

SEQUENTIAL ALLOCATIONS THAT REDUCE RISK FOR MULTIPLE COMPARISONS

Stephen E. Chick
Koichiro Inoue

Department of Industrial and Operations Engineering
The University of Michigan
1205 Beal Avenue
Ann Arbor, Michigan 48109-2117, U.S.A.

ABSTRACT

We consider how to efficiently allocate computing resources in order to infer the best of a finite set of simulated systems, where best means that the system has the maximal expected performance measure. Commonly-used frequentist procedures that are based on the indifference zone and ‘worst possible configuration’ tend to suggest an inefficiently large number of replications in practice. Recent work suggests that simulating likely competitors for the ‘best’ may lead to an order of magnitude improvement in computing effort for simulations. Much of that work, however, makes strong assumptions that might not be seen in practice, such as known variance, or the same cost of running a replication for each system. This paper discusses the problem of allocating computer resources to identify the best simulated system while relaxing general conditions, including different cost per replication for each system, both opportunity cost (linear loss) and 0-1 loss, and known or unknown variance for populations whose samples are normally distributed.

1 INTRODUCTION

How to select the best of a finite set of simulated systems is an important problem in discrete-event simulation (e.g., see Law and Kelton 1991; Goldsman and Nelson 1994). Based on an initial set of simulations of each system, a measure of evidence that a given system is best is desired (some authors suggest P-values and the Bonferroni inequality). If the measure of evidence is insufficient, an indication of how many additional computer simulation replications (or data samples) is required in order to reduce the risk of potentially selecting an inferior choice. The problem is also important in medicine, pharmacology, quality engineering, business, computer design, and a number of other fields (e.g., see Bechhofer, Hayter, and

Tamhane 1991; Gupta, Nagel, and Panchapakesan 1979; Gupta 1965; Hsu 1996).

Commonly-used frequentist procedures (Dudewicz and Dalal 1975; Rinott 1978) based on the ‘worst possible configuration’ are known suggest an inefficiently large number of additional replications in practice. Some Bayesian approaches (Gupta and Miescke 1994; Gupta and Miescke 1996; Chen, Chen, and Dai 1996) suggest that preferring to simulate likely competitors for the ‘best’ may lead to significant improvement in computing effort for simulations. Much of this work, however, makes strong assumptions that might not be seen in practice, such as known variance and the same cost of running a replication for each system.

In this paper, we discuss the problem of allocating computer resources to identify the best simulated system under much more general conditions, based on initial modeling work of Chick (1997).

Suppose that there are K different simulated stochastic systems, and that the system with the largest mean performance measure is to be identified. One system might be selected immediately, or one may observe the results of a finite number of additional independent simulation replications $x_{k,i}$ (for $k = 1, \dots, K; i = 1, \dots$) before selecting the ‘best’ system. As is commonly practice, we suppose the $x_{k,i}$ to have joint Gaussian distribution, conditional on the mean vector $\mathbf{w} = (w_1, \dots, w_K)$ and precision vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K)$. This assumption is often reasonable in the simulation context, where output often satisfies a central limit theorem. The mean is assumed unknown. Common practice is to consider the case of known precision (sometimes assumed identical for each system). Here we consider both known (and potentially unequal) and unknown precision. We also allow for different costs per replication, $\mathbf{c} = (c_1, \dots, c_K)$ for each system, and consider for both opportunity cost (or linear loss)

$$\mathcal{L}_{o.c.}(k, w) = \max_j w_j - w_k \quad (1)$$

and the 0-1 loss

$$\mathcal{L}_{0-1}(k, w) = \begin{cases} 0 & \text{when } w_k \geq \max_j w_j \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

For the one-stage problem, one of the K systems is to be selected after observing $r_k \geq 0$ output samples $\mathbf{x}_{r_k} = (x_{k,1}, \dots, x_{k,r_k})$ for $k = 1, \dots, K$. Set $\mathbf{r} = (r_1, \dots, r_K)$ and $\mathbf{x}_{\mathbf{r}} = (\mathbf{x}_{r_1}, \dots, \mathbf{x}_{r_K})$.

One is interested in determining the allocation of replications, \mathbf{r} , which minimizes the sum of the expected loss and the cost of observing the systems. The two-stage problem can be expressed in this form as well by letting the second-stage allocation use the posterior distribution after the first stage as a prior for the second stage. We also consider the cost-constrained allocation problem, where $\mathbf{c}\mathbf{r}^T = B$ for some computational budget $B > 0$. The cost-constrained problem can be used to develop a sequential decision procedure, where B is spent on sampling at each step of the procedure. Although formally \mathbf{r} is a vector of non-negative integers, our analysis follows a standard approximation that \mathbf{r} is real-valued. For steady-state simulations, fractional observations have a natural interpretation (e.g. simulate 5.32 years of a stochastic process). For terminating simulations, \mathbf{r} must be considered an approximation and rounded to integers. Throughout we write vectors in boldface, random variables in upper case, and their values in lower case.

Gupta and Miescke (1994) show that when $K = 2$; $c_1 = c_2$; $r_1 + r_2 = m$, the optimal policy for both $\mathcal{L}_{o.c.}(\cdot, \mathbf{w})$ and $\mathcal{L}_{0-1}(\cdot, \mathbf{w})$ is to minimize the absolute difference of the posterior precision for the mean of each system. Gupta and Miescke (1996) discuss the general decision problem, and provide an optimal allocation for a single replication for the case of $K \geq 3$, opportunity cost, known variance, and $c_1 = \dots = c_K$. Also see Gupta and Miescke (1988). A closed-form solution has proven difficult to obtain for other cases.

