: increasing the performance of a bad alternative until it is good will not increase the probability that any other bad alternative is selected, and, thus, the probability of making a good selection will not decrease. Running this argument in reverse says that decreasing the performance of a good alternative will not increase the PGS. Hence, for an arbitrary configuration θ in the indifference zone, iteratively reducing the performances of all good (but not best) alternatives until they

are at $\theta_{[k]} - \delta$ will not increase the PGS. The PGS in configuration θ will, therefore, be no less than the PGS of configuration θ^* , which is at least $1 - \alpha$.

Unlike Condition 1, Condition 2 is difficult to verify since deriving an expression for the probability that a given alternative is selected is often complicated, if not futile. Indeed, Hayter (1994) remarks that even formulating stronger conditions that imply Condition 2 is challenging. Despite the difficulty of verifying Condition 2, one might expect it to hold for reasonable selection procedures, even those that use screening. All else being equal, increasing the performance of an alternative *should* increase the probability of that alternative being selected, at the expense of every other alternative. In other words, there is no apparent reason that another alternative would benefit from having a stronger competitor. Unfortunately, screening procedures involve complicated dependencies among elimination decisions that can undermine this intuition. That is, changing the performance of an alternative can change the times of elimination decisions involving that alternative, and in turn affect future elimination decisions.

Hayter (1994) provides a counter-example of a selection procedure with screening that does not satisfy Condition 2. Suppose there are three alternatives with performances $\theta_1 < \theta_2 < \theta_3$, i.e., Alternative 3 is the best, and consider the following simple procedure that takes observations independently across alternatives:

1. Take n_0 observations from each alternative.
2. Eliminate all but the two alternatives with the highest estimated performances.
3. Take n_1 additional observations from the two surviving alternatives.
4. Select the alternative with the higher overall estimated performance.

We will explain how for this procedure, increasing θ_1 may – for some parameter settings – *increase* the probability that Alternative 2 is selected, thereby violating Condition 2.

Fix the initial sample size $n_0 \geq 1$ and consider the extreme cases where $n_1 = 0$ and $n_1 = \infty$. When $n_1 = 0$, the procedure selects the alternative with the highest estimated first-stage performance, since there is no second stage. Therefore, increasing θ_1 will increase the probability that Alternative 1 is selected while decreasing the probabilities that Alternatives 2 and 3 are selected.

When $n_1 = \infty$, the procedure exhaustively simulates the two alternatives that receive second-stage sampling and so always selects the better of the two. Although an infinite simulation budget is impractical, this case is intended to reflect extensive sampling of alternatives that survive screening. In this case, the probability of selecting Alternative 1 is zero since it will lose any pairwise comparison if it survives screening. In contrast, Alternative 3 is always selected when it survives screening, meaning that the probability of selecting Alternative 3 is exactly the probability that it survives screening. This probability decreases as θ_1 is increased, but not past θ_2 . From these two relationships, it follows that the probability of selecting Alternative 2 must increase as θ_1 is increased. It is possible to find a large, finite value of n_1 for which this selection procedure violates Condition 2.

This counter-example demonstrates how using screening in a selection procedure leads to complicated dependencies between alternatives that may result in non-monotone changes in the probabilities of selecting alternatives. Fairweather (1968) is an example of a selection procedure with a PCS guarantee over the preference zone that might violate Condition 2. Likewise, the NSGS procedure of Nelson et al. (2001) could possibly violate Condition 2 if a high value of α_0 is used in the screening stage and then a small (conservative) value of α_1 is used in the selection stage.

Some researchers have attempted to directly prove PGS guarantees for procedures that use screening and deliver IZ-inspired PCS guarantees, with mixed results. For example, Osogami (2009) uses Anderson's probability bound (Anderson 1960) to show that a PGS guarantee holds for some parameter settings of his sequential procedure. In addition, Fan et al. (2016) show that one of their indifference-zone-free procedures has an asymptotic PGS guarantee. On the other hand, the purported proofs of PGS guarantees in Kim and Nelson (2001) and Pichitlamken et al. (2006) are incorrect. Another approach is to adapt sequential selection procedures with IZ-inspired PCS guarantees by pushing out the boundaries of the continuation

regions, effectively weakening the elimination rule (Kao and Lai 1980; Jennison et al. 1982). However, these modifications result in overly conservative procedures.

