Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

EXPECTED OPPORTUNITY COST GUARANTEES AND INDIFFERENCE ZONE SELECTION PROCEDURES

Stephen E. Chick

Technology Management Area
INSEAD
Boulevard de Constance
77305 Fontainebleau CEDEX, FRANCE

ABSTRACT

Selection procedures help identify the best of a finite set of simulated alternatives. The indifference-zone approach focuses on the probability of correct selection, but the expected opportunity cost of a potentially incorrect decision may make more sense in business contexts. This paper provides the first selection procedure that guarantees an upper bound for the expected opportunity cost, in a frequentist sense, of a potentially incorrect selection. The paper therefore bridges a gap between the indifference-zone approach (with frequentist guarantees) and the Bayesian approach to selection procedures (which has considered the opportunity cost). An expected opportunity cost guarantee is provided for all configurations of the mean, and need not rely upon an indifference zone parameter to determine a so-called least favorable configuration. Further, we provide expected opportunity cost guarantees for two existing indifference zone procedures that were designed to provide probability of correct selection guarantees.

1 INTRODUCTION

Statistical selection procedures provide a mechanism to identify the best of a finite set of simulated alternatives, where best is defined in terms of the maximum (or minimum) expected value of each alternative. A sample application is the selection of one of several design proposals for a supply chain when simulation is used to evaluate the performance of each alternative.

Many selection procedures identify the best system by running an initial stage of simulations of each system to get a rough estimate of the mean and variance of each system, then additional sampling occurs before making a final decision. The additional sampling can occur all at once in a second stage, or sequentially. Because simulation output has randomness, the best alternative cannot be selected with certainty. Instead, selection procedures provide some measure of the quality of a selection.

This note focuses on two stage sampling procedures where the simulation output is normally distributed. The main contribution is to provide an alternate guarantee for the quality of a correct selection. Much of the extensive frequentist ranking and selection and multiple comparisons literature provides results to guarantee that the probability of a correct selection, P(CS), exceeds some prespecified threshold P^*, subject to the condition that the best system be at least δ^* better than the other systems (Dudewicz and Dalal 1975, Rinott 1978, Bechhofer, Santner, and Goldsman 1995). More recent results provide guarantees for the probability of a good selection, P(GS), the probability that the selected system is within some specified δ^* of the best (Matejcik and Nelson 1995, Nelson, Swann, Goldsman, and Song 2001), for all configurations of the unknown means. In both cases, the probability statements are made with respect to repeated applications of the procedure.

These guaranteed lower bounds on P(CS) do not, however, reflect how poor a potentially incorrect selection might be. The expected opportunity cost does penalize particularly bad choices. For example, it may be better to be wrong 99% of the time if the penalty for being wrong is $1 (an expected opportunity cost of 0.99×$1 = $0.99) rather than being wrong only 1% of the time if the penalty is $1,000 (an expected opportunity cost of 0.01×$1,000 = $10).

New selection procedures that are based on a Bayesian decision-theoretic foundation provide a mechanism to allocated second stage samples to reduce the expected opportunity cost (Chick and Inoue 2001a, Chick and Inoue 2001b). That work also provides a measure of the posterior probability of correct selection, or posterior expectation of the potential opportunity costs, based on a single application of the procedure. But frequentist P(CS) guarantees have not yet been provided for those procedures.

This paper merges ideas from both the indifference zone and decision-theoretic literature to provide what is believed to be the first selection procedure with a frequentist expected opportunity cost guarantee. We indicate that the least favorable configuration (LFC) for the procedure is not necessarily the slippage configuration for a given δ^*. In fact an expected opportunity cost guarantee can be provided for all configurations of the mean without reference to indifference zone parameter δ^*. We also provide expected opportunity cost guarantees for the procedure in Rinott (1978) and a procedure in Nelson and Banerjee (2001).

2 SELECTION PROCEDURES AND OPPORTUNITY COST

This section recalls the formal description of the selection problem from a classical, indifference-zone perspective, then presents a new procedure that provides a guaranteed upper bound on the expected opportunity cost of potentially selecting the wrong system.

2.1 Setup for Selection Problem

The best of k simulated alternatives is to be identified using a two-stage selection procedure. The simulation output $x_{i,j}$ for replication j of system i is is presumed independent and normally distributed for $i = 1, 2, \ldots, k$; $j = 1, 2, \ldots$, with unknown means $\mu = (\mu_1, \mu_2, \ldots, \mu_k)$ and variances $\sigma^2 = (\sigma_1^2, \sigma_2^2, \ldots, \sigma_k^2)$. Best here is defined by the system with the maximal mean (the minimum is handled similarly). Let $\mu_{[1]} \le \mu_{[2]} \le \ldots \le \mu_{[k-1]} < \mu_{[k]}$ be the unknown ordering, so system $[k]$ is best.

A correct selection occurs when the system selected by the procedure, system d, is the same as the true best, $[k]$. The system selected is commonly the system with the highest overall sample mean (e.g. Rinott 1978, Nelson, Swann, Goldsman, and Song 2001, although see Dudewicz and Dalal 1975). Since the output is variable, d is the realization of a random variable D that identifies the selected system. The probability of correct selection, P(CS), is the probability that $D = [k]$, the probability taken over repeated applications of the procedure to the same problem:

$$\mathrm{P(CS)} = E\left[\mathbf{1}_{(D=[k])}\right],$$

where $\mathbf{1}_{(\cdot)}$ is the indicator function (1 if the argument is true, 0 otherwise), and the distribution of the selected system, D, is determined by structure of the selection procedure. The validity of an indifference zone procedure is established by showing that a bound on the P(CS) is available, given some conditions. For example, the well-known two-stage procedure of Rinott (1978), which we call Procedure R, takes two parameters to specify a minimum desired probability of correct selection P^*, and the minimum difference δ^* worth detecting between the best and the others:

$$\min_{\mu \in \Omega(\delta^*)} \mathrm{P(CS)} \ge P^*$$

where $\Omega(\delta^*) = \{\mu : \mu_{[k]} - \delta^* \ge \mu_i, i \ne [k]\}$. Procedure R provides this guarantee by allocating a total number (both stages) of replications that is proportional to the first stage sample variance.

