
New Two-Stage and Sequential Procedures for Selecting the Best Simulated System

Author(s): Stephen E. Chick and Koichiro Inoue

Source: *Operations Research*, Vol. 49, No. 5 (Sep. - Oct., 2001), pp. 732-743

Published by: INFORMS

Stable URL: <http://www.jstor.org/stable/3088571>

Accessed: 09/10/2009 08:57

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=informs>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Operations Research*.

NEW TWO-STAGE AND SEQUENTIAL PROCEDURES FOR SELECTING THE BEST SIMULATED SYSTEM

STEPHEN E. CHICK
KOICHIRO INOUE

*Department of Industrial and Operations Engineering, The University of Michigan,
1205 Beal Avenue, Ann Arbor, Michigan, 48109-2117, sechick@engin.umich.edu*

(Received July 1998; revision received June 1999; accepted May 2000)

Standard “indifference-zone” procedures that allocate computer resources to infer the best of a finite set of simulated systems are designed with a statistically conservative, least favorable configuration assumption consider the probability of correct selection (but not the opportunity cost) and assume that the cost of simulating each system is the same. Recent Bayesian work considers opportunity cost and shows that an average case analysis may be less conservative but assumes a known output variance, an assumption that typically is violated in simulation. This paper presents new two-stage and sequential selection procedures that integrate attractive features of both lines of research. They are derived assuming that the simulation output is normally distributed with unknown mean and variance that may differ for each system. We permit the reduction of either opportunity cost loss or the probability of incorrect selection and allow for different replication costs for each system. The generality of our formulation comes at the expense of difficulty in obtaining exact closed-form solutions. We therefore derive a bound for the expected loss associated potentially incorrect selections, then asymptotically minimize that bound. Theoretical and empirical results indicate that our approach compares favorably with indifference-zone procedures.

An important use of stochastic simulation is the identification of the best of several alternative systems, where “best” is defined in terms of the maximum (or minimum) mean value of simulation output (Law and Kelton 1991, Banks et al. 1996). There is therefore interest in statistical selection procedures to identify the best system (Matejcek and Nelson 1995, Bechhofer et al. 1995, Goldsman and Nelson 1998). Typically, a few simulation replications for each system are run in a preliminary stage, and the decision-maker measures the evidence that a given system is best. If the evidence is insufficient, additional replications are run to obtain more information about the identity of the best system.

Well-known indifference-zone procedures determine the number of additional replications based on a statistically conservative, least favorable configuration (LFC) assumption (Rinott 1978, Bechhofer et al. 1995). Some Bayesian approaches to the problem of selecting the best system (Chen 1996; Gupta and Miescke 1994, 1996) incorporate first-stage sample mean information that is ignored by indifference-zone procedures. They suggest that an average case analysis may lead to a significant reduction in computing effort, relative to indifference-zone procedures. However, these Bayesian formulations formally assume a known output variance, an assumption that is likely to be violated in simulation practice.

This paper derives new two-stage and sequential selection procedures that integrate attractive features of both lines of research, and it provides additional flexibility to address practical concerns. First, the new procedures are derived with a Bayesian, average case analysis to avoid

the statistically conservative LFC assumption. Second, we assume that the variance for each system may be different and unknown, a feature of some indifference-zone procedures that is not formally afforded by previous Bayesian developments. Further, we allow the analyst either to increase the probability of correct selection or to reduce the expected opportunity cost of potentially incorrect selections. Indifference-zone procedures consider only the probability of correct selection. Finally, we allow for a computing budget constraint and permit the cost of each system’s simulations to be different. None of the work mentioned above explicitly accounts for differing simulation costs.

Medicine and agriculture also employ statistical procedures to select the best system. (“Which population has the highest average blood cholesterol?” “Which crop treatment improves harvest the most?”) This paper applies to those fields as well, although the relevant terminology differs (replace “replication” with “sample,” and “system” with “treatment”).

In §1 we recall two indifference-zone procedures (Rinott 1978, Goldsman and Nelson 1998) that serve as references for comparison. We motivate our general Bayesian formulation and explain how it differs from the indifference-zone approach in §2. Closed-form solutions of certain special cases (e.g., known variance, common sampling costs, allocation of a single sample) are known (Gupta and Miescke 1994, 1996; Berger 1988), but a closed-form solution to the general problem is unknown at present. Here we develop suboptimal allocations of additional replications for the general problem by asymptotically minimizing a bound on

Subject classifications: Simulation, statistical analysis: selecting the best simulated system. Statistics, Bayesian: Bayesian inference for ranking and selection.
Area of review: SIMULATION.

the expected loss (cost of replications plus expected loss for potentially selecting an inferior system). New two-stage procedures that reduce the expected opportunity cost and the probability of incorrect selection are presented and justified in §3.1 and §3.2, respectively. We provide sequential analogs of those procedures in §3.3, in an attempt to provide further sampling efficiencies. The new procedures compare favorably with indifference-zone procedures in an empirical study in §4.

1. BACKGROUND

The indifference-zone formulation has strongly influenced the design of procedures for identifying the best of k systems. The idea is to identify the best system with a guaranteed lower bound P^* on the probability of correct selection (PCS), given the assumption that the best is better than the other systems by at least δ^* , the smallest difference worth detecting in the opinion of a decision-maker. Indifference-zone procedures are statistically conservative because second-stage replications are allocated to guarantee the probability of correct selection when the unknown means are in the LFC.

To benchmark the new procedures in §3, we present two indifference-zone procedures that provide PCS guarantees when the simulation output is normally distributed.

Rinott (1978) proposes a well-known two-stage procedure that we call Procedure \mathcal{R} . First, δ^* , P^* , and the first-stage sample size r_0 are specified. Independent replications $x_{i,1}, \dots, x_{i,r_0}$ are run for each system, $i = 1, \dots, k$, and the first-stage sample mean $\bar{x}_i = \sum_{j=1}^{r_0} x_{i,j}/r_0$ and sample variance $\hat{\sigma}_i^2 = \sum_{j=1}^{r_0} (x_{i,j} - \bar{x}_i)^2/(r_0 - 1)$ are computed. The number of second-stage replications $\mathbf{r} = (r_1, \dots, r_k)$ depends on δ^* , $\hat{\sigma}_i^2$, r_0 , and the solution $g = g(k, P^*, r_0 - 1)$ to Rinott's integral (e.g., see Bechhofer et al. 1995),

$$r_i = \max \left\{ 0, \left\lceil \left(\frac{g}{\delta^*} \right)^2 \hat{\sigma}_i^2 \right\rceil - r_0 \right\}. \quad (1)$$

After the second stage, the system with the highest overall sample mean is selected as best.

The total number of replications $r_0 + r_i$ per system is proportional to $\hat{\sigma}_i^2$ if r_0 is small, regardless of the first-stage sample mean. Simulating systems with clearly inferior first-stage sample means may be unnecessary. This motivates the combined screening and selection procedure of Goldsman and Nelson (1998), called Procedure \mathcal{C} here, which attempts to improve sampling efficiency by paring out noncompetitive systems using a subset selection technique. Nelson et al. (1999) present empirical evidence that Procedure \mathcal{C} outperforms Procedure \mathcal{R} over a range of settings.

PROCEDURE \mathcal{C} , A COMBINED SCREENING AND SELECTION PROCEDURE.

1. Specify the indifference-zone parameter δ^* , probability of correct selection guarantee $P^* = 1 - \alpha$, and first-stage sample size $r_0 \geq 2$. Set $t = t_{1-(1-\alpha/2)^{1/(k-1)}, r_0-1}$,

and let $g_{\alpha/2} = g(k, 1 - \alpha/2, r_0 - 1)$ solve Rinott's integral.

2. Take independent replications $x_{i,1}, \dots, x_{i,r_0}$ for each system, $i = 1, \dots, k$.
3. For each system, compute the first-stage sample mean $\bar{x}_i = \sum_{j=1}^{r_0} x_{i,j}/r_0$ and sample variance $\hat{\sigma}_i^2 = \sum_{j=1}^{r_0} (x_{i,j} - \bar{x}_i)^2/(r_0 - 1)$.
4. Calculate the quantity $\omega_{ij} = t((\hat{\sigma}_i^2 + \hat{\sigma}_j^2)/r_0)^{1/2}$ for all $i \neq j$. Form the screening subset I , containing each system i such that

$$\bar{x}_i \geq \bar{x}_j - (\omega_{ij} - \delta^*)^+ \text{ for all } j \neq i.$$

5. If I contains a single index, then stop and return that system as the best. Otherwise, for all $i \in I$, compute the second-stage sample sizes

$$r_i = \max \{ 0, \lceil (g_{\alpha/2}/\delta^*)^2 \hat{\sigma}_i^2 \rceil - r_0 \}.$$

6. Take r_i additional observations from all systems $i \in I$, independent of all other replications for all systems.
7. Estimate the means with the overall sample means $\bar{\bar{x}}_i = \sum_{j=1}^{r_0+r_i} x_{i,j}/(r_0 + r_i)$.
8. Select the system with the largest $\bar{\bar{x}}_i$ as best.

2. AN ALTERNATE FORMULATION

While Procedure \mathcal{C} improves on Procedure \mathcal{R} by incorporating some first-stage sample information about the unknown mean performance of each system, it does not incorporate all such information. The Bayesian decision-theoretic formulation presented here attempts to improve second-stage sampling efficiency by incorporating additional first-stage information into the allocation decision. We also describe useful flexibility that the Bayesian approach offers that indifference-zone approach has not yet offered.

