Estimating the Probability That a Simulated System Will Be the Best

J.O. Miller,[1] Barry L. Nelson,[2] Charles H. Reilly[3]

[1]*Department of Operational Sciences, Air Force Institute of Technology, Wright-Patterson AFB, Ohio 45433-7765*

[2]*Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208-3119*

[3]*Department of Industrial Engineering and Management Systems, University of Central Florida, Orlando, Florida 32816-2450*

Received July 2001; revised November 2001; accepted 26 November 2001

Abstract: Consider a stochastic simulation experiment consisting of v independent vector replications consisting of an observation from each of k independent systems. Typical system comparisons are based on mean (long-run) performance. However, the *probability* that a system will actually be the best is sometimes more relevant, and can provide a very different perspective than the systems' means. Empirically, we select one system as the best performer (i.e., it wins) on each replication. Each system has an unknown constant probability of winning on any replication and the numbers of wins for the individual systems follow a multinomial distribution. Procedures exist for selecting the system with the largest probability of being the best. This paper addresses the companion problem of estimating the probability that each system will be the best. The maximum likelihood estimators (MLEs) of the multinomial cell probabilities for a set of v vector replications across k systems are well known. We use these same v vector replications to form v^k unique vectors (termed pseudo-replications) that contain one observation from each system and develop estimators based on AVC (*A*ll *V*ector *C*omparisons). In other words, we compare every observation from each system with every combination of observations from the remaining systems and note the best performer in each pseudo-replication. AVC provides lower variance estimators of the probability that each system will be the best than the MLEs. We also derive confidence intervals for the AVC point estimators, present a portion of an extensive empirical evaluation and provide a realistic example. © 2002 Wiley Periodicals, Inc. Naval Research Logistics 49: 341–358, 2002; Published online in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/nav.10019

Keywords: multinomial; ranking and selection; simulation

1. INTRODUCTION

Suppose we have $k \geq 2$ independent populations, denoted $\pi_1, \pi_2, \ldots, \pi_k$. In the stochastic simulation context each "population" is a simulated system. We consider the problem of estimating the probability that each system will be the best system in a single comparison (trial)

Correspondence to: J.O. Miller

for all of the systems. Often—but not always—the best performer in terms of mean or long-run average performance will correspond to the system with the largest probability of being the best. But maximizing the probability of being the best answers a fundamentally different question than maximizing the long-run average performance does, and can provide a very different perspective on system performance. We provide examples in this paper.

Let X_{ji} represent the ith replication of some performance measure from system j. Note that the X_{ji} are independent both within and across systems. Each system has an unknown constant probability, denoted $p_j, j = 1, 2, \ldots, k$, of having the largest value of the performance measure on any replication containing one observation from each system. Our performance measure is a continuous random variable, so that a tie between systems occurs with probability 0. We define the best system as the system most likely to have the largest performance measure (i.e., it wins) in any comparison across all systems.[1] Such a comparison corresponds to a multinomial trial, where one and only one system can win in any given trial; therefore, the numbers of wins for each individual system in v independent trials follows a multinomial distribution. Our objective is to provide estimates for the unknown multinomial success probabilities, $p_j, j = 1, 2, \ldots, k$.

Closely related to the estimation of the p_j's is the problem of determining which of the systems has the largest probability of being the best system. This is known as the multinomial selection problem (MSP). Bechhofer and Sobel [2] introduced the use of multinomial selection procedures to find the system most likely to produce the largest observation on a given vector-trial. Goldsman [4] first suggested the more general use of this type of procedure to find the system most likely to produce the "most desirable" observation on a given vector-trial, where "most desirable" can be almost any criterion of goodness.

An MSP approach is appropriate when we are interested in selecting the system expected to yield the highest benefit in a single trial, rather than the system with the largest average response over a large number of trials. Simulation examples include selecting the best set of tactical or strategic military actions to achieve maximum damage in a single strike; selecting the bridge design most likely to require the least costly repairs after a one-time catastrophic event, such as an earthquake; or selecting the computer system with the highest probability of running the longest without failure. For the type of problem considered in this paper, we require a quantitative measure of system performance so that each system in each trial can be compared with other systems across any or all of the remaining trials.

When comparing systems, it is often argued that common random numbers (CRN) should be used to evaluate the different systems under the same experimental conditions, where our experimental conditions are the random variates used to determine various system reliabilities, probabilities of target damage, etc. The use of CRN implies that the X_{ji}'s are no longer independent across systems in each replication. When estimating the difference in mean system performance, CRN tends to reduce the variance of the difference by inducing a positive correlation across the systems' responses; it does this without changing the true mean response from each system. On the other hand, if we estimate the probability that one system's performance will be better than the others', then the value of this probability is altered by using CRN. In other words, the MSP with and without CRN are actually different problems.

Miller and Bauer [10] showed the use of CRN in MSP may not change the *identity* of best system—the system with the largest probability of having the best performance—even though it does change the *value* of that probability. They also showed that the probability of correctly

[1]Throughout this paper we assume that a larger value of the performance measure is better, but neither the analysis nor the results change if smaller performance is better.

selecting the best system can increase or decrease with the use of CRN, implying that CRN may help or hurt in the selection problem.

This paper considers *estimating* the probability that each system will be the best; thus, it matters which problem we intend to solve (with or without CRN). The analyst must decide whether or not to use CRN based on whether it is logically appropriate, rather than whether or not it will reduce variance. When the systems react to the same set of experimental conditions in the same ways, then simulating them using CRN may be a faithful representation of the physical process. But when the systems react in different ways to the experimental conditions, or cause those conditions to be changed, then simulating them independently is logically correct. The new estimator introduced in this paper is only appropriate when the systems are simulated independently (no CRN), so from here on we assume independence across systems.

Let $\mathbf{X_i} = (X_{1i}, X_{2i}, \ldots, X_{ki})$ represent the ith replication across all k systems. Let $Y_{ji} = 1$ if $X_{ji} > X_{\ell i}$, for $\ell = 1, 2, \ldots, k$, but $\ell \neq j$; and let $Y_{ji} = 0$ otherwise. In other words, $Y_{ji} = 1$ if X_{ji} is the largest observation in $\mathbf{X}_i$. Suppose that there are v independent replications across all systems, and let $Y_j = \sum_{i=1}^{v} Y_{ji}$ represent the number of times system j wins out of these v replications. For simplicity assume that ties have probability 0, and let $p_j = \Pr\{X_{ji} > X_{\ell i}, \forall \ell \neq j\}$, where $0 < p_j < 1$ and $\sum_{j=1}^{k} p_j = 1$. Then $\sum_{j=1}^{k} Y_j = v$, and the k-variate discrete random variable $\mathbf{Y} = (Y_1, Y_2, \ldots, Y_k)$ follows a multinomial distribution with v trials and success probabilities $\mathbf{p} = (p_1, p_2, \ldots, p_k)$. The usual point estimators for the multinomial success probabilities are the maximum likelihood estimators (MLEs) given by

$$\hat{p}_j = \frac{Y_j}{v} \tag{1}$$

for $j = 1, 2, \ldots, k$.

