

SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK

TECHNICAL REPORT NO. 596

November, 1983

ON SELECTING THE BEST OF K
SIMULATED SYSTEMS:
AN EXPOSITORY SURVEY

by

David Goldman
and
Lee Schruben

Abstract

We investigate the problem of selecting the 'best' one of k arbitrary systems or alternatives. Consider one observation from each of the k systems. By 'best,' we mean that system which has the highest probability of yielding the 'most desirable' of the k observations. The term 'most desirable' is defined according to some criterion of goodness determined by the experimenter. We show that this problem can be formulated as a multinomial selection problem. Hence, multinomial selection procedures are, in a sense, nonparametric procedures. An up-to-date survey of such multinomial procedures is given. Further, we describe how some of these procedures can be adapted for use in the simulation environment.

1. Introduction

Consider k different competing populations (or systems or alternatives, etc). A natural question to ask is: Which of these k systems is 'best?' By 'best' system, we could informally mean, e.g.

- that one of k inventory policies which maximizes profit,
- that one of k scales which is the most precise, or
- that one of k computer systems which has the greatest availability.

Thus, 'best' can take on a variety of meanings depending on the practical problem at hand.

Denote the k populations (sources of observations) as $\Pi_1, \Pi_2, \dots, \Pi_k$, respectively. Suppose we take independent vector-observations (X_1, X_2, \dots, X_k) , where X_i is from Π_i , $i=1, \dots, k$. Further, for $i=1, \dots, k$, denote

$$p_i = P\{X_i \text{ is the 'most desirable' of } X_1, X_2, \dots, X_k\}.$$

The term 'most desirable' must be defined according to some criterion of goodness determined by the experimenter. Assume that nothing is known beforehand concerning the values of the p_i 's. Obviously, that Π_i associated with the largest of the p_i 's is the population which has the *highest probability of yielding the 'most desirable' observation* (of those observations from the k -vector). In this paper, our goal will be to find that Π_i associated with the largest of the p_i 's. We refer to that Π_i as the 'best' population.

In order to motivate this definition, consider a simple example. Let A and B be two (s, S) inventory policies. Profit is taken to be the criterion of desirability. Suppose that

$$\begin{array}{llll} \text{Profit from } A & = & 1000 & \text{with probability } 0.001 \\ & = & 0 & \text{" " " } 0.999 \end{array}$$

and

Profit from B = 0.999 with probability 1.

Clearly, $E(\text{Profit from A}) = 1 > 0.999 = E(\text{Profit from B})$; i.e., A gives the higher average profit. However, $P\{\text{Profit from B} > \text{Profit from A}\} = 0.999$; therefore, B gives the higher profit *almost all of the time*. For this reason, the experimenter might justifiably consider policy B to be better than policy A.

Hence, it is meaningful to consider as 'best' the policy which will most likely produce the 'most desirable' observation.

This goal of finding the 'best' population can be viewed as that of finding that cell of a k -nomial distribution with the largest underlying probability. Suppose that we take one observation from each of the k populations. Award a one (a 'success') to the Π_i corresponding to the 'most desirable' of these k observations (use randomization if necessary.) Award a zero to the remaining $k-1$ Π_i 's. Clearly, this is the same as taking an observation from a multinomial distribution with cell probabilities p_1, \dots, p_k .

Thus, the problem of finding the 'best' one of k arbitrary populations can be formulated as a problem of finding that one category of a k -nomial distribution with the highest underlying 'success' probability. This implies that any procedure which finds the multinomial cell associated with the largest probability is a *nonparametric* procedure. Since most real-life systems do not follow one of the 'usual' probability distributions, such nonparametric procedures are seen to be very useful. We group these nonparametric procedures under the heading of *multinomial selection procedures*. Additional motivation for the above arguments can be found in Bechhofer and Sobel (1958).

In Section 2 of this paper, we give a brief summary of the

pertinent notation and terminology. In Section 3, some of the existing selection procedures are presented. Section 4 is concerned with applications to simulation.

2. Background

We now introduce notation and terminology which will be useful for investigating the problem of finding the multinomial cell which has the largest cell probability. Suppose that we take independent observations sequentially from a k -nomial distribution with cell probabilities p_1, p_2, \dots, p_k , $p_i \geq 0$, $\sum p_i = 1$, until some stopping criterion (several of which will be given in the sequel) is met. Most of the procedures which we will study take observations (up to a limit, perhaps) until one cell has 'significantly more' successes than the other cells. In this case, the stopping criteria call for the termination of the procedures.