Because closed form solutions have proven difficult to obtain, we proceed by deriving natural upper and lower bounds on the expected risk for selecting the system with the highest expected mean. Motivated by an analogy with the Bonferroni inequality, we determine an optimal asymptotic allocation to minimize the lower bound on expected risk (loss plus cost of experiment). We then provide an allocation when the computer budget for running simulation replications is constrained, and the allocation is consistent with our asymptotic result. We then develop a sequential procedure, where at each step of a finite budget is allocated for sampling. Calculating these allocations often during a sequential procedure does not pose a significant computational burden, and therefore may

be an easy-to-implement efficiency improvement technique for selecting the best system. Our results also indicate that for $K > 2$, the DT condition is not optimal.

2 EXPECTED OPPORTUNITY COST

The expected opportunity cost can be expressed in terms of the truncated linear loss function of Eq. 1. Here we assume that the precision λ_k for the output of each system is known, and that uncertainty regarding the unknown mean w_k is expressed as a normally-distributed random variable with mean μ_k and precision τ_k .

$$W_k \sim \mathcal{N}(\mu_k, \tau_k)$$

Values of τ_k close to 0 indicate a high degree of uncertainty about the value of the unknown mean. We first consider the case of no budget constraint, with linear cost for systems, and then introduce a budget constraint.

Define $\Psi(s) = \int_s^\infty (x - s)\phi(x)dx = \phi(s) - s(1 - \Phi(s))$, where $\phi(s)$ and $\Phi(s)$ are the density function and cumulative distribution function, respectively, of the standard normal random variable. Intuitively, $\Psi(s)$ is the expected linear loss associated with asserting that the value of a standard normal random variable is less than s . Note that $\partial\Psi(s)/\partial s = \Phi(s) - 1$.

It is useful to define the permutation $[i]$ such that $\mu_{[1]} \geq \dots \geq \mu_{[K]}$ is a non-increasing ordering of the prior means for the w_i . Define $Z_i \stackrel{\text{def}}{=} E[W | \mathbf{X}_{r_i}] = \mu_i(\mathbf{X}_{r_i})$ to be the unknown posterior mean, a random variable which depends on the number of samples, r_i . Then the prior distribution of $W_{[i]} - W_{[j]}$ is $\mathcal{N}(\mu_{[i]} - \mu_{[j]}, \tau_{i,j})$, where $\tau_{i,j} = \left(\frac{\tau_{[i]}\tau_{[j]}}{\tau_{[i]} + \tau_{[j]}}\right)$. Further, the prior predictive distribution of $Z_i - Z_j$ is

$$Z_i - Z_j = \mu_{[i]}(\mathbf{X}_{r_i}) - \mu_{[j]}(\mathbf{X}_{r_j}) \sim \mathcal{N}(\mu_i - \mu_j, \tau_{\{i,j\}}),$$

$$\text{where } \tau_{\{i,j\}} = \left(\frac{\tau_{[i]}\lambda_{[i]}}{\tau_{[i]}[\tau_{[i]} + r_{[i]}\lambda_{[i]}} + \frac{\tau_{[j]}\lambda_{[j]}}{\tau_{[j]}[\tau_{[j]} + r_{[j]}\lambda_{[j]}}\right)^{-1}.$$

Theorem 1. Assume that the $x_{k,i}$ are jointly independent with $\mathcal{N}(w_k, \lambda_k)$ distribution, given w_k and known λ_k . Denote by ζ the prior distribution of $\mathbf{W} = (W_1, \dots, W_K)$, with $W_k \sim \mathcal{N}(\mu_k, \tau_k)$ jointly independent. Let δ_k represent the decision w_k is maximal, let the linear loss function $\mathcal{L}_{o.c.}(\delta_k, w)$ be as in Eq. 1, and let $G_{\mathbf{r}}$ be the prior predictive distribution of the posterior mean, \mathbf{Z} . Let $\mathcal{A}_k = \{\mathbf{z} \mid z_k = \max_j z_j\}$ be the event that the posterior mean of system k is maximal. Then:

- The optimal (Bayes) decision policy $\delta^*(\mathbf{r})$ is the natural decision rule $\delta^N(\mathbf{r})$,

$$\delta^N(\mathbf{r}) = \delta_k, \text{ when } z_k \geq \max_j z_j.$$

- The expected total risk $\rho_{o.c.}^*(\zeta, \mathbf{r})$,

$$\begin{aligned} \rho_{o.c.}^*(\zeta, \mathbf{r}) &= \mathbf{c}\mathbf{r}^T + E_{\zeta}[\max_j w_j - w_{[1]}] \\ &\quad - \sum_{k=2}^K p_{G_{\mathbf{r}}}(\mathcal{A}_{[k]}) E_{G_{\mathbf{r}}|\mathcal{A}_{[k]}}[z_{[k]} - z_{[1]}] \end{aligned}$$

is bounded above by:

$$\begin{aligned} &\mathbf{c}\mathbf{r}^T + E_{\zeta}[\max_j w_j - w_{[1]}] \\ &\quad - \max_{k \neq 1} \left\{ \tau_{\{1,k\}}^{-1/2} \Psi \left[\tau_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] \right\} \end{aligned}$$

and below by:

$$\begin{aligned} \hat{\rho}_{o.c.}^*(\zeta, \mathbf{r}) &\stackrel{\text{def}}{=} \mathbf{c}\mathbf{r}^T + E_{\zeta}[\max_j w_j - w_{[1]}] \\ &\quad - \sum_{k=2}^K \tau_{\{1,k\}}^{-1/2} \Psi \left[\tau_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] \end{aligned}$$