Due to convention and convenience when comparing simulated system responses from independent systems, the responses are typically grouped by replication, corresponding to a trial in a physical experiment. Grouping system responses in this fashion is arbitrary (unless there is some attempt to synchronize comparisons across systems, like the use of CRNs), and since our simulated responses are quantitative, we can compare any observation from one system with any observation from each of the remaining systems. This means that a single observation from system 1 can be grouped into a vector comparison with any one of the v observations from system 2, and with any one of the v observations from system 3, and so on, up to and including any one of the v observations from system k. Since there are v observations from system 1 as well, this gives us a total of v^k vector comparisons (trials) that can be formed with v independent observations from each of the k systems. We incorporate this setup, which we call AVC for *A*ll *V*ector *C*omparisons, to construct new point estimators for the p_j's. Our estimators turn out to be k-sample U-statistics [13]. When we arbitrarily perform only the v vector comparisons implied by the MLEs, we disregard the information available from the remaining $v^k - v$ comparisons. AVC exploits this additional information.

In this paper we prove that the variance of the AVC estimator is no larger than the variance of the MLE, and is typically much smaller. We use a specific example to demonstrate the magnitude of the variance reduction that can be expected. We also derive the asymptotic variance of the AVC estimator as an approximation to its true variance. An estimate of the asymptotic variance is used to form confidence intervals that are shorter than their MLE counterparts.

The paper is organized as follows: We first motivate the need to look at performance measures other than the mean, and then briefly review the MSP and introduce terminology used in our point-estimation problem. This leads to a description of AVC. Properties and empirical evaluation of the

AVC estimators follow, including the construction of confidence intervals. We conclude with an example illustrating the use of these procedures in selecting the best performer among a number of highly reliable computer systems.

2. MOTIVATION AND BACKGROUND

Comparisons based on mean or long-run average performance are so common that this choice of performance measure is rarely questioned. Provided that long-run performance of a stable system is what matters, comparisons based on the mean make sense. However, when risk is important, or when one-shot performance matters, then measures such as the probability of being the best are relevant.

In some situations the system with the best mean performance may not be the one that is most likely to be the best. Goldsman [5] gives the following simple example: Let A and B be two inventory policies. Profit from A is 1000 with probability 0.001 and 0 with probability 0.999. However, profit from B is certain to be 0.999. In this situation the expected profit from A, which is 1, is greater than the expected profit from B, which is 0.999. However, the $\Pr\{\text{Profit from B} > \text{Profit from A}\} = 0.999$, so that B is almost certain to give higher profit. In this case the system with the best expected performance is not the same as the one that is most likely to be the best. This example also illustrates that in our context we are not interested in the magnitude of the difference in performance among competing systems in any given trial (replication).

Even when the best system as determined by mean performance or the probability of being the best is the same, our perception of how much better one system is than another can be quite different depending on which performance measure we choose. For example, suppose now that the Profit from A is distributed $N(11, \sigma^2)$, while the Profit from B is N(10, 1). Thus, A is 10% better than B in terms of mean performance for all values of σ^2. Figure 1 shows the $\Pr\{\text{Profit from A} > \text{Profit from B}\}$ as a function of σ^2 for $0.1 \leq \sigma^2 \leq 100$. Notice that $\Pr\{\text{Profit from A} > \text{Profit from B}\}$ ranges from 0.83 (quite certain) to 0.54 (a bit better than a coin toss) depending on the variance of the return from A. In this example knowledge of the probability of being the best adds to our understanding of the relationship between A and B beyond what is provided by their means.

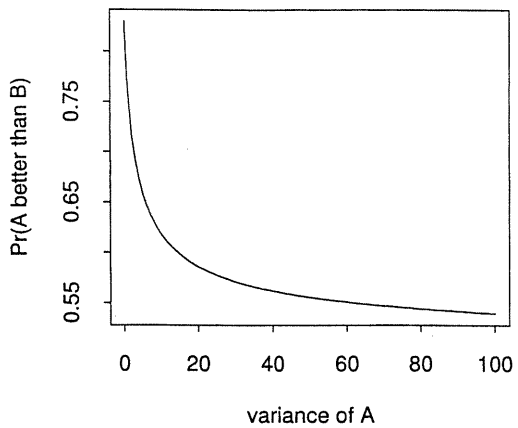

Figure 1. The probability that policy A has higher profit than policy B as a function of the variance of its return.

One thesis of this paper is that probability of being the best is a useful complement to mean performance, and AVC provides a statistically efficient way to estimate it from the same data used to estimate the mean.

Because of the close relationship between the MSP and the estimation of multinomial success probabilities, we briefly introduce a classical solution procedure for the MSP and define some terminology we will use throughout our discussion.

Bechhofer, Elmaghraby, and Morse [1] describe a single-stage procedure, Procedure BEM, for selecting the multinomial event (population or system) which has the largest success probability. Procedure BEM requires the specification of P^* (where $1/k < P^* < 1$), a minimum probability of correctly identifying the population with the largest success probability subject to an indifference zone constraint, and $\theta^* > 1$, the indifference zone constraint defined as the minimum ratio of the largest success probability to the second largest success probability that we want to be able to detect. The probability of correct selection (PCS) is a property of the procedure and provides no information about the values of the p_j's. If we desire estimates of the p_j's, then the natural set of estimators based on the same data as BEM are the MLEs given in (1). It is well known that $\mathrm{E}[\hat{p}_j] = p_j$ and

$$\mathrm{Var}(\hat{p}_j) = \frac{p_j(1-p_j)}{v}. \tag{2}$$

3. AVC ESTIMATORS

We propose a new method to provide point estimators for the multinomial success probabilities with smaller variances than the corresponding MLEs using the same replications $\mathbf{X}_i, i = 1, 2, \ldots, v$. Rather than comparing the ith replication for each system with the ith replication of the other systems, consider a total of v^k pseudo-replications formed by associating each X_{ji} $(j = 1, 2, \ldots, k; i = 1, 2, \ldots, v)$, with all possible combinations of the remaining $X_{\ell h}$ $(\ell = 1, 2, \ldots, k; \ell \neq j; h = 1, 2, \ldots, v)$. Each such pseudo-replication contains one observation from each system. Notice that the v^k pseudo-replications include the v independent replications from which the pseudo-replications are formed.

Assuming ties are not possible, define

$$Z_j = \sum_{a_1=1}^{v} \sum_{a_2=1}^{v} \cdots \sum_{a_k=1}^{v} \prod_{\ell=1;\ell\neq j}^{k} \phi(X_{ja_j} - X_{\ell a_\ell}) \tag{3}$$

for $j = 1, 2, \ldots, k$, with

$$\phi(a) = \begin{cases} 1, & a > 0 \\ 0, & a < 0. \end{cases}$$

Thus, Z_j represents the number of times out of v^k pseudo-replications that system j wins and $\sum_{j=1}^{k} Z_j = v^k$. Our new point estimator is then

$$\bar{p}_j = \frac{Z_j}{v^k}, \tag{4}$$

the fraction of wins out of v^k pseudo-replications for system j. We refer to the estimator in (4) as an AVC estimator, or to $\bar{\mathbf{p}} = (\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_k)$ collectively as our AVC estimator. Notice that

if ties are possible, then a tie for the best must be broken by randomly selecting one of the tied systems as being the best *independently for each pseudo-replication*. When ties can occur, a simple representation of the number of wins such as (3) is not possible. Throughout the remainder of the paper we restrict attention to the case in which the X_{ji} are such that ties have probability 0.