Assume that the simulation output $x_{i,j}$ for system i is normally distributed with an unknown mean w_i and variance σ_i^2 ($i = 1, \dots, k; j = 1, \dots$), and that the $x_{i,j}$ are jointly independent. It will be easier at times to refer to the precision $\lambda_i = 1/\sigma_i^2$ instead of the variance. Set $\mathbf{w} = (w_1, \dots, w_k)$ and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_k)$. Throughout, we write vectors in boldface, random variables in upper case, and their realizations in lower case.

We incorporate first-stage information by assigning a prior distribution and using Bayes' rule to infer the values of the unknown mean W_i and precision Λ_i . Allocations are then made with respect to the *expected* probability of correct selection rather than the worst-case LFC. If we assume a noninformative prior distribution, it is straightforward (de Groot 1970) to show that after observing the first-stage output, the joint distribution for W_i and Λ_i is a normal-gamma distribution,

$$\Lambda_i \sim \mathcal{G}((r_0 - 1)/2, \hat{\sigma}_i^2(r_0 - 1)/2), \quad (2)$$

$$W_i | \Lambda_i \sim \mathcal{N}(\bar{x}_i, \Lambda_i^{-1}/r_0), \quad (3)$$

where the gamma distribution $\mathcal{G}(\alpha, \beta)$ has mean α/β and variance α/β^2 , and the normal distribution $\mathcal{N}(\mu, \sigma^2)$ has mean μ and variance σ^2 . Further, the marginal distribution and frequentist confidence interval for the unknown mean are both described by the same three-parameter Student-t distribution,

$$W_i \sim \text{St}(\bar{x}_i, r_0/\hat{\sigma}_i^2, r_0 - 1), \quad (4)$$

where the three-parameter Student-t distribution $\text{St}(\mu, \kappa, \nu)$ has mean μ , precision κ , and degrees of freedom ν . When $\nu > 2$, the variance is $\kappa^{-1}\nu/(\nu - 2)$.

The probability of correct selection is an essential feature of indifference-zone procedures. We therefore consider the 0-1 loss function associated with selecting i as best when the means are actually \mathbf{w} ,

$$\mathcal{L}_{0-1}(i, \mathbf{w}) = \begin{cases} 0 & \text{when } w_i = \max_j w_j, \\ 1 & \text{otherwise.} \end{cases} \quad (5)$$

The expectation $E_{\mathbf{w}}[\mathcal{L}_{0-1}(i, \mathbf{w})]$ is the probability that i is not the best system.

In many business and engineering applications, however, the expected opportunity cost is more relevant to a decision maker. For example, a 90% chance of incorrect selection with a loss of \$1 may be preferable to a 1% chance of losing \$1 million. We therefore also consider the *opportunity cost* (sometimes called the linear loss)

$$\mathcal{L}_{o.c.}(i, \mathbf{w}) = \max_j w_j - w_i. \quad (6)$$

The loss is 0 when the best system is correctly selected—otherwise the loss is the difference in means between the best and the selected system.

In practice, there may be a deadline for completing a simulation analysis, and the CPU times per replication $\mathbf{c} = (c_1, \dots, c_k)$ may differ for each system. For instance, conclusions from a simulation study may be needed in 24 hours, and simulations of systems with complex algorithms for factory floor control may take longer to complete than for systems with simple control algorithms. This motivates a budget-constrained allocation, with $\mathbf{c}\mathbf{r}^T = b$. Further, if \mathbf{c} can be converted to units that match the simulation output (say, dollars), then a trade-off can be made between additional replications and the expected value of information from those replications (an unconstrained budget allocation). The total number of replications can be constrained to be b , regardless of the CPU time for each replication, by setting $\mathbf{c} = (1, \dots, 1)$.

The goal is to determine the second-stage allocation \mathbf{r} that minimizes the cost of replications plus the expected loss to a decision maker after all replications have been run. Let $\mathbf{x}_{r_i} = (x_{i, r_0+1}, \dots, x_{i, r_0+r_i})$ denote the second-stage output for system i , and let $\mathbf{x}_{\mathbf{r}} = (\mathbf{x}_{r_1}, \dots, \mathbf{x}_{r_k})$ denote all second-stage output. The number of stages of data used for an estimate is given by the number of bars or hats, so that the overall sample mean is

$$\bar{x}_i = \sum_{j=1}^{r_0+r_i} x_{i,j} / (r_0 + r_i),$$

and $\hat{\sigma}_i^2 = \sum_{j=1}^{r_0+r_i} (x_{i,j} - \bar{x}_i)^2 / (r_0 + r_i - 1)$ is the overall sample variance. Then the posterior distribution of the unknown mean, conditional on \mathbf{x}_{r_i} , is (de Groot 1970)

$$W_i | \mathbf{x}_{r_i} \sim \text{St}(\bar{x}_i, (r_0 + r_i) / \hat{\sigma}_i^2, r_0 + r_i - 1). \quad (7)$$

Let $d^N(\mathbf{x}_{\mathbf{r}}) = \max_i \bar{x}_i$ be the system with the highest overall sample mean. Given $\mathbf{x}_{\mathbf{r}}$, the expected loss for loss function \mathcal{L} is $E_{\mathbf{w}|\mathbf{x}_{\mathbf{r}}}[\mathcal{L}(d^N(\mathbf{x}_{\mathbf{r}}), \mathbf{w}) | \mathbf{x}_{\mathbf{r}}]$.

Because \mathbf{r} is chosen before the second stage, we take the expectation with respect to the random variable $\mathbf{X}_{\mathbf{r}}$. The cost of replications plus the expected loss for selecting the system with the maximum overall sample mean is therefore

$$\rho(\mathbf{r}) \stackrel{\text{def}}{=} \mathbf{c}\mathbf{r}^T + E_{\mathbf{x}_{\mathbf{r}}}[E_{\mathbf{w}|\mathbf{x}_{\mathbf{r}}}[\mathcal{L}(d^N(\mathbf{X}_{\mathbf{r}}), \mathbf{w}) | \mathbf{X}_{\mathbf{r}}]]. \quad (8)$$

This leads to the following optimization problem:

$$\min_{\mathbf{r}} \rho(\mathbf{r}), \text{ s.t. } r_i \geq 0 \text{ for } i = 1, \dots, k. \quad (9)$$

Special cases of the problem in Equation (9) have been solved elsewhere. Gupta and Miescke (1994) show that when $k = 2$, $c_1 = c_2$, and the budget is constrained, the optimal second-stage allocation minimizes the absolute difference of the posterior precision for the mean of each system, regardless of whether the 0-1 loss or opportunity cost is used. For the opportunity cost, $k \geq 2$, $c_1 = \dots = c_k = 1$, and known precision, Gupta and Miescke (1996) provide an optimal allocation assuming that only one observation can be made ($b = 1$). Repeatedly allocating a single replication at a time leads to a one-at-a-time sequential procedure that may not be optimal amongst all sequential procedures (Berger 1985).

3. NEW TWO-STAGE AND SEQUENTIAL PROCEDURES

The general formulation in §2 allows flexibility for allocating replications that is not allowed by the indifference-zone approach, but there is currently no known closed-form solution for the general problem in Equation (9). Further, numerical determination of the optimal allocation may be so computationally demanding that time might be better spent running more simulations rather than determining an optimal allocation.

This section derives good allocations of replications that are much simpler to compute than the optimal allocation. The allocations asymptotically minimize a bound on the expected loss. The bound is obtained by examining the $k - 1$ pairwise comparisons between the system with the highest first-stage sample mean and each other system. We use this idea to derive procedures that reduce the expected opportunity cost in §3.1, and procedures that improve the probability of correct selection in §3.2. Both subsections present two procedures. The first applies when there is no budget constraint on the number of replications, and the second applies when the budget is constrained.

Sequential extensions of the two-stage procedures are presented in §3.3.

3.1. Two-Stage Procedures: Opportunity Cost

Procedure \mathcal{LL} is a new two-stage procedure that reduces the expected opportunity cost loss when the budget is not constrained. The name \mathcal{LL} (for linear loss) is used rather than \mathcal{OC} (for opportunity cost) to avoid confusion with the OCBA procedure, a different approach explored by Chen (1996).

Procedure \mathcal{LL} is similar to Procedures \mathcal{R} and \mathcal{C} in that r_0 independent replications of each system are observed during a first stage, a second stage further distinguishes the performance of each system, and the system with the highest overall sample mean is selected as best. The novelty of Procedure \mathcal{LL} is that the number of replications is determined by balancing the cost of additional replications against the reduction in expected opportunity cost obtained from second-stage output. Procedure \mathcal{LL} is based on the approximation of Theorem 1 that assumes small replication costs, as occurs when the budget is essentially unconstrained.

PROCEDURE \mathcal{LL} , FOR OPPORTUNITY COST (LINEAR LOSS) AND UNCONSTRAINED BUDGET.

