

PERFORMANCE MEASURES FOR RANKING AND SELECTION PROCEDURES

Rolf Waeber
Peter Frazier
Shane Henderson

Department of Operations Research and Information Engineering
Cornell University
Ithaca, NY 14850, USA

ABSTRACT

To efficiently compare Ranking and Selection procedures, we present a three-layer performance evaluation process. The two most popular formulations, namely the Bayes and Indifference Zone formulations, have a common representation analogous to convex risk measures used in quantitative risk management. We study decision makers' acceptance sets via an axiomatic approach and introduce a new performance measure using computational cost.

1 MOTIVATION

In ranking and selection (R&S), we use stochastic simulation to determine which of several simulated systems is the best. This need arises, for example, if we have a simulation of a call center that we would like to use to choose among several possible staffing policies. We would first decide how many samples of the customer waiting time to take using each staffing policy, and would then use these samples to decide which system has the most desirable waiting time distribution. The crucial decision in R&S is choosing how many samples to take from each system. If we make this decision intelligently, we can find a good system more reliably and in less time than can be found using naive sampling strategies.

An enormous number of R&S procedures have been introduced in the literature, and two opposing theoretical frameworks have been introduced for evaluating the performance of these procedures. The first framework is the Indifference Zone (IZ) framework, which considers the worst-case performance of an R&S policy over a specific set of configurations. The second framework is the Bayesian framework, which considers average performance under a prior probability distribution. Each framework has advantages and disadvantages. The IZ framework provides a statistical guarantee on performance that provides peace of mind, but is often extremely conservative in the number of samples that it requires an R&S procedure to take. The Bayesian framework, in contrast, results in R&S procedures that perform well in many cases, but might perform badly in certain configurations. The two frameworks seem to be at odds with each other, presenting radically different approaches to the R&S problem.

In this paper, we introduce a new and more general theoretical framework that includes both the IZ and Bayesian formulations of the R&S problem, as well as a continuum of other frameworks that strike a balance between the conservatism of IZ and the Bayesian emphasis on average-case performance. This allows choosing, in a cohesive way, among the enormous number of available R&S procedures. Surprisingly, this theoretical framework is structurally similar to that of convex risk measures studied in mathematical finance.

We begin by investigating how the performance of selection policies can be studied using a three layer performance evaluation process. We then, inspired by convex risk measures, introduce a new performance measure based solely on the number of simulation iterations required for the policy to become *acceptable* to a decision maker. A decision maker determines what performance standards an R&S policy should satisfy, and the set of policies that meet these standards then comprise her acceptance set. When performance can be controlled by the total number of simulation iterations budgeted, the smallest simulation budget for which a policy's performance falls within the acceptance set defines a new performance measure for R&S policies.

The existing literature on R&S policies is extensive. R&S problems have been heavily studied in sequential statistical testing and experiment design, see [Santner and Tamhane \(1984\)](#) and [Bechhofer, Santner, and Goldsman \(1995\)](#) for overviews. In simulation literature, selection procedures are popular tools for simulation optimization

problems (Boesel, Nelson, and Kim 2003) and discrete-event simulation problems (Swisher, Jacobson, and Yücesan 2003). Malone (2004) provides an overview for the Bernoulli R&S problem. For the normal R&S problem, there exist mainly two research streams: one is the IZ approach (Kim and Nelson 2006), the other uses Bayesian methods (Chick 2006). This paper focuses on the underlying performance evaluation for selection procedures and uses similar methods as in statistical decision theory, see Berger (1985) and DeGroot (1970). By far the most popular performance measure for selection procedures is the *probability of correct selection* (PCS), resp. its counterpart the *probability of incorrect selection* (PICS). Chick (1997), Chick and Inoue (2001) and Chen, He, Fu, and Lee (2008) also consider other performance measures. Branke, Chick, and Schmidt (2007) provided a comparison of existing R&S policies for the normal R&S problem based on an extensive set of examples. Our idea is that performance measures used in quantitative finance, namely convex risk measures, could prove effective for comparison of selection procedures, see Föllmer and Schied (2004) and McNeil, Frey, and Embrechts (2005) for an introduction to risk measures used in finance, and Rockafellar (2007) for use of convex risk measures in optimization.

The outline of this paper is as follows. Section 2 introduces the R&S problem and outlines the three-layer performance evaluation process. Section 3 shows the unified representation for existing performance measures analogous to convex risk measures. Section 4 introduces acceptance sets and the new performance measure based on computational cost. In Section 5 the new performance measure is used to evaluate three popular R&S policies for independent normal systems with unknown means and variances. Section 6 presents conclusions and future research directions.

2 PERFORMANCE EVALUATION FOR RANKING AND SELECTION PROCEDURES

In this section, we introduce the R&S problem and necessary notation. We then outline the performance evaluation process for R&S policies and show how the Bayes and the IZ formulation fit into a decision theoretic framework.

2.1 Mathematical Framework

We consider a setting with k independent systems. Each system i corresponds to a simulation experiment producing iid random outcomes Y_i . The distribution of these outcomes is described by a vector ψ_i , which is also called the *configuration* of system i . From each configuration a corresponding *performance parameter* $\theta_i \in \mathbb{R}$ may be computed. In general, θ_i is unknown, and all or part of ψ_i is unknown as well. The performance parameters across all systems are summarized in a vector $(\theta_1, \dots, \theta_k) = \theta \in \Theta \subset \mathbb{R}^k$, where Θ denotes the corresponding parameter space, and the underlying configurations are summarized $(\psi_1, \dots, \psi_k) = \psi \in \Psi$. We are interested in identifying the system with the *highest* performance parameter θ_i , i.e., we want to select system i^* where $\theta_{i^*} = \max_{i \in \{1, \dots, k\}} \theta_i$. For example, if sampling distributions are normal, then we could take $\psi_i = (\mu_i, \sigma_i^2)$ to contain the mean and the variance of samples from system i , and $\theta_i = \mu_i$ to be the mean. Although θ_i is often taken to be the mean of the sampling distribution in the R&S literature, our framework allows other choices as well. For example, in finance, one is often interested in a portfolio composition with the highest Sharpe ratio, which can be estimated by $\theta_i = \mu_i / \sigma_i$. Because θ depends on ψ , we sometimes write $\theta(\psi)$ to emphasize this dependence. More often, however, we simply write θ when the context is clear.