For convenience in comparing our AVC estimators with the MLEs, we have used the same number of observations for each of the k systems. However, a potential advantage of the AVC estimators is the ability to easily define them for differing number of observations among the k systems. Let v_j represent the number of independent observations from system $j, j = 1, 2, \ldots, k$. We then have

$$Z_j = \sum_{a_1=1}^{v_1} \sum_{a_2=1}^{v_2} \cdots \sum_{a_k=1}^{v_k} \prod_{\ell=1;\ell\neq j}^{k} \phi(X_{ja_j} - X_{\ell a_\ell}) \tag{5}$$

with $\phi(a)$ defined as above, so that $\sum_{j=1}^{k} Z_j = \prod_{\ell=1}^{k} v_\ell$. Then our AVC point estimator becomes $\bar{p}_j = Z_j / \prod_{\ell=1}^{k} v_\ell$. For simplicity, we assume from here on that we have v replications from each system.

The basic AVC *selection* procedure was introduced in [11] and applied to the MSP. In [11] we show that many fewer vector replications are required to solve the MSP using an AVC-based procedure, rather than procedure BEM, and we show how to determine the required number of replications. We also define our AVC point estimator for the best system, show that it is a k-sample U statistic, and use an asymptotic argument to compute and compare the probability of correct selection with procedure BEM. In this paper we examine the performance of AVC-based point estimators for *all* competing systems, derive confidence intervals for them and compare these estimators with the MLEs obtained using BEM.

Once the data are collected, computing all of the AVC point and variance estimators is an $O(kv \ln kv)$ calculation. This compares to an $O(kv)$ calculation for the MLEs. Thus, obtaining the AVC point estimators may not be worth the extra effort if the computing cost per additional replication is trivial. However, in typical systems simulation problems the computation per replication is orders of magnitude more than the additions, multiplications and sorting required to produce the AVC estimators, suggesting that the reduction in variance is not offset by the additional effort. An efficient algorithm for AVC calculation is given in the Appendix.

Clearly $\mathrm{E}[\bar{p}_j] = p_j$, but the $\mathrm{Var}(\bar{p}_j)$ is more complex to calculate than $\mathrm{Var}(\hat{p}_j)$. To find the variances for the individual $\bar{p}_j$'s we will show that our AVC estimator is a k-sample U-statistic and specialize a general expression for its variance.

Let F_j denote the cumulative distribution function (cdf) of the random variable X_{ji}. From [13], we say $\mathbf{p}$ is *estimable* of degree $(1, 1, \ldots, 1)$ for cdfs $(F_1, F_2, \ldots, F_k)$ of $(X_{1i}, X_{2i}, \ldots, X_{ki})$, in some family of distributions $\mathcal{F}$, if $(X_{1i}, X_{2i}, \ldots, X_{ki})$ is the smallest sample size (one observation from each system) for which there exists an unbiased estimator of $\mathbf{p}$ for every $(F_1, F_2, \ldots, F_k) \in \mathcal{F}$. Formally stated

$$E_{(F_1,\ldots,F_k)}[h^{(j)}(X_{11}, \ldots, X_{k1})] = p_j$$

for $j = 1, 2, \ldots, k$, for a k-sample symmetric kernel $h^{(j)}(\cdot)$. In our case this kernel is

$$h^{(j)}(X_{1i^*}, \ldots, X_{ki^*}) = \prod_{\ell=1;\ell\neq j}^{k} \phi(X_{ji^*} - X_{\ell i^*}), \qquad i^* \in 1, 2, \ldots, v^k, \tag{6}$$

where i^* represents the index of an arbitrary pseudo-replication. Thus, $h^{(j)}(X_{1i^*}, \ldots, X_{ki^*}) = 1$ if the observation from the jth system is the largest in any vector comparison across all systems. A k-sample U-statistic is the average value of such a kernel over all vectors of observations with one observation from each system, which is precisely (3) divided by v^k. Therefore, $\bar{p}_j, j = 1, 2, \ldots, k$ are each k-sample U-statistics, and we can use well-known results for U-statistics to calculate their variances [9]. In Section 4.1. we prove that the variance of the AVC estimator is always less than or equal to the variance of the MLEs. Unfortunately, the variance formula becomes increasingly complex, and its value increasingly difficult to estimate, as k increases. In Section 5, we present a simpler asymptotic approximation to the variance of the AVC estimator which forms the basis for confidence intervals.

4. SMALL-SAMPLE PROPERTIES OF AVC ESTIMATORS

The following analytical study illustrates a number of important properties of our AVC estimator. First, we provide a general proof that the variance of our AVC estimator is no larger than the variance of the MLE. We then illustrate the size of this variance reduction for a specific small-sample case and discuss the weak dependence of this reduction on the underlying distributions of the X_{ji}. Finally, we address the difficulty in obtaining analytical results for even a small number of systems and observations, and thus provide motivation for our asymptotic variance approximation.

4.1. Smaller Variance of the AVC Estimator

To show that AVC provides a point estimator with a smaller variance than the MLE, suppose we have simulation results for k systems with v observations from each system, $X_{ji}, j = 1, 2, \ldots, k, i = 1, 2, \ldots, v$.

Let $\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k$ represent the empirical cumulative distribution functions (edf) for each system based on these data, and suppose we wish to estimate p_1. From Eqs. (1) and (2) the MLE is then $\hat{p}_1 = Y_1/v$ with variance $\mathrm{Var}(\hat{p}_1) = p_1(1 - p_1)/v$.

To arrive at our AVC estimator let

$$p_1(\hat{\mathbf{F}}) = \Pr\{X_{1i^*} > X_{ji^*}, \quad \forall j \neq 1 | \hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k\},$$

where i^* represents the index of an arbitrary pseudo-replication, so that $p_1(\hat{\mathbf{F}})$ is the probability system 1 is the best when the data are distributed as $\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k$. Under $\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k$, each pseudo-replication occurs with an equal probability of $1/v^k$. Thus, we can write

$$p_1(\hat{\mathbf{F}}) = \frac{(\text{the number of pseudo-replications where system 1 wins})}{v^k},$$

implying that the numerator is Z_1 from (3). Therefore, an equivalent representation of our AVC estimator is $\bar{p}_1 = p_1(\hat{\mathbf{F}})$.

Using standard definitions for conditional expectation and variance [3], we know that

$$\begin{aligned}
\mathrm{Var}(\hat{p}_1) &= \mathrm{Var}[\mathrm{E}(\hat{p}_1 | \hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k)] + \mathrm{E}[\mathrm{Var}(\hat{p}_1 | \hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k)] \\
&= \mathrm{Var}[p_1(\hat{\mathbf{F}})] + \mathrm{E}[\mathrm{Var}(\hat{p}_1 | \hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k)] \\
&= \mathrm{Var}[\bar{p}_1] + \mathrm{E}[\mathrm{Var}(\hat{p}_1 | \hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k)].
\end{aligned}$$

Since $\mathrm{E}[\mathrm{Var}(\hat{p}_1|\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_k)] \geq 0$, we have shown that $\mathrm{Var}(\bar{p}_1) \leq \mathrm{Var}(\hat{p}_1)$. In the following section we quantify the size of this reduction in estimator variance for a specific case.

4.2. Example

Without loss of generality, let π_k be the best system. We will assume that all the remaining systems, $\pi_1, \pi_2, \ldots, \pi_{k-1}$ are identically distributed, and let X represent an observation from π_k and W represent an observation from any of the inferior systems.