Denote $x_{i,t}$ as the number of observations from cell i after t multinomial observations (or 'stages') have been taken, $i=1, \dots, k$; $t=1, 2, \dots$. Further, denote $p_{[1]} \leq p_{[2]} \leq \dots \leq p_{[k]}$ as the ordered p_i 's and $x_{[1],t} \leq \dots \leq x_{[k],t}$ as the ordered $x_{i,t}$'s. Assume that we have no a priori knowledge as to how the $p_{[i]}$'s are paired with the multinomial cells.

Our goal is to select as best that cell which is associated with $p_{[k]}$, the largest probability. If the cell corresponding to $p_{[k]}$ is actually chosen, we say that a *correct selection* (CS) has been made. Also, it is desired that the *probability of correct selection* ($P\{CS\}$) be at least P^* whenever $\theta^* p_{[k-1]} \leq p_{[k]}$, where $\{P^*, \theta^*\}$ is pre-specified by the user (with $1 < \theta^* < \infty$ and $1/k < P^* < 1$). Define $\Omega_{\theta^*} = \{p | \theta^* p_{[k-1]} \leq p_{[k]}\}$. We call Ω_{θ^*} the *preference zone* and $\Omega_{\theta^*}^c$ the *indifference-zone*. [Multinomial procedures such as those to be considered below fall under the classification of so-called *indifference-zone selection procedures*. Another rich family of selection procedures employs the so-called *subset approach*; this

approach will not be emphasized here. The reader should refer to Gupta and Panchapakesan (1979) for material concerning the indifference-zone and subset methodologies.]

We will consider the following configuration of $p_{[i]}$'s as a benchmark for comparison among procedures:

$$p_{[k]} = \theta^* p_{[i]}, \quad i=1, \dots, k-1 \quad (\text{SC})$$

I.e., $p_{[i]} = (k-1+\theta^*)^{-1}$, $i=1, \dots, k-1$; $p_{[k]} = \theta^* (k-1+\theta^*)^{-1}$. SC stands for *slippage configuration* (with slippage factor θ^*). For some sampling procedures (cf: Bechhofer, Elmaghraby, and Morse (1959)), this configuration of $p_{[i]}$'s minimizes the $P\{CS\}$ over $p \in \Omega_{\theta^*}$. In this case, the SC is called the *least-favorable configuration* (LFC). Informally, the LFC can be viewed as a 'worst case' configuration (given that $p \in \Omega_{\theta^*}$). It is not known whether the SC is the LFC for all of the multinomial procedures to be presented in the sequel. However, this is a reasonable conjecture; we shall treat the SC as if it is the LFC. Since we desire $P\{CS\} \geq P^*$ for all configurations $p \in \Omega_{\theta^*}$, then (assuming the conjecture to be true) we can equivalently require that $P\{CS | p=SC\} \geq P^*$. Hence, it is meaningful to investigate the SC.

Another interesting configuration is the *equal probability configuration* (EPC), where $p_i = 1/k$ for all i . Of course, the term 'correct selection' is now meaningless; but the EPC is useful as another benchmark in that we would expect such a configuration to maximize a multinomial procedure's expected sample size (i.e., the expected number of multinomial observations needed before the termination criterion is met). Denote the sample size for a procedure P as S_P . $E(S_P)$ is the expected sample size.

Ideally, we wish to find a procedure which guarantees $P\{CS|p=SC\} \geq P^*$ but which is also parsimonious with observations; that is, $E(S_p|p=SC)$ and $E(S_p|p=EPC)$ should be 'low.'

3. Some Multinomial Procedures

In this section, we concentrate on indifference-zone procedures for selecting the multinomial cell which has the largest probability. Recall that when using the indifference-zone approach, the experimenter must pre-specify two constants, P^* and θ^* . The procedures to be discussed below insure that

$$P\{CS|\theta^* P_{[k-1]} \leq P_{[k]}\} \geq P^*, \quad (PR)$$

where PR stands for *probability requirement*. For all of these procedures, we establish the following conventions:

- All observations are independent multinomial observations.
- T is defined to be the stage at which the procedure in question terminates sampling. T may be a random variable.
- We will choose as best that cell corresponding to $x_{[k],T}$ (using randomization if necessary).

3.1 A single-sample procedure

The first procedure we consider is that of Bechhofer, Elmaghraby, and Morse (1959), denoted as P_{BEM} .