We denote by N the *total* number of measurements budgeted across all systems. For the moment, we assume that the simulation budget N is given at the beginning of the experiment. In Section 4, we relax this assumption and use N to control the performance of the policy. The goal of a selection procedure is to allocate these N measurements to the k systems to best support the ultimate selection decision.

An R&S policy π consists of an allocation rule, which determines which systems to simulate, and a selection rule, which determines which system to select based on the simulation results. We allow adaptive allocation rules. Formally, for $n = 1, \dots, N$, we define $x_n \in \{1, \dots, k\}$ to be the system simulated at time n and $y_n(x_n) \in \mathbb{R}$ the corresponding outcome. The decision x_n depends upon the previously observed outcomes, i.e., $x_n(x_{1:n-1}, y_{1:n-1})$, although we suppress this in the notation. Define $i_\pi \in \{1, \dots, k\}$ to be the selection made once all N simulations have been carried out. This decision is a function of all the observed outcomes $y_{1:N}$, i.e., $i_\pi(x_{1:N}, y_{1:N})$. A policy π can be written as $\pi = (x_{1:N}, i_\pi)$, where we use the notation $x_{1:n} := (x_1, \dots, x_n)$. Finally, let Π be the set of all possible policies π .

The essential question in R&S is this: which policy π should we use? The performance evaluation analysis in the next section discusses the issues that arise in answering this question.

2.2 Three-Layer Performance Evaluation

When we use an R&S procedure, we hope that it selects the best alternative. Inevitably, however, it will sometimes fail to do so. In this section, we discuss ways to measure the risk associated with this failure to select the best. Together, there are three layers of analysis that determine this risk: the loss of the decision, the configuration-specific risk, and the overall risk. In this section, we discuss and give examples of these layers of risk.

2.2.1 Loss of the Decision (Outcome-Specific Risk)

For a known θ , we can assign a quantity $L(i, \theta) \in \mathbb{R}$ to decision i , which reflects the *loss* (cost) associated with choosing system i given the true parameter θ . Examples of such loss functions L include:

1. $L(i, \theta) := \mathbb{1}_{\{\theta_i \neq \theta_{i^*}\}}$. This is the 0-1 loss function, used throughout the R&S literature. It penalizes all incorrect selections equally. The corresponding expected loss is the *probability of incorrect selection* (PICS).
2. $L(i, \theta) := \mathbb{1}_{\{\theta_i \leq (\theta_{i^*} - \delta)\}}$, for $\delta > 0$. This loss function only penalizes incorrect selection when the selected alternative is more than δ worse than the best. This reflects the idea that we may be indifferent to differences in performance of δ or less. This notion is similar to the IZ approach in R&S.
3. $L(i, \theta) := \theta_{i^*} - \theta_i$. This is the linear loss function used in [Chick and Inoue \(2001\)](#), and is also referred to as *regret*. It is equal to the difference in performance between the best system and the selected system.
4. $L(i, \theta) := c - \theta_i$, where $c \in \mathbb{R}$ is a constant. This can be used when our loss depends not on the difference in performance between the selected and the best, but between the selected and some known threshold. In some cases, e.g., the Bayesian framework using expected loss, this induces the same preference order as the linear loss function. In other cases, it does not.
5. $L(i, \theta) := f(\theta_{i^*} - \theta_i)$, where f is a convex and increasing function on \mathbb{R} . This generalizes linear loss, with the function f modeling the risk aversion of the decision-maker. Possible choices for f include $f(x) = \mathbb{1}_{\{x \geq 0\}} x^p$ for some constant $p \geq 1$ or $f(x) = \exp(x)$.

2.2.2 Configuration-Specific Risk

The loss $L(i, \theta)$ quantifies the risk associated with a single outcome of the selection decision i_π . However, this decision $i_\pi(x_{1:N}, y_{1:N})$ made by an R&S policy π depends on random simulation samples, and so $L(i_\pi, \theta)$ is a random variable. Thus, an additional source of risk is this randomness induced by stochastic simulation, given a single fixed configuration ψ . We call this *configuration-specific risk*, and refer to it as $R(\pi, \psi)$.

To quantify the configuration-specific risk, one possibility is to use the expected loss, $R(\pi, \psi) = \mathbb{E}[L(i_\pi, \theta(\psi))]$. This is a reasonable measure when the policy is used repeatedly for R&S in a large number of similar situations, or by a decision-maker with little aversion to risk. However, when selection errors are costly, or if a policy is only to be used once or a small number of times (for example in medical trials), more conservative functionals of the loss distribution, such as quantiles or expected loss above a given quantile, should be considered ([Chick 1997](#)). Generally, we write the configuration-specific risk as

$$R(\pi, \psi) := r(L(i_\pi, \theta(\psi))), \quad (1)$$

where r is a functional of the distribution of the random variable $L(i_\pi, \theta(\psi))$, given a fixed ψ .

2.2.3 Overall Risk

The configuration-specific risk $R(\pi, \psi)$ defined in (1) is still dependent on the *unknown* underlying configuration $\psi \in \Psi$. There is also risk associated with the fact that this configuration is unknown. A risk-averse decision-maker would prefer a policy π that has moderately low $R(\pi, \psi)$ across all possible configurations ψ . In contrast, a decision-maker indifferent to risk might prefer a policy π with extremely low $R(\pi, \psi)$ in most configurations, but high $R(\pi, \psi)$ in a certain few problematic configurations.

To quantify this risk, we use a functional ρ on the function $\psi \mapsto R(\pi, \psi)$. With a slight abuse of notation, we write $\rho(\pi) = \rho(R(\pi, \cdot)) \in \mathbb{R}$. This performance mapping ρ depends only on the policy π , and not on the configuration ψ . It can then be used to compare different policies and induces a preference order on Π .