Consider $X \sim \exp(\lambda)$ and $W \sim \exp(\mu)$ and let $\lambda < \mu$, where $\lambda > 0$ and $\mu > 0$ are exponential rates. This particular example was suggested by Goldsman [7]. For $k = 2$ and $v = 2$, we have $p_2 = \Pr\{X > W\} = \mu/(\lambda + \mu)$ and $p_1 = \Pr\{X < W\} = \lambda/(\lambda + \mu)$. From (2) we have $\mathrm{Var}(\hat{p}_2) = \frac{\lambda\mu}{2(\lambda+\mu)^2}$. In a similar fashion using the derivation in [9] we have

$$\mathrm{Var}(\bar{p}_2) = \frac{3}{4} \frac{\lambda\mu(\lambda^2 + 3\lambda\mu + \mu^2)}{(2\lambda + \mu)(\lambda + 2\mu)(\lambda + \mu)^2}. \tag{7}$$

To illustrate the variance reduction achieved by the AVC estimator, we display the ratio of $\mathrm{Var}(\bar{p}_2)/\mathrm{Var}(\hat{p}_2)$ plotted against $\theta = p_2/p_1$ in Figure 2. We give results for exponential systems with $k = 2$ and $v = 2$ as described above, and also include results for continuous uniform systems from [9] to illustrate the weak dependence of the variance reduction on the distribution of the X_{ji}. A ratio less than 1 indicates a variance reduction with AVC. The reduction in variance is on the order of roughly 20% over a range of θ between 1 and 4. This covers most of the practical range of θ included in standard tables for BEM and used by experimenters. Clearly, as θ increases, both estimators approach 1 and the associated variances approach 0. Asymptotic results in [9] show that the variance reduction will be roughly $(k-1)/(2k-1) \times 100\%$ when the number of replications v is large, implying that as much as a 50% variance reduction can be achieved for the cost of a different calculation of the point estimator.

The weak dependence of the variance reduction on the underlying distribution led us to consider the asymptotic approximation for the variance of an AVC estimator that is presented below.

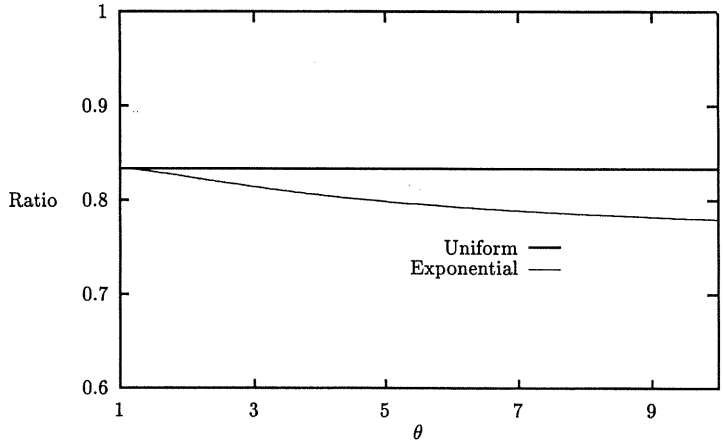

Figure 2. Ratio of $\mathrm{Var}(\bar{p}_2)/\mathrm{Var}(\hat{p}_2)$ for exponential and uniform systems with $k = 2$, $v = 2$.

5. LARGE-SAMPLE PROPERTIES OF AVC ESTIMATORS

The results presented so far (and others in [9]) for small k and v show that the variance of the AVC estimator appears to be weakly distribution dependent; it is also tedious to compute. Here we derive the asymptotic variance of the AVC point estimator, which is much easier to estimate.

Recall that $p_j = \Pr\{X_{ji} > X_{\ell i}, \forall \ell \neq j\}$. Define

$$\begin{aligned} Y_j(v) &= \text{number of wins for system } j \text{ under MLE for fixed } v, \\ Z_j(v) &= \text{number of wins for system } j \text{ under AVC for fixed } v, \end{aligned}$$

which gives us MLE and AVC point estimators $\hat{p}_j(v) = Y_j(v)/v$ and $\bar{p}_j(v) = Z_j(v)/v^k$. Notice that these are the same estimators we defined in (1) and (4), respectively, except that we have explicitly included the dependence on v. Our approach is based on the fact that standardized versions of $\hat{p}$ and $\bar{p}$ are asymptotically normal.

5.1. Asymptotic Variance of MLE and AVC Estimators

Consider the asymptotic behavior of the MLEs. It is well known that as $v \to \infty$

$$\sqrt{v}(\hat{p}_j - p_j) \stackrel{\mathcal{D}}{\Rightarrow} \mathrm{N}(0, p_j(1-p_j)). \tag{8}$$

Thus, the asymptotic variance of $\hat{p}_j$ is $p_j(1-p_j)$.

Similarly, consider the asymptotic behavior of our AVC estimator as v goes to infinity. Let $N = kv$. We will use the following theorem due to Lehmann [8, p. 964], with some of the notation simplified for our context.

THEOREM 1: Let $\bar{p}_a = U_a(X_{11}, \ldots, X_{1v}; \ldots; X_{k1}, \ldots, X_{kv})$ be a k-sample U-statistic based on symmetric kernal $h^{(a)}$ for the parameter p_a of degree $(1, 1, \ldots, 1)$. If $\lim_{N\to\infty}(v/N) = \lambda_i, 0 < \lambda_i < 1$, for $i = 1, \ldots, k$, and if $\mathrm{E}[\{h^{(a)}(X_{11}, \ldots, X_{k1})\}^2] < \infty$, then $\sqrt{N}(\bar{p}_a - p_a)$ has a limiting normal distribution as $N \to \infty$ with mean 0 and variance

$$\sigma_a^2 = \sum_{i=1}^{k} \frac{\xi_{0,\ldots,0,1,0,\ldots,0}}{\lambda_i}, \tag{9}$$

(where we define $\xi_{0,\ldots,0,1,0,\ldots,0}$ below) provided that $\sigma_a^2 > 0$.

Proof: See Lehmann [8].

We now give an expression for $\xi_{0,\ldots,0,1,0,\ldots,0}$. Following Randles and Wolfe [13, pp. 105–106], let i be an integer such that $1 \leq i \leq k$ and define $H_{i1}^{(a)} = h^{(a)}(X_{1\alpha_{11}}, \ldots, X_{k\alpha_{k1}})$ and $H_{i2}^{(a)} = h^{(a)}(X_{1\beta_{11}}, \ldots, X_{k\beta_{k1}})$ where $\alpha_{j1} = \beta_{j1}$ if $j = i$, and $\alpha_{j1} \neq \beta_{j1}$ if $j \neq i$. Then define the covariance terms

$$\begin{aligned} \xi^{(a,a)}_{0,\ldots,0,1,0,\ldots,0} &= \mathrm{Cov}[H_{i1}^{(a)}, H_{i2}^{(a)}] \\ &= \mathrm{E}[H_{i1}^{(a)} H_{i2}^{(a)}] - p_a^2, \end{aligned} \tag{10}$$

where the only 1 in the subscript of $\xi^{(a,a)}_{0,\ldots,0,1,0,\ldots,0}$ is in the ith position and $a \in \{1, 2, \ldots, k\}$. In this notation a represents the system whose parameter we are estimating, and i represents the

only system with a common observation in $H_{i1}^{(a)}$ and $H_{i2}^{(a)}$. Using our kernel from (6), we have for $a = i$

$$\mathrm{E}[H_{a1}^{(a)} \quad H_{a2}^{(a)}] = \Pr\left\{ X_{a\alpha_a} > \max_{\ell \neq a} \{X_{\ell\alpha_\ell}, X_{\ell\beta_\ell}\} \right\} \tag{11}$$

and for $a \neq i$

$$\mathrm{E}[H_{i1}^{(a)} \quad H_{i2}^{(a)}] = \Pr\left\{ \begin{array}{l} X_{a\alpha_a} > X_{i\alpha_i}, X_{a\alpha_a} > \max_{\ell \neq a,i} \{X_{\ell\alpha_\ell}\}; \\ X_{a\beta_a} > X_{i\alpha_i}, X_{a\beta_a} > \max_{\ell \neq a,i} \{X_{\ell\beta_\ell}\} \end{array} \right\}, \tag{12}$$

where $\alpha_\ell \neq \beta_\ell$. This leaves us with just two covariance expressions for each system

$$\mathcal{L}_a^{(a)}(1) = \mathrm{Cov}[H_{a1}^{(a)}, H_{a2}^{(a)}] \tag{13}$$

and

$$\mathcal{L}_i^{(a)}(2) = \mathrm{Cov}[H_{i1}^{(a)}, H_{i2}^{(a)}], \quad a \neq i. \tag{14}$$