The new procedure below, Procedure OC_f, is similar in structure to Procedure R, but differs in that the expected opportunity cost of a potentially incorrect selection is guaranteed to be less than some user-specified upper bound, Δ. The name Procedure OC_f comes from this **f**requentist expected **o**pportunity **c**ost guarantee. Recall that the opportunity cost of selecting system d is

$$OC = \mu_{[k]} - \mu_d.$$

If the best system is correctly selected, then $OC = 0$. If not, then OC increases with the difference in the mean performance of the best and the mean performance of the selected system. The expected opportunity cost E[OC] of a selection procedure is a frequentist measure of the expected opportunity cost associated with selecting a system, where

$$E[OC] = E\left[\mu_{[k]} - \mu_D\right]$$

is the expectation of OC taken over repeated applications of the procedure.

Procedure OC_f

1. Specify the expected opportunity cost guarantee Δ, and first-stage sample size $n_0 \ge 3$. Set $g = g(\Delta, k, n_0)$ as described in Section 4.
2. First stage. Observe the output $x_{i,1}, x_{i,2}, \ldots, x_{i,n_0}$ of independent simulation runs for $i = 1, 2, \ldots, k$.
3. For each system, compute the first-stage sample mean $\bar{x}_i = \sum_{j=1}^{n_0} x_{i,j}/n_0$ and sample variance, $s_i^2 = \sum_{j=1}^{n_0}(x_{i,j} - \bar{x}_i)^2/(n_0 - 1)$.
4. Compute the total number of runs for each system,
$$n_i = \max\left\{n_0, \left\lceil g^2 s_i^2 \right\rceil\right\}.$$
5. Second stage. Obtain independent simulation output $x_{i,n_0+1}, x_{i,n_0+2}, \ldots, x_{i,n_i}$ for $i = 1, 2, \ldots, k$.
6. Compute the overall sample means for each system, $\bar{\bar{x}}_i = \sum_{j=1}^{n_i} x_{i,j}/n_i$.
7. Select system $d = \arg\max_i \bar{\bar{x}}_i$ as best.

Procedure OC_f does not need an indifference zone parameter, and therefore has one less parameter in Step 1 than Procedure R, using Δ rather than P^* and δ^*. The choice of Δ should be tied to the economic value associated with the simulated systems. If simulation is used to select the best of a set of manufacturing system designs, for example, the simulation output can be taken to be realizations of the revenue minus the cost over the usable lifetime of those systems. Smaller values of Δ require more replications. The cost of more replications can be traded off with a lower upper bound for $E[OC]$. This is similar to the tradeoff between increasing P^* and running more replications.

3 ANALYSIS

This section shows how to select the parameter g in Step 4 of Procedure OC_f so that the expected opportunity cost is bounded above by the maximum desired expected opportunity cost, $E[OC] \leq \Delta$, for *all* configurations of the means. First we simplify the case of selecting from $k \geq 2$ systems to the case of $k = 2$ systems by considering the $k-1$ pairwise comparisons between systems i and $[k]$ (*not* the $k(k-1)/2$ comparisons between all pairs).

3.1 Opportunity Cost Bounds

Let $\bar{X} = (\bar{X}_1, \bar{X}_2, \ldots, \bar{X}_k)$ and $S^2 = (S_1^2, S_2^2, \ldots, S_k^2)$ be the vector of first stage sample means and variances. Random quantities are written in upper case and realizations are in lower case. The selected system, D, depends upon the overall sample means, $\bar{\bar{X}} = (\bar{\bar{X}}_1, \bar{\bar{X}}_2, \ldots, \bar{\bar{X}}_k)$.

$$D = \arg \max_{j \in \{1,2,\ldots,k\}} \bar{\bar{X}}_j$$

Let $D_i = \arg\max_{j \in \{i,[k]\}} \bar{\bar{X}}_j$ be the random variable that selects a system in a pairwise comparison between systems i and $[k]$, and define $\delta_i = \mu_{[k]} - \mu_i$ to be the difference in the means of the best system and system i, for all $i \neq [k]$. The $E[OC]$ for k systems is bounded from above by the sum of the expected losses from the $k-1$ pairwise comparisons (the best versus each alternative) because $\mu_{[k]} - \mu_D \leq \sum_{i=1}^{k-1} \mu_{[k]} - \mu_{D_i}$ for each realization.

$$\begin{aligned} E[OC] &= E_{\bar{\bar{X}}, S^2}\left[\mu_{[k]} - \mu_D\right] \\ &\leq E_{\bar{\bar{X}}, S^2}\left[\sum_{i=1}^{k-1} \mu_{[k]} - \mu_{D_i}\right] \qquad (1) \\ &= \sum_{i: i \neq [k]} E_{\bar{\bar{X}}_i, \bar{\bar{X}}_{[k]}, S_i^2, S_{[k]}^2}\left[\delta_i \mathbf{1}_{(\bar{\bar{X}}_i > \bar{\bar{X}}_{[k]})}\right] \\ &\stackrel{def}{=} \sum_{i: i \neq [k]} E[OC_i] \end{aligned}$$

The implication is that if we can guarantee $E[OC_i] \leq \Delta/(k-1)$ in each of $k-1$ pairwise comparisons, then an overall $E[OC]$ guarantee is provided. We proceed along lines analogous to Rinott (1978). For all $i \neq [k]$, define

$$\begin{aligned} Z_i &= \frac{(\bar{\bar{X}}_i - \mu_i) - (\bar{\bar{X}}_{[k]} - \mu_{[k]})}{\left[\frac{\sigma_i^2}{n_i} + \frac{\sigma_{[k]}^2}{n_{[k]}}\right]^{1/2}}, \\ R_i &= \frac{\delta_i}{\left[\frac{\sigma_i^2}{n_i} + \frac{\sigma_{[k]}^2}{n_{[k]}}\right]^{1/2}}. \end{aligned}$$

Condition on the first stage sample variances $S^2 = s^2$ (Rinott 1978), so that the n_i can be considered fixed. Then $\mathbf{1}_{(\bar{\bar{X}}_i > \bar{\bar{X}}_{[k]})} = \mathbf{1}_{(Z_i > R_i)}$ and Z_i is a standard normal random variable. Note that $n_i \geq g^2 s_i^2$ implies $R_i \geq Q_i$, where

$$Q_i = \frac{g\delta_i}{\left[\frac{\sigma_i^2}{s_i^2} + \frac{\sigma_{[k]}^2}{s_{[k]}^2}\right]^{1/2}}$$

for all $i \neq [k]$. Therefore $\mathbf{1}_{(Z_i > R_i)} \leq \mathbf{1}_{(Z_i > Q_i)}$. Let $\Phi(\cdot)$ be the cumulative distribution function of a standard normal random variable, and let $\phi(\cdot)$ be its probability density function. Then

$$\begin{aligned} E[OC_i] &\leq E[E[\delta_i \mathbf{1}_{(Z_i \geq Q_i)} \mid S_i^2, S_{[k]}^2]] \\ &= E\left[\delta_i \Phi\left(\frac{-g\delta_i}{\left[\frac{\sigma_i^2}{S_i^2} + \frac{\sigma_{[k]}^2}{S_{[k]}^2}\right]^{1/2}}\right)\right] \\ &= E\left[\delta_i \Phi\left(\frac{-g\delta_i}{\left[(n_0 - 1)\left(\frac{1}{Y_i} + \frac{1}{Y_{[k]}}\right)\right]^{\frac{1}{2}}}\right)\right] \qquad (2) \end{aligned}$$

where the final expectation is taken with respect to the independent random variables $Y_j = (n_0 - 1)S_j^2/\sigma_j^2$ which are known to have a $\chi^2_{n_0-1}$ distribution.

Choose g so that the right hand side of Equation (2) equals $\Delta/(k-1)$ to guarantee an expected loss $E[OC_i] \leq \Delta/(k-1)$ for a comparison with two systems. Inequality (1) in turn implies that the overall expected opportunity cost is less than Δ.

3.2 Finding g Given Δ

Section 3.1 might suggest that we identify a parameter g for a given maximum expected opportunity cost Δ, along with a value of δ_i that corresponds to an indifference-zone parameter. Here, we turn the problem around and show that for a given g, we can determine the *least favorable* $\delta_i = \mu_{[k]} - \mu_i$. The least favorable δ_i then determines a bound on $E[OC_i]$ for each pairwise comparison. Since g determines $E[OC_i]$, we take a given bound Δ for the expected opportunity cost, and find the smallest g that delivers the guarantee. The guarantee is delivered for *all* configuration of the means. That explains why Procedure OC_f does not have an indifference-zone parameter.

Let $f(\delta_i)$ be the upper bound in the right hand side of Equation (2), viewed as a function of δ_i. Then

$$\begin{aligned} W &= \frac{-g\delta_i}{((n_0-1)(1/Y_1+1/Y_2))^{1/2}} \\ f(\delta_i) &= E[\delta_i \Phi(W)]. \end{aligned} \qquad (3)$$

The derivative of f with respect to δ_i gives a first order optimality condition for the least favorable δ_i.

$$\frac{df}{d\delta_i} = E[\Phi(W) + W\phi(W)] = 0 \qquad (4)$$

Both the distribution of the random variable W and the optimality condition are *invariant to transformations* of (g, δ_i) to any $(g\alpha, \delta_i/\alpha)$, where $\alpha > 0$ is real-valued. The implication is that knowing the LFC for any one value of g immediately leads to a knowledge of the LFC for any other value of g. Similarly, the expected opportunity cost bound for the comparison scales with δ_i (see Equation (2)). Since W depends upon the Y_i, and the distribution of the Y_i changes as a function of n_0, numerical solutions for δ_i may be required for different values of n_0.

Figure 1 illustrates the scaling properties transformations from (g, δ_i) to $(g\alpha, \delta_i/\alpha)$. Doubling g from 0.5 to 1 halves the least favorable δ_i from 2.66 to 1.33, and the maximum pairwise loss drops from 0.576 to 0.288. As $\delta_i \to 0$ the probability of an incorrect selection approaches 1/2, but the penalty for choosing the wrong system goes to zero with δ_i. As $\delta_i \to \infty$ the penalty for an incorrect selection grows without bound, but the probability of incorrect selection goes to 0 so fast that expected opportunity cost approaches 0. This discussion justifies the following characterization of the least favorable δ_i, relative to Equation (2).

Lemma 1 *Choose $\beta, \alpha > 0$. Let $\delta_i = \mu_{[k]} - \mu_i$ be the LFC for a comparison between systems i and $[k]$ when $g = \beta$, relative to Equation (2), so the resulting bound on the expected opportunity cost is $f(\delta_i)$. Then the LFC if $g = \beta\alpha$ is δ_i/α and the corresponding expected opportunity cost bound is $f(\delta_i)/\alpha$.*

The least favorable δ_i for a given g can be determined numerically. We did this for several values of n_0, and report the results below in Section 4. Section 4 also indicates how to use a table to choose g for Step 4 of Procedure OC_f to guarantee that the overall expected opportunity cost $E[OC]$ is less than the specified bound, Δ.

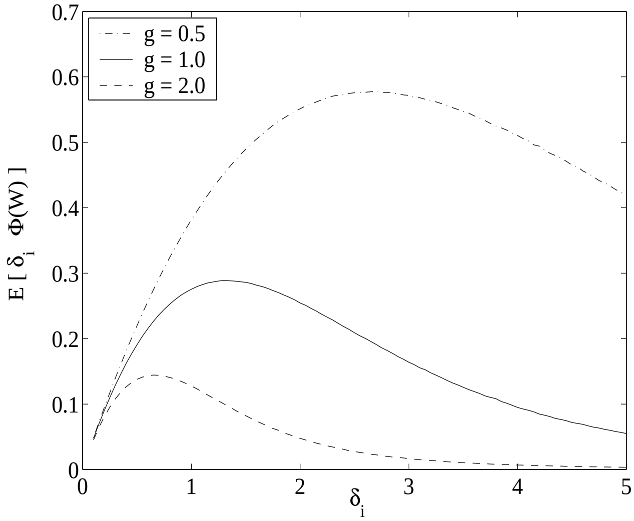

Figure 1: The Upper Bound $E[\delta_i \Phi(W)]$ for the Expected Opportunity Cost, Given Several Values of the Allocation Parameter g ($k = 2$, $n_0 = 5$)

3.3 Opportunity Cost And Existing Procedures

This section shows that two previously proposed indifference-zone procedures implicitly provide expected opportunity cost guarantees, even though they were initially designed to provide guarantees for the selection probability.

Like Procedure OC_f, Procedure R of Rinott (1978) allocates a total number of observations that is proportional to the first-stage sample variance and selects the system with the best overall sample mean. This suggests that Procedure R also provides an expected opportunity cost guarantee.

Lemma 2 *Procedure R provides an expected opportunity cost guarantee for all configurations of the means, not just those in the indifference-zone.*

Proof: Recall that Procedure R obtains a minimum probability of correct selection of P^* when the best system is at least δ^* better than the others by setting

$$n_i = \max\left\{n_0, \left\lceil \left(\frac{h}{\delta^*}\right)^2 s_i^2 \right\rceil\right\}, \qquad (5)$$

where $h = h(k, P^*, n_0 - 1)$ is the solution to Rinott's integral (e.g., see Bechhofer, Santner, and Goldsman 1995). Set $g = h/\delta^*$. By the results in Section 3, the choice of δ^* and P^* implicitly determine an $E[OC]$ guarantee for *all* configurations of the means. Conversely, the choice of an $E[OC]$ guarantee determines the value of g, which in turn determines an implied indifference-zone parameter δ^* for each desired P^*. □

Nelson and Banerjee (2001) present procedures that provide a probability of *good* selection, P(GS), guarantee, meaning that a lower bound for the probability of selecting a system within δ^* of the best is provided for any configuration of the mean. One of them, Procedure S, allocates a total number of samples proportional to the first stage sample mean, and selects the system with the largest overall sample mean as best. Therefore Procedure S also provides an implicit expected opportunity cost guarantee.

Lemma 3 *Procedure S (Nelson and Banerjee 2001) provides an expected opportunity cost guarantee for all configurations of the means.*

Proof: Same as for Lemma 2, except $g = \sqrt{2}v^{P^*}_{n_0-1}/\delta^*$ is the analogous factor for Procedure S. □

Interestingly, the LFC for Procedure R and Procedure S with respect to the probability of correct selection *may* or *may not* be the same as the LFC for those procedures with respect to the expected opportunity cost. Section 5.4 explores this issue in more detail.

Procedure R and Procedure S may require more replications than are needed because it they are statistically conservative. Both procedures ignore first stage information about the sample mean. All systems have a total number of observations proportional to the first stage sample variance, even if the first stage sample means of some systems are significantly inferior (with high probability). Procedure OC_f suffers from the same criticism. Nelson, Swann, Goldsman, and Song (2001) proposed combined screening and selection procedures to address the statistical conservatism of indifference zone procedures. The idea is to screen out systems whose first stage sample mean and variance indicate that they are not likely contenders for the 'best'. A second stage then allocates samples proportional to the sample variances of the remaining systems in a way that guarantees P(CS) $\geq P^*$. An area for future work is the development of a combined screening/selection procedure that can provide an expected opportunity cost guarantee with potentially fewer replications.

4 TABLE FOR PROCEDURE OC_f

Table 1 gives the LFC for a comparison between $k = 2$ systems, $[k]$ and i, as a function of the number of first stage replications n_0 of each system. The table presumes $g = 1$, and defines the LFC to be the value of $\delta_i = \mu_{[k]} - \mu_i$ that maximizes the upper bound $E[OC_i]$ for the expected opportunity cost of a comparison in Equation (2). The table also gives the value of that bound when the means are in the LFC. If $g \neq 1$, Lemma 1 says that the corresponding values of the LFC and $E[OC_i]$ are obtained by dividing the appropriate values in the table by g.

Table 1 is straightforward to use. Suppose that there are $k = 5$ systems, that $n_0 = 10$ replications are to be run for each system during the first stage, and that the maximum acceptable $E[OC]$ is $\Delta = 0.4$. The correct factor g for Procedure OC_f to guarantee $E[OC] \leq \Delta$ is determined as follows. There are $k - 1 = 4$ comparisons between each alternate and the best, so the acceptable maximal loss per paired comparison is $\Delta/(k-1) = 0.1$. The opportunity cost entry in Table 1 corresponding to $n_0 = 10$ is $E[OC_i] = 0.2602$. By setting

$$g = \frac{E[OC_i]}{\Delta/(k-1)} = \frac{0.2602}{0.1} = 2.602,$$

we can guarantee the $E[OC]$ bound.

Table 1: Least Favorable $\delta_i = \mu_{[k]} - \mu_i$ for Procedure OC_f and Loss for Given Number of First Stage Replications, n_0, Assuming $g = 1$, $k = 2$

n_0	LF $\hat{\delta}_i$	$E[\hat{OC}_i]$	$\hat{SE}_{E[OC_i]}$
3	1.767	0.3524	0.0004
4	1.449	0.3078	0.0003
5	1.330	0.2886	0.0002
6	1.267	0.2779	0.0002
7	1.227	0.2711	0.0001
8	1.201	0.2663	0.0001
9	1.182	0.2629	0.0001
10	1.167	0.2602	0.0001
11	1.156	0.2581	0.0001
12	1.147	0.2564	0.0001
13	1.139	0.2550	0.0001
14	1.133	0.2539	0.0001
15	1.128	0.2529	0.0001
16	1.123	0.2520	0.0001
17	1.119	0.2513	0.0001
18	1.116	0.2506	0.0001
19	1.113	0.2500	0.0001
20	1.110	0.2495	0.0001

The $E[OC_i]$ estimates in Table 1 were determined by generating 200,000 values of (Y_1, Y_2) for Equation (3) using CRN across the values of n_0 and have a standard error given by $\hat{SE}_{E[OC_i]}$. The least favorable $\delta_i = \mu_{[k]} - \mu_i$ for each n_0 was determined by sample path optimization with the `fminsearch` function of Matlab. The estimation process was repeated several times for the smaller values of n_0, which have a larger standard error. The estimated values of the least favorable δ_i varied within ± 0.003 of the values reported in the table.

5 EXPERIMENTAL RESULTS

This section examines Procedure OC_f with simulation.

5.1 Tightness of Bound and the Variance

The LFC in Table 1, relative to the opportunity cost bound, does not depend upon the variance of each system. This raises the question of how tight the bound may be as a function of the size of the variances of each system. As the

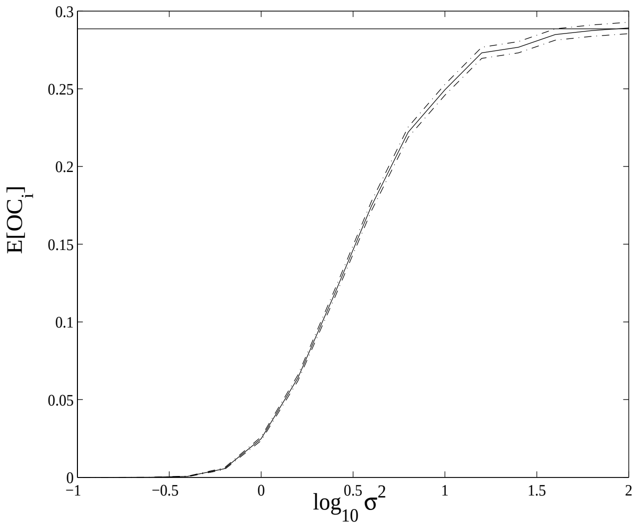

Figure 2: Expected Opportunity Cost as Function of a Common Variance σ^2 ($k = 2$, $\delta_i = 1.33$, $n_0 = 5$, $g = 1$)

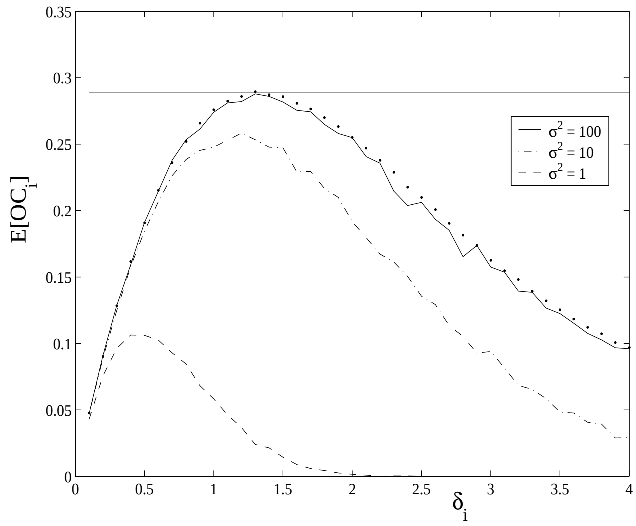

Figure 3: Expected Opportunity Cost in Comparison, Compared with Bound, Common σ^2 ($k = 2$, $n_0 = 5$, $g = 1$)

variances go to 0, the chance of an incorrect selection and the expected opportunity cost go to 0, because the n_0 first stage replications will accurately identify the relative order of the means, and additional replications will be required extremely rarely. The bound is expected to be tighter if all systems have larger variances, as the χ^2 approximation is better when the total number of replications is less likely to be constrained by the first stage sample size ($n_i = \max\{n_0, \lceil g^2 s_i^2 \rceil\}$).

Figure 2 illustrates the tightness of the bound improves as the variance increases for a comparison of $k = 2$ systems with $g = 1$, $n_0 = 5$, $\delta_i = 1.33$ (the least favorable δ_i relative to the bound for losses in Equation (3)). The graph presumes a common variance $\sigma^2 = \sigma^2_{[k]} = \sigma^2_i$, and was generated with 50,000 applications of the selection procedure for each value of σ^2. The 90% CI is demarcated with the dotted lines. For small values of the variance the expected loss well below the bound, but when the variance is on the order of 100 the expected opportunity cost is quite close to the theoretical upper bound. The majority of the loss increases occurs as σ^2 rises in the range from 1 to 10. The same qualitative phenomenon was observed in numerical experiments with $n_0 = 10$.

5.2 Bound: δ_i and Variance Interaction

The bound on $E[OC_i]$ is relatively tight for large variances (e.g. on the order of 100 or more) when the difference in means is close to the least favorable δ_i. Figure 3 shows how tight the bound for $E[OC_i]$ is to the actual opportunity cost lost in a comparison ($k = 2$) over a range of δ_i for 'small', 'medium' and 'large' values of the variance of both systems (25000 applications of the selection procedure for each combination of δ_i and σ^2).

Figure 3 suggests that a smaller value of g may be used in a variation of Procedure OC_f that requires the variance of each system be less than or equal to a threshold (an indifference-zone like constraint on the variances rather than the means). Losses appear to be smaller for smaller values of the variance, and the least favorable δ_i seems to decrease as the maximum allowable variance is decreased.

5.3 Slack in Bound for $k > 2$

The upper bound for $E[OC]$ in Equation (1) may be loose when more than two systems are being compared because the bound adds the sum of maximum losses from $k - 1$ comparisons with 2 systems. That bound is therefore analogous to a P(CS) bound with the Bonferroni inequality. This section examines how loose the bound is with a numerical experiment for $k = 2, 3, 5$ and 10 systems.

Figure 4 illustrates the expected opportunity cost as a function of the difference between the best system and the performance of each of the other systems (50,000 applications of the selection procedure for each combination of δ_i and k). The variance of each system is presumed to be 100. The figure suggests that the value of δ_i for the least favorable slippage configuration (each of the nonbest systems has the same mean, so the δ_i are the same for all $i \neq [k]$) appears to increase as a function of k. A proof of this, and a proof of the conjecture that the LFC is a slippage configuration when the variance of each system is the same, is an area for future work. If the conjecture is true, tables for the least favorable δ_i and the worst case $E[OC]$ would be straightforward to construct.

Figure 4 also confirms the looseness in the bound due to the Bonferroni-like summation of losses. The maximum

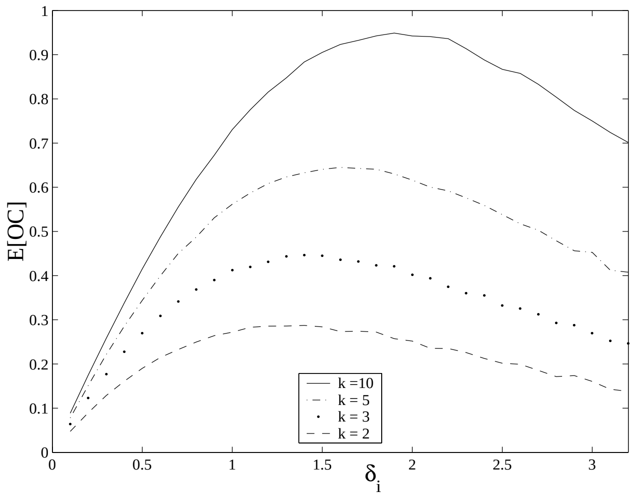

Figure 4: Estimate of Expected Opportunity Cost with $k \geq 2$ Simulated Systems (For Common $\sigma^2 = 100$, $n_0 = 5$, $g = 1$)

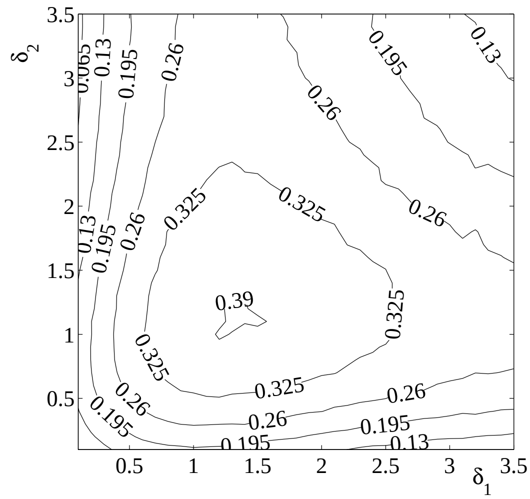

Figure 5: Expected Opportunity Cost with $k = 3$ Simulated Systems (Different σ_i^2, $n_0 = 5$, $g = 1$)

loss when $k = 3$ systems are compared (two comparisons of two systems) is approximately 0.45, and is less than two times the loss from one comparison of two systems, $2 \times 0.288 \approx 0.576$. More starkly, the worst-case when $k = 10$, which has nine comparisons with 2 systems, is certainly less than nine times the worst case loss with $k = 2$. Slepian's inequality is a tool for improving upon the Bonferonni inequality for P(CS) bounds. The development of better bounds when $k > 2$ for more general loss functions, like the expected opportunity cost, is an area of future work.

The analysis above presumes that each system has the same variance. Figure 5 describes one situation when the variances are different in a comparison of $k = 3$ systems (40,000 applications of the selection procedure for each combination of δ_i, with $\sigma_1^2 = 5$, $\sigma_2^2 = 10$, $\sigma_3^2 = 10$). The figure presents a contour plot of the estimated expected opportunity cost lost as a function of the differences between the best system. The LFC still appears to be near the slippage configuration. Analytically one would suspect this when the variances are all large, since the 'max' operator for the second stage allocation would have less effect, so the χ^2 distribution approximation in Equation (2) improves. When one variance is particularly small the second stage allocation will most likely be zero. This may cause the LFC to not be a slippage configuration if the variance of different systems are constrained to be small with different upper bounds.

5.4 P(CS) and δ^* versus E[OC]

Specifying a P(CS) of P^* and indifference zone parameter δ^* for Procedure R leads to an expected opportunity cost of $\delta^*(1 - P^*)$ when the means are in the LFC for Procedure R. But this is not the LFC for Procedure OC_f, so the $E[OC]$ may differ from $\delta^*(1 - P^*)$.

Consider the case $k = 2$. The second stage alloction of Procedure R with $n_0 = 5$ and $P^* = 0.95$ depends upon the constant $h = h(2, 0.95, 5 - 1) = 3.107$ (see Bechhofer, Santner, and Goldsman 1995, p. 63). If $\delta^* = 1$, that corresponds to a value of $g = h/\delta^* = 3.107$. The LFC for Procedure R under these assumptions is $\delta_i = 1$, with expected opportunity cost $1 \times (1 - 0.95) = 0.05$. But the LFC for Procedure OC_f with $g = 3.107$ is $\delta_i = 1.330/g = 0.428$. The associated $E[OC]$ bound is $0.2886/g = 0.0929$, which is greater than the $E[OC]$ of 0.05 that occurs in the LFC for Procedure R. More generally, the LFC for Procedure OC_f may lead to a more severe E[OC] than the E[OC] associated with the LFC of Procedure R. The implication is that specifying P^* and δ^* with Procedure R does *not* give an accurate bound for an expected opportunity cost guarantee Δ.

On the other hand, if P^* and δ^* are selected by a decision maker with the idea that a maximum expected opportunity cost of $\Delta = \delta^*(1 - P^*)$ is tolerable, then Procedure OC_f can be used with that Δ. Typically this will require more replications than required by an indifference-zone procedure with parameters P^*, δ^*, because $E[OC] \leq \Delta$ is guaranteed over *all* configurations of the means, not just for the slippage configuration determined by δ^*.

6 IS THE LFC FOR PROCEDURE OC_f A SLIPPAGE CONFIGURATION?

The above sections provide some analytical results, some empirical observations, and several questions. One question is whether or not the LFC for Procedure OC_f is a slippage configuration when $k > 2$. This section provides a preliminary analysis that indicates that the slippage configuration satisfies first order optimality conditions for being a LFC when $k = 3$ at least one of two special assumptions is true—if the variances of all nonbest systems are the same (a weak form of homoscedasticity), or when the variances all approach infinity together (an asymptotic argument).

Since $E[OC]$ is not convex even if $k = 2$ (see Figure 3), using convexity in the δ_i to show that the LFC is a SC is not viable. The approach taken here is to find the LFC subject to the constraint that the δ_i lie on a given simplex, $\sum_{i \neq [k]} \delta_i = c$, with $\delta_i \geq 0$ for $i \neq [k]$. If the LFC for each simplex is a SC, then the LFC over all $\delta_i \geq 0$ is then the least favorable SC.

Suppose that there are $k = 3$ systems, and without loss of generality (WLOG) suppose that the best is system $[k] = 3$. We proceed by first examining the loss contributed by systems 1 and 2, and try to find the least favorable δ_i for those two systems, subject to the constraint $\delta_1 + \delta_2 = c$. Condition on the values of $\bar{\bar{x}}_3$ and each s_i^2. A loss of δ_1 is incurred when the overall sample mean of system 1 exceeds that of all other systems. The probability of that event can be determined by conditioning on whether or not system 2 exceeds $\bar{\bar{x}}_3$. Set $a_i = (\sigma_i^2/n_i)^{-1/2}$. Recall that $\delta_2 = c - \delta_1$ and that the $\bar{\bar{X}}_i$ are conditionally independent and normally distributed, given the s_i^2, to obtain

$$\begin{aligned} & \Pr(\bar{\bar{X}}_1 > \bar{\bar{x}}_3, \bar{\bar{X}}_1 > \bar{\bar{X}}_2) \\ = \; & E[\mathbf{1}_{(\bar{\bar{X}}_1 > \bar{\bar{x}}_3 > \bar{\bar{X}}_2)} + \mathbf{1}_{(\bar{\bar{X}}_1 > \bar{\bar{X}}_2 > \bar{\bar{x}}_3)}] \\ = \; & \Phi[-(\bar{\bar{x}}_3 + \delta_1)a_1]\Phi[(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ & + E_{\bar{\bar{X}}_2}[\mathbf{1}_{(\bar{\bar{X}}_2 > \bar{\bar{x}}_3)} \Phi[-(\bar{\bar{X}}_2 + \delta_1)a_1]] \end{aligned} \tag{6}$$

and

$$\begin{aligned} & \Pr(\bar{\bar{X}}_2 > \bar{\bar{x}}_3, \bar{\bar{X}}_2 > \bar{\bar{X}}_1) \\ = \; & \Phi[(\bar{\bar{x}}_3 + \delta_1)a_1]\Phi[-(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ & + E_{\bar{\bar{X}}_1}[\mathbf{1}_{(\bar{\bar{X}}_1 > \bar{\bar{x}}_3)} \Phi[-(\bar{\bar{X}}_1 + c - \delta_1)a_2]]. \end{aligned} \tag{7}$$

The expected loss $H(\bar{\bar{x}}_3, \delta_1)$, conditional on $\bar{\bar{x}}_3$ and the $a_i = (\sigma_i^2/n_i)^{-1/2}$, due to systems 1 and 2 is

$$\begin{aligned} H(\bar{\bar{x}}_3, \delta_1) \; = \; & \delta_1 \Pr(\bar{\bar{X}}_1 > \bar{\bar{x}}_3, \bar{\bar{X}}_1 > \bar{\bar{X}}_2) \\ & + (c - \delta_1) \Pr(\bar{\bar{X}}_2 > \bar{\bar{x}}_3, \bar{\bar{X}}_2 > \bar{\bar{X}}_1), \end{aligned} \tag{8}$$

so the unconditional loss contributed by systems 1 and 2 is

$$E_{\bar{\bar{X}}_3, S_i^2}[H(\bar{\bar{x}}_3, \delta_1) \mid \bar{\bar{X}}_3 = \bar{\bar{x}}_3, S_i^2 = s_i^2]. \tag{9}$$

The first order optimality condition for the LFC, given $\delta_1 + \delta_2 = c$, is $\partial E[H]/\partial \delta_1 = 0$. The function H is sufficiently 'nice' to interchange the derivative and expectation, resulting in the condition $E[\partial H/\partial \delta_1] = 0$, where

$$\begin{aligned} \frac{\partial H}{\partial \delta_1} \; = \; & \Phi[-(\bar{\bar{x}}_3 + \delta_1)a_1]\Phi[(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ - \; & a_1\delta_1\phi(-(\bar{\bar{x}}_3 + \delta_1)a_1)\Phi[(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ - \; & a_2\delta_1\Phi(-(\bar{\bar{x}}_3 + \delta_1)a_1)\phi[(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ + \; & E_{\bar{\bar{X}}_2}[\mathbf{1}_{(\bar{\bar{X}}_2 > \bar{\bar{x}}_3)}\Phi[-(\bar{\bar{X}}_2 + \delta_1)a_1]] \\ - \; & a_1\delta_1 E_{\bar{\bar{X}}_2}[\mathbf{1}_{(\bar{\bar{X}}_2 > \bar{\bar{x}}_3)}\phi[-(\bar{\bar{X}}_2 + \delta_1)a_1]] \\ - \; & \Phi[(\bar{\bar{x}}_3 + \delta_1)a_1]\Phi[-(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ + \; & a_1(c - \delta_1)\phi((\bar{\bar{x}}_3 + \delta_1)a_1)\Phi[-(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ + \; & a_2(c - \delta_1)\Phi((\bar{\bar{x}}_3 + \delta_1)a_1)\phi[-(\bar{\bar{x}}_3 + c - \delta_1)a_2] \\ - \; & E_{\bar{\bar{X}}_1}[\mathbf{1}_{(\bar{\bar{X}}_1 > \bar{\bar{x}}_3)}\Phi[-(\bar{\bar{X}}_1 + c - \delta_1)a_2]] \\ + \; & a_2(c - \delta_1)E_{\bar{\bar{X}}_1}[\mathbf{1}_{(\bar{\bar{X}}_1 > \bar{\bar{x}}_3)}\phi[-(\bar{\bar{X}}_1 + c - \delta_1)a_2]]. \end{aligned} \tag{10}$$

Lemma 4 *Suppose $k = 3$. If the a_i are independent and identically distributed (*i.i.d.*) for all $i \neq [k]$, then the slippage configuration satisfies the first order optimality conditions for being a LFC of Procedure OC_f.*

Proof: Suppose that the a_i are independent and identically distributed, for all $i \neq [k]$, and let $c \geq 0$ be arbitrary. WLOG assume $[k] = 3$. If $\delta_1 = \delta_2 = c/2$, the first and sixth terms of Equation (10) cancel when the expectation over a_1, a_2 is taken (which is equivalent to taking the expectation over s_1^2, s_2^2). Similarly the fourth and ninth terms cancel, and the fifth and tenth terms cancel in expectation. Some algebra indicates that the sum of the remaining four terms is also 0 in expectation. This is true for all $\bar{\bar{x}}_3$, and therefore this first order optimality condition also holds when the expectation over $\bar{\bar{x}}_3$ is taken. □

Further inspection indicates that $E[H(\bar{\bar{x}}_3, \delta_1) \mid \bar{\bar{x}}_3] = E[H(\bar{\bar{x}}_3, c - \delta_1) \mid \bar{\bar{x}}_3]$, so $E[H(\bar{\bar{x}}_3, \delta_1) \mid \bar{\bar{x}}_3]$ is symmetric about $\delta_1 = c/2$ when the a_i are *i.i.d.*. We hypothesize but do not prove further optimality properties for $k \geq 2$.

Corollary 5 *Suppose $k = 3$. The slippage configuration satisfies the first order condition for the LFC of Procedure OC_f if the variance of each system is the same, i.e. $\sigma_i^2 = \sigma^2$, for all $i \neq [k]$.*

Proof: Assume that $\sigma_i^2 = \sigma^2$ for all $i \neq [k]$. Then the distributions of the s_i^2 are the same, so the distribution of the $n_i = \max\{n_0, \lceil g^2 s_i^2 \rceil\}$ are the same (for all $i \neq [k]$), as are the distributions of the $a_i = (\sigma_i^2/n_i)^{-1/2}$. □

The conclusions of Corollary 5 are therefore also true for the homoscedastic case, when *all* systems, including the best, have the same variance. The conclusions of Corollary 6 *asymptotically* cover the heteroscedastic case.

Corollary 6 *Suppose $k = 3$. The slippage configuration satisfies the first order conditions to be a LFC of Procedure OC_f asymptotically as the variances of all* *systems go to infinity together, i.e. when* $Var[X_{i,j}] = \kappa\sigma_i^2$ *for each* i*, in the limit* $\kappa \to \infty$.

Proof: Under the assumption of the hypothesis, as $\kappa \to \infty$, the distribution of $a_i^2(n_0-1)/g^2$ approaches a $\chi^2_{n_0-1}$ distribution, for all i (Section 3 or Rinott 1978). □

If the slippage configuration eventually is shown to be the LFC for $k = 3$ (not just satisfying first order conditions) and eventually for $k > 3$, subject to the constraint $\sum_{i \neq [k]} \delta_i = c$, as we hypothesize, this transforms a $k-1$ dimensional problem involving all δ_i into a one dimensional problem of finding the least favorable slippage configuration. Tables for the least favorable δ_i as a function of n_0 and k could be constructed on that basis. Checking second order optimality conditions is an area for future work.

7 CONCLUSIONS

This paper appears to provide the first selection procedure that guarantees an upper bound for the expected opportunity cost, taken in a frequentist sense. The analysis shows that the indifference-zone parameter δ^* can be dispensed with for this procedure. Further, we proved that some existing indifference-zone procedures have implicit expected opportunity cost guarantees for all configurations of the means, even though the LFC with respect to the probability of correct selection may not be the same as the LFC for the expected opportunity cost.

The paper identified several potential areas for future research in the area of frequentist style guarantees for the expected opportunity cost or other loss functions. A proof that the LFC is a slippage configuration would facilitate the development of tables when $k > 2$ that would improve sampling efficiency. Further improvements would be anticipated if combined screening/selection procedures could be developed. Extensions to handle common random numbers are also of interest. Alternately, developing variations of Bayesian decision-theoretic procedures in order to provide frequentist guarantees is an alternate route for cross-fertilization. Allocating replications proportional to the first stage sample variances appears to simplify the analysis due to the independence of the sample variance and sample mean. The Bayesian allocations do not have that property, so a frequentist analysis for a Bayesian allocation appears to be more challenging.

REFERENCES

Bechhofer, R. E., T. J. Santner, and D. M. Goldsman. 1995. *Design and analysis for statistical selection, screening, and multiple comparisons*. New York: John Wiley & Sons, Inc.

Chick, S. E., and K. Inoue. 2001a. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49 (5): 732–743.

Chick, S. E., and K. Inoue. 2001b. New procedures for identifying the best simulated system using common random numbers. *Management Science* 47 (8): 1133–1149.

Dudewicz, E. J., and S. R. Dalal. 1975. Allocation of observations in ranking and selection with unequal variances. *Sanhkya* B37 (1): 28–78.

Matejcik, 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., 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.

Rinott, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics* A7 (8): 799–811.

AUTHOR BIOGRAPHY

STEPHEN E. CHICK is Associate Professor of Technology Management at INSEAD in Fontainebleau, France. He is on leave from the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. In addition to stochastic simulation, his research focuses on Bayesian statistics, public health, decision analysis and computational methods in statistics. His research is motivated by problems in manufacturing, operations, and health care. He can be contacted at <stephen.chick@insead.edu>.