Depending on the risk tolerance of the decision-maker there are different ways to define ρ . Three popular choices are:

1. *Worst-Case Performance:*

$$\rho_{WC}(\pi) := \sup_{\psi \in \Psi} R(\pi, \psi),$$

This performance measure prefers policies which perform well in the worst underlying configuration of the experiment. This is the most conservative performance measure.

2. *Indifference Zone:*

$$\rho_{IZ}(\pi) := \sup_{\psi \notin IZ} R(\pi, \psi), \quad (2)$$

This formulation is closely related to the worst-case performance measure. A set $IZ \subset \Psi$, the so-called *Indifference zone*, is defined as a subset of the set of configurations. The IZ performance measure prefers policies that perform well in the worst underlying configuration *outside* the Indifference zone. The intuition behind this approach is that some configurations of the experiment have performance measures so close to each other that it is not worth detecting the difference. This formulation is still quite conservative, but not as restrictive as the worst-case analysis.

3. *Bayes Risk:* The Bayesian approach incorporates a prior distribution P_0 on the configuration $\psi \in \Psi$. This prior reflects a *risk weighting* on ψ and does not necessarily correspond to the prior belief regarding ψ or θ . In the Bayesian approach, ψ is considered to be a random vector so $R(\pi, \psi)$ is a random variable. Different functionals of the distribution of $R(\pi, \psi)$ with respect to P_0 can be used to quantify the performance of a policy. The most popular approach is to use the expected value, i.e.,

$$\rho_{Bayes}(\pi) := \mathbb{E}_{P_0}[R(\pi, \psi)].$$

In contrast to the worst-case and the indifference zone formulation, the Bayes formulation focuses on an *average* performance rather than a conservative worst-case performance. When the risk weighting P_0 is rather diffuse, such an average case analysis might be too optimistic. One could introduce more risk aversion by using a different loss function L , a more conservative risk weighting P_0 or a different functional of $R(\pi, \psi)$ with respect to P_0 , e.g. quantiles or expected shortfall.

2.2.4 Decision-Theoretic Formulation for R&S Procedures

We briefly summarize the process needed to induce a preference order on the set of policies Π . To compare different policies, three quantities must be defined:

1. The loss of a decision $L(i, \theta(\psi))$;
2. The configuration-specific risk $R(\pi, \psi)$;
3. The overall risk $\rho(\pi)$.

Figure 1 provides an overview of these three concepts and how they interact. There is no single best way to define these quantities. Rather, they should be chosen carefully according to the decision-maker's risk tolerance for the particular problem instance addressed.

3 REPRESENTATION THEOREM FOR PERFORMANCE MEASURES

In the previous section, we described the process that underlies the performance evaluation of selection procedures. We now focus on the connections between the three popular performance measures defined in Section 2.2.3. These performance measures have a common representation analogous to *convex risk measures* studied in finance. This representation then suggests other performance measures in-between the measures given in Section 2.2.3 that may potentially overcome the drawbacks of these popular performance measures.

We first define convex risk measures as they are used in finance.

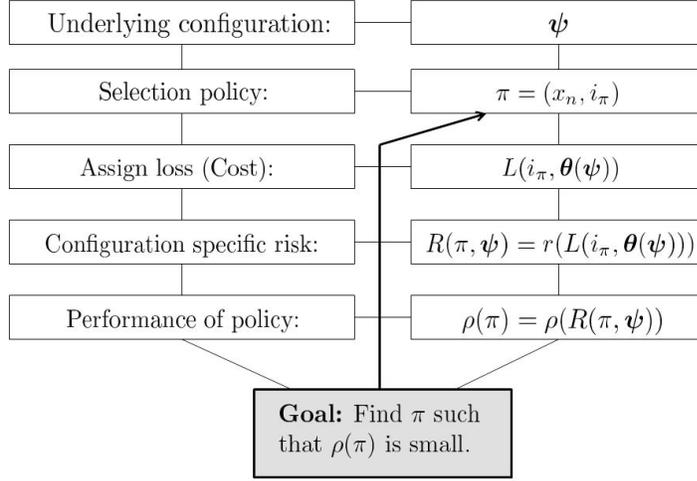


Figure 1: Quantities used to define a preference order on the set of policies.

Definition 1. A mapping $\rho : L^2 \rightarrow \mathbb{R}$ is called a convex risk measure if it satisfies the following conditions for all $X, Y \in L^2$ (L^2 denotes the space of random variables with finite variance).

1. Monotonicity: If $X \leq Y$, then $\rho(X) \geq \rho(Y)$.
2. Cash Invariance: If $m \in \mathbb{R}$, then $\rho(X + m) = \rho(X) - m$.
3. Convexity: $\rho(\lambda X + (1 - \lambda)Y) \leq \lambda \rho(X) + (1 - \lambda)\rho(Y)$, for $0 \leq \lambda \leq 1$.

Here the restriction to the L^2 space is sufficient but not necessary. At a first glance, this definition of convex risk measures does not seem related to our problem setting. However, a strong connection can be identified through the following representation property of convex risk measures.

Theorem 1. Denote by \mathcal{Q} the set of all finitely additive probability measures on a measurable space (Ω, \mathcal{F}) . Let $\alpha : \mathcal{Q} \rightarrow \mathbb{R} \cup \{\infty\}$ be a functional with $\inf_{Q \in \mathcal{Q}} \alpha(Q) \in \mathbb{R}$. Then the mapping

$$\rho(X) := \sup_{Q \in \mathcal{Q}} (\mathbb{E}_Q[-X] - \alpha(Q)), \quad (3)$$

defines a convex risk measure on L^2 .

With further technical assumptions the reverse direction also holds, i.e., every convex risk measure has a representation of the form (3). See Föllmer and Schied (2004) for more details and a proof of Theorem 1. We now state a similar representation theorem for the performance measures used for R&S procedures.

Theorem 2. The Worst-Case, Indifference zone and Bayes formulations for Ranking and Selection procedures defined in Section 2.2.3 can each be represented as

$$\rho(\pi) = \sup_{Q \in \mathcal{Q}} (\mathbb{E}_Q[R(\pi, \psi)] - \alpha(Q)), \quad (4)$$

where \mathcal{Q} is some appropriate set of probability measures on Ψ and its associated Borel sigma-field, and $\alpha : \mathcal{Q} \rightarrow \mathbb{R} \cup \{\infty\}$ is some penalty function.