Substituting (13) and (14) into (9), we obtain

$$\sigma_j^2 = k \left[\mathcal{L}_j^{(j)}(1) + \sum_{i \neq j} \mathcal{L}_i^{(j)}(2) \right], \qquad j = 1, 2, \ldots, k. \tag{15}$$

Therefore, from Theorem 1 we have

$$\sqrt{N}(\bar{p}_j - p_j) \overset{\mathcal{D}}{\Rightarrow} \mathrm{N}(0, \sigma_j^2), \qquad j = 1, 2, \ldots, k \tag{16}$$

as $N \to \infty$, where $N = kv$ and σ_j^2 is defined in (15).

In comparing σ_j^2 with the asymptotic MLE variance in (8), we notice that the MLE expression is defined as $v \to \infty$ while the AVC expression is defined as $N = kv \to \infty$. In order to compare these asymptotic variances with the exact variance of $\bar{p}_k$ for any k, we divide the asymptotic variance in (8) by v and the asymptotic variance in (15) by N to give

$$\mathrm{Var}_{\mathcal{A}}(\hat{p}_j) = \frac{1}{v} p_j(1 - p_j), \tag{17}$$

$$\mathrm{Var}_{\mathcal{A}}(\bar{p}_j) \approx \frac{1}{v} \left[\mathcal{L}_j^{(j)}(1) + \sum_{i \neq j} \mathcal{L}_i^{(j)}(2) \right]. \tag{18}$$

Notice that the asymptotic MLE variance from (17) is equal to the true MLE variance. The asymptotic AVC variance from (18), however, is an approximation.

Figure 3 plots $v \cdot \mathrm{Var}_{\mathcal{A}}(\bar{p}_2)$, $v \cdot \mathrm{Var}_{\mathcal{A}}(\hat{p}_2)$, and v times the exact variance found in (7) against v for $k = 2$ exponential systems with $\theta = 1.2$. This figure illustrates some important facts about the relationship among the variances of our estimators. First, notice that the AVC variance is

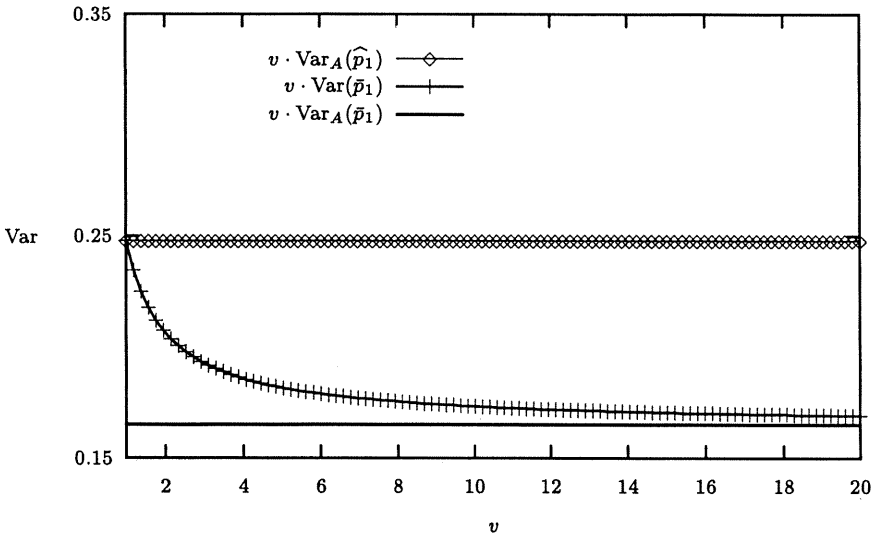

Figure 3. $v \cdot \mathrm{Var}(\hat{p}_2)$, $v \cdot \mathrm{Var}(\bar{p}_2)$ and $v \cdot \mathrm{Var}_{\mathcal{A}}(\bar{p}_2)$ for $k = 2$ exponential systems with $\theta = 1.2$.

smaller than the MLE variance. Also notice how quickly the exact AVC variance approaches the AVC asymptotic approximation. At $v = 15$, the difference is only about 3%. This indicates that the asymptotic approximation for the AVC variance is quite good at relatively small values of v ($v \geq 15$) for $\theta = 1.2$. Calculations for larger values of θ at $v = 15$ show the difference between the exact and approximate AVC variance still about 3% for $\theta = 2.0$ and increasing to a difference of roughly 4% at $\theta = 3.0$. This shows that the $\mathrm{Var}_{\mathcal{A}}(\bar{p}_2)$ provides a better approximation to $\mathrm{Var}(\bar{p}_2)$ when the difference between systems is smaller.

In the following section we describe our methodology for extracting an estimate of (18) from a simulation and using this estimate to calculate confidence intervals for the $p_j s$.

6. VARIANCE ESTIMATORS AND CONFIDENCE INTERVALS

We seek a confidence interval based on the AVC point estimator that is analogous to the asymptotically valid normal-theory confidence interval for the MLE given by

$$\hat{p}_j \pm \mathcal{Z}_{1-\alpha/2} \sqrt{\widehat{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j)}, \tag{19}$$

where $\widehat{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j) = \hat{p}_j(1-\hat{p}_j)/v$ and $\mathcal{Z}_{1-\alpha/2}$ denotes the $1-\alpha/2$ quantile of the standard normal distribution. Thus, we propose an interval of the form

$$\bar{p}_j \pm \mathcal{Z}_{1-\alpha/2} \sqrt{\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)}, \tag{20}$$

where $\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)$ is an estimator of (18). This variance estimator is formed by directly estimating each of the covariance terms (13) and (14). We estimate these covariance terms by first counting the pairs of vectors that meet the conditions from either (11) or (12), and then dividing by the

total number of vector pairs with a common element from one system. The final step involves subtracting the square of the estimate for p_a as shown in (10).

Specifically, we estimate $\mathrm{E}[H_{j1}^{(j)}\ H_{j2}^{(j)}]$ by

$$\hat{\mathrm{E}}[H_{j1}^{(j)}\ \ H_{j2}^{(j)}] = \frac{2}{v^k(v-1)^{k-1}} \sum \mathcal{I}\left\{X_{j\alpha_j} > \max_{\ell \neq j}\{X_{\ell\alpha_\ell}, X_{\ell\beta_\ell}\}\right\},$$

where $\mathcal{I}$ is the indicator function, and the sum is over all indices such that $\alpha_\ell \neq \beta_\ell$. Similarly, we estimate $\mathrm{E}[H_{i1}^{(j)}\ H_{i2}^{(j)}]$ for $i \neq j$ by

$$\hat{\mathrm{E}}[H_{i1}^{(j)}\ \ H_{i2}^{(j)}] = \frac{2}{v^k(v-1)^{k-1}} \sum \mathcal{I}\left\{ \begin{array}{l} X_{j\alpha_j} > X_{i\alpha_i}, X_{j\alpha_j} > \max_{\ell \neq j,i}\{X_{\ell\alpha_\ell}\}; \\ X_{j\beta_j} > X_{i\alpha_i}, X_{j\beta_j} > \max_{\ell \neq j,i}\{X_{\ell\beta_\ell}\} \end{array} \right\}.$$

And finally, we estimate p_j by $\bar{p}_j$. Estimating all the $\bar{p}_j$'s and their associated variances using the efficient algorithm in the Appendix is an $O(kv \ln kv)$ calculation. Notice, however, that $\widehat{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j)$ could be used to provide a conservative confidence interval for the AVC estimators if avoiding this variance calculation is desirable.