Procedure P_{BEM} :

1. Specify k , P^* , and θ^* .
2. Take N_{BEM} observations, where $N_{BEM} = N_{BEM}(k, P^*, \theta^*)$ is to be found in the tables of BEM (1959). N_{BEM} is the number of multinomial observations which must be taken in order to satisfy the PR. //

Remarks 3.1:

1. Kesten and Morse (1959) prove that the SC is the LFC.
2. In P_{BEM} , the number of observations we take is *fixed* at N_{BEM} . For this reason, the procedure is said to be a *fixed-sample* or *single-sample procedure*.

Example 3.1.1:

Suppose that $k=3$ and that we specify $P^*=0.75$ and $\theta^*=3$.

Use the following table (abstracted from BEM (1959)) in order to find N_{BEM} .

N_{BEM}	θ^*	1.1	1.5	2.0	3.0
1		.355	.429	.500	.600
2		.355	.429	.500	.600
3		.362	.464	.563	.696
4		.367	.484	.594	.734
5		.370	.496	.617	.769
6		.374	.515	.646	.804

Table 1 (for P_{BEM}):
 $P\{CS|k=3, p=LFC\}$ for selected θ^* and N_{BEM}

Reading down the $\theta^*=3.0$ column, we see that $N_{BEM}=5$ is the smallest value of N_{BEM} which achieves the PR (Note that owing to the discrete nature of the multinomial distribution, P_{BEM} overshoots slightly the desired $P^*=0.75$.) Hence, if we take

5 observations, the PR will be guaranteed. //

3.2 P_{BK} , an improved version of P_{BEM}

By considering the following example, it becomes apparent that P_{BEM} is sometimes wasteful with observations.

Example 3.2.1:

Suppose that $k=2$, $N_{BEM}=7$, and $x_5=(x_{1,5}, x_{2,5})=(4,1)$.

Obviously, it is impossible to terminate sampling with $x_{1,T} < x_{2,T}$.

In other words, there is no chance for cell 2 to be chosen.

Since cell 1 is guaranteed to be the victor regardless of the remaining two observations, we should stop sampling at $T=5$. //

With this example in mind, we compare two procedures, the latter due to Bechhofer and Kulkarni (1983).

Procedure P_{BEM} :

1. Specify k and N .
2. Take N observations. //

Procedure P_{BK} :

1. Specify k and N .
2. Take observations until either
 - A. The stage $t=N$ or
 - B. $x_{[k],t} - x_{[k-1],t} = N-t$ (Stop sampling if the cell(s) with the second largest number of observations can only tie the cell corresponding to $x_{[k],t}$, even if the remaining $N-t$

observations are taken.) //

Remarks 3.2:

1. Note that P_{BK} is a *sequential* procedure.
2. It is clear that $E(S_{P_{BK}}) \leq E(S_{P_{BEM}})$.
3. Bechhofer and Kulkarni show that $P\{CS|P_{BEM}\} = P\{CS|P_{BK}\}$.
Thus, P_{BK} preserves the $P\{CS\}$ of the less parsimonious procedure, P_{BEM} . Hence, we can use the more efficient P_{BK} with no loss of $P\{CS\}$.

Example 3.2.2:

Let $k=3$, $P^*=0.75$, and $\theta^*=3$. Then $E(S_{P_{BEM}}) = N_{BEM} = 5$. It is straightforward (but tedious) to show that $E(S_{P_{BK}}) = 3.95$ in the LFC. //

3.3 A sequential procedure due to Ramey and Alam (1979)

Procedure P_{RA} :

1. Specify k , P^* , θ^* .
2. Take observations until either
 - A. $x_{[k],t} = N$ or
 - B. $x_{[k],t} - x_{[k-1],t} = r$, where r and N are determined by k , P^* , and θ^* , and are to be found in tables for certain k , P^* , and θ^* (NB: See Remarks below.) //

Remarks 3.3:

1. Ramey and Alam's tables actually contain a number of errors;

- the user is advised to consult Bechhofer and Goldsman (1984a).
2. The number of observations which P_{RA} takes is *bounded* by $kN-k+1$.
 3. It is not known whether the SC is the LFC for all k for P_{RA} , but we will make the reasonable assumption that this is the case.
 4. r and N are determined in such a way that the PR is satisfied and $E(S_{P_{RA}} | p=LFC)$ is minimized over the (r, N) grid.
 5. P_{RA} is not directly comparable to P_{BK} . However, for most choices of k , P^* , and θ^* , it seems that P_{RA} requires fewer observations (on the average) than P_{BK} .