3.3 Multiple Comparisons

Another classical approach to proving PGS guarantees comes from multiple comparisons, a field of statistics concerned with constructing simultaneous confidence statements about the performances of alternatives. In the multiple comparisons literature, three events related to good selection are

$$\begin{aligned}\mathcal{A} &= \{Y_{[i]} - Y_{[j]} - (\theta_{[i]} - \theta_{[j]}) < \delta, \forall i \neq j\}, \\ \mathcal{B} &= \{Y_{[i]} - Y_{[k]} - (\theta_{[i]} - \theta_{[k]}) < \delta, \forall i \neq k\}, \text{ and} \\ \mathcal{C} &= \{Y_{[i]} - Y_{[k]} - (\theta_{[i]} - \theta_{[k]}) < \max(\delta, \theta_{[k]} - \theta_{[i]}), \forall i \neq k\}.\end{aligned}$$

Event \mathcal{A} is called multiple comparisons all pairwise and states that the differences in performances between all pairs of alternatives have all been estimated within δ of their true values. Similarly, event \mathcal{B} is called multiple comparisons with the best and states that the differences in performances between each alternative and the best are within δ of their true values. Event \mathcal{C} is a slight relaxation of event \mathcal{B} , and it can be seen that $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{C}$.

Each of the three events independently implies good selection, with proofs given by Hsu (1996) and Nelson and Banerjee (2001) and reproduced in Lee and Nelson (2017). As an example, event \mathcal{C} implies good selection since for any bad alternative $[i]$, i.e., $\theta_{[i]} \leq \theta_{[k]} - \delta$, event \mathcal{C} implies that $Y_{[i]} < Y_{[k]}$, and, therefore, Alternative $[i]$ will not be selected. If a selection procedure can guarantee that any of the events \mathcal{A} , \mathcal{B} , or \mathcal{C} occurs with a probability greater than or equal to $1 - \alpha$, then the procedure has a PGS guarantee with the same confidence level. For example, Event \mathcal{B} is used by Nelson and Matejcek (1995) to show that some procedures with IZ-inspired PCS guarantees also deliver PGS guarantees. It is also used to prove the PGS guarantee for procedure GSP' of Ni et al. (2017). Meanwhile, event \mathcal{C} is used by Nelson and Banerjee (2001) to produce lower confidence bounds on the PGS for a given selection procedure. The three events are also used in the asymptotic PGS guarantees of the bootstrapping procedures of Lee and Nelson (2017), discussed later in Section 4.1. We direct the interested reader to Hsu (1996) for other examples of classical procedures that incorporate these ideas.

4 RELATED PGS GUARANTEES

In this section, we discuss variants of the frequentist PGS guarantee studied in Sections 2 and 3, namely the asymptotic PGS guarantee that arises from bootstrapping and the posterior PGS guarantee under a Bayesian framework.

4.1 Bootstrapping

Recently, selection procedures have been designed that leverage bootstrapping. Lee and Nelson (2016) and Lee and Nelson (2017) introduce a family of procedures that make no distributional assumptions about the observations from the alternatives and allow for the use of common random numbers. The procedures take an equal number of observations, n , from each alternative and use bootstrapping to estimate the probability of events \mathcal{A} , \mathcal{B} , and \mathcal{C} discussed in Section 3.3. Depending on whether common random numbers are used, the bootstrap samples are drawn from either the marginal or joint empirical distributions. Since all three events imply good selection, the procedures terminate when the bootstrap probability of an event exceeds $1 - \alpha$ and select the alternative with the best estimated performance. This selection rule is used because the procedures are determining a common sample size for which the bootstrap probability that the alternative with the best estimated performance is a good alternative exceeds $1 - \alpha$.

These bootstrapping procedures estimate the frequentist probabilities that events \mathcal{A} , \mathcal{B} , or \mathcal{C} occur, given the observed data. In other words, bootstrapping treats the empirical distributions as the true

distributions of the observations from the alternatives and estimates the probability that a given event occurs if n observations were taken from each alternative. In contrast, as we will see in Section 4.2, Bayesian R&S procedures calculate the posterior probability of an event, given the collected observations *and* a prior distribution.