The proof that the AVC confidence intervals are asymptotically valid as v goes to infinity requires proving that the variance estimator $\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)$ is consistent using a standard, but very tedious, argument presented in [9]. In the next section we provide empirical results to evaluate the variance estimator and associated confidence interval at realistic sample sizes.

7. EMPIRICAL STUDY

For our empirical study we focus on the variance reduction achieved by AVC, and the performance of the approximate AVC variance estimator and confidence interval, all relative to the performance of the MLE. All results are computed for $p_k = p_{[k]}$ (the largest p_j) using exponential system distributions as described in Section 4.2. This is a small portion of an extensive empirical study that is presented in [9].

We performed $M = 10{,}000$ macroreplications of kv observations each, for $k = 2, 3$ and $v = 20, 50, 100$ and 200. Let $\hat{p}_{j\ell}, \bar{p}_{j\ell}, \widehat{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j)_\ell$, and $\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)_\ell$ represent the point and variance estimates from the ℓth macroreplication.

Since the sample variance is an unbiased estimator of the true variance, we estimate the true variance of the AVC estimator by the sample variance across the macroreplications, specifically $\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_j) = \frac{1}{M-1}\sum_{\ell=1}^{M}(\bar{\bar{p}}_j - \bar{p}_{j\ell})^2$, where $\bar{\bar{p}}_j = \frac{1}{M}\sum_{\ell=1}^{M}\bar{p}_{j\ell}$. We also calculate averages of the variance estimates across macroreplications as $\overline{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j) = \frac{1}{M}\sum_{\ell=1}^{M}\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)_\ell$ and $\overline{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j) = \frac{1}{M}\sum_{\ell=1}^{M}\widehat{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_j)_\ell$. We then estimate the relative bias of our AVC variance estimator by

$$\widehat{\mathrm{RelativeBias}}(\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j)) = \frac{\overline{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j) - \widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_j)}{\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_j)}$$

and the variance reduction by $\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_k)/\overline{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_k)$.

Table 1 presents the variance results for each method for $k = 2, 3$ systems and a number of different values of v. We find a significant negative bias in our AVC variance estimator at $v = 20$ for all k. At $v = 50$ this bias drops to about 6% for $k = 2$ and less than 10% for $k = 3$. Moving up to $v = 100$, the bias for both $k = 2$ and 3 drops below 7%. At $v = 200$ the bias effectively goes

Table 1. AVC and MLE variance for p_k.

k	θ^*	v	$\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_k)$	$\overline{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_k)$	$\overline{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_k)$	Relative Bias $(\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_k))$	$\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_k)/\overline{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_k)$
2	1.2	20	.0086	.0074	.0118	−.1395	.7288
	1.2	50	.0034	.0032	.0049	−.0588	.6939
	1.2	100	.0017	.0016	.0025	−.0588	.6800
	1.2	200	.0008	.0008	.0012	.0000	.6667
2	2.0	20	.0076	.0065	.0105	−.1447	.7238
	2.0	50	.0030	.0028	.0044	−.0667	.6818
	2.0	100	.0015	.0014	.0022	−.0667	.6818
	2.0	200	.0007	.0007	.0011	.0000	.6364
3	1.2	20	.0075	.0061	.0111	−.1867	.6757
	1.2	50	.0030	.0027	.0046	−.1000	.6522
	1.2	100	.0015	.0014	.0023	−.0667	.6522
	1.2	200	.0007	.0007	.0012	.0000	.5833
3	2.0	20	.0085	.0069	.0119	−.1882	.7143
	2.0	50	.0034	.0031	.0049	−.0882	.6939
	2.0	100	.0017	.0016	.0025	−.0588	.6800
	2.0	200	.0008	.0008	.0012	.0000	.6667

to zero for all k. These results indicate that we may underestimate $\mathrm{Var}(\bar{p}_k)$ by using $\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_k)$ for $v \leq 50$. The last column of Table 1 shows the ratio $\widehat{\mathrm{Var}}_{\mathcal{S}}(\bar{p}_k)/\overline{\mathrm{Var}}_{\mathcal{A}}(\hat{p}_k)$ to illustrate the variance reduction with the AVC estimator. Focusing on the results for $v = 200$ at $\theta^* = 1.2$, we see roughly a 33% reduction at $k = 2$ and roughly a 42% reduction at $k = 3$.

Table 2. AVC and MLE 95% confidence intervals for p_k.

			AVC		MLE	
k	θ^*	v	Percentage coverage	Average width	Percentage coverage	Average width
2	1.2	20	.92	.34	.92	.42
	1.2	50	.94	.22	.93	.27
	1.2	100	.94	.16	.94	.19
	1.2	200	.95	.11	.94	.14
2	2.0	20	.91	.31	.92	.40
	2.0	50	.94	.21	.95	.26
	2.0	100	.94	.15	.96	.18
	2.0	200	.95	.10	.95	.13
3	1.2	20	.90	.31	.94	.41
	1.2	50	.93	.20	.94	.27
	1.2	100	.94	.15	.95	.19
	1.2	200	.95	.10	.94	.13
3	2.0	20	.91	.32	.96	.43
	2.0	50	.93	.22	.93	.27
	2.0	100	.94	.16	.94	.20
	2.0	200	.95	.11	.95	.14

Table 3. MLE and AVC individual 95% confidence intervals for airline reservation problem.

	v	MLE	MLE width	AVC	AVC width	AVC reduction
p_1	667	[0.35869, 0.43291]	0.07422	[0.36188, 0.42046]	0.05858	21%
p_2	667	[0.29561, 0.36705]	0.07144	[0.30495, 0.36047]	0.05552	22%
p_3	667	[0.23906, 0.30667]	0.06761	[0.25080, 0.30145]	0.05065	25%
p_1	1101	[0.35815, 0.41569]	0.05754	[0.35847, 0.40375]	0.04527	21%
p_2	1101	[0.30371, 0.35932]	0.05561	[0.31321, 0.35649]	0.04328	22%
p_3	1101	[0.25499, 0.30813]	0.05313	[0.26401, 0.30406]	0.04004	25%

To compare confidence intervals, we constructed 95% confidence intervals for p_k from each macroreplication using (19) and (20). We then counted how many of the M intervals formed with each method captured p_k and computed the average confidence interval width. Looking at the 95% confidence interval results in Table 2, we see undercoverage and relatively large intervals for $v \leq 50$ for both MLE and AVC. This indicates that the normal approximation is not particularly good for either method at small values of v. At $v = 100$ both MLE and AVC coverages jump up to 94% or 95% in nearly all cases, with MLE coverage slightly better than the AVC coverage. However, at $v = 200$ the AVC coverage slightly exceeds the MLE coverage, with both methods achieving 95% coverage in almost all cases. In all cases the average AVC interval width shows a 20–25% reduction over the average MLE interval width.