- In the limit $c_k \rightarrow 0$ (small cost of replications), the optimal system sizes r_k^* to minimize $\hat{\rho}_{o.c.}^*(\zeta, \mathbf{r})$ are asymptotically:

$$\tilde{r}_{[k]}^* = \left(\frac{\phi \left[(\tau_{1,k})^{1/2} (\mu_{[k]} - \mu_{[1]}) \right]}{2(\tau_{1,k})^{-1/2} c_{[k]} \lambda_{[k]}} \right)^{1/2} - \frac{\tau_{[k]}}{\lambda_{[k]}} \quad (3)$$

for $k \neq [1]$, and

$$\tilde{r}_{[1]}^* = \left(\sum_{k=2}^K \frac{\phi \left[(\tau_{1,k})^{1/2} (\mu_{[k]} - \mu_{[1]}) \right]}{2(\tau_{1,k})^{-1/2} c_{[1]} \lambda_{[1]}} \right)^{1/2} - \frac{\tau_{[1]}}{\lambda_{[1]}}.$$

Proof. See Chick and Inoue (1998). \square

The lower bound $\hat{\rho}_{o.c.}^*(\zeta, \mathbf{r}) \leq \rho_{o.c.}^*(\zeta, \mathbf{r})$ is obtained by summing of expected opportunity costs for each system that is better than the system with highest prior mean. The upper bound is derived by looking at each pairwise comparison individually, and taking the the most extreme of the expected pairwise losses. Special cases of these bounds (e.g., common known precision) were provided by Miescke (1979). The derivation of $\tilde{r}_{[k]}^*$ does not require that $E_{\zeta}[\max_j w_j - w_{[1]}]$ be calculated, as it does not depend on $\tilde{r}_{[k]}^*$. Further the same allocation is obtained by attempting to minimize the surrogate function:

$$\begin{aligned} \mathbf{c}\mathbf{r}^T &+ \sum_{k=2}^K \tau_{1,k}^{-1/2} \Psi \left[\tau_{1,k}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] \\ &- \sum_{k=2}^K \tau_{\{1,k\}}^{-1/2} \Psi \left[\tau_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right], \end{aligned}$$

which has the sum of expected pairwise losses between the ‘current’ best and the others. Since the Bonferroni inequality is also a sum of pairwise losses (albeit for the

0-1 loss), the minimization of the lower bound for Bayes risk that leads to Eq. 3 corresponds in some sense to a Bonferroni-type optimization.

Note that our allocations account for both the prior means of each system, as well as the precision. Widely-used techniques due to Dudewicz and Dalal (1975) and Rinott (1978) only consider the latter.

3 PROBABILITY OF CORRECT SELECTION

The probability of correct selection P(CS) when system k is selected can justifiably be defined as $p(w_k = \max_{i=1, \dots, K} w_i) = 1 - E[\mathcal{L}_{0-1}(k, \mathbf{w})]$. Recall that the Bayes decision to maximize P(CS) is *not* δ^N , but we evaluate the expected loss for δ^N to design an appropriate budget allocation for a simulation experiment. In this section it is useful to define $\tilde{\tau}_{1,2} = \left(\frac{1}{\tau_{[1]} + r_{[1]} \lambda_{[1]}} + \frac{1}{\tau_{[2]} + r_{[2]} \lambda_{[2]}} \right)^{-1}$ to be the posterior precision of $W_{[1]} - W_{[k]}$, given \mathbf{x} .

The case of $K = 2$ systems is considered first.

Lemma 2. Represent by ζ the prior distribution of $\mathbf{W} = (W_1, W_2)$, with $W_k \sim \mathcal{N}(\mu_k, \tau_k)$ jointly independent; and assume that the $x_{k,i}$ are jointly independent with $\mathcal{N}(w_k, \lambda_k)$ distribution, given w_k and a known precision λ_k . Let $G_{\mathbf{r}}$ be the prior predictive distribution for the unknown posterior mean \mathbf{Z} , as in Theorem 1. WLOG, assume $\mu_1 \geq \mu_2$. Consider the 0 – 1 loss as in Eq. 2. Then the expected total risk $\rho_{0-1}(\zeta, \delta^N(\mathbf{r}))$ is:

$$\begin{aligned} \mathbf{c}\mathbf{r}^T &+ \Phi \left[\tau_{1,2}^{1/2} (\mu_2 - \mu_1) \right] \\ &+ \Phi \left[\tau_{\{1,2\}}^{1/2} (\mu_2 - \mu_1) \right] \\ &\cdot \left(1 - 2E_{G_{\mathbf{r}}|z_2 \geq z_1} \left[\Phi \left[\tilde{\tau}_{1,2}^{1/2} (z_2 - z_1) \right] \right] \right) \end{aligned} \quad (4)$$