Proof. The representation (4) is not unique. We show two ways in which the performance measures in Section 2.2.3 can be represented according to (4). In the first method, we take \mathcal{Q} to be the set of all probability measures on Ψ and adjust the penalty function α accordingly. An unconstrained optimization problem must be solved to determine $\rho(\pi)$. In the second method, which may be more intuitive, we restrict the set of probability measures \mathcal{Q} considered.

This approach has the advantage that the penalty function α is equal to zero, but it requires solving a constrained optimization problem.

We first show the theorem using the first method. Let \mathcal{Q} be the set of all probability measures on Ψ .

- If the penalty function α is

$$\alpha_{WC}(Q) = \begin{cases} 0, & \text{if } Q \text{ is a Dirac point measure for } \psi \in \Psi, \\ +\infty, & \text{otherwise,} \end{cases}$$

then ρ is the Worst-Case performance measure, i.e., $\rho(\pi) = \rho_{WC}(\pi)$.

- If the penalty function α is

$$\alpha_{IZ}(Q) = \begin{cases} 0, & \text{if } Q \text{ is a Dirac point measure for } \psi \in \Psi \setminus IZ, \\ +\infty, & \text{otherwise,} \end{cases}$$

then ρ is the IZ performance measure, i.e., $\rho(\pi) = \rho_{IZ}(\pi)$.

- If the penalty function α is

$$\alpha_{Bayes}(Q) = \begin{cases} 0, & \text{if } Q = P_0; \\ +\infty, & \text{otherwise,} \end{cases}$$

then ρ is the Bayes performance measure, i.e., $\rho(\pi) = \rho_{Bayes}(\pi)$.

We now show the theorem using the second method. Let $\alpha \equiv 0$. Then

- the set $\mathcal{Q}_{WC} = \{\text{all Dirac point measures for } \psi \in \Psi\}$ yields $\rho(\pi) = \rho_{WC}(\pi)$;
- the set $\mathcal{Q}_{IZ} = \{\text{all Dirac point measures for } \psi \in \Psi \setminus IZ\}$ yields $\rho(\pi) = \rho_{IZ}(\pi)$;
- the set $\mathcal{Q}_{Bayes} = \{P_0\}$ yields $\rho(\pi) = \rho_{Bayes}(\pi)$.

□

Theorem 2 suggests that the class of convex risk measures is helpful in analyzing the performance of R&S policies. For example, if a decision-maker is interested in analyzing the performance of a policy under Bayes and IZ formulations she could consider $\rho(\pi)$ with $\alpha \equiv 0$ and $\mathcal{Q} = \mathcal{Q}_{IZ} \cup \mathcal{Q}_{Bayes}$. Another decision-maker might be interested only in the performance of a policy under the Bayes formulation, but be unsure about how robust the performance is with respect to the risk weighting P_0 . She could then use $\mathcal{Q}_C = \{P : P = (1 - \varepsilon)P_0 + \varepsilon C, C \in \mathcal{C}\}$, for $0 < \varepsilon < 1$ and a class of possible 'contaminations' \mathcal{C} . Such a class of 'contaminated' priors is studied in Bayesian robustness theory, see [Berger \(1985\)](#). A third decision-maker might put emphasis on average performance, but also be somewhat concerned about worst-case performance. He might take $\mathcal{Q} = \mathcal{Q}_{WC} \cup \mathcal{Q}_{Bayes}$, $\alpha(Q) = 0$ for $Q \in \mathcal{Q}_{Bayes}$, and $\alpha(Q)$ equal to a large strictly positive constant c for $Q \in \mathcal{Q}_{WC}$. With this large value of the penalty function on \mathcal{Q}_{WC} , the performance measure $\rho(\pi)$ would be equal to the Bayes performance unless the worst-case performance was extremely bad, when $\rho(\pi)$ would equal the worst-case performance minus c .

4 ACCEPTANCE SETS AND A NEW PERFORMANCE MEASURE

In the previous sections we described how, given a fixed simulation budget N , the performance of a policy π can be evaluated via a performance measure ρ . Another important aspect of R&S is *controlling* performance. By relaxing the assumption of a fixed simulation budget we can use the "necessary" number of simulation iterations as a performance control of the policy. To do this, we must first define what it means for a policy to be *acceptable*. We then quantify the performance of a policy as the smallest simulation budget necessary to make the policy acceptable.

We use the representation of ρ in Theorem 2 to define an acceptable policy. While the IZ and Bayes formulations are special cases of (4), this representation can be used to define much more general performance measures. Each probability measure $Q \in \mathcal{Q}$ corresponds to a *scenario* of the underlying experiment configuration Ψ and \mathcal{Q} is the set of all scenarios considered. Examples of \mathcal{Q} are \mathcal{Q}_{WC} , \mathcal{Q}_{IZ} , \mathcal{Q}_{P_0} and \mathcal{Q}_C as defined in Section 3. For this paper we assume that $|\mathcal{Q}| = m < \infty$. The extension to infinitely many scenarios is mathematically more challenging and left open for future research.

For a given loss function L , a risk function $R(\pi, \psi)$, and a number of measurements N , the risk of a policy π_N that takes N measurements can be represented as a mapping $f_{\pi_N} : \mathcal{Q} \rightarrow \mathbb{R}, Q \mapsto \mathbb{E}_Q[R(\pi_N, \psi)]$. The set of *all* risks generated by arbitrary combinations of L , $R(\pi, \psi)$ and N corresponds to the functional space $\mathcal{G} = \{f : \mathcal{Q} \rightarrow \mathbb{R}\}$, which can be identified with \mathbb{R}^m . The risk of a particular policy π_N can then be represented as a risk vector

$$f(\pi_N) = \begin{pmatrix} f_{\pi_N}(Q_1) \\ \vdots \\ f_{\pi_N}(Q_m) \end{pmatrix} = \begin{pmatrix} \mathbb{E}_{Q_1}[R(\pi_N, \theta)] \\ \vdots \\ \mathbb{E}_{Q_m}[R(\pi_N, \theta)] \end{pmatrix} \in \mathbb{R}^m.$$