1. Specify the first-stage sample size r_0 .
2. Take independent replications $x_{i,1}, \dots, x_{i,r_0}$, for each system, $i = 1, \dots, k$.
3. Compute the first-stage sample mean $\bar{x}_i = \sum_{j=1}^{r_0} x_{i,j}/r_0$ and sample variance

$$\hat{\sigma}_i^2 = \frac{\sum_{j=1}^{r_0} (x_{i,j} - \bar{x}_i)^2}{r_0 - 1}.$$

4. Determine the first-stage order statistics, $\bar{x}_{[1]} \leq \dots \leq \bar{x}_{[k]}$. Compute the precision $\lambda_{i,k} = r_0/(\hat{\sigma}_{[k]}^2 + \hat{\sigma}_{[i]}^2)$ of the difference $W_{[i]} - W_{[k]}$ in unknown means, for $[i] \neq [k]$.
5. Estimate the cost per replication c_i for each system, based on the average run times during first stage sampling (express in same units as simulation output).
6. Compute the number of additional replications (round up if necessary),

$$r_{[i]} = \max \left\{ 0, \left\lceil \left(\left(\hat{\sigma}_{[i]}^2 (\lambda_{i,k})^{1/2} \frac{r_0 - 1 + \lambda_{i,k} (\bar{x}_{[k]} - \bar{x}_{[i]})^2}{(r_0 - 1) - 1} \right. \right. \right. \\ \left. \left. \left. \phi_{r_0-1}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] \right) / 2c_{[i]} \right)^{1/2} \right\rceil - r_0 \right\},$$

for $[i] \neq [k]$ and

$$r_{[k]} = \max \left\{ 0, \left\lceil \left(\left(\hat{\sigma}_{[k]}^2 \sum_{i=1}^{k-1} (\lambda_{i,k})^{1/2} \frac{r_0 - 1 + \lambda_{i,k} (\bar{x}_{[k]} - \bar{x}_{[i]})^2}{(r_0 - 1) - 1} \right. \right. \right. \\ \left. \left. \left. \phi_{r_0-1}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] \right) / 2c_{[k]} \right)^{1/2} \right\rceil - r_0 \right\}.$$

7. Take r_i additional observations from system i , independent of all other replications for all systems, for $i = 1, \dots, k$.

8. Compute the overall sample means $\bar{\bar{x}}_i = \sum_{j=1}^{r_0+r_i} x_{i,j}/(r_0 + r_i)$.
9. Select the system with the largest $\bar{\bar{x}}_i$ as best.

The computing requirements for Procedures \mathcal{R} , \mathcal{C} , and \mathcal{LL} are comparable. All compute the sample mean and variance of each system. Procedure \mathcal{LL} needs a few additional math operations per system, and the others require the solution of Rinott's integral.

The total number of replications of all three procedures is quite sensitive to subjectively specified parameters (c for Procedure \mathcal{LL} ; δ^* and P^* for Procedures \mathcal{R} and \mathcal{C}). The number of replications is readily controlled, however, with a budget constraint, $\mathbf{c}^T = b$. Procedure $\mathcal{LL}(\mathcal{B})$ is a budget-constrained version of Procedure \mathcal{LL} and is a consequence of Corollary 1.

PROCEDURE $\mathcal{LL}(\mathcal{B})$, FOR OPPORTUNITY COST (LINEAR LOSS) WITH BUDGET CONSTRAINT b .

1. Complete Steps 1–4 of Procedure \mathcal{LL} .
2. Estimate c_i with the average CPU time per replication for system i for $i = 1, \dots, k$; select b for the budget constraint $\mathbf{c}^T = b$ for second-stage replications; and initialize the set of systems considered for second-stage replications, $\mathcal{S} = \{1, \dots, k\}$.
3. Compute a tentative number of additional replications for each system $[i] \in \mathcal{S}$,

$$r_{[i]} = \frac{b + \sum_{j \in \mathcal{S}} r_0 c_j}{\sum_{j \in \mathcal{S}} \left(\frac{c_j c_{[i]} \hat{\sigma}_{[j]}^2 \eta_j}{\hat{\sigma}_{[i]}^2 \eta_{[i]}} \right)^{1/2}} - r_0,$$

where

$$\eta_{[i]} = \begin{cases} (\lambda_{i,k})^{1/2} \frac{r_0 - 1 + \lambda_{i,k} (\bar{x}_{[k]} - \bar{x}_{[i]})^2}{(r_0 - 1) - 1} & \text{for } [i] \neq [k] \\ \phi_{r_0-1}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] & \text{for } [i] = [k] \end{cases}$$

4. If all allocations r_i are nonnegative, then continue to Step 5. Otherwise, remedy the nonnegativity constraint violation: (a) For each $[i] \in \mathcal{S}$ such that $r_{[i]} < 0$, remove $[i]$ from \mathcal{S} and set $r_{[i]} = 0$, (b) For each $[i]$, if $[i] \in \mathcal{S}$ or $[k] \in \mathcal{S}$ then reassign

$$\lambda_{i,k} = \begin{cases} r_0/(\hat{\sigma}_{[k]}^2 + \hat{\sigma}_{[i]}^2) & \text{if } [i], [k] \in \mathcal{S} \\ r_0/\hat{\sigma}_{[k]}^2 & \text{if } [i] \notin \mathcal{S}; [k] \in \mathcal{S} \\ r_0/\hat{\sigma}_{[i]}^2 & \text{if } [i] \in \mathcal{S}; [k] \notin \mathcal{S}, \end{cases}$$

and (c) return to Step 3.

5. Round the r_i to an integer number of replications (see the note after the procedure).
6. Take r_i additional observations from system i , independent of all other replications for all systems, for $i = 1, \dots, k$.
7. Compute the overall sample means $\bar{\bar{x}}_i = \sum_{j=1}^{r_0+r_i} x_{i,j}/(r_0 + r_i)$.
8. Select the system with the largest $\bar{\bar{x}}_i$ as best.

The rounding in Step 5 may violate the constraint $\mathbf{c}\mathbf{r}^T = b$. Budget overruns are not likely to be a concern if the amount of rounding is small. If the budget is truly tight, a practical approach for overruns is to solve for \mathbf{r} again with a slightly decreased budget.

The theorem and corollary that justify these procedures rely on some additional notation. The state of uncertainty after the first stage (Equations (2) and (3)) is denoted by $\mu_i = \bar{x}_i$, $n_i = r_0$, $\alpha = (r_0 - 1)/2$, and $\beta_i = \hat{\sigma}_i^2(r_0 - 1)/2$. Define $[i]$ so that $\mu_{[1]} \leq \dots \leq \mu_{[k]}$ is nondecreasing. Ties occur here with probability 0. Denote by $\lambda_{i,j}$ the precision of the difference in unknown means $W_{[i]} - W_{[j]}$,

$$\lambda_{i,j} = \alpha \left(\frac{\beta_{[i]}}{n_{[i]}} + \frac{\beta_{[j]}}{n_{[j]}} \right)^{-1}. \quad (10)$$

Denote the expected value z_i of the unknown mean of system i , given \mathbf{x}_{r_i} , by

$$z_i \stackrel{\text{def}}{=} E[W_i | \mathbf{x}_{r_i}] = \bar{x}_i. \quad (11)$$

After the first stage but prior to observing \mathbf{x}_{r_i} , the unknown overall output mean

$$Z_i = E[W_i | \mathbf{X}_{r_i}] \quad (12)$$

has $\text{St}(\bar{x}_i, r_0(r_0 + r_i)/(r_i \hat{\sigma}_i^2), r_0 - 1)$ distribution (de Groot 1970). Denote by $\lambda_{[i,j]}$ the predictive precision for the difference in unknown posterior means $Z_{[i]} - Z_{[j]}$,

$$\lambda_{[i,j]} = \alpha \left(\frac{r_{[i]}\beta_{[i]}}{n_{[i]}(n_{[i]} + r_{[i]})} + \frac{r_{[j]}\beta_{[j]}}{n_{[j]}(n_{[j]} + r_{[j]})} \right)^{-1}. \quad (13)$$

Let $\phi_\nu(s)$ and $\Phi_\nu(s)$ be the density function and cumulative distribution function of the standard Student-t random variable with ν degrees of freedom, and let $\Psi_\nu(s) = \int_s^\infty (x-s)\phi_\nu(x)dx$. Bracken and Schleifer (1964) indicate that $\Psi_\nu(s) = \frac{\nu+s^2}{\nu-1}\phi_\nu(s) - s(1-\Phi_\nu(s))$. The theory below presumes that the allocations r_i are continuous.

THEOREM 1. Assume that the $x_{i,j}$ are jointly independent with $\mathcal{N}(w_i, \lambda_i^{-1})$ distribution. Denote by ζ the joint prior distribution of \mathbf{W}, Λ , with $\Lambda_i \sim \mathcal{G}(\alpha, \beta_i)$, $W_i | \lambda_i \sim \mathcal{N}(\mu_i, \lambda_i^{-1}/n_i)$ given λ_i , and (W_i, Λ_i) independent for $i = 1, \dots, k$. Let the loss function be the opportunity cost of Equation 6, and let $\mathbf{c}, [i]$, $\lambda_{[i,j]}$, and $\lambda_{i,j}$ be as above. Then:

• The objective function $\rho(\mathbf{r})$ of Equation 8 is bounded below by

$$\rho_{o.c.}^*(\mathbf{r}) = \mathbf{c}\mathbf{r}^T + E_\zeta[\max_j w_j - w_{[k]}] - \sum_{i=1}^{k-1} \lambda_{[i,k]}^{-1/2} \Psi_{2\alpha}[\lambda_{[i,k]}^{1/2}(\mu_{[i]} - \mu_{[k]})], \quad (14)$$

and the bound is tight when $k = 2$.

• If the c_i are sufficiently small for all i , the $r_{[i]}$ that minimize $\rho_{o.c.}^*(\mathbf{r})$ are approximately

$$\tilde{r}_{[i]}^* = \left(\frac{(\lambda_{i,k})^{1/2} \frac{2\alpha + \lambda_{i,k}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{i,k})^{1/2}(\mu_{[k]} - \mu_{[i]})]}{2c_{[i]}(\alpha/\beta_{[i]})} \right)^{1/2} - n_{[i]}, \quad (15)$$

for $[i] \neq [k]$, and

$$\tilde{r}_{[k]}^* = \left(\frac{\sum_{i=1}^{k-1} (\lambda_{i,k})^{1/2} \frac{2\alpha + \lambda_{i,k}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{i,k})^{1/2}(\mu_{[k]} - \mu_{[i]})]}{2c_{[k]}(\alpha/\beta_{[k]})} \right)^{1/2} - n_{[k]}.$$

PROOF. See the appendix.