Proof. See Chick and Inoue (1998). \square

Insight from Eq. 4 can be obtained by examining the expected loss (excluding the cost of replications) $\rho = \rho_{0-1}(\zeta, \delta^N(\mathbf{r})) - \mathbf{c}\mathbf{r}^T$ at extremes. When there are no additional replications ($\mathbf{r} \rightarrow (0, 0)$), then $\tilde{\tau}_{1,2} \rightarrow \tau_{1,2}$, $\tau_{\{1,2\}} \rightarrow \infty$, and

$$\begin{aligned} \rho &= \Phi \left[\tau_{1,2}^{1/2} (\mu_2 - \mu_1) \right] + 0(1 - 2\gamma) \\ &= \text{prior probability 2 is best} \\ &= 1 - \text{prior } P(\text{CS}), \end{aligned}$$

for some $\gamma \in (1/2, 1]$, as desired. If the number of replications grows without bound ($\mathbf{r} \rightarrow (\infty, \infty)$), then $\tilde{\tau}_{1,2} \rightarrow \infty$, $\tau_{\{1,2\}} \rightarrow \tau_{1,2}$, $G_{\mathbf{r}} \rightarrow \zeta$, and

$$\begin{aligned} \rho &\rightarrow \Phi \left[\tau_{1,2}^{1/2} (\mu_2 - \mu_1) \right] \\ &+ \Phi \left[\tau_{1,2}^{1/2} (\mu_2 - \mu_1) \right] (1 - 2 \cdot 1) \\ &= 0, \end{aligned}$$

as expected, given the value of perfect information.

Because $E_{G_{\mathbf{r}}|z_2 \geq z_1} \left[\Phi \left[\tilde{\tau}_{1,2}^{-1/2} (z_2 - z_1) \right] \right] \rightarrow 1$ when the c_k are small (large r_k), a reasonable lower bound for $\rho_{0-1}(\zeta, \delta^N(\mathbf{r}))$ is therefore given by

$$\hat{\rho}_{0-1}^*(\zeta, \mathbf{r}) \stackrel{\text{def}}{=} \mathbf{c}\mathbf{r}^T + \Phi \left[\tau_{1,2}^{1/2} (\mu_2 - \mu_1) \right] - \Phi \left[\tau_{\{1,2\}}^{1/2} (\mu_2 - \mu_1) \right].$$

To minimize this lower bound, choose:

$$\tilde{r}_k^* = \left(\frac{\tau_{1,2}^{3/2} (\mu_2 - \mu_1) \phi \left[\left(\frac{\tau_{1,2}}{\tau_{\{1,2\}}} \right)^{1/2} (\mu_2 - \mu_1) \right]}{2c_k \lambda_k} \right)^{1/2} - \frac{\tau_k}{\lambda_k}.$$

With a budget constraint $\mathbf{c}\mathbf{r}^T = B$ (unconstrained \mathbf{r}) and the lower bound, choose same as for linear loss:

$$\tilde{r}_{k,B}^* = \frac{B + \sum_{j=1}^2 \frac{c_j \tau_j}{\lambda_j}}{\sum_{j=1}^2 \left(\frac{c_j \lambda_k}{\lambda_j} \right)^{1/2}} - \frac{\tau_k}{\lambda_k}.$$

These results are special cases of the following general theorem for $K \geq 2$.

Theorem 3. Represent by ζ the prior distribution of $\mathbf{W} = (W_1, \dots, W_K)$, with $W_k \sim \mathcal{N}(\mu_k, \tau_k)$ jointly independent; and assume that the $x_{k,i}$ are jointly independent with $\mathcal{N}(w_k, \lambda_k)$ distribution with known precision λ_k . Consider the 0–1 loss as in Eq. 2, and let $G_{\mathbf{r}}, \mathcal{A}_{[k]}$ be as in Theorem 1. Set $\tilde{\tau}_{1,k} = \left(\frac{1}{\tau_{[1]} + r_{[1]} \lambda_{[1]}} + \frac{1}{\tau_{[k]} + r_{[k]} \lambda_{[k]}} \right)^{-1}$ to be the posterior precision of $W_{[1]} - W_{[k]}$, given \mathbf{x} . Then:

- The expected total loss $\rho_{0-1}(\zeta, \delta^N(\mathbf{r}))$ of the natural decision rule:

$$\rho_{0-1}(\zeta, \delta^N(\mathbf{r})) = \mathbf{c}\mathbf{r}^T + (1 - p_{\zeta}(\mathcal{A}_{[1]})) + \sum_{k=2}^K \left\{ p_{G_{\mathbf{r}}}(\mathcal{A}_{[k]}) \left(1 - 2E_{G_{\mathbf{r}}|\mathcal{A}_{[k]}} \left[\Phi \left[\tilde{\tau}_{1,k}^{-1/2} (z_{[k]} - z_{[1]}) \right] \right] \right) \right\}$$

is bounded below by:

$$\hat{\rho}_{0-1}^*(\zeta, \mathbf{r}) \stackrel{\text{def}}{=} \mathbf{c}\mathbf{r}^T + (1 - p_{\zeta}(\mathcal{A}_{[1]})) - \sum_{k=2}^K \Phi \left[\tau_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] \quad (5)$$