Example 3.3.1:

Again, let $k=3$, $P^*=0.75$, and $\theta^*=3$. We abstract a small portion of the necessary (corrected) tables for P_{RA} from Bechhofer and Goldsman (1984a).

P^*	θ^*	r	N	$P\{CS\}$	$E(S)$
.75	3.0	2	3	.796	3.68
.75	2.4	2	5	.760	4.70
.75	2.0	4	5	.756	8.80
.75	1.6	4	12	.757	18.24

Table 2 (for P_{RA}):
 $P\{CS|k=3, p=LFC\}$, $E(S|.)$ for various P^* , θ^*

We see that if $r=2$ and $N=3$ are chosen, a $P\{CS\}$ of 0.796 will be achieved in the conjectured LFC. The overshoot of the $P\{CS|p=LFC\}$ (0.796 vs. $P^*=0.75$) is again due to the discrete nature of the problem. Further, *in this example*,

$$E(S_{P_{RA}}|p=LFC)=3.68 < 3.95=E(S_{P_{BK}}|p=LFC). //$$

3.4 An unbounded sequential procedure, P_{BKS}

Bechhofer, Kiefer, and Sobel (1968) give an *unbounded* (or *open*) sequential procedure which satisfies the PR.

Procedure P_{BKS} :

1. Specify k , P^* , θ^* .
2. Take observations until

$$\sum_{i=1}^{k-1} (1/\theta^*)^{x[k],t} \leq (1-P^*)/P^*. //$$

Remark 3.4:

1. BKS show that the SC is the LFC for this procedure.

Example 3.4.1:

Let $k=3$, $P^*=0.75$, $\theta^*=3$. Consulting the appropriate tables in Bechhofer and Goldsman (1984b), we immediately find that $P\{CS|p=LFC\} = 0.842$ (.0004) and $E(S_{P_{BKS}}|p=LFC) = 4.526$ (.051). These results are Monte Carlo estimates obtained via simulation; the entries in parentheses are the accompanying standard errors. The results are nearly exact, as can be seen by the small standard errors. //

3.5 P_{BG} , an improved version of P_{BKS}

As in the above example, it turns out that P_{BKS} frequently yields $P\{CS|p=LFC\} \gg P^*$. This extra $P\{CS\}$ is at the cost of unnecessary observations. Therefore, Bechhofer and Goldsman (1984b) give a procedure which decreases the attained $P\{CS\}$ to a level slightly greater than P^* , but which also saves observations.

Procedure P_{BG} :

1. Specify k , P^* , θ^* .
2. Take observations until either
 - A. $\sum_{i=1}^{k-1} (1/\theta^*)^{x_{[k],t} - x_{[i],t}} \leq (1-P^*)/P^*$ or
 - B. the stage $t=N_{BG}$, where N_{BG} is determined by k , P^* , θ^* , and is to be found in Bechhofer and Goldsman's tables for certain values of k , P^* , θ^* . //

Remarks 3.5:

1. N_{BG} is chosen as the smallest upper bound on the total number of observations such that the PR is satisfied.
2. Unlike P_{BKS} , P_{BG} is *bounded*.
3. It is not known whether the SC is the LFC for this procedure, but we so conjecture.
4. P_{BG} is neither directly comparable to P_{BK} nor P_{RA} . For many choices of k , P^* , θ^* , it seems that P_{BG} requires fewer observations (on the average) than P_{BK} . The authors feel that the user should consult the relevant tables when designing an experiment.

Example 3.5.1:

Let $k=3$, $P^*=0.75$, $\theta^*=3$. We now abstract a small portion of the necessary tables for P_{BG} from Bechhofer and Goldsman (1984b).

P^*	θ^*	N_{BG}	$P\{CS\}$	$E(S)$
.75	3.0	5	.757	3.48
.75	2.4	8	.760	5.59
.75	2.0	13	.751	8.18
.75	1.6	32	.752	17.80

Table 3 (for P_{BG}):
 $P\{CS|k=3, p=LFC\}$, $E(S|.)$ for various P^* , θ^*

We see that we must choose $N_{BG}=5$ with the resulting

$$P\{CS|p=LFC\}=0.757 \text{ and } E(S_{P_{BG}}|p=LFC)=3.48. //$$

3.6 P_{A1} , an augmented version of P_{BG}

We now employ the same device which was used in P_{BK} ; viz., we stop sampling when the cell in second place only has a chance to tie.