A limitation of such bootstrapping procedures is that only asymptotic PGS guarantees can be proved due to two consequences of finite sample sizes. First, the empirical distributions will not match the true distributions. Second, in the bootstrapping setting, the definitions of events \mathcal{B} and \mathcal{C} rely on an *estimate* of the identity of the best alternative and its performance. Thus, for finite sample sizes, the estimated best alternative may not be the true best alternative. The procedures estimating the probabilities of events \mathcal{A} and \mathcal{B} deliver asymptotic PGS guarantees as $\delta \downarrow 0$, whereas the procedure estimating the probability of event \mathcal{C} involves an additional limit requiring that the relative spacing of the performances is maintained as the configuration contracts.

Another drawback of these bootstrapping procedures is that they must take an equal number of observations from all alternatives to maintain their asymptotic validity. This means that alternatives that appear to be clearly inferior cannot be eliminated early. Therefore, when the number of alternatives is large, these procedures can be inefficient.

4.2 Bayesian PGS

The frequentist R&S problem discussed in Sections 2 and 3 can also be examined from a Bayesian perspective (Berger and Deely 1988; Gupta and Miescke 1996). Under the Bayesian interpretation, the problem instance is treated as a random variable, and expectations are taken over a posterior distribution, after incorporating observed evidence and a prior distribution. Accordingly, the parameters of the joint distribution F are viewed as random variables drawn from distributions with known families but unknown parameters. For example, in the Bayesian R&S literature, it is commonly assumed that the vector of observations (X_{1j}, \dots, X_{kj}) comes from a multivariate normal distribution. When alternatives are simulated independently, the observations from Alternative i are distributed as $X_{ij} \sim \mathcal{N}(W_i, \sigma_i^2)$ where the mean W_i and variance σ_i^2 are random and for computational convenience, the normal-gamma conjugate pair on $(W_i, 1/\sigma_i^2)$ is often used (Chick and Inoue 2001b). On the other hand, when alternatives are simulated using common random numbers, the normal-Wishart conjugate pair can be used for the mean vector and precision matrix (Chick and Inoue 2001a). In either case, the configuration is represented by the random mean vector $W = (W_1, \dots, W_k)$, in contrast to the frequentist view of the problem instance as being unknown but fixed.

To use a Bayesian selection procedure, a decision-maker begins with prior distributions on the unknown parameters of F . After taking observations from the alternatives, the standard Bayesian machinery is used to update the hyperparameters of the posterior distributions. The posterior distributions on the parameters of F encapsulate the decision-maker's beliefs on the unknown problem instance, given the observed evidence. For instance, when observations are normally distributed as described above, the posterior distribution of W (often a multivariate t -distribution) can be used to calculate the *posterior* probability of an event involving the alternatives' performances. Specifically, the posterior PGS for a given alternative, i , is defined as

$$\text{PGS}_i^{\text{Bayes}} = P(W_i > W_j - \delta \quad \forall j \neq i),$$

where the probability is taken with respect to the posterior distribution of W , given the observed evidence.

The posterior PGS for a given alternative amounts to an integral of the posterior distribution of the mean vector over a polyhedron defined by $k - 1$ inequalities of the form $W_i - W_j > -\delta$ for all $j \neq i$. This polyhedron corresponds to the event that the performance of Alternative i is no worse than δ below the maximum of the others. When the number of alternatives is small, this integral might be computed numerically, and when the number of alternatives is large, it can be estimated using Monte Carlo methods, although the latter introduces a frequentist error estimate that deserves consideration. Alternatively, Slepian's inequality (Slepian 1962) can be used to obtain a lower bound on the posterior PGS of a given alternative that can

be easily computed (Branke et al. 2005). This lower bound can be used as a stopping rule for expected value of information (VIP) and optimal computing budget allocation (OCBA) procedures (Chick and Inoue 2001b; Chen et al. 2000). A similar approach is taken to bound the PCS in the enhanced two-stage selection (ETSS) procedure (Chen and Kelton 2005).