- As costs $c_k \rightarrow 0$, the optimal system sizes r_k^* to minimize $\hat{\rho}_{0-1}^*(\zeta, \mathbf{r})$ approach:

$$\tilde{r}_{[k]}^* = \left(\frac{\phi \left[\tau_{1,k}^{1/2} (\mu_{[1]} - \mu_{[k]}) \right]}{2c_{[k]} \lambda_{[k]} / (\tau_{1,k}^{3/2} (\mu_{[1]} - \mu_{[k]}))} \right)^{\frac{1}{2}} - \frac{\tau_{[k]}}{\lambda_{[k]}} \quad (6)$$

for $[k] \neq [1]$, and

$$\tilde{r}_{[1]}^* = \left(\sum_{k=2}^K \frac{\phi \left[\tau_{1,k}^{1/2} (\mu_{[1]} - \mu_{[k]}) \right]}{2c_{[1]} \lambda_{[1]} / (\tau_{1,k}^{3/2} (\mu_{[1]} - \mu_{[k]}))} \right)^{\frac{1}{2}} - \frac{\tau_{[1]}}{\lambda_{[1]}}.$$

Proof. See Chick and Inoue (1998). \square

Minimizing the following surrogate function

$$\mathbf{c}\mathbf{r}^T + \sum_{k=2}^K \left\{ \Phi \left[\tau_{1,k}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] - \Phi \left[\tau_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right] \right\} \quad (7)$$

gives the same solution $\tilde{r}_{[k]}^*$ of Eq. 6 for minimizing Eq. 5. Neglecting the cost of replications, note that this surrogate function represents sum of approximations of pairwise probabilities incorrect selection, and that the Bonferroni inequality used in frequentist selection procedures is also based upon sums of pairwise probabilities of incorrect selection. For the special case of no replications, the sum in Eq. 7 becomes $\sum_{k=2}^K \Phi \left[\tau_{1,k}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right]$, an upper bound on the probability of incorrect selection, and is exactly the Bonferroni inequality applied to the prior probability of correct selection.

4 UNKNOWN VARIANCE

It is typically unreasonable to assume that the precision of the output is known. We therefore revisit the analysis of Sec. 3 under the assumption that both the mean and precision of the output are unknown. In particular, we represent uncertainty about $\mathbf{w}, \boldsymbol{\lambda}$ by assuming that $\Lambda_k \sim \mathcal{G}(\alpha_k, \beta_k)$; $W_k \sim \mathcal{N}(\mu_k, n_k \lambda_k)$, given λ_k ; and the vectors (w_k, λ_k) for $k = 1, \dots, K$ are assumed independent. The gamma distribution for Λ_k is parameterized so that $f(\lambda_k | \alpha_k, \beta_k) = \frac{(\beta_k)^{\alpha_k}}{\Gamma(\alpha_k)} \lambda_k^{\alpha_k - 1} e^{-\beta_k \lambda_k}$. Alternate prior distributions can be selected, but the analytic advantage of using conjugate prior distributions will be lost.

Given these assumptions, it is well-known (de Groot 1970) that the marginal distribution of W_k has Student-t distribution $\mathbf{St}(\mu_k, n_k \alpha_k / \beta_k, 2\alpha_k)$, where a random variable X is said to have $\mathbf{St}(\mu, \lambda, \nu)$ distribution when the density function $f(x | \mu, \lambda, \nu)$ is

$$f(x | \mu, \lambda, \nu) = c \left[1 + \frac{(x - \mu)^2 \lambda (x - \mu)}{\nu} \right]^{-\frac{\nu+1}{2}},$$

where $\mu = E[X]$, and $\text{Var}[X] = \lambda^{-1}\nu/(\nu - 2)$, when these integrals exist, and c is a normalizing constant. λ is commonly called the precision parameter.

The posterior distribution of W_k , given \mathbf{x}_k , is $\text{St}\left(\mu_k(\mathbf{x}_k), \frac{(n_k+r_k)(\alpha_k+r_k/2)}{\beta_{\mathbf{x}_k}}, 2\alpha_k+r_k\right)$, where $\mu_k(\mathbf{x}_k) = \frac{n_k\mu_k+r_k\bar{\mathbf{x}}_k}{n_k+r_k}$ is the posterior mean, $\bar{\mathbf{x}}_k = \sum_{i=1}^r x_{i,k}/r_k$ is the sample mean, and $\beta_{\mathbf{x}_k} = \beta_k + \frac{n_k r_k}{n_k+r_k}(\mu_k - \bar{\mathbf{x}}_k)^2 + \sum_{i=1}^{r_k} (x_{i,k} - \bar{\mathbf{x}}_k)^2/2$. Also, the pre-posterior distribution $G_{\mathbf{r}}$ for $\mathbf{Z} = E[W | \mathbf{X}_{\mathbf{r}}]$ is $\text{St}\left(\mu, \frac{n(n+r)\alpha}{r\beta}, 2\alpha\right)$.