Procedure P_{A1} :

1. Specify k, P^*, θ^* .

2. Take observations until

$$A. \sum_{i=1}^{k-1} (1/\theta^*)^{x_{[k],t} - x_{[i],t}} \leq (1-P^*)/P^* \text{ or}$$

$$B. t=N_{A1}=N_{BG}, \text{ where } N_{BG} \text{ is from } P_{BG} \text{ or}$$

$$C. x_{[k],t} - x_{[k-1],t} = N_{A1}^{-t}. //$$

Remarks 3.6:

$$1. \text{ Clearly, } E(S_{P_{A1}}) \leq E(S_{P_{BG}}).$$

2. By reasoning similar to that given in Bechhofer and Kulkarni (1983), $P\{CS|P_{A1}\} = P\{CS|P_{BG}\}$. That is, no $P\{CS\}$ is lost between the two procedures.

3. Tables for P_{A1} are currently being prepared. See Remark 3.6.2 above for information concerning the $P\{CS\}$.

Example 3.6.1:

Again, let $k=3, P^*=0.75, \theta^*=3$. Then $N_{A1}=5$ and $P\{CS|p=LFC\}=0.757$ as before. Now, $E(S_{P_{A1}}|p=LFC)=3.24 < 3.48=E(S_{P_{BG}}|p=LFC). //$

3.7 General remarks

We have seen procedures which follow a poset of sorts in terms of sampling efficiency. P_{BEM} leads to the more efficient P_{BK} . Similarly, P_{BKS} leads to P_{BG} which, in turn, leads to P_{A1} . P_{RA} stands alone. We note that augmentations may be made to P_{RA} , but this makes our search for the optimal combination of r , N , etc., intractable.

$$P_{BEM} \longrightarrow P_{BK}$$

$$P_{BKS} \longrightarrow P_{BG} \longrightarrow P_{A1}$$

$$P_{RA}$$

In lieu of work currently in progress, the authors recommend use of P_{RA} or P_{A1} when these procedures are applicable to the situation at hand.

4. Applications and Augmentations in Simulation

We are now interested in the more general problem of determining which of k arbitrary populations Π_1, \dots, Π_k , is the 'best.' Suppose X_i is an independent observation from Π_i , $i=1, \dots, k$. Recall that we can correspond each of the k Π_i 's with a cell of a k -nomial distribution with cell probabilities p_1, \dots, p_k , where $p_i = P\{X_i \text{ is the 'most desirable' of } X_1, \dots, X_k\}$. Hence, the multinomial procedures described in the last section are *nonparametric*. This fact is of tremendous importance for simulators since the underlying distributions of the Π_i 's (i.e., k simulated systems) are frequently unknown.

4.1 An example

Suppose that we wish to choose that one of k different (s, S) inventory policies which will have the highest probability of yielding the maximum profit for a small company. Here, profit is taken to be the criterion of desirability. It is assumed that the financial affairs of the company are complicated enough such that an analytic solution of this problem is not possible. Thus, it is necessary to resort to the use of simulation and multinomial selection techniques.

For the sake of simplicity, suppose that $k=3$, $P^*=0.75$, and $\theta^*=3$. That is, we must choose among three (s, S) policies; it is desired that $P\{CS!p_{[k]} \geq \theta^* p_{[k-1]}\}$, where p_i is the probability that the i -th policy yields the highest profit in a given k -vector observation. We will use procedure P_{A1} .

We simulate each of the three (s, S) policies (with different

pseudo-random number sequences) to obtain vector observations $\underline{Y}_1, \underline{Y}_2, \dots$. Let $Y_{j,t}$ = the profit from policy j on the t -th simulation run. $\underline{Y}_t = (Y_{1,t}, Y_{2,t}, Y_{3,t})$, $t=1, \dots, T$, where T is the stage of sampling at which P_{A1} terminates. After the t -th stage of sampling is completed, identify the policy which yields the highest profit among $(Y_{1,t}, Y_{2,t}, Y_{3,t})$. If necessary, use randomization to break ties. Increment the count in the corresponding multinomial cell by one.