8. A COMPUTER SYSTEM DESIGN PROBLEM

To illustrate the MSP approach to selecting the best system, and the AVC estimators of the probability that each system will be the best, we use a modification of the airline-reservation system example used by Goldsman, Nelson, and Schmeiser [6] to illustrate finding the system with the best mean performance.[2] The example consists of $k = 3$ different airline-reservation systems where the single measure of performance is the time to failure, TTF, so that larger is better. Each system works if either of two computers work. Rare computer failures and short repair times typically result in a large TTF for each system. The differences in the three systems result from varying parameters affecting the time-to-failure and time-to-repair distributions. From experience we know that the average TTFs are roughly 100,000 minutes (about 70 days) for all three systems. We are interested in which airline reservation system is most likely to have the longest TTF; in other words, which system is most likely to go the longest without a failure. Thus, p_j is the probability that system j has the longest TTF, for $j = 1, 2, 3$.

Suppose that we want to design our experiment so that we are 95% certain of finding the system with the largest p_j when $\theta = p_{[3]}/p_{[2]} \geq 1.15$; that is, we want to find the system with the largest probability of having the longest TTF when that probability is at least 15% larger than the second largest probability. Using FORTRAN code provided by Goldsman [7], we find that $v = 1101$ vector observations are required to guarantee a PCS of 95% ($P^* = 0.95$) using procedure BEM. Based on the large-sample approximation in [11], only $v = 667$ vector observations are required by AVC to provide the same PCS. Of course, in a real application we would only take 667 vector

[2]Our modification consists of dropping system 2 from the original example, a system which was shown to be clearly inferior.

Table 4. MLE and AVC simultaneous 95% confidence intervals for airline reservation problem.

	v	MLE	MLE width	AVC	AVC width	AVC reduction
p_1	667	[0.35017, 0.44143]	0.09127	[0.35515, 0.42718]	0.07203	21%
p_2	667	[0.28741, 0.37526]	0.08785	[0.29857, 0.36685]	0.06828	22%
p_3	667	[0.23130, 0.31443]	0.08313	[0.24498, 0.30726]	0.06228	25%
p_1	1101	[0.35155, 0.42230]	0.07075	[0.35328, 0.40895]	0.05567	21%
p_2	1101	[0.29732, 0.36571]	0.06839	[0.30824, 0.36146]	0.05322	22%
p_3	1101	[0.24890, 0.31423]	0.06533	[0.25942, 0.30865]	0.04923	25%

observations and use AVC, but to illustrate the benefits of AVC we also perform the larger number of replications here.[3]

For either sample size, and using either estimator $\hat{p}_j$ or $\bar{p}_j$, system 1 has the largest sample probability of being the best, and would therefore be selected. Using the results in this paper, we can also provide confidence intervals for the p_js. Confidence intervals for each p_j give the decision maker a clear quantitative measure of uncertainty in these estimates and present this information in a manner to identify the statistical and/or practical difference in the probabilities of competing systems to provide the best performance. Tables 3 and 4 give individual and simultaneous 95% confidence intervals, respectively, for both sample sizes, where the simultaneous intervals were constructed by applying Bonferroni's inequality with an equal coverage probability $1 - \alpha/3 = 1 - 0.05/3 \approx 0.98$ for each individual interval (implying an overall coverage probability greater than or equal to 0.95). At over 600 vector observations the normal approximation should be adequate for both types of intervals.

The widths of the AVC intervals are shorter than the MLE intervals, from 21–25% shorter when the two estimators are based on the same sample size. Notice that the individual confidence intervals (Table 3) based on the MLEs have overlapping endpoints at both sample sizes, while the individual intervals based on AVC do not overlap. For the Bonferroni simultaneous intervals (Table 4) both the MLE and AVC intervals have overlapping endpoints, but the AVC intervals are more nearly disjoint and would clearly become disjoint at a much smaller sample size than required for the MLE intervals. Thus, AVC can provide conclusive ranking of the systems with less data than MLE.

9. CONCLUSIONS

We have presented an efficient method for comparing performance among competing simulated systems when the probability of being the best is a more appropriate measure than long-run average performance. Our results prove the variance of the AVC estimator is no larger than the variance of the MLE and we demonstrate the magnitude of the reduction in variance for a specific example. We expand results presented in [11] to derive the asymptotic variance of the AVC estimator. This AVC asymptotic variance is shown to be significantly smaller than the equivalent MLE asymptotic variance directly, and through the construction of confidence intervals for a realistic example. These results illustrate how AVC makes more efficient use of the data already available to provide a more precise set of estimators for the multinomial success probabilities.

[3]For the smaller sample size we simply used the first 667 vectors from the larger sample.

APPENDIX

In this appendix we present formulas for the AVC point and variance estimators that are suitable for implementation in computer algorithms. The formulas, which are appropriate when the output of interest is such that ties are not possible, are presented without proof; the proofs are tedious and available from the authors on request.

Let X_{ji} be the ith observation from system j, for $i = 1, 2, \ldots, v$ and $j = 1, 2, \ldots, k$. As above, we assume that bigger is better. Define

$$N_{\ell j}(i) = \#\{\alpha : 1 \le \alpha \le v \text{ and } X_{\ell\alpha} < X_{ji}\}$$

for $i = 1, 2, \ldots, v, j, \ell = 1, 2, \ldots, k$, but $\ell \neq j$. In other words, $N_{\ell j}(i)$ is the number of observations from system ℓ that are beaten by X_{ji}.

Our AVC point and variance estimators can be written as

$$\bar{p}_j = \frac{1}{v^k} \sum_{h=1}^{v} \prod_{\ell \neq j} N_{\ell j}(h) = \frac{1}{v^k} A_j$$

and

$$\begin{aligned}
&\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_h) \\
&= \frac{1}{v^{k+1}(v-1)^{k-1}} \left\{ \sum_{h=1}^{v} \prod_{\ell \neq j} N_{\ell j}(h)(N_{\ell j}(h) - 1) \right. \\
&\quad \left. + 2 \sum_{i \neq j} \left[\sum_{h=1}^{v-1} \sum_{m=h+1}^{v} N_{ij}(m) \prod_{\ell \neq i,j} N_{\ell j}(m)(N_{\ell j}(h) - 1) \right] \right\} - \frac{k}{v} \bar{p}^2 \\
&= \frac{1}{v^{k+1}(v-1)^{k-1}} \{B_j + 2C_j\} - \frac{k}{v} \bar{p}^2
\end{aligned}$$

These estimators can be computed in $O(kv \ln kv)$ time using the algorithm presented below. The algorithm is based on sorting all kv observations X_{ji} *together*, but maintaining an indicator for which system generated each observation. Appropriate sorting algorithms include *Heapsort* or *Quicksort* [12]. After sorting the data, only one pass through it is required to compute all k point and variance estimates.

Key to developing this efficient algorithm is the fact that B_j and C_j can be written as

$$B_j = \sum_{h=1}^{v} \left\{ \prod_{\ell \neq j} N_{\ell j}(h) \right\} \left\{ \prod_{\ell \neq j} (N_{\ell j}(h) - 1)^+ \right\}$$

and

$$C_j = \sum_{h=1}^{v-1} \left\{ \sum_{i \neq j} \prod_{\ell \neq i,j} (N_{\ell j}(h) - 1)^+ \right\} \left\{ \sum_{m=h+1}^{v} \prod_{\ell \neq j} N_{\ell j}(m) \right\}$$

where $y^+ = \max\{0, y\}$. The proof involves nothing more than tedious algebra.