Since we are interested in comparisons between systems, the distribution of the difference of two Student-t variables is of interest. However, the distribution of the difference is also Student-t only when the degrees of freedom match. In particular, if $\alpha = \alpha_i = \alpha_j$, then the prior and preposterior precision parameters for $W_i - W_j$ are $\lambda_{i,j} = \alpha \left(\frac{\beta_{[i]}}{n_{[i]}} + \frac{\beta_{[j]}}{n_{[j]}} \right)^{-1}$ and $\lambda_{\{i,j\}} = \alpha \left(\frac{r_{[i]}\beta_{[i]}}{n_{[i]}(n_{[i]}+r_{[i]})} + \frac{r_{[j]}\beta_{[j]}}{n_{[j]}(n_{[j]}+r_{[j]})} \right)^{-1}$, respectively, and the degrees of freedom is 2α . The following theorem studies the case $\nu/2 = \alpha = \alpha_1 = \dots = \alpha_K$.

Theorem 4. Assume that the $x_{k,i}$ are jointly independent with $\mathcal{N}(w_k, \lambda_k)$ distribution. Denote by ζ the prior distribution of $\mathbf{W}, \mathbf{\Lambda} = (W_1, \Lambda_1, \dots, W_K, \Lambda_K)$, with $W_k | \lambda_k \sim \mathcal{N}(\mu_k, n_k \lambda_k)$, $\Lambda_k \sim \mathcal{G}(\alpha, \beta_k)$, and (w_k, λ_k) independent for $k = 1, \dots, K$. (Note that all systems have the same α .) Suppose that the 0–1 loss function is as in Eq. 2. Then:

- The total risk is bounded below by:

$$\begin{aligned} \hat{\rho}_{0-1}^*(\zeta, \mathbf{r}) &= \mathbf{c}\mathbf{r}^T + E_{\zeta}[\text{max}_j w_j - w_{[1]}] \\ &- \sum_{k=2}^K \Phi_{2\alpha} \left[\lambda_{\{1,k\}}^{1/2} (\mu_{[k]} - \mu_{[1]}) \right], \end{aligned}$$

and the bound is tight for $K = 2$.

- In the limit $c_k \rightarrow 0$ (small cost of replications), the optimal system sizes r_k^* to minimize $\hat{\rho}_{0-1}^*(\zeta, \mathbf{r})$ are asymptotically:

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

for $[k] \neq [1]$, and

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

Proof. See Chick and Inoue (1998). \square

Although the assumption of identical α_k seems rather restrictive, this situation arises naturally in certain contexts. For instance, if R_0 replications for each system are taken in phase 1 of a two-stage procedure for multiple comparisons, and a reference prior (Bernardo and Smith 1994) is used, then $\alpha = (R_0 - 1)/2$ for all systems at the start of the second stage. Sec. 5 discusses this point further.

When the α_k are not all equal, it might be reasonable to use some approximation formulas for the difference of two Student-t random variables with different degrees of freedom. For instance when $\alpha_1 \neq \alpha_k$, one might approximate terms involving $\phi_{2\alpha}, \lambda_{1,k}$ via with $\phi_{2\tilde{\alpha}_k}, \tilde{\lambda}_{1,k}$ via Welsh's approximation (e.g., see Law and Kelton 1991) or Patil's approximation (e.g., see Box and Tiao 1973) for the Behrens-Fisher distribution. Terms such as $\alpha/\beta_{[k]}$ should be replaced with $\alpha_{[k]}/\beta_{[k]}$.

Corollary 5. The solution to $\min \hat{\rho}_{0-1}^*(\zeta, \mathbf{r})$ subject to $\mathbf{c}\mathbf{r}^T = B$ for large B is approximately:

$$\tilde{r}_{k,B}^* = \frac{B + \sum_{j=1}^K c_j n_j}{\sum_{j=1}^K \left(\frac{c_j c_k \beta_j \gamma_{j,\alpha}}{\beta_k \gamma_{k,\alpha}} \right)^{1/2}} - n_k, \quad (8)$$

where

$$\gamma_{[k],\alpha} = \lambda_{1,k}^{3/2} (\mu_{[1]} - \mu_{[k]}) \phi_{2\alpha} \left[\lambda_{1,k}^{1/2} (\mu_{[1]} - \mu_{[k]}) \right]$$

for $[k] \neq [1]$, and $\gamma_{[1],\alpha} = \sum_{j|j \neq [1]} \gamma_{j,\alpha}$.

Proof. See Chick and Inoue (1998). \square

What happens when the constraint $r_k \geq 0$ is added to the optimization problem of Corollary 5? Suppose that one or more $\tilde{r}_{k,B}^* < 0$, as determined by Eq. 8, and let $S = \{k | \tilde{r}_{k,B}^* \geq 0\}$. One cannot make a negative number of replications, so we reset $\tilde{r}_{k,B}^* = 0$ for those $k \notin S$. Resetting $\tilde{r}_{k,B}^* = 0$, however, causes the total sampling budget to exceed B , which indicates the number of systems for other systems should decrease. Eq. 8 should therefore be applied again to recalculate $\tilde{r}_{k,B}^*$ for all $k \in S$. This, in turn, might cause the recalculated $\tilde{r}_{k,B}^*$ to be negative for some $k \in S$. Remove such k from S and recalculating again is appropriate, as given by algorithm of Table 1.

Corollary 5 and the algorithm of Table 1 use the approximation $\tau_{i,j} \approx \tau_{\{i,j\}}$ (which assumes $\tilde{r}_i^*, \tilde{r}_j^*$ are large), via the $\gamma_{[k],\alpha}$. Direct calculation indicates that the approximation also holds when one of \tilde{r}_i^* or \tilde{r}_j^* is zero and the other is large, which further justifies the use of this approximation in the algorithm, even when some $\tilde{r}_i^* = 0$.