When $k = 2$, $E_\zeta[\max_j w_j - w_{[k]}] = \lambda_{1,2}^{-1/2} \Psi_{2\alpha}[\lambda_{1,2}^{1/2}(\mu_{[1]} - \mu_{[2]})]$. Excluding replication costs, the expected opportunity cost of a pairwise comparison is therefore exactly

$$\lambda_{1,2}^{-1/2} \Psi_{2\alpha}[\lambda_{1,2}^{1/2}(\mu_{[1]} - \mu_{[2]})] - \lambda_{1,2}^{-1/2} \Psi_{2\alpha}[\lambda_{1,2}^{1/2}(\mu_{[1]} - \mu_{[2]})].$$

When $k > 2$, the simulation costs plus the expected losses from $k - 1$ pairwise comparisons between the system with the highest first-stage sample mean and each other system is then

$$\mathbf{c}\mathbf{r}^T + \sum_{i=1}^{k-1} \lambda_{i,k}^{-1/2} \Psi_{2\alpha}[\lambda_{i,k}^{1/2}(\mu_{[i]} - \mu_{[k]})] - \lambda_{i,k}^{-1/2} \Psi_{2\alpha}[\lambda_{i,k}^{1/2}(\mu_{[i]} - \mu_{[k]})]. \quad (16)$$

Because the Bonferroni inequality (e.g., see Law and Kelton 1991) can be considered to be a sum of pairwise losses (albeit for the 0-1 loss), the allocation that minimizes Equation (14) also minimizes the Bonferroni-like approximation of Equation (16).

The following corollary justifies the second-stage allocations for Procedure $\mathcal{LL}(\mathcal{B})$. This asymptotic result presumes that b is large relative to the c_i .

COROLLARY 1. The solution to $\min \rho_{o.c.}^*(\mathbf{r})$ subject to $\mathbf{c}\mathbf{r}^T = b$ for asymptotically large b is

$$\tilde{r}_{[i],b}^* = \frac{b + \sum_{j=1}^k c_j n_j}{\sum_{j=1}^k \left(\frac{c_j c_{[i]} \beta_j \eta_j}{\beta_{[i]} \eta_{[i]}} \right)^{1/2}} - n_{[i]}, \quad (17)$$

where

$$\eta_{[i]} = \begin{cases} (\lambda_{i,k})^{1/2} \frac{2\alpha + \lambda_{i,k}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{i,k})^{1/2}(\mu_{[k]} - \mu_{[i]})] & \text{for } [i] \neq [k] \\ \sum_{j=1}^{k-1} \eta_{[j],\alpha} & \text{for } [i] = [k] \end{cases} \quad (18)$$

PROOF. See the appendix.

When b is small, some of the allocations of Equation (17) may violate the nonnegativity constraint for replications. Suppose that one or more $\tilde{r}_{[i],b}^* < 0$, and let $\mathcal{S} = \{[i] | \tilde{r}_{[i],b}^* > 0\}$. To comply with the nonnegativity constraint, set $\tilde{r}_{[j],b}^* = 0$ for all $[j] \notin \mathcal{S}$. But this change violates the budget constraint, so $\sum_{[i] \in \mathcal{S}} \tilde{r}_{[i],b}^*$ must decrease. Further, the precision $\lambda_{i,k}$ of the difference $Z_{[i]} - Z_{[k]}$ is no longer well approximated by $\lambda_{i,k}$, but by

$$\lambda_{i,k} \approx \begin{cases} \lambda_{i,k} & \text{if } \tilde{r}_{[i],b}^*, \tilde{r}_{[k],b}^* \text{ are both large,} \\ \alpha n_{[k]}/\beta_{[k]} & \text{if } \tilde{r}_{[i],b}^* = 0 \text{ and } \tilde{r}_{[k],b}^* \text{ is large,} \\ \alpha n_{[i]}/\beta_{[i]} & \text{if } \tilde{r}_{[i],b}^* \text{ is large and } \tilde{r}_{[k],b}^* = 0. \end{cases} \quad (19)$$

We therefore replace $\lambda_{i,k}$ in Equation (18) with the appropriate approximation for $\lambda_{[i,k]}$ from Equation (19) when recalculating $\tilde{r}_{[i],b}^*$ for $[i] \in \mathcal{S}$, using $[i] \in \mathcal{S}$ as a criterion for indicating if $\tilde{r}_{[i],b}^*$ is large. An approximation for $\lambda_{i,k}$ is not needed when $[i], [k] \notin \mathcal{S}$.

Procedure $\mathcal{LL}(\mathcal{B})$ implements this process iteratively until the nonnegativity constraint is guaranteed. At most $k-1$ recalculations are needed, as $|\mathcal{S}|$ is decreased by at least one for each added iteration. The worst-case computational complexity is therefore $O(k^2)$ for small b . For sufficiently large b , the worst-case complexity is $O(k)$ because all systems have nonnegative allocations on the first pass.

3.2. Two-Stage Procedures: 0-1 Loss

The two-stage procedures presented in §3.1 reduce the expected opportunity cost of a selection. This section presents analogous procedures that reduce the expected 0-1 loss, or equivalently the probability of incorrect selection. Procedure 0-1, the 0-1 loss function analog of Procedure \mathcal{LL} , is identical to Procedure \mathcal{LL} , except that Step 6 becomes

6'. Compute the number of additional replications,

$$r_{[i]} = \max \left\{ 0, \left\lceil \left(\hat{\sigma}_{[i]}^2 (\lambda_{i,k})^{3/2} (\bar{x}_{[k]} - \bar{x}_{[i]}) \cdot \phi_{r_{0-1}}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] / 2c_{[i]} \right)^{1/2} \right\rceil - n_{[i]} \right\},$$

for $[i] \neq [k]$, and

$$r_{[k]} = \max \left\{ 0, \left\lceil \left(\sum_{i=1}^{k-1} \hat{\sigma}_{[k]}^2 (\lambda_{i,k})^{3/2} (\bar{x}_{[k]} - \bar{x}_{[i]}) \cdot \phi_{r_{0-1}}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] / 2c_{[k]} \right)^{1/2} \right\rceil - n_{[k]} \right\}.$$

Procedure 0-1 differs from Procedure \mathcal{R} mainly because an approximation to the expected posterior probability of correct selection, rather than the LFC, drives the allocation of second-stage replications. Theorem 2 justifies Procedure 0-1.

THEOREM 2. *Assume that the $x_{i,j}$ are jointly independent with $\mathcal{N}(w_i, \lambda_i^{-1})$ distribution. Denote by ζ the joint prior distribution of $\mathbf{W}, \mathbf{\Lambda}$, with $\Lambda_i \sim \mathcal{G}(\alpha, \beta_i)$, $W_i | \lambda_i \sim \mathcal{N}(\mu_i, \lambda_i^{-1}/n_i)$ given λ_i , and (W_i, Λ_i) independent for $i = 1, \dots, k$. Let the loss function be the 0-1 loss of Equation (5), and let $\mathbf{c}, [i], \lambda_{i,j}$, and $\lambda_{[i,j]}$ be as above. Then:*

• The objective function $\rho(\mathbf{r})$ of Equation (8) is bounded below by

$$\rho_{0-1}^* \mathbf{r} = \mathbf{c} \mathbf{r}^T + E_{\zeta} \left[\max_j w_j - w_{[k]} \right] - \sum_{i=1}^{k-1} \Phi_{2\alpha} \left[\lambda_{[i,k]}^{1/2} (\mu_{[i]} - \mu_{[k]}) \right]. \quad (20)$$

• For sufficiently small c_i for all i , the $r_{[i]}$ that minimize $\rho_{0-1}^*(\mathbf{r})$ are approximately

$$\tilde{r}_{[i]}^* = \left(\frac{(\lambda_{i,k})^{3/2} (\mu_{[k]} - \mu_{[i]}) \phi_{2\alpha}[(\lambda_{i,k})^{1/2} (\mu_{[k]} - \mu_{[i]})]}{2c_{[i]} (\alpha/\beta_{[i]})} \right)^{1/2} - n_{[i]} \quad (21)$$

for $[i] \neq [k]$, and

$$\tilde{r}_{[k]}^* = \left(\frac{\sum_{i=1}^{k-1} (\lambda_{i,k})^{3/2} (\mu_{[k]} - \mu_{[i]}) \phi_{2\alpha}[(\lambda_{i,k})^{1/2} (\mu_{[k]} - \mu_{[i]})]}{2c_{[k]} (\alpha/\beta_{[k]})} \right)^{1/2} - n_{[k]}. \quad (22)$$

PROOF. See the appendix.

Unlike the bound in Equation (14), the bound in Equation (20) is not tight for $k=2$ because of an approximation in the proof. The bound becomes tighter as the r_i increase.

Selecting the system with the highest sample mean does not necessarily maximize the probability of correct selection (e.g., see Gupta and Miescke 1994). However, this selection rule has intuitive appeal and is used implicitly by indifference-zone procedures, and we conjecture that our procedures tend to avoid situations where the rule is suboptimal.

Procedure 0-1(\mathcal{B}), a budget-constrained version of Procedure 0-1, is obtained by modifying Step 3 of Procedure $\mathcal{LL}(\mathcal{B})$ to account for the 0-1 loss function.