Example: If $\underline{Y}_1 = (356, 422, 297)$, then the highest profit (for this vector observation) is realized by Π_2 . Thus, the count $\underline{x}_1 = (x_{1,1}, x_{2,1}, x_{3,1}) = (0, 1, 0)$. //

Take 3-vector simulated observations until P_{A1} calls for the termination of sampling. Recall from Section 3.6 that P_{A1} terminates when

1. $\sum_{i=1}^{k-1} (1/\theta^*)^{x_{[k],t} - x_{[i],t}} \leq (1-P^*)/P^* \quad (= 1/3) \text{ or}$
2. $t = N_{A1} \quad (= 5) \text{ or}$
3. $x_{[k],t} - x_{[k-1],t} = N_{A1} - t \quad (= 5-t)$

In the table below, we continue the example. The first column gives the sampling stage - i.e., the number of 3-vector observations which have been taken. In the next three columns, the 3-vectors of simulated data are given. These are followed by the corresponding multinomial cell $x_{i,t}$'s.

Stage t	$Y_{1,t}$	$Y_{2,t}$	$Y_{3,t}$	$x_{1,t}$	$x_{2,t}$	$x_{3,t}$
1	356	422	297	0	1	0
2	411	378	314	1	1	0
3	374	393	380	1	2	0
4	368	374	378	1	2	1

At stage $t = 4$, P_{A1} calls for procedure termination since $x_{[k],t} - x_{[k-1],t} = N_{A1} - t$. We choose policy two as 'best,' since that is the policy corresponding to $x_{[3],T}$ //

4.2 Pseudo-observations

We discuss an augmentation of P_{A1} that eliminates populations which seem to be 'inferior.' The augmentation takes advantage of the possibility that in the course of sampling, some of the Π_i 's will have no chance of 'winning' (being chosen as 'best').

For instance, in the example of Section 4.1,

$$\underline{x}_3 = (x_{1,3}, x_{2,3}, x_{3,3}) = (1, 2, 0).$$

Claim: Given that $\underline{x}_3 = (1, 2, 0)$, it is impossible for Π_3 to win (in this example).

Proof:

Case 1: If $\underline{x}_4 = (2, 2, 0)$, then only Π_1 and Π_2 can win (since $N_{A1} = 5$).

Case 2: If $\underline{x}_4 = (1, 3, 0)$, then sampling terminates

and Π_2 wins (since $\sum_{i=1}^{k-1} (1/\theta^*)^{x_{[k],t} - x_{[i],t}} \leq (1-P^*)/P^*$).

Case 3: If $\underline{x}_4 = (1, 2, 1)$, then Π_2 wins (since

$$x_{[k],t} - x_{[k-1],t} = N_{A1} - t). //$$

Thus, in this example, it is *pointless to sample from* Π_3 given that $\underline{x}_3 = (1,2,0)$.

With this example in mind, consider the following augmented procedure, P_* , which no longer takes observations from Π_3 : Suppose that *before* the next vector observation is taken, a $U(0,1)$ probability die is rolled. Let the outcome of the roll be $0 \leq u \leq 1$. Since the PR must be satisfied, assume that $(p_{[1]}, p_{[2]}, p_{[3]}) = (p, p, \theta^* p)$, where $p = 1/(\theta^* + 2)$. That is, the underlying configuration of p_i 's is the SC (the conjectured LFC). If $u \leq p$, award a 'success' to multinomial cell 3 (i.e., increment cell 3's count by one: $x_{3,4} = x_{3,3} + 1$) *without actually taking vector observation* \underline{y}_4 . In this case, we have generously awarded Π_3 a 'free success,' and we call this non-observation a *pseudo-observation*. If $u > p$, define $\underline{y}_4 = (y_{1,4}, y_{2,4})$. In this case, we only sample from Π_1 and Π_2 . Increment as usual the count of the cell corresponding to the 'more desirable' of the two observations. Take observations in this manner until any of the stopping criteria from P_{A1} are met (where the $x_{i,t}$'s are defined as above).

With the example still in mind, let B be the event that we are using procedure P_{A1} , the underlying configuration of the p_i 's is the SC, and $\underline{x}_3 = (1,2,0)$. Define C similarly except that P_{A1} is to be replaced by P_* .

Claim: $P\{CS|B\} = P\{CS|C\}$.

Proof: Since we operate in the SC, $p_3 = p$ or $\theta^* p$.

Case 1: If $p_3 = \theta^* p$, then cell 3 is the correct cell (since $\theta^* > 1$). However, it is clear that

$$P\{CS|p_3=\theta^*p, B\} = P\{CS|p_3=\theta^*p, C\} = 0. \quad //$$

Case 2: Suppose $p_3 