The posterior PGS for a given alternative can be interpreted as the probability that the problem instance is one for which that alternative is good. If Alternative i were selected by a procedure with a Bayesian posterior PGS guarantee, then based on the evidence, with probability exceeding $1 - \alpha$, the *random* problem instance is one for which the *fixed* Alternative i has a mean within δ of the best. In contrast, for a selection procedure with a frequentist PGS guarantee, with probability exceeding $1 - \alpha$, the *random* alternative selected by the procedure has a mean within δ of the best for this *fixed* problem instance. In other words, if the selection procedure were run many times on the same problem instance, on at least $(1 - \alpha) \times 100\%$ of those runs, the procedure would select a good alternative. A potential drawback of this interpretation is that in practice, the decision-maker is likely to run the procedure only once.

Compared to its frequentist counterpart, the Bayesian concept of posterior PGS has several advantages. Foremost among these is that it can be recalculated as more observations are collected without affecting a procedure's validity, while avoiding issues with repeatedly looking at the data that arise in the frequentist framework. As a consequence, Bayesian selection procedures can use the posterior PGS for the alternatives as a stopping condition, terminating when the posterior PGS of an alternative exceeds $1 - \alpha$ and selecting that alternative (Branke et al. 2005; Russo 2016).

Another appealing aspect of posterior PGS is that Bayesian selection procedures typically use only the desired confidence level, $1 - \alpha$, in setting a stopping condition. Thus, allocation rules that control how many observations are taken from each alternative can be designed without regard for $1 - \alpha$, and are, therefore, free of any need to consider a least favorable configuration. For frequentist selection procedures, $1 - \alpha$ must be specified in advance and the procedures are then designed to deliver a PGS guarantee for any configuration, making them inherently conservative. Unsurprisingly, empirical comparisons have shown that selection procedures designed to deliver Bayesian fixed-confidence guarantees are often more efficient than ones with frequentist guarantees (Branke et al. 2007).

5 CONCLUSIONS

We reviewed PGS guarantees as they appear in the R&S and multi-armed bandit literature and surveyed several proof techniques, including concentration inequalities and multiple comparisons. We also argued that PGS guarantees are superior to the popular IZ-inspired PCS guarantees and demonstrated how – under certain conditions – the latter can imply the former. In our discussion, we illustrated how the analysis of selection procedures that use screening is complicated by the dependencies among alternatives and their estimators. It remains an open question whether such selection procedures that deliver IZ-inspired PCS guarantees, e.g., the fully sequential procedures of Paulson (1964) and Kim and Nelson (2001), also deliver PGS guarantees. Verifying Condition 2 for these procedures may be a promising approach. Lastly, we examined some of the relative advantages and disadvantages of PGS guarantees for bootstrapping and Bayesian selection procedures.

The idea of delivering a fixed-confidence guarantee for selecting a δ -optimal alternative can also be applied to problems for which the number of alternatives is infinite. As an example, an intriguing area for future research would be stochastic optimization problems with continuous domains. In this setting, one might explore how to design algorithms that guarantee with high probability to return a solution whose objective function value is within δ of a local optimum. Aside from regularity conditions on the error associated with estimating the objective function, some structural assumptions about the objective function would be essential.

Rather than having a single alternative be selected for them, decision-makers may prefer to obtain a small subset of alternatives from which they can make a final decision based on some additional criteria.

For this reason, more research is needed on subset-selection procedures with fixed-confidence guarantees on the quality of the selected alternatives.

ACKNOWLEDGMENTS

We thank Steve Chick, Jeff Hong, and Barry Nelson for valuable conversations. We also thank the reviewers for their detailed suggestions for improving this tutorial. This work is supported by the National Science Foundation under grants DGE-1650441 and CMMI-1537394 and by the Army Research Office under grant W911NF-17-1-0094.