3'. Compute a tentative number of additional replications for each system $[i] \in \mathcal{S}$,

$$r_{[i]} = \frac{b + \sum_{j \in \mathcal{S}} r_{0-1} c_j}{\sum_{j \in \mathcal{S}} \left(\frac{c_j \gamma_{[i]} \hat{\sigma}_{[i]}^2 \gamma_j}{\sigma_{[i]}^2 \gamma_{[i]}} \right)^{1/2}} - r_{0-1},$$

where

$$\gamma_{[i]} = \begin{cases} (\lambda_{i,k})^{3/2} (\bar{x}_{[k]} - \bar{x}_{[i]}) & \text{for } [i] \neq [k] \\ \phi_{r_{0-1}}[(\lambda_{i,k})^{1/2} (\bar{x}_{[k]} - \bar{x}_{[i]})] & \text{for } [i] = [k] \end{cases}$$

This allocation is a consequence of Corollary 2.

COROLLARY 2. *The solution to $\min \hat{\rho}_{0-1}^T(\mathbf{r}) = \rho_{0-1}(\mathbf{r}) + \mathbf{c} \mathbf{r}$ subject to $\mathbf{c} \mathbf{r}^T = b$ for asymptotically large b is*

$$\tilde{r}_{[i],b}^* = \frac{b + \sum_{j=1}^k c_j n_j}{\sum_{j=1}^k \left(\frac{c_j \gamma_{[i]} \beta_j \gamma_j}{\beta_{[i]} \gamma_{[i]}} \right)^{1/2}} - n_{[i]},$$

where

$$\gamma_{[i]} = \begin{cases} (\lambda_{i,k})^{3/2} (\mu_{[k]} - \mu_{[i]}) & \text{for } [i] \neq [k] \\ \times \phi_{2\alpha}[(\lambda_{i,k})^{1/2} (\mu_{[k]} - \mu_{[i]})] & \text{for } [i] = [k] \end{cases}$$

PROOF. The proof of this parallels that for Corollary 1, except that $\gamma_{[i]}$ replaces $\eta_{[i]}$. \square

3.3. Sequential Procedures

Chen (1996) and Bechhofer et al. (1995) describe sequential procedures that improve sampling efficiency based on different assumptions than the decision theoretic approach here. Sequential procedures allow a modeler to run a certain number of replications, observe the output, and repeat until a suitable stopping condition is achieved. A potential advantage of sequential procedures is that sampling efficiency may be improved at each stage by incorporating information from all earlier stages. A second advantage is that the stopping rule might be used to provide a Bayesian PCS guarantee. Additional stages can be run until a pre-specified posterior PCS can be claimed.

To develop sequential variations on our new procedures, a technical issue requires attention. Because the number of replications $r_{0,i}$ seen so far for each system may differ, the shape parameter $\alpha_i = (r_{0,i} - 1)/2$ for the unknown precision may differ as well. The results in §§3.1 and 3.2 require a common shape parameter α . The problem with different α_i is that the differences $W_{[i]} - W_{[j]}$ and $Z_{[i]} - Z_{[j]}$ do not have three parameter Student-t distributions because of a mismatch in the degrees of freedom.

One approach to this problem is to approximate the distributions of these differences. A sequential procedure, Procedure $\mathcal{LL}(\mathcal{S})$, is obtained with the following modifications to Procedure $\mathcal{LL}(\mathcal{B})$. The changes incorporate the standard Welch approximation for the difference of two Student-t random variables with different degrees of freedom (e.g., see Law and Kelton 1991).

- Allow for Steps 2–7 to be repeated until a user-specified condition is met (e.g., allocate τ at each stage until a total budget of $b \geq \tau$ is exhausted).

- Calculate $\mu_i = \bar{x}_i$, $\alpha_i = (r_{0,i} - 1)/2$, $\beta_i = \hat{\sigma}_i^2(r_{0,i} - 1)/2$, $n_i = r_{0,i}$, based on all $r_{0,i}$ replications seen so far.

- In Step 4, the precision parameter $\lambda_{i,k}$ becomes

$$\lambda_{i,k} = \begin{cases} (\hat{\sigma}_{[i]}^2/r_{0,[i]} + \hat{\sigma}_{[k]}^2/r_{0,[k]})^{-1} & \text{if } [i], [k] \in \mathcal{S}, \\ r_{0,[k]}/\hat{\sigma}_{[k]}^2 & \text{if } [i] \notin \mathcal{S}; [k] \in \mathcal{S}, \\ r_{0,[i]}/\hat{\sigma}_{[i]}^2 & \text{if } [i] \in \mathcal{S}; [k] \notin \mathcal{S}. \end{cases} \quad (23)$$

- Replace all occurrences of $r_0 - 1$ with the Welch approximation for the degrees of freedom, $\nu_{[i],[k]}$, where

$$\nu_{[i],[k]} = \begin{cases} \frac{(\hat{\sigma}_{[i]}^2/r_{0,[i]} + \hat{\sigma}_{[k]}^2/r_{0,[k]})^2}{\left(\frac{\hat{\sigma}_{[i]}^2}{r_{0,[i]}}\right)^2/(r_{0,[i]} - 1) + \left(\frac{\hat{\sigma}_{[k]}^2}{r_{0,[k]}}\right)^2/(r_{0,[k]} - 1)} & \text{if } [i], [k] \in \mathcal{S}, \\ r_{0,[k]} - 1 & \text{if } [i] \notin \mathcal{S}; [k] \in \mathcal{S}, \\ r_{0,[i]} - 1 & \text{if } [i] \in \mathcal{S}; [k] \notin \mathcal{S}. \end{cases} \quad (24)$$

- Replace all other occurrences of r_0 with $r_{0,[i]}$.

Similar changes to Procedure 0-1(\mathcal{B}) result in the sequential Procedure 0-1(\mathcal{S}).

4. AN EMPIRICAL COMPARISON OF THE PROCEDURES

This section presents some results from a small empirical study that assess the effectiveness of the procedures

described above to identify the best system. Each procedure's effectiveness is evaluated with respect to four performance measures described in §4.1. The selection problems in §4.2 are artificially designed to provide some indication of the ability of Procedures \mathcal{C} , $\mathcal{LL}(\mathcal{B})$ and 0-1(\mathcal{B}) to allocate replications to competitive systems and avoid simulating non-competitive systems. Section 4.3 uses a less artificial selection problem to analyze the performance of Procedures \mathcal{R} , \mathcal{C} , $\mathcal{LL}(\mathcal{B})$ and 0-1(\mathcal{B}), as a function of the total second-stage sampling. We use $c_i = 1$ to constrain the total number of replications for our new procedures. A thorough exploration of the potential benefits of allowing differing costs is beyond the scope of this paper.

4.1. Performance Measures for the Procedures

The effectiveness of each procedure is measured with respect to four performance measures, each of which is relevant to the justification of at least one of the procedures.

The first performance measure to evaluate the procedures is the PCS, estimated as the empirical fraction of correct selections (EFCS) of the best system. The derivations of Procedures \mathcal{C} and \mathcal{R} both provide PCS guarantees.

The second performance measure to evaluate the procedures is the empirical fraction of selections that are within δ^* of the best system. Nelson and Matejcek (1995) show that many indifference-zone procedures, including Procedure \mathcal{R} , select a system within δ^* of the best for *all* configurations of the means.

The third performance measure is the expected value of the Bonferroni bound on the probability of correct selection (EBPCS), given the output obtained from an application of a selection procedure. By design, Procedure 0-1(\mathcal{B}) allocates replications to asymptotically maximize EBPCS. The Bonferroni bound for the probability of correct selection (BPCS) is estimated at the end of the procedure, when the system $d^N(\mathbf{x}_r)$ with the largest overall sample mean $\bar{x}_{d^N(\mathbf{x}_r)}$ is selected. Let $\tilde{\lambda}_{i,d^N(\mathbf{x}_r)}$ and $\nu_{[i],d^N(\mathbf{x}_r)}$ be the approximate precision and degrees of freedom for $W_{[i]} - W_{d^N(\mathbf{x}_r)}$ using the Welch approximation, based on *all* output. The BPCS is then

$$\max \left\{ 0, 1 - \sum_{[i] \neq d^N(\mathbf{x}_r)} \Phi_{\nu_{[i],d^N(\mathbf{x}_r)}} \left[\tilde{\lambda}_{i,d^N(\mathbf{x}_r)}^{1/2} (\bar{x}_{[i]} - \bar{x}_{d^N(\mathbf{x}_r)}) \right] \right\} \quad (25)$$

(Reverse signs, $\bar{x}_{d^N(\mathbf{x}_r)} - \bar{x}_{[i]}$, for a minimization problem.) The Welch approximation is used even when the degrees of freedom match for numerical stability. Inoue and Chick (1998) further discuss the relationship between P-values and the BPCS.

The fourth performance measure is the expected value of the Bonferroni-like bound on opportunity cost (EBOC), the opportunity cost analog of Equation (25),

$$\sum_{[i] \neq d^N(\mathbf{x}_r)} \tilde{\lambda}_{i,d^N(\mathbf{x}_r)}^{-1/2} \Psi_{\nu_{[i],d^N(\mathbf{x}_r)}} \left[\tilde{\lambda}_{i,d^N(\mathbf{x}_r)}^{-1/2} (\bar{x}_{[i]} - \bar{x}_{d^N(\mathbf{x}_r)}) \right].$$

By design, Procedure $\mathcal{LL}(\mathcal{B})$ allocates replications to asymptotically minimize EBOC.