To introduce the concept of acceptance sets, we take the viewpoint of a decision-maker. If the decision-maker wants to use an R&S procedure, she should specify an *acceptance* set $\mathcal{A} \subset \mathbb{R}^m$ such that a policy π is *acceptable* if and only if $f(\pi) \in \mathcal{A}$. The acceptance set represents the decision-maker's risk attitude towards different risk scenarios Q . The set \mathcal{A} should be defined in a coherent way, i.e., all reasonable acceptance sets should satisfy certain axioms *after* being appropriately normalized. We assume that after normalization, $f_\pi(Q) \in [0, 1]$ for all Q in \mathcal{Q} , where $f_\pi(Q) = 0$ corresponds to the minimal risk and $f_\pi(Q) = 1$ to the maximal risk. Further, we assume that the policy that randomly chooses a system has risk $1 - 1/k$. Such a normalization is adapted to the PICS formulation. We propose that *any* coherent acceptance set \mathcal{A} should satisfy the following axioms:

- Axiom 1: $0 \in \mathcal{A}$. A policy with no risk under any scenario should always be acceptable.
- Axiom 2: $\left(\frac{k-1}{k} \cdot 1, 1\right]^m \cap \mathcal{A} = \emptyset$, where 1 represents a vector of ones in \mathbb{R}^m . Any acceptable policy should be as good as randomly identifying a system as the best. Clearly, this axiom depends on our choice of normalization.
- Axiom 3: For any point $q \in \mathcal{A}$, $[0, q]^m \subseteq \mathcal{A}$. All policies with risks smaller than some acceptable policy must also be acceptable.

Note that vector notation is used to define subsets of \mathbb{R}^m , e.g. for two vectors $p, q \in \mathbb{R}^m$, we define the set $[p, q]^m := [p_1, q_1] \times [p_2, q_2] \times \dots \times [p_m, q_m] \subseteq \mathbb{R}^m$.

Once a decision-maker has identified the acceptance set according to her risk tolerance, we can check whether or not a given policy π_N is acceptable. Now, instead of taking N as a predetermined simulation budget, we can treat N as a *control* chosen by the decision-maker. This interpretation motivates the new performance measure

$$\rho_{\mathcal{A}}(\pi) := \inf \{n \in \mathbb{N} \mid f(\pi_n) \in \mathcal{A}\}, \quad (5)$$

which corresponds to the minimum number of simulation iterations required to make the policy acceptable to the decision-maker.

The intuition behind this performance measure stems from the fact that the *cost* of a simulation can be measured as the number of iterations required to achieve a certain confidence level. If an infinite number of simulations is allowed, then *any* reasonable policy can be satisfactory. Here a reasonable policy refers to an algorithm that converges to the right answer, i.e., $f(\pi_N) \rightarrow 0$ as $N \rightarrow \infty$. A performance measure for an R&S policy should depend on the convergence rate of that policy.

Figure 2 shows conceptually how the risk measure $\rho_{\mathcal{A}}$ can be interpreted for two policies $\pi^{(1)}$ and $\pi^{(2)}$, where the grey area \mathcal{A} represents a possible acceptance set. We consider only two scenarios, Q_1 and Q_2 . Assume that the decision-maker wants to compare policy $\pi^{(1)}$ and $\pi^{(2)}$. As N varies between 0 and ∞ , each policy traces out a trajectory in \mathbb{R}^2 indicated by the dashed arrows. When the trajectory crosses the boundary of the acceptance set \mathcal{A} for the first time, the corresponding number of iterations N is precisely the performance measure $\rho_{\mathcal{A}}$.

The IZ and Bayesian formulations can be combined effectively in this framework. Consider scenarios $Q_1 = LFC$ and $Q_2 = P_0$, where LFC represents the *least favorable configuration* of the experiment where the supremum in (2) is attained. It is not clear whether such a configuration can always be identified and in many cases the LFC may depend on the policy itself. For illustrative purposes, we can think of the LFC as one single *challenging* configuration. To be more rigorous one should use \mathcal{Q}_{IZ} defined in the proof to Proposition 2. Figure 3 shows the acceptance set \mathcal{A}_1 of a decision-maker using the Bayes formulation, as well as the acceptance set \mathcal{A}_2 of a decision-maker applying the IZ formulation. A decision-maker who wants to satisfy both the IZ and the Bayesian formulation would use the intersection of \mathcal{A}_1 and \mathcal{A}_2 as her acceptance set. The dashed arrows indicate possible trajectories of two R&S policies.

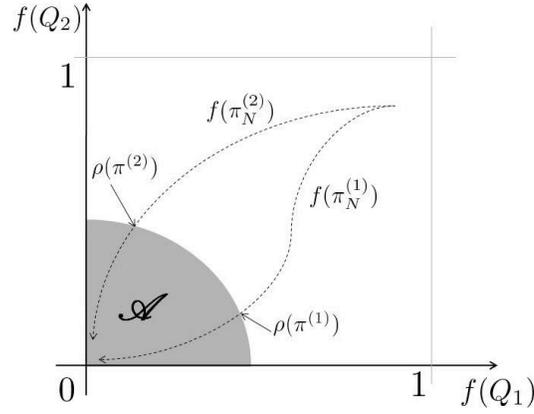


Figure 2: Possible acceptance set \mathcal{A} and trajectories of two policies $\pi^{(1)}$ and $\pi^{(2)}$. The trajectories outline the function $f(\pi_N)$ for $N \rightarrow \infty$. Note that the trajectories are not continuous curves as $N \in \mathbb{N}$, but we drew them continuously here for simplicity.