REFERENCES

- Anderson, T. W. 1960. "A Modification of the Sequential Probability Ratio Test to Reduce the Sample Size". *Annals of Mathematical Statistics* 31(1):165–197.
- Bechhofer, R. E. 1954. "A Single-Sample Multiple Decision Procedure for Ranking Means of Normal Populations with Known Variances". *Annals of Mathematical Statistics* 25(1):16–39.
- Berger, J. O., and J. Deely. 1988. "A Bayesian Approach to Ranking and Selection of Related Means with Alternatives to Analysis-of-Variance Methodology". *Journal of the American Statistical Association* 83(402):364–373.
- Blanchet, J., J. Liu, and B. Zwart. 2008. "Large Deviations Perspective on Ordinal Optimization of Heavy-Tailed Systems". In *Proceedings of the 2008 Winter Simulation Conference*, edited by S. J. Mason et al., 489–494. Piscataway, New Jersey: IEEE.
- Branke, J., S. E. Chick, and C. Schmidt. 2005. "New Developments in Ranking and Selection: An Empirical Comparison of the Three Main Approaches". In *Proceedings of the 2005 Winter Simulation Conference*, edited by M. E. Kuhl et al., 708–717. Piscataway, New Jersey: IEEE.
- Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a Selection Procedure". *Management Science* 53(12):1916–1932.
- Broadie, M., M. Han, and A. Zeevi. 2007. "Implications of Heavy Tails on Simulation-based Ordinal Optimization". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson et al., 439–447. Piscataway, New Jersey: IEEE.
- Bubeck, S., R. Munos, and G. Stoltz. 2011. "Pure Exploration in Finitely-armed and Continuous-armed Bandits". *Theoretical Computer Science* 412(19):1832–1852.
- Chen, C.-H., S. E. Chick, L. H. Lee, and N. A. Pujowidianto. 2015. "Ranking and Selection: Efficient Simulation Budget Allocation". In *Handbook of Simulation Optimization*, edited by M. Fu, 45–80. New York: Springer.
- 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, E. J., and W. D. Kelton. 2005. "Sequential Selection Procedures: Using Sample Means to Improve Efficiency". *European Journal of Operational Research* 166(1):133–153.
- Chen, P. 1982. *An Alternative Definition of Correct Selection in Ranking and Selection Problems*. Ph. D. thesis, UC Santa Barbara, Santa Barbara, California.
- Chick, S. E., and K. Inoue. 2001a. "New Procedures to Select the Best Simulated System Using Common Random Numbers". *Management Science* 47(8):1133–1149.
- Chick, S. E., and K. Inoue. 2001b. "New Two-Stage and Sequential Procedures for Selecting the Best Simulated System". *Operations Research* 49(5):732–743.
- Chick, S. E., and Y. Wu. 2005. "Selection Procedures with Frequentist Expected Opportunity Cost Bounds". *Operations Research* 53(5):867–878.
- Chiu, W. K. 1974a. "A Generalized Selection Goal for a Ranking Problem". *South African Statistical Journal* 8(1):45–48.