Let $\mathbf{X}$ be a $kv \times 1$ vector of all the data, sorted from largest to smallest, and let $\mathbf{I}$ be a $kv \times 1$ vector such that $\mathbf{I}(\ell) = j$ if $\mathbf{X}(\ell)$ is an observation from system j. As the main loop of the algorithm progress from smallest to largest observation in $\mathbf{X}$, the variable e_j counts the number of observations from system j that have been encountered thus far; therefore, at any index r in the main loop, e_j gives the number of observations from system j that are beaten by (smaller than) $\mathbf{X}(r)$, which is an observation from system $\mathbf{I}(r)$. Also, at every index r we update the product

$$\Delta = \prod_{\ell=1}^{k} e_\ell$$

so that

$$\frac{\Delta}{e_{\mathbf{I}(r)}} = \prod_{\ell \neq \mathbf{I}(r)} e_\ell = \prod_{\ell \neq \mathbf{I}(r)} N_{\ell,\mathbf{I}(r)}(v - e_{\mathbf{I}(r)} + 1).$$

The quantity Δ^- maintains a similar product of $N_{\ell,\mathbf{I}(r)}(v - e_{\mathbf{I}(r)} + 1) - 1$ terms. These products are essentially all that we require to compute A_j, B_j and C_j.

`Algorithm AVC Calc`

1. sort the $kv \times 2$ array $\mathbf{W} = [\mathbf{X}, \mathbf{I}]$ from largest to smallest by $\mathbf{X}$
2. set $e_j = A_j = B_j = C_j = D_j = 0$ for $j = 1, 2, \ldots, k$
   set $S = 0$ and $\Delta = \Delta^- = 1$
3. for $r = kv$ to 1 by -1 do
   (i) $e_{\mathbf{I}(r)} = e_{\mathbf{I}(r)} + 1$
   (ii) if $e_{\mathbf{I}(r)} > 1$ then
   $\Delta = \Delta \times \frac{e_{\mathbf{I}(r)}}{e_{\mathbf{I}(r)-1}}$
   if $e_{\mathbf{I}(r)} = 2$ then $S = S + 1$
   endif
   (iii) if $e_{\mathbf{I}(r)} > 2$ then
   $\Delta^- = \Delta^- \times \frac{e_{\mathbf{I}(r)-1}}{e_{\mathbf{I}(r)-2}}$
   $S = S - \frac{1}{e_{\mathbf{I}(r)-2}} + \frac{1}{e_{\mathbf{I}(r)-1}}$
   endif
   (iv) if $e_j > 0$ for all j then
   $D_{\mathbf{I}(r)} = A_{\mathbf{I}(r)}$
   $A_{\mathbf{I}(r)} = A_{\mathbf{I}(r)} + \frac{\Delta}{e_{\mathbf{I}(r)}}$
   endif
   (v) if $e_j > 1$ for all $j \neq \mathbf{I}(r)$ then
   $B_{\mathbf{I}(r)} = B_{\mathbf{I}(r)} + \frac{\Delta}{e_{\mathbf{I}(r)}} \times \frac{\Delta^-}{\max\{1, e_{\mathbf{I}(r)-1}\}}$
   endif
   (vi) if $e_j \geq 1$ for all $j, e_{\mathbf{I}(r)} > 1$ and $\#\{\ell : \ell \neq \mathbf{I}(r), e_\ell = 1\} \leq 1$ then
   $C_{\mathbf{I}(r)} = C_{\mathbf{I}(r)} + \frac{\Delta^-}{e_{\mathbf{I}(r)-1}} \times (S - \frac{1}{e_{\mathbf{I}(r)-1}}) \times D_{\mathbf{I}(r)}$
   elseif $e_j \geq 1$ for all $j, e_{\mathbf{I}(r)} = 1$ and $\#\{\ell : \ell \neq \mathbf{I}(r), e_\ell = 1\} \leq 1$ then
   $C_{\mathbf{I}(r)} = C_{\mathbf{I}(r)} + \Delta^- \times S \times D_{\mathbf{I}(r)}$
   endif
   loop
4. for $j = 1$ to k do
   $\bar{p}_j = \frac{A_j}{v^k}$
   $\widehat{\mathrm{Var}}_{\mathcal{A}}(\bar{p}_j) = \frac{1}{v^{k+1}(v-1)^{k-1}}\{B_j + 2C_j\} - \frac{k}{v}\bar{p}_j^2$
   loop

REMARK: `Algorithm AVC Calc` is written for ease of exposition rather than for direct implementation. In particular, the tests in steps 3(iv), 3(v), and 3(vi) can be implemented in such a way that only one comparison, rather than k comparisons, are required to determine if the condition is satisfied.

ACKNOWLEDGMENTS

This research was partially supported by National Science Foundation Grant No. DMI-9622065, and by Rockwell Software and Symix/Pritsker Corporation. The authors gratefully acknowledge the assistance of Leonardo Lopes in developing the efficient implementation of the AVC algorithm.

DISCLAIMER

The views expressed in this article are those of the authors and do not reflect the official policy of the U.S. Air Force, Department of Defense, or the U.S. Government.

REFERENCES

[1] R.E. Bechhofer, S. Elmaghraby, and N. Morse, A single-sample multiple-decision procedure for selecting the multinomial event which has the highest probability, Ann Math Stat 30 (1959), 102–119.

[2] R.E. Bechhofer and M. Sobel, Non-parametric multiple-decision procedures for selecting that one of k populations which has the highest probability of yielding the largest observation (preliminary report), Abstract, Ann Math Stat 29 (1958), 325.

[3] G. Casella and R.L. Berger, Statistical inference, Duxbury Press, Pacific Grove, California, 1990.

[4] D. Goldsman, A multinomial ranking and selection procedure: Simulation and applications, Proceedings of the 1984 Winter Simulation Conference, S. Shepard, U.W. Pooch, and C.D. Pegden (Editors), The Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1984, pp. 259–264.

[5] D. Goldsman, On selecting the best of K systems: An expository survey of indifference-zone multinomial procedures, Proceedings of the 1984 Winter Simulation Conference, S. Shepard, U.W. Pooch, and C.D. Pegden (Editors), The Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1984, pp. 107–112.

[6] D. Goldsman, B.L. Nelson, and B. Schmeiser, Methods for selection the best system, Proceedings of the 1991 Winter Simulation Conference, B. Nelson, W.D. Kelton, and G. Clark (Editors), The Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1991, pp. 177–186.

[7] D. Goldsman, personal communication, 1995.

[8] E.L. Lehmann, Robust estimation in analysis of variance, Ann Math Stat 34 (1963), 957–966.

[9] J.O. Miller, Efficient multinomial selection in simulation, unpublished Ph.D. Thesis, Department of Industrial, Welding, and Systems Engineering, The Ohio State University, Columbus, 1997.

[10] J.O. Miller and K.W. Bauer, How common random numbers affect multinomial selection, Proceedings of the 1997 Winter Simulation Conference, S. Andradottir, K.J. Healy, D.H. Withers, and B.L. Nelson, The Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1997, pp. 342–347.

[11] J.O. Miller, B.L. Nelson, and C.H. Reilly, Efficient multinomial selection in simulation, Nav Res Logistics 45 (1998), 459–482.

[12] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, Numerical recipes: The art of scientific computing (FORTRAN version), Cambridge University Press, Cambridge, 1989.

[13] H.R. Randles and D.A. Wolfe, Introduction to the theory of nonparametric statistics, Wiley, New York, 1979.