Table 1: Computation of optimal allocation $\tilde{r}_{k,B}^*$ to minimize lower bound on total risk, for 0-1 loss, unknown precision, with fixed budget B , as in Corollary 5.

1. Determine a budget B . Initialize $S = \{1, \dots, K\}$.
2. For all $k \in S$, calculate

$$\tilde{r}_{k,B}^* = \frac{B + \sum_{j \in S} c_j n_j}{\sum_{j \in S} \left(\frac{c_j c_k \beta_j \gamma_{j,\alpha}}{\beta_k \gamma_{k,\alpha}} \right)^{1/2}} - n_k.$$

3. Set redo to false.
4. For all $k \in S$, If $\tilde{r}_{k,B}^* < 0$ Then: (a) Set $S = S \setminus \{k\}$, (b) Set $\tilde{r}_{k,B}^* = 0$, (c) Set redo to true.
5. If redo is true, Then go to Step 2.
6. Suggest that $\tilde{r}_{k,B}^*$ systems be observed for system k , $k = 1, \dots, K$.

5 COMPARISON FOR TWO-STAGE

If we consider a special case of the above development (variance is unknown, and may be different from system to system, same cost of replications for each system, 0-1 loss), we can make a comparison with classical two-stage indifference zone procedures. We select the procedure due to Rinott (1978), but other procedures can be compared similarly.

Two-stage classical procedures typically require each of K systems to be simulated with R_0 (≈ 10 to 20) replications during the first stage. Based on the sample mean and variance of each system, an additional number $n_{k,Rin}^*$ of replications during the second stage are to be run for each system, $k = 1, \dots, K$. The number of second-stage replications typically depends on K , and indifference zone Δ , and the desired level P^* for the probability of correct selection, often taken to be 90 to 95%. For instance, Rinott (1978) suggests:

$$n_{k,Rin}^* = \max\left\{0, \left\lceil \frac{h^2 S_k^2}{\Delta^2} \right\rceil - R_0 \right\} \quad (9)$$

where $S_k^2 = \sum_{j=1}^{R_0} (x_{j,k} - \bar{x}_{R_0})^2 / (R_0 - 1)$ is the sample variance of the first stage output of system k , and h solves Rinott's integral for R_0, K , and P^* (e.g., see Wilcox 1984).

An analogous approach based on the arguments in the above sections would be to similarly run $r_k = R_0$ replications of each system during a first stage. The posterior from the first stage could then be used as a prior distribution for the second stage to determine the number of

additional replications, using the results of Theorem 4. If a reference prior for each system were used (Chick 1997), then the appropriate prior for the second stage would have $\mu_k = \bar{x}_k$, $n_k = R_0$, $\alpha = (R_0 - 1)/2$, $\beta_k = \frac{R_0 - 1}{2} S_k^2$, and $\lambda_{1,k} = R_0 \left(S_{[1]}^2 + S_{[k]}^2 \right)^{-1}$.

To compare the classical and present allocations, one could evaluate the expected number of replications required to reach a given P(CS). Alternately, one can evaluate the expected P(CS) obtainable for a given computer budget. We choose to do the latter. For the procedure of Rinott (1978), this requires selecting Δ so that the correct total number of computer replications is achieved.

We simulated the 5 inventory systems described in Law and Kelton (1991) in the section on comparing more than two systems (p. 591) of their widely-used text. All parameters and simulation design are identical to those used by Law and Kelton (1991) except that different uniform random number streams were used. The goal is to identify the inventory policy which leads to the minimal average monthly cost. (This minimization is easily converted to a maximization problem, as required by Sec. 4.) Based on a first-stage with $R_0 = 10$ replications for each system, we obtained the output summarized in Fig. 2. Also included in Fig. 2 are the allocations determined by the Rinott and Bayesian approaches.

Table 2: Stage One (10 replications) Statistics and Stage Two Allocation ($B = 100$) for Five (s, S) Inventory Policies of Law and Kelton (1991)

System	s	S	\bar{x}_k	S_k	Rinott	Bayes
1	20	40	123.86	4.11	48	42
2	20	80	122.70	2.01	3	30
3	40	60	124.48	2.54	12	28
4	40	100	132.36	2.80	16	0
5	60	100	145.19	3.03	21	0

We ran a second stage, and observed the posterior Bonferroni P(CS) measure $\sum_{k=1}^{K-1} \Phi \left[\tilde{r}_{k,K}^{1/2} (z_{[k]} - z_{[K]}) \right]$, which measures the evidence of correct selection with the Bonferroni inequality and pairwise comparisons (P-value for frequentists and posterior probability of correct selection for a Bayesian). We repeated the second stage 500 times in order to estimate

$$E \left[\sum_{k=1}^{K-1} \Phi \left[\tilde{r}_{k,K}^{1/2} (Z_{[k]} - Z_{[K]}) \right] \right], \quad (10)$$

which we call the expected posterior Bonferroni P(CS), the expectation taken from the information available prior to running the replications.