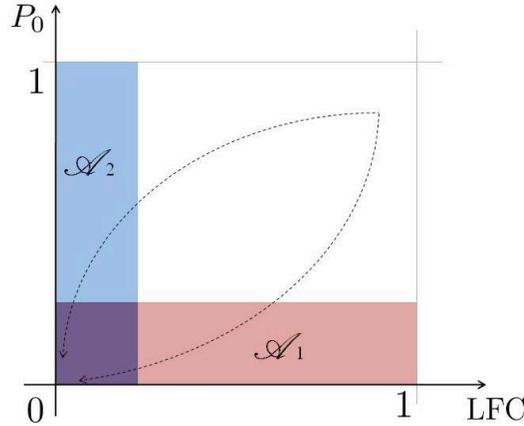


Figure 3: Acceptance sets corresponding to Bayes and IZ formulations.

It is instructive to analyze the connection between the performance measure $\rho(\pi) = \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[R(\pi, \theta)] - \alpha(Q)$ defined in Theorem 2 and the new performance measure $\rho_{\mathcal{A}}(\pi) = \inf \{n \in \mathbb{N} | f(\pi_n) \in \mathcal{A}\}$ introduced in (5). For simplicity, consider the case where \mathcal{Q} encompasses all relevant scenarios so the penalty function in (4) can be ignored, i.e., $\rho(\pi) = \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[R(\pi, Q)]$. An acceptance set can then be defined directly by setting a bound on $\rho(\pi)$, i.e., $\rho(\pi) \leq c$ for some $c \in \mathbb{R}$, and this corresponds to a hypercube in \mathbb{R}^m with side-length c . On the other hand, if an acceptance set \mathcal{A} is defined independently of ρ , then $\rho(\pi_{\rho_{\mathcal{A}}})$ corresponds to the *maximal* risk among all scenarios as the policy trajectory enters the acceptance set \mathcal{A} .

Remark 1. $\rho_{\mathcal{A}}$ does not directly induce a preference order on the set of policies. In R&S algorithms, not only the rate of convergence needs to be considered, but also the cost of determining the measurement allocation rule. For example, complex sequential R&S procedures require extensive computation at each allocation step, whereas a simple equal allocation does not require any computation between simulations. If simulation of the systems is relatively inexpensive compared to the computation of the allocation rule, an equal allocation policy is most likely preferred. If it is possible to determine a cost c_{π} for an allocation decision of policy π , then a preference order can be formed via the quantity $c_{\pi} \rho_{\mathcal{A}}(\pi)$, i.e., a policy π_1 is preferred to a policy π_2 if and only if $c_{\pi_1} \cdot \rho_{\mathcal{A}}(\pi_1) \leq c_{\pi_2} \cdot \rho_{\mathcal{A}}(\pi_2)$. If all the allocation costs are equal, or negligible compared to simulation costs, then the risk measure $\rho_{\mathcal{A}}$ itself induces a preference order on the set of policies.

The cost of simulating each individual system should also be considered. If the systems require different computational budget, different allocations should be considered and policies that spend too much time on expensive systems should be penalized. This aspect has been considered in some policies in [Chick and Inoue \(2001\)](#).

Remark 2. *In order to extend the present framework so that policies with random stopping times can also be compared, the assumption of a fixed simulation budget N needs to be relaxed. This could be done by considering the expected running time. Although it is possible to determine whether or not a policy is acceptable to a decision-maker based on the expected running time, it is not obvious how these policies can be controlled. For policies with fixed simulation budgets, the parameter N provides this important control and can be used as a performance measure in (5). For policies with random runlength, a different control such as a confidence parameters can be used.*

Remark 3. *R&S procedures can also be useful for infinitely many systems, especially in continuous simulation optimization problems. In this case, the convergence $f(\pi_N) \rightarrow 0$ as $N \rightarrow \infty$ is often not guaranteed. However, our framework can still be used by setting $\rho_{\mathcal{A}} = \infty$ if the trajectory $f(\pi_N)$ does not reach the acceptance set \mathcal{A} .*

5 SIMULATION EXAMPLE

In this section, we show via a small simulation example how the performance measure $\rho_{\mathcal{A}}$ can be used to induce a preference order on a set of policies for the independent normal R&S problem with unknown means and variances. Since explicit calculations of $\rho_{\mathcal{A}}$ can be quite challenging, we rely on estimation procedures.

We consider $k = 10$ systems with independent normal output $Y_i \sim N(\mu_i, \sigma_i^2)$, $i = 1, \dots, k$, where the underlying configuration of the experiment is $\psi_i = (\mu_i, \sigma_i^2)$ and we are interested in identifying the system with the highest mean, i.e., $\theta_i = \mu_i$. This problem has been studied extensively in the R&S literature and there exist many competing policies. Branke, Chick, and Schmidt (2007) compared different policies for this problem. We follow their outline but use the new performance measure introduced in (5) to induce a preference order on Π .

An extensive comparison of existing policies would go beyond the scope of this paper. Instead, we compare three popular policies, namely the equal allocation policy, the optimal computing budget allocation policy (OCBA), and the $(0-1)(S)$ expected value of information policy. The equal allocation policy simply allocates the allowed measurements equally among the k systems, OCBA is a policy introduced in Chen (1996) that dynamically allocates simulation replications, thereby essentially screening out bad systems, and $(0-1)(S)$ is a Bayesian policy introduced in Chick and Inoue (2001). For detailed descriptions of these policies, see Branke, Chick, and Schmidt (2007). We explore whether or not our framework can give insights regarding their relative performance.

First, it is necessary to define the considered set of scenarios \mathcal{Q} . For illustrative purposes we only consider two scenarios, the IZ formulation (Q_1) and the Bayes formulation (Q_2). As mentioned earlier, infinitely many scenarios must be considered in order to properly capture the IZ formulation. Instead we only focus on one *challenging* configuration of the experiment and refer to it as LFC even though it is not necessarily the least favorable configuration. We define the two scenarios as

$$\begin{aligned} Q_1 &:= \text{A point mass at } \{\psi_1 = (0.4, 1), \psi_i = (0, 1) \text{ for } i = 2, \dots, 10\}, \text{ and} \\ Q_2 &:= \{\psi_i = (\mu_i, \sigma_i), \text{ where } \mu_i \text{ (iid)} \sim U(0, 2), \sigma_i \text{ (iid)} \sim U(0.2, 1) \text{ for } i = 1, \dots, 10\}. \end{aligned}$$

Scenario Q_2 is approximated by 100 independent random replications of ψ . We choose the 0-1 loss function $L(i, \theta) := \mathbb{1}_{\{\theta_i \neq \theta_{i^*}\}}$ and the risk mapping $R(\pi, \psi) = \mathbb{E}[L(i_\pi, \theta(\psi))]$ which corresponds to the PICS for a single configuration ψ .