4.2. A Stylized Selection Problem

Stylized selection problems can provide insight into the performance of selection procedures in controlled environments. Here we use the monotone decreasing means (MDM) configuration, which evenly spaces the means of each system,

$$w_i = w_1 - (i-1)\zeta\delta^*, \quad \text{for } i = 2, \dots, k,$$

for some ζ . The extensive study of Nelson et al. (1999) indicates that Procedures \mathcal{C} requires a smaller number of replications than Procedures \mathcal{R} in the MDM configuration when the number of systems is large ($k \geq 10$ in one experiment). Because Procedures \mathcal{R} and \mathcal{C} are already compared elsewhere, this section focuses on the ability of Procedures \mathcal{C} , $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$ to effectively allocate replications to identify the best system, as a function of the number of systems ($k = 2, 5, 10, 100$).

The procedures are compared by running a common first stage. The total second-stage allocation of each procedure is made the same by letting the total number of second-stage replications b for Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$ be the same as the total second-stage allocation of Procedure \mathcal{C} . Independent replications are then run for the second stage of each procedure, and the system selected to be best by each procedure is recorded. This process tests whether the new procedures can more effectively allocate the number of replications suggested by Procedure \mathcal{C} . Three thousand macroreplications are run so that the error for estimating a probability of correct selection of 0.95 gives two-decimal-place accuracy ($\sqrt{0.95 \times (1-0.95)/3000} \approx 0.004$).

We initially set $\zeta = 1/2$, $r_0 = 10$, $P^* = 0.95$, $\delta^* = 1/\sqrt{r_0}$ and $\sigma_i^2 = 4$ for $i = 1, \dots, k$. System 1 therefore performs best, and System 2 performs within δ^* of the best. Table 1 presents the average number of second-stage replications

(ANR) per system for each procedure, as well as the performance measures.

Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$ perform at least as well as Procedure \mathcal{C} for all 4 performance measures, for each of $k = 2, 5, 10$, and 100. More than half the comparisons indicate a statistically significant difference (95% confidence with a paired t test), but several comparisons are within error bounds. Procedure $\mathcal{LL}(\mathcal{B})$ also performs at least as well as Procedure $0-1(\mathcal{B})$, or within error bounds, for all performance measures and values of k tested. Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$ outperform Procedure \mathcal{C} for EBPCS and EBOC for each k tested, and outperform Procedure \mathcal{C} for the empirical fraction of correct selections (EFCS) of the best system for larger numbers of systems ($k = 10, 100$). The use of additional first-stage information therefore improves the second-stage allocations in this experiment. Procedures $\mathcal{LL}(\mathcal{B})$, $0-1(\mathcal{B})$, and \mathcal{C} perform more similarly on this experiment when measured by the empirical fraction of selections within δ^* of the best. For this MDM experiment, then, the new Bayesian procedures more frequently identify the true best system, particularly when the number of systems is large, but when Procedure \mathcal{C} does not identify the best system, it often identifies a very good system.

Somewhat surprisingly, Procedure $0-1(\mathcal{B})$ does not perform best with respect to EBPCS, a measure that the procedure was designed to improve. This is because the derivation of Procedure $0-1(\mathcal{B})$ requires one more asymptotic approximation that the derivation of Procedure $\mathcal{LL}(\mathcal{B})$. Procedure $\mathcal{LL}(\mathcal{B})$ is designed to improve EBOC, and the procedure performs best with respect to this measure. Further, its performance improves relative to the other procedures as the number of systems increases. In this experiment, then, the Bonferroni-like bound and asymptotic

Table 1. Performance of Procedures \mathcal{C} , $0-1(\mathcal{B})$, and $\mathcal{LL}(\mathcal{B})$ for the monotone decreasing means (MDM) experiment of §4.2.

Figure of Merit	Procedure	Number of Systems, k			
		2	5	10	100
ANR	Procedures \mathcal{C} , $0-1(\mathcal{B})$, $\mathcal{LL}(\mathcal{B})$	738	3,429	8,784	42,862
Empirical fraction of correct selections (EFCS)	Procedure \mathcal{C}	0.8363	0.9140	0.9323	0.9763
	Procedure $0-1(\mathcal{B})$	0.8527*	0.9117	0.9480*	0.9937*
	Procedure $\mathcal{LL}(\mathcal{B})$	0.8500	0.9293*	0.9660*	0.9987*
Empirical fraction of selections within δ^* of the best	Procedure \mathcal{C}	1.0000	0.9943	0.9990	1.0000
	Procedure $0-1(\mathcal{B})$	1.0000	0.9953	0.9977	0.9997
	Procedure $\mathcal{LL}(\mathcal{B})$	1.0000	0.9953	0.9993	1.0000
Expected Bonferroni bound on PCS (EBPCS)	Procedure \mathcal{C}	0.8336	0.8379	0.8649	0.9318
	Procedure $0-1(\mathcal{B})$	0.8446*	0.8339	0.8821*	0.9717*
	Procedure $\mathcal{LL}(\mathcal{B})$	0.8470*	0.8462*	0.9022*	0.9842*
Expected Bonferroni bound on opportunity cost (EBOC)	Procedure \mathcal{C}	0.0176	0.0138	0.0104	0.0037
	Procedure $0-1(\mathcal{B})$	0.0157*	0.0128*	0.0075*	0.0012*
	Procedure $\mathcal{LL}(\mathcal{B})$	0.0154*	0.0110*	0.0055*	0.0005*

The symbol * indicates a statistically significant difference in performance when compared with Procedure \mathcal{C} .

approximation used in the derivation of Procedure $\mathcal{LL}(\mathcal{B})$ do not appear to deleteriously affect its performance.

We repeated the experiment with small modifications to the parameters ($P^* = 0.90$ and $\sigma_i^2 = 1$ for $i = 1, \dots, k$) and obtained the same qualitative results (data not shown). This study did not investigate the relative performance of the procedures when the variance for each system differs.

4.3. An Inventory Policy Selection Problem

We also consider the performance of the selection procedures with respect to a problem that has less structure to the configuration of the means and variances of each system. The systems considered are the five (s, S) inventory policies introduced by Koenig and Law (1985) and analyzed later by Nelson and Matejcek (1995). The best system is the policy that has the minimum expected cost per period, evaluated over 60 periods.


In this experiment we use a somewhat different comparison mechanism to evaluate the performance of the procedures as a function of the total second-stage budget. We constrain the total number of replications after the first stage, rather than letting Procedure \mathcal{C} determine the second-stage allocations for Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$. For Procedures \mathcal{R} and \mathcal{C} , this means changing the proportionality constant for the number of second-stage replications for each system to obtain a specified total number of replications. Let Procedures $\mathcal{R}(\mathcal{B})$ and $\mathcal{C}(\mathcal{B})$ denote budget constrained version of these indifference-zone procedures. Although no PCS guarantee is claimed, replications are still allocated proportional to the first-stage sample variance.

Suppose that for a *specific* first-stage, $r_0 = 5$, $\bar{x} = (122.9, 121.4, 126.0, 132.4, 145.9)$ and the sample variances are $(13.7, 4.0, 11.4, 6.6, 4.9)$. If the second-stage budget is constrained to $b = 30$, Procedure $\mathcal{R}(\mathcal{B})$ allocates $\mathbf{r} = (14, 0, 10, 4, 2)$; Procedure $\mathcal{C}(\mathcal{B})$ screens out Systems 4 and 5 ($\delta^* = 1$) and allocates $\mathbf{r} = (16, 1, 13, 0, 0)$; Procedure $0-1(\mathcal{B})$ allocates $\mathbf{r} = (15, 8, 7, 0, 0)$; and Procedure $\mathcal{LL}(\mathcal{B})$ allocates $\mathbf{r} = (16, 8, 6, 0, 0)$. The new procedures allocate $r_4 = r_5 = 0$, because more replications for Systems 4 and 5 are unlikely to appreciably change the posterior probability that either is best. They *might* be selected, however, if the additional replications for the other systems indicate that Systems 4 or 5 have the best overall sample mean after both stages are completed. On the other hand, once Procedure \mathcal{C} screens out a system, that system may no longer be selected as best.

We estimate the *average* performance of each two-stage procedure by running 3000 independent macroreplications of first ($r_0 = 5$) and second stages. Each procedure is then compared with respect to EPCS, EBPCS and EBOC, as a function of the total number of second-stage replications.

We tested several variations of Procedure $\mathcal{C}(\mathcal{B})$, corresponding to different levels of δ^* on the interval $[0, 1]$. Once the total number of second-stage replications is fixed for $\mathcal{C}(\mathcal{B})$, the effect of δ^* is to change the number of systems screened out. A smaller δ^* corresponds to a smaller

Figure 1. Empirical fraction of correct selections (EFCS) as a function of the total number b of second-stage replications.



number of screened systems. Although Procedure \mathcal{C} cannot be run with $\delta^* = 0$ because a division by δ^* is required, Procedure $\mathcal{C}(\mathcal{B})$ can be run with $\delta^* = 0$, because second-stage replications are allocated proportional to the sample variance of unscreened systems so that a budget constraint is satisfied. Nelson et al. (1999) remark that when $\delta^* = 0$, the screening in Procedure $\mathcal{C}(\mathcal{B})$ is a generalization of the subset selection procedure of Gupta (1965). Screening with $\delta^* = 0$ results in a better empirical performance than with $\delta^* = 1$ for this experiment, so we present results for $\delta^* = 0$.

The EFCS is displayed as a function of the number of additional replications in Figure 1. (System 2 is considered “best” based on many thousands of replications). Procedure $\mathcal{C}(\mathcal{B})$ outperforms Procedure $\mathcal{R}(\mathcal{B})$ for EFCS because it does not allocate replications to screened systems that are apparently inferior. Both are outperformed by Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$. To reach the same EFCS obtained from $b = 200$ replications for Procedure $\mathcal{R}(\mathcal{B})$, Procedure $\mathcal{LL}(\mathcal{B})$ requires 110 replications and Procedure $\mathcal{C}(\mathcal{B})$ requires 160 replications. The better performance of the new procedures can be attributed to their inclusion of additional first-stage information when allocating second-stage replications.