- Chiu, W. K. 1974b. “The Ranking of Means of Normal Populations for a Generalized Selection Goal”. *Biometrika* 61(3):579–584.
- Clark, G. M., and W.-N. Yang. 1986. “A Bonferroni Selection Procedure When Using Common Random Numbers with Unknown Variances”. In *Proceedings of the 1986 Winter Simulation Conference*, edited by J. R. Wilson et al., 313–315. New York: IEEE.
- Dudewicz, E. J., and S. R. Dalal. 1975. “Allocation of Observations in Ranking and Selection with Unequal Variances”. *Sankhyā: The Indian Journal of Statistics* 37(1):28–78.
- Eckman, D. J., and S. G. Henderson. 2018a. *Fixed-Confidence, Fixed-Tolerance Guarantees for Selection-of-the-Best Procedures*. Working paper, Cornell University, School of Operations Research and Information Engineering, <http://people.orie.cornell.edu/dje88/mypdfs/PGS-overview-paper.pdf>.
- Eckman, D. J., and S. G. Henderson. 2018b. “Reusing Search Data in Ranking and Selection: What Could Possibly Go Wrong?”. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 28(3):18:1–18:15.
- Even-Dar, E., S. Mannor, and Y. Mansour. 2002. “PAC Bounds for Multi-Armed Bandit and Markov Decision Processes”. In *Fifteenth Annual Conference on Computational Learning Theory (COLT)*, edited by J. Kivinen and R. H. Sloan, 255–270. Berlin: Springer.
- Even-Dar, E., S. Mannor, and Y. Mansour. 2006. “Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems”. *Journal of Machine Learning Research* 7:1079–1105.
- Fabian, V. 1962. “On Multiple Decision Methods for Ranking Population Means”. *Annals of Mathematical Statistics* 33(1):248–254.
- Fairweather, W. R. 1968. “Some Extensions of Somerville’s Procedure for Ranking Means of Normal Populations”. *Biometrika* 55(2):411–418.
- Fan, W., L. J. Hong, and B. L. Nelson. 2016. “Indifference-Zone-Free Selection of the Best”. *Operations Research* 64(6):1499–1514.
- Feigin, P. D., and I. Weissman. 1981. “On the Indifference Zone Approach to Selection – A Consistency Result”. *Annals of the Institute of Statistical Mathematics* 33(1):471–474.
- Frazier, P. 2014. “A Fully Sequential Elimination Procedure for Indifference-Zone Ranking and Selection with Tight Bounds on Probability of Correct Selection”. *Operations Research* 62(4):926–942.
- 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 et al., 577–585. Piscataway, New Jersey: IEEE.
- Glynn, P., and S. Juneja. 2015. “Selecting the Best System, Large Deviations, and Multi-armed Bandits”. *ArXiv e-prints*. Preprint arXiv:1507.04564v2, <http://arxiv.org/abs/1507.04564v2>.
- Guiard, V. 1996. “Different Definitions of Δ -Correct Selection for the Indifference Zone Formulation”. *Journal of Statistical Planning and Inference* 54(2):175–199.
- Gupta, S. S., and K. J. Miescke. 1996. “Bayesian Look Ahead One-Stage Sampling Allocations for Selection of the Best Population”. *Journal of Statistical Planning and Inference* 54(2):229–244.
- Hayter, J. 1994. “On the Selection Probabilities of Two-Stage Decision Procedures”. *Journal of Statistical Planning and Inference* 38(2):223–236.
- Ho, Y.-C., C. G. Cassandras, C.-H. Chen, and L. Dai. 2000. “Ordinal Optimisation and Simulation”. *Journal of the Operational Research Society* 51(4):490–500.
- Hsu, J. 1996. *Multiple Comparisons: Theory and Methods*. New York: CRC Press.
- Hunter, S. R., and R. Pasupathy. 2010. “Large-Deviation Sampling Laws for Constrained Simulation Optimization on Finite Sets”. In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson et al., 995–1002. Piscataway, New Jersey: IEEE.
- Jennison, C., I. M. Johnstone, and B. W. Turnbull. 1982. “Asymptotically Optimal Procedures for Sequential Adaptive Selection of the Best of Several Normal Means”. In *Statistical Decision Theory and Related Topics III*, edited by S. S. Gupta and J. O. Berger, Volume 2, 55–86. New York: Elsevier.