Based on these assumptions, we first analyze the performance under IZ and Bayes formulations. Figure 4 shows the estimated trajectories of the three policies as N increases from 40 to 530. Throughout we use a batchsize of 10, i.e., the policy is updated after 10 simulation runs. Batching speeds up the simulations since some allocation rules require significant computation and updating the allocation after each single observation would be too time consuming. The dashed line indicates the acceptance set for a decision-maker adopting an IZ approach and the dotted line indicates the acceptance set for a decision-maker adopting a Bayes approach. The confidence level is set to be $c = 0.22$, i.e., a policy is acceptable if and only if $\rho_{IZ}(\pi) \leq c$ resp. $\rho_{Bayes}(\pi) \leq c$. A preference order among these three policies cannot be deduced based on the trajectories alone. Table 1 gives the estimation results for the performance measure $\rho_{\mathcal{A}}$. Figure 5 shows the convergence towards the acceptance set for each of the two scenarios.

Let us now consider a decision-maker who seeks good performance in both scenarios Q_1 and Q_2 . Figure 6 shows four possible acceptance sets that combine the two scenarios. The set A_1 is most conservative, where the decision-maker requires that both the IZ and the Bayes performances are bounded by a certain confidence level c . The set A_4 is the least conservative, where the decision-maker requires that at least one performance measure

Table 1: Estimated performance measures ρ_{sd} (with standard deviations) for a decision-maker with IZ or Bayes acceptance set. We used bootstrapping techniques to estimate ρ_{sd} and its standard deviations based on a sample of 100 trajectories. The standard deviation for the Bayes scenario is larger than that for the LFC. This is due to the fact that the slope at the entry point of the trajectory in the Bayes scenario is much smaller than the slope for the LFC scenario, see Figure 5.

| | LFC | Bayes |
|--------|-----------|----------|
| Equal | 360 (1.0) | 464 (63) |
| OCBA | 278 (4.6) | 207 (33) |
| 0-1(S) | 323 (4.8) | 259 (49) |

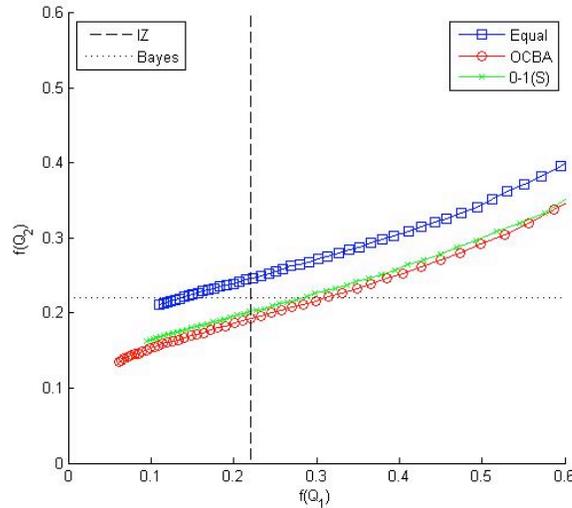


Figure 4: Trajectories for the Equal, OCBA and VIP policies, along with the IZ and Bayes acceptance sets.

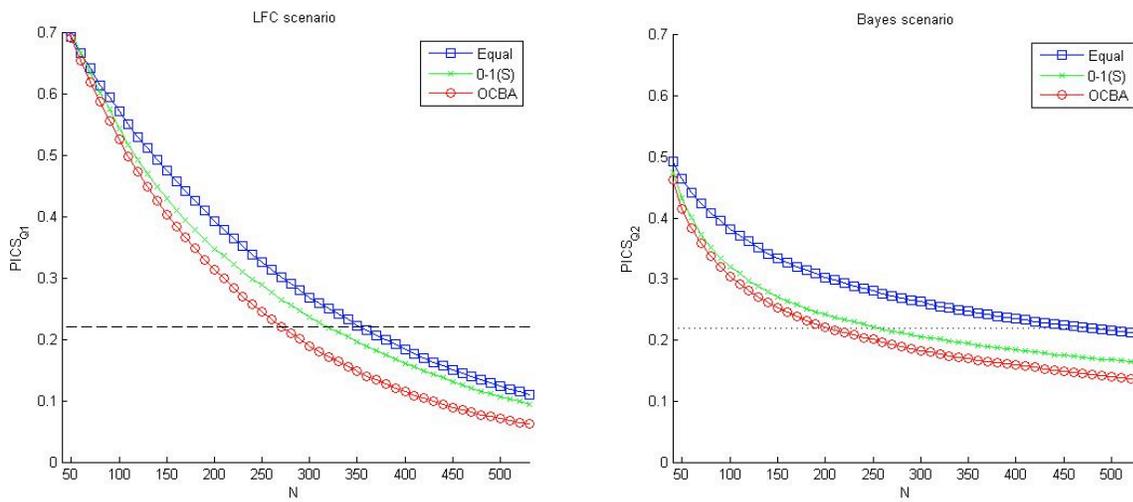


Figure 5: Convergence in terms of N for each scenario separately.

satisfies a given confidence level c . Table 2 summarizes the estimation results for the performance measure ρ_{sd} defined in (5) and the corresponding *estimated* performance measure $\rho = \rho(\pi_{\rho_{sd}})$ defined in (4) from Proposition 2.