Table 3: Expected Bonferroni P(CS) Estimates, Using Allocations in Table 2

Allocation	P(CS)
Prior ($B = 0$)	0.6307
Rinott ($B = 100$)	0.9687
Bayes ($B = 100$)	0.9993

The prior Bonferroni P(CS), and estimates of the expected posterior Bonferroni P(CS) for both the Rinott and Bayesian allocation are summarized in Table 3. The Bayesian allocation provides a much better measure of evidence, in expectation, than the Rinott allocation. Additional numerical experimentation (results not shown) indicates that the Bayesian allocation with 50 second stage replications gives an expected posterior Bonferroni P(CS) of 0.9736, somewhat higher than the Rinott with 100 replications.

6 CONCLUSIONS

In many areas of stochastic simulation, computational efficiency improvement has received a great deal of attention. For instance, numerous analytic and empirical studies for Monte Carlo integration evaluate the trade-offs between the cost of generating and the variance of independent samples that can be used to estimate an integral.

The incorporation of sampling costs for the efficient selection of the best simulated system, however, has not previously been studied by the standard treatments of multiple selection problems. This paper proposes to incorporate sampling costs explicitly, and further, uses value of information arguments to improve the computational efficiency of identifying the best system. We explicitly handle known or unknown variance, heteroscedastic simulation responses, and both opportunity cost and 0-1 loss. Our general approach is to (i) determine the expected total risk, (ii) derive a natural Bonferroni-like lower bound, (iii) determine asymptotically optimal allocations as the cost of replications gets arbitrarily small, (iv) establish allocations when there are budget constraints that are asymptotically optimal, for large budgets.

The asymptotic budget allocations may be used to develop sequential sampling strategies for simulation experiments, by allocating a small budget at each stage of the sequential design. The cost of calculating these allocations at each stage does not represent an onerous burden, as allocations can be determined with simple mathematical operations (sum, multiply, square root, exponentiation), and do not need numerical root finding.

REFERENCES

- Bechhofer, R. E., A. J. Hayter, and A. C. Tamhane. 1991. Designing experiments for selecting the largest normal mean when the variances are known and unequal: Optimal sample size allocation. *Journal of Statistical Planning and Inference* 28:271–289.
- Bernardo, J. M. and A. F. M. Smith. 1994. *Bayesian theory*. Chichester, UK: Wiley.
- Box, G. E. P. and G. C. Tiao. 1973. *Bayesian inference in statistical analysis*. New York: John Wiley and Sons.
- Chen, C.-H., H.-C. Chen, and L. Dai. 1996. A gradient approach for smartly allocating computing budget for discrete event simulation. In *Proceedings of the Winter Simulation Conference*, ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain, 398–405. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Chick, S. E. 1997. Selecting the best system: A decision-theoretic approach. In *Proceedings of the Winter Simulation Conference*, ed. S. Andradóttir, K. Healy, D. Withers, and B. Nelson, 326–333. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Chick, S. E. and K. Inoue. 1998. How to more effectively allocate computing resources to identify the best simulated alternative. In review.
- de Groot, M. H. 1970. *Optimal statistical decisions*. New York: McGraw-Hill, Inc.
- Dudewicz, E. J. and S. R. Dalal. 1975. Allocation of observations in ranking and selection with unequal variances. *Sanhkya* B37:28–78.
- Goldsman, D. and B. L. Nelson. 1994. Ranking, selection, and multiple comparisons in computer simulation. In *Proceedings of the Winter Simulation Conference*, ed. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 192–199. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Gupta, S. S. 1965. On some multiple decision (selection and ranking) rules. *Technometrics* 7:225–245.
- Gupta, S. S. and K. J. Miescke. 1988. On the problem of finding the largest normal mean under heteroscedasticity. In *Statistical decision theory and related topics IV*, Vol. 2, ed. S. S. Gupta and J. O. Berger. New York: Springer-Verlag.
- Gupta, S. S. and K. J. Miescke. 1994. Bayesian look ahead one stage sampling allocations for selecting the largest normal mean. *Statistical Papers* 35:169–177.
- Gupta, S. S. and K. J. Miescke. 1996. Bayesian look ahead one-stage sampling allocations for selecting the best population. *Journal of Statistical Planning and Inference* 54:229–244.

- Gupta, S. S., K. Nagel, and S. Panchapakesan. 1979. *Multiple decision procedures: Theory and methodology of selecting and ranking populations*. New York: Wiley.
- Hsu, J. C. 1996. *Multiple comparisons: Theory and methods*. New York: Chapman & Hall.
- Law, A. M. and W. D. Kelton. 1991. *Simulation Modeling & Analysis*. 2d ed. New York: McGraw-Hill, Inc.
- Miescke, K. J. 1979. Bayesian subset selection for additive and linear loss functions. *Commun. Statist.-Theor. Meth.* A8 (12): 1205–1226.
- Rinott, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics* A7:799–811.
- Wilcox, R. R. 1984. A table for Rinott's selection procedure. *Journal of Quality Technology* 16:97–100.

AUTHOR BIOGRAPHIES

STEPHEN E. CHICK is an assistant professor of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. In addition to stochastic simulation, his research interests include Bayesian statistics, reliability, decision analysis, computational methods in statistics, and applications to health care, particularly epidemiology. His work experience includes several years of material handling system design for the automotive industry using simulation analysis.

KOICHIRO INOUE is a Ph.D. candidate in the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. His dissertation work focuses on the application of Bayesian statistics to simulation output analysis. He is also interested in multi-attribute decision analysis.