Figure 2 indicates that the procedures have approximately the same relative ranking for EBPCS as observed above for EFCS. The one difference is that Procedure $\mathcal{C}(\mathcal{B})$ outperforms Procedure $0-1(\mathcal{B})$, when the number of second-stage replications is small. The asymptotic approximations in the derivation of Procedure $0-1(\mathcal{B})$ apparently therefore have a negative effect when the total budget is particularly small.

Procedure $\mathcal{R}(\mathcal{B})$ achieves an EBPCS of 0.9714 with $b = 200$ replications during the second stage. Procedures $\mathcal{LL}(\mathcal{B})$ and $0-1(\mathcal{B})$, on the other hand, require approximately $b = 125$ and $b = 145$, respectively, to obtain the same EBPCS. This means that the new procedures tend

Figure 2. Expected Bonferroni bound for the probability of correct selection (EBPCS) as a function of the total number b of second-stage replications.





Figure 3. Expected Bonferroni bound for opportunity cost (EBOC) as a function of the total number b of second-stage replications.



to provide more evidence for correct selection, on the basis of the simulation output observed at the completion of the procedure, than the indifference-zone procedures.

Figure 3 indicates that Procedures 0-1(\mathcal{B}) and $\mathcal{C}(\mathcal{B})$ perform similarly with respect to EBOC. Procedure $\mathcal{LL}(\mathcal{B})$ is designed to reduce EBOC, so there is no surprise that it performs best. Procedure $\mathcal{R}(\mathcal{B})$ performs worst because it ignores the most first-stage information.

Changing the systems slightly by doubling the number of simulated months per replication does not modify the qualitative nature of the results displayed in the figures.

Sequential allocation provides a distinct advantage for this problem. If Procedure 0-1(\mathcal{S}) is used with $\tau = 5$ replications allocated per stage until a total of $b = 100$ replica-

tions are observed, the EBPCS is 0.9748, an increase from 0.9486 from the two-stage Procedure 0-1(\mathcal{B}). Sequential allocation also improves EFCS and EBOC.

The CPU time required to compute the second-stage allocation is not included in the sampling budget constraint, but may influence the computational efficiency of the procedures if the cost of computing the allocation is nontrivial relative to the sampling budget b . Here the average CPU time per allocation computation for Procedure $\mathcal{R}(\mathcal{B})$ is 0.0020 second, compared with 0.018 second for Procedure 0-1(\mathcal{S}). Procedure $\mathcal{R}(\mathcal{B})$ ($b = 100$) results in an EBPCS of 0.925, and requires 15.1 seconds on average (CPU time for allocation plus 100 replications). Procedure 0-1(\mathcal{S}) ($b = 50$ replications allocated $\tau = 5$ at a time) results in an EBPCS of 0.926, and requires 9.39 seconds on average. Procedure 0-1(\mathcal{S}) therefore needs $9.39/15.1 \approx 0.62$ as much effort as Procedure $\mathcal{R}(\mathcal{B})$ to provide the same EBPCS. The allocations of the new procedures require more time than for the indifference-zone procedures. The penalty for the increased computation time for the new procedures is more than offset by the improved efficiency for identifying the best system in this example. In general, the benefit increases as simulation run times become longer.

5. DISCUSSION

Many two-stage indifference-zone procedures ignore a fair amount of first-stage sampling information. Because Procedure \mathcal{C} can screen after the first stage, it incorporates much more first-stage information than Procedure \mathcal{R} . Procedure \mathcal{C} therefore outperforms Procedure \mathcal{R} when several systems are screened out (Nelson et al. 1999).

Procedures $\mathcal{LL}(\mathcal{B})$ and 0-1(\mathcal{B}) both use more first stage information than Procedures \mathcal{R} and \mathcal{C} . The new procedures are justified by (i) deriving a Bonferroni-like approximation for the total expected loss, (ii) determining an allocation that asymptotically minimizes that approximation as the cost of replications gets arbitrarily small, and (iii) establishing budget-constrained allocations that are asymptotically optimal as the budget gets large. The asymptotic and Bonferroni-like approximations cause the procedures to be suboptimal. However, the improved use of first-stage information when allocating additional replications seems to outweigh the deleterious effect of the approximations, even when the number of systems is $k = 100$.

Procedure $\mathcal{LL}(\mathcal{B})$ seems to perform best among the four procedures considered here, for both the MDM experiments and the inventory selection problem, with respect to four measures of effectiveness. Additional experiments (data not shown) indicate that modifying r_0 does not change this ranking. Procedure 0-1(\mathcal{B}) also shows significant improvements over Procedure \mathcal{R} , but the performance improvement is less impressive than for Procedure $\mathcal{LL}(\mathcal{B})$. The extra asymptotic approximation in the derivation of Procedure 0-1(\mathcal{B}) is the likely culprit for the performance degradation. An open question is whether other bounds for the expected loss result in more efficient procedures.

In this paper we did not consider common random numbers, an important efficiency improvement technique for distinguishing the performance of systems. This presents an avenue for further research.

6. CONCLUSIONS

The new procedures presented here provide flexibility that was previously unavailable to a decision-maker that uses simulation to select the best system. Whereas indifference-zone procedures are concerned with the probability of correct selection, we allow for a decision-maker to improve the probability of correct selection or the expected opportunity cost associated with a potentially incorrect decision. A budget constraint on the number of replications arises naturally in practice, and is readily incorporated into the decision-theoretic framework. Although our budget-constrained versions of the indifference-zone procedures are not great leaps, budget constraints are not found in standard treatments of indifference-zone procedures.

Our central tenet is that computer replications should be allocated so that the expected value of information from the experiments is maximized. To accomplish this, our procedures incorporate more first-stage information about sample means than indifference-zone procedures. Although the procedures are based on suboptimal asymptotic approximations, the deleterious effect of the approximations seems to be outweighed by the use of additional first-stage information. Empirical results indicate that Procedure $\mathcal{LL}(\mathcal{B})$ seems to be a particularly effective two-stage procedure. The performance of these new procedures for a broader range of selection problems warrants further investigation.

APPENDIX

PROOF OF THEOREM 1 (OPPORTUNITY COST). We first determine the expected total loss $\rho(\mathbf{r})$. Consider the modified loss function:

$$\mathcal{L}_{o.c.}^*([i], \mathbf{w}) = \mathcal{L}_{o.c.}([i], \mathbf{w}) - \mathcal{L}_{o.c.}([k], \mathbf{w}) = w_{[k]} - w_{[i]}.$$

It is well known that adding $-\mathcal{L}_{o.c.}([k], \mathbf{w})$ to the loss function does not change the optimal decision (e.g., see de Groot 1970). Thus $E[\mathcal{L}_{o.c.}^*([i], \mathbf{w}) | \mathbf{x}] = z_{[k]} - z_{[i]}$, the difference of the posterior means. Further, $[i]$ is selected only if the event $\mathcal{A}_{[i]} = \{\mathbf{z}: z_{[i]} = \max_j z_j\}$ occurs. Take the expectation over second-stage outcomes \mathbf{X}_r , add in $E_{\mathcal{L}}[\mathcal{L}_{o.c.}([k], \mathbf{w})]$, to compensate for subtracting $-\mathcal{L}_{o.c.}([k], \mathbf{w})$ earlier, and add the cost $\mathbf{c}\mathbf{r}^T$ of the experiment to obtain $\rho(\mathbf{r}) = \mathbf{c}\mathbf{r}^T + E_{\mathcal{L}}[\max_j w_j - w_{[k]}] - \sum_{i=1}^{k-1} p_{\mathbf{z}}(\mathcal{A}_{[i]})E[Z_{[i]} - Z_{[k]} | \mathcal{A}_{[i]}]$.

Now consider the loss for a pairwise comparison between systems $[i]$ and $[k]$. Then

$$E[\mathcal{L}_{o.c.}^*([i], \mathbf{w}) | \mathbf{x}] \leq \begin{cases} z_{[k]} - z_{[i]} & \text{when } z_{[k]} < z_{[i]}, \\ 0 & \text{otherwise,} \end{cases} \quad (26)$$

for each \mathbf{x} , $[i] \neq [k]$. Because $Z_{[k]} - Z_{[i]}$ has a Student-t distribution (Equation (12)), a result of Bracken and Schleifer

(1964) implies that the expected value of this pairwise loss is

$$\lambda_{\{i,k\}}^{-1/2} \Psi_{2\alpha}[\lambda_{\{i,k\}}^{1/2}(\mu_{[i]} - \mu_{[k]})].$$

Define $\mathcal{B}_i = \{\mathbf{z} | z_i \geq z_{[k]}\}$ so that $\mathcal{A}_i \subseteq \mathcal{B}_i$, and let $\chi_{\mathcal{A}_i}(\cdot)$ be the indicator function. The lower bound Equation (14) is obtained by noting:

$$\begin{aligned} \sum_{i=1}^{k-1} p_{\mathbf{z}}(\mathcal{A}_{[i]})E[Z_{[i]} - Z_{[k]} | \mathcal{A}_{[i]}] &= \sum_{i=1}^{k-1} E[\chi_{\mathcal{A}_i}(\mathbf{Z})(Z_{[i]} - Z_{[k]})] \\ &\leq \sum_{i=1}^{k-1} E[\chi_{\mathcal{B}_i}(\mathbf{Z})(Z_{[i]} - Z_{[k]})] \\ &= \sum_{i=1}^{k-1} p_{\mathbf{z}}(\mathcal{B}_{[i]})E[Z_{[i]} - Z_{[k]} | \mathcal{B}_{[i]}] \\ &= \sum_{i=1}^{k-1} \tau_{\{i,k\}}^{-1/2} \Psi[\tau_{\{i,k\}}^{1/2}(\mu_{[i]} - \mu_{[k]})]. \end{aligned}$$