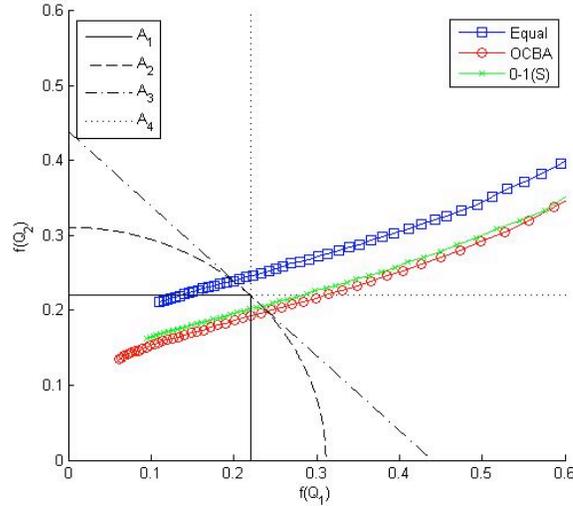


Figure 6: Trajectories for Equal, OCBA and VIP policy and different acceptance sets A_1 through A_4 .

Table 2: Estimated risk measures ρ_{sd} (along with standard deviations) for different acceptance sets. We used bootstrapping techniques to estimate ρ_{sd} and its standard deviations based on a sample of 100 trajectories. The corresponding risk measures ρ defined in (4) are also given. We omit standard deviations for ρ because they were small ($< 5\%$).

| | A_1 | | A_2 | | A_3 | | A_4 | |
|--------|--------------|--------|--------------|--------|--------------|--------|--------------|--------|
| | ρ_{A_1} | ρ | ρ_{A_2} | ρ | ρ_{A_3} | ρ | ρ_{A_4} | ρ |
| Equal | 473 (68) | 0.22 | 385 (20) | 0.24 | 382 (17) | 0.24 | 358 (10) | 0.25 |
| OCBA | 280 (5.4) | 0.21 | 259 (9) | 0.24 | 258 (10) | 0.24 | 207 (31) | 0.31 |
| 0-1(S) | 332 (14) | 0.22 | 308 (13) | 0.23 | 306 (14) | 0.23 | 256 (42) | 0.29 |

From these results we may conclude that OCBA outperforms the other policies on these scenarios and acceptance sets. Having said that, our primary goal is to indicate how a comprehensive performance analysis might be carried out.

6 CONCLUSIONS AND FURTHER RESEARCH

We presented a performance-analysis process for Ranking and Selection procedures. This process states all assumptions needed to incorporate a preference order on the set of policies. The three most popular approaches have a common representation analogous to convex risk measures used in mathematical finance. Based on the performance evaluation process, we introduced a new performance measure using computational cost and the definition of acceptable policies. This easy-to-interpret performance measure can be used to compare selection procedures for different acceptance sets specified by the decision-maker.

Extensions to this framework involve the treatment of infinitely many scenarios in the scenario approach and the analysis of policies with random simulation lengths. Analogous to [Branke, Chick, and Schmidt \(2007\)](#), a comprehensive comparison of existing R&S policies using different acceptance sets would be helpful in identifying robust selection procedures. This paper provides a framework regarding how such a comparison could be carried out, but we have yet to identify good or optimal policies for specific acceptance sets. The computation required for the illustrative example was nontrivial, which suggests that efficient computation of these performance measures should also be explored.

ACKNOWLEDGMENTS

This research was supported, in part, by National Science Foundation Grant CMMI-0800688.

REFERENCES

- Bechhofer, R., T. Santner, and D. Goldsman. 1995. *Design and analysis of experiments for statistical selection, screening, and multiple comparisons*. Wiley series in probability and statistics.
- Berger, J. 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer.
- Boesel, J., B. Nelson, and S. Kim. 2003. Using Ranking and Selection to “Clean Up” after Simulation Optimization. *Operations Research* 51 (5): 814–825.
- Branke, J., S. Chick, and C. Schmidt. 2007. Selecting a selection procedure. *Management Science* 53 (12): 1916–1932.
- Chen, C. 1996. A lower bound for the correct subset-selection probability and its application to discrete-event system simulations. *IEEE transactions on automatic control* 41 (8): 1227–1231.
- Chen, C., D. He, M. Fu, and L. Lee. 2008. Efficient simulation budget allocation for selecting an optimal subset. *INFORMS Journal on Computing* 20 (4): 579.
- Chick, S. 1997. Selecting the best system: A decision-theoretic approach. In *Proceedings of the 29th conference on Winter simulation*, 326–333. IEEE Computer Society.
- Chick, S. 2006. Subjective probability and Bayesian methodology. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science. Elsevier.
- Chick, S., and K. Inoue. 2001. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49 (5): 732–743.
- DeGroot, M. 1970. *Optimal Statistical Decision*. McGraw-Hill, New York.
- Föllmer, H., and A. Schied. 2004. *Stochastic Finance, An Introduction in Discrete Time*. Walter de Gruyter & Co., Berlin.
- Kim, S., and B. Nelson. 2006. Selecting the best system. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science. Elsevier.
- Malone, G. 2004. *Ranking and Selection Procedures for Bernoulli and Multinomial Data*. Ph. D. thesis, Georgia Technical Institute, Atlanta GA.
- McNeil, A., R. Frey, and P. Embrechts. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton university press, Princeton, NJ.
- Rockafellar, T. 2007. Coherent approaches to risk in optimization under uncertainty. In *Tutorials in Operations Research: OR Tools and Applications. Glimpses of Future Technologies*. INFORMS.
- Santner, T., and A. Tamhane. 1984. *Design of experiments: ranking and selection*. Marcel Dekker, inc.
- Swisher, J., S. Jacobson, and E. Yücesan. 2003. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 13 (2): 134–154.

AUTHOR BIOGRAPHIES

ROLF WAEBER is a Ph.D. student in the School of Operations Research and Information Engineering at Cornell University. His research interest is in quantitative risk management. His e-mail address can be found via www.orie.cornell.edu.

PETER I. FRAZIER is an assistant professor in the School of Operations Research and Information Engineering at Cornell University. He received a Ph.D. in Operations Research and Financial Engineering from Princeton University in 2009. His research interest is in the optimal acquisition of information, with applications in simulation, medicine and operations management. His web page can be found via www.orie.cornell.edu.

SHANE G. HENDERSON is a professor in the School of Operations Research and Information Engineering at Cornell University. He is the simulation area editor at Operations Research, and an associate editor for Management Science. He co-edited the Proceedings of the 2007 Winter Simulation Conference. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services. His web page can be found via www.orie.cornell.edu.