- Kao, S. C., and T. L. Lai. 1980. "Sequential Selection Procedures Based on Confidence Sequences for Normal Populations". *Communications in Statistics – Theory and Methods* A9(16):1657–1676.
- Karnin, Z., T. Koren, and O. Somekh. 2013. "Almost Optimal Exploration in Multi-armed Bandits". In *Proceedings of the 30th International Conference on Machine Learning*, edited by S. Dasgupta and D. McAllester, 1238–1246. Stroudsburg, Pennsylvania: IMLS.
- Kiekhäfer, A. P. 2011. *Simulation Ranking and Selection Procedures and Applications in Network Reliability Design*. Ph. D. thesis, University of Iowa, Iowa City, Iowa.
- Kim, S.-H., and A. B. Dieker. 2011. "Selecting the Best by Comparing Simulated Systems in a Group of Three". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain et al., 3987–3997. Piscataway, New Jersey: IEEE.
- Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-Zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11(3):251–273.
- Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, 501–534. Amsterdam: Elsevier.
- Koenig, L. W., and A. M. Law. 1985. "A Procedure for Selecting a Subset of Size m Containing the l Best of k Independent Normal Populations, with Applications to Simulation". *Communications in Statistics – Simulation and Computation* 14(3):719–734.
- Lau, T. W. E., and Y. C. Ho. 1997. "Universal Alignment Probabilities and Subset Selection for Ordinal Optimization". *Journal of Optimization Theory and Applications* 93(3):455–489.
- Lee, S., and B. L. Nelson. 2016. "General-Purpose Ranking and Selection for Computer Simulation". *IIE Transactions* 48(6):555–564.
- Lee, S., and B. L. Nelson. 2017. *An Omnibus Ranking-and-Selection Procedure*. Working Paper, Northwestern University, Department of Industrial Engineering & Management Sciences.
- Ma, S., and S. G. Henderson. 2017. "An Efficient Fully Sequential Selection Procedure Guaranteeing Probably Approximately Correct Selection". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan et al., 2225–2236. Piscataway, New Jersey: IEEE.
- Mannor, S., and J. N. Tsitsiklis. 2004. "The Sample Complexity of Exploration in the Multi-armed Bandit Problem". *Journal of Machine Learning Research* 5:623–648.
- Matejčík, F. J., and B. L. Nelson. 1995. "Two-Stage Multiple Comparisons with the Best for Computer Simulation". *Operations Research* 43(4):633–640.
- Nelson, B. L., and S. Banerjee. 2001. "Selecting a Good System: Procedures and Inference". *IIE Transactions* 33(3):149–166.
- Nelson, B. L., and F. Matejčík. 1995. "Using Common Random Numbers for Indifference-Zone Selection and Multiple Comparisons in Simulation". *Management Science* 41(12):1935–1945.
- Nelson, B. L., J. Swann, D. Goldsman, and W. Song. 2001. "Simple Procedures for Selecting the Best Simulated System When the Number of Alternatives Is Large". *Operations Research* 49(6):950–963.
- Ni, E., D. Ciocan, S. Henderson, and S. Hunter. 2017. "Efficient Ranking and Selection in Parallel Computing Environments". *Operations Research* 65(3):821–836.
- Osogami, T. 2009. "Finding Probably Best Systems Quickly via Simulations". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 19(3):12:1–12:19.
- Parnes, M., and R. Srinivasan. 1986. "Some Inconsistencies and Errors in the Indifference Zone Formulation of Selection". *Sankhyā: The Indian Journal of Statistics* 48(1):86–97.
- Paulson, E. 1964. "A Sequential Procedure for Selecting the Population with the Largest Mean from k Normal Populations". *Annals of Mathematical Statistics* 35(1):174–180.
- 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.
- Pichitlamken, J., B. L. Nelson, and L. J. Hong. 2006. "A Sequential Procedure for Neighborhood Selection-of-the-Best in Optimization via Simulation". *European Journal of Operational Research* 173(1):283–298.

- Rinott, Y. 1978. “On Two-Stage Selection Procedures and Related Probability Inequalities”. *Communications in Statistics – Theory and Methods* 7(8):799–811.
- Russo, D. 2016. “Simple Bayesian Algorithms for Best-Arm Identification”. *ArXiv e-prints*. Preprint arXiv:1602.08448v4, <http://arxiv.org/abs/1602.08448v4>.
- Slepian, D. 1962. “The One-Sided Barrier Problem for Gaussian Noise”. *Bell Labs Technical Journal* 41(2):463–501.
- Sobel, M., and M. J. Huyett. 1957. “Selecting the Best One of Several Binomial Populations”. *Bell Labs Technical Journal* 36(2):537–576.
- Zhong, Y., and L. J. Hong. 2017. “A New Framework of Designing Sequential Ranking-and-Selection Procedures”. In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan et al., 2237–2244. Piscataway, New Jersey: IEEE.
- Zhou, Y., X. Chen, and J. Li. 2014. “Optimal PAC Multiple Arm Identification with Applications to Crowdsourcing”. In *Proceedings of the 31st International Conference on Machine Learning*, edited by E. P. Xing and T. Jebara, 217–225. Stroudsburg, Pennsylvania: IMLS.

AUTHOR BIOGRAPHIES

DAVID J. ECKMAN is a Ph.D. candidate in the School of Operations Research and Information Engineering at Cornell University. His research interests deal with simulation optimization with a focus on ranking and selection and exploratory multi-armed bandits. His email address is dje88@cornell.edu and his website is <https://people.orie.cornell.edu/dje88/>.

SHANE G. HENDERSON is a professor in the School of Operations Research and Information Engineering at Cornell University. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services and bike sharing applications. He co-edited the Proceedings of the 2007 Winter Simulation Conference. His email address is sg9@cornell.edu and his web page is <http://people.orie.cornell.edu/~shane>.