To minimize $\rho_{o.c.}^*(\mathbf{r})$, consider r_i to be continuous and take partial derivatives. Noting that $\partial \Psi_{\nu}(s)/\partial s = \Phi_{\nu}(s) - 1$ leads to optimality conditions

$$c_{[i]} = \frac{(\lambda_{\{i,k\}})^{1/2} \frac{2\alpha + \lambda_{\{i,k\}}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{\{i,k\}})^{1/2}(\mu_{[i]} - \mu_{[k]})]}{2(n_{[i]} + r_{[i]})^2(\alpha/\beta_{[i]})},$$

for $[i] \neq [k]$, and

$$c_{[k]} = \sum_{i=1}^{k-1} \frac{(\lambda_{\{i,k\}})^{1/2} \frac{2\alpha + \lambda_{\{i,k\}}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{\{i,k\}})^{1/2}(\mu_{[i]} - \mu_{[k]})]}{2(n_{[k]} + r_{[k]})^2(\alpha/\beta_{[k]})}.$$

In the limit $c_i \rightarrow 0$ for all i , we have large r_i , and $\lambda_{\{i,k\}} \rightarrow \lambda_{i,k}$. Substitute $\lambda_{i,k}$ for $\lambda_{\{i,k\}}$ in the optimality conditions to obtain the stated number of replications. \square

PROOF OF COROLLARY 1. Let θ be a Lagrange multiplier for the constraint $\mathbf{c}\mathbf{r}^T = b$. Take derivatives to obtain optimality conditions

$$\theta c_{[i]} + \frac{(\lambda_{\{i,k\}})^{1/2} \frac{2\alpha + \lambda_{\{i,k\}}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{\{i,k\}})^{1/2}(\mu_{[i]} - \mu_{[k]})]}{2(n_{[i]} + r_{[i]})^2(\alpha/\beta_{[i]})} = 0,$$

for $[i] \neq [k]$, and

$$\theta c_{[k]} + \sum_{i=1}^{k-1} \frac{(\lambda_{\{i,k\}})^{1/2} \frac{2\alpha + \lambda_{\{i,k\}}(\mu_{[k]} - \mu_{[i]})^2}{2\alpha - 1} \phi_{2\alpha}[(\lambda_{\{i,k\}})^{1/2}(\mu_{[i]} - \mu_{[k]})]}{2(n_{[k]} + r_{[k]})^2(\alpha/\beta_{[k]})} = 0.$$

When b is large, the number of replications of each system is large, so that $\lambda_{\{i,k\}} \rightarrow \lambda_{i,k}$. Let $\eta_{[k]}$ be as in Equation (18), substitute the limiting value $\lambda_{i,k}$ for $\lambda_{\{i,k\}}$, and replace r_i with \tilde{r}_i^* to indicate that an asymptotic approximation is being made, to obtain

$$\theta c_{[i]} + \frac{\eta_{[i]}}{2(\eta_{[i]} + \tilde{r}_{[i]}^*)^2(\alpha/\beta_{[i]})} = 0.$$

As a consequence, $\tilde{r}_{[j]}^*$ is therefore related to $\tilde{r}_{[i]}^*$ by

$$\frac{(\eta_{[j]} + \tilde{r}_{[j]}^*)^2}{(\eta_{[i]} + \tilde{r}_{[i]}^*)^2} = \frac{\beta_{[j]}\eta_{[j]}/c_{[j]}}{\beta_{[i]}\eta_{[i]}/c_{[i]}}.$$

Recall that $\mathbf{c}\mathbf{r}^T = b$ and solve for $\tilde{r}_{[i]}^*$ to obtain the desired allocation. \square

PROOF OF THEOREM 2 (0-1 loss). The proof parallels that of Theorem 1, except that an extra approximation is made for the expected loss. Consider the *pairwise* loss between systems $[i]$ and $[k]$ in order to develop a bound on $\rho(\mathbf{r})$. Add $-\mathcal{L}_{0-1}([k], \mathbf{w})$ to the loss function to obtain

$$\begin{aligned} \mathcal{L}_{0-1}^*([i], \mathbf{w}) &= \mathcal{L}_{0-1}([i], \mathbf{w}) - \mathcal{L}_{0-1}([k], \mathbf{w}) \\ &= \begin{cases} 0 & \text{if } [i] = [k] \text{ or neither } [i] \text{ nor } [k] \\ & \text{is best,} \\ -1 & \text{if } [i] \neq [k] \text{ and } [i] \text{ is best,} \\ 1 & \text{if } [i] \neq [k] \text{ and } [k] \text{ is best,} \end{cases} \end{aligned}$$

$$E[\mathcal{L}_{0-1}^*([k])|\mathbf{x}] = 0,$$

$$E[\mathcal{L}_{0-1}^*([i])|\mathbf{x}] = p(w_{[k]} \text{ best } |\mathbf{x}) - p(w_{[i]} \text{ best } |\mathbf{x}) \geq -1.$$

This expected pairwise loss is incurred whenever system $[i]$ has a larger overall sample mean than system $[k]$, and this occurs with probability $\Phi[\tau_{[i,k]}^{1/2}(\mu_{[i]} - \mu_{[k]})]$. As $r_{[i]}, r_{[k]}$ grow without bound, $E[\mathcal{L}_{0-1}^*([i])|\mathbf{x}]$, system $[i]$ is best $\rightarrow -1$ almost surely (correct selection is assured by the perfect information obtained from an infinite number of replications). In numerical tests, this limit was typically approached rather quickly.

The lower bound of Equation (20) is obtained by adding the expected pairwise losses, approximating $E[\mathcal{L}_{0-1}^*([i])|\mathbf{x}]$ by -1 for each pairwise comparison, then adding the simulation costs and $E_{\xi}[\mathcal{L}_{0-1}([k], \mathbf{w})]$. Differentiate the lower bound and approximate $\lambda_{[i,k]}$ as in the proof of Theorem 1 to obtain the stated number of replications. \square

ACKNOWLEDGMENTS

The authors are pleased to acknowledge the financial support of the Engineering Research Center for Reconfigurable Machining Systems (NSF grant # EEC-9529125), and the assistance from the members of the ERC (both industry and academia). Feedback from Steve Pollock, Barry Nelson, Dave Goldsman, and the referees is also greatly appreciated.

REFERENCES

- Banks, J., J. S. Carson, B. L. Nelson. 1996. *Discrete-Event System Simulation* (2nd ed.). Prentice-Hall, Inc., Upper Saddle River, NJ.

- Bechhofer, R. E., T. J. Santner, D. M. Goldsman. 1995. *Design and Analysis for Statistical Selection, Screening, and Multiple Comparisons*. John Wiley & Sons, Inc., New York.
- Berger, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis* (2nd ed.). Springer-Verlag, New York.
- . 1988. A Bayesian approach to ranking and selection of related means with alternatives to analysis-of-variance methodology. *J. Amer. Statist. Assoc.* **83**(402) 364–373.
- Bracken, J., A. Schleifer. 1964. Tables for Normal Sampling with Unknown Variance. Graduate School of Business Administration, Harvard University, Cambridge, M.A.
- Chen, C.-H. 1996. A lower bound for the correct subset-selection probability and its application to discrete event simulations. *IEEE Trans. Auto. Control* **41**(8) 1227–1231.
- de Groot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, Inc., New York.
- Goldsman, D., B. L. Nelson. 1998. Statistical screening, selection, and multiple comparisons in computer simulation. *Proc. Winter Simulation Conference*. D. J. Madeiros, E. F. Watson, J. S. Carson, M. S. Manivannan (eds.). Institute of Electrical and Electronics Engineers, Inc., Piscataway, NJ. 159–166.
- Gupta, S. S. 1965. On some multiple decision (selection and ranking) rules. *Technometrics* **7** 225–245.
- , K. J. Miescke. 1994. Bayesian look ahead one stage sampling allocations for selecting the largest normal mean. *Statist. Papers* **35** 169–177.
- , —. 1996. Bayesian look ahead one-stage sampling allocations for selecting the best population. *J. Statist. Planning and Inference* **54** 229–244.
- Inoue, K., S. E. Chick. 1998. Comparison of Bayesian and frequentist assessments of uncertainty for selecting the best system. *Proc. Winter Simulation Conference*. D. J. Medeiros, E. J. Watson, M. Manivannan, J. Carson (eds.). Institute of Electrical and Electronics Engineers, Inc., Piscataway, NJ. 727–734.
- Koenig, L. W., 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. *Commun. Statist.-Simulation and Computation* **14**(3) 719–734.
- Law, A. M., W. D. Kelton. 1991. *Simulation Modeling & Analysis* (2nd ed.). McGraw-Hill, Inc., New York.
- Matejcik, F. J., B. L. Nelson. 1995. Two-stage multiple comparisons with the best for computer simulation. *Oper. Res.* **43** 633–640.
- Nelson, B. L., F. J. Matejcik. 1995. Using common random numbers for indifference-zone selection and multiple comparisons in simulation. *Management Sci.* **41** 1935–1945.
- , J. Swann, D. Goldsman, W. Song. 1999. Simple procedures for selecting the best simulated system when the number of alternatives is large. Technical Report, Northwestern University, Department of Industrial Engineering and Management Science, Evanston, IL.
- Rinott, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Comm. Statist.* **A7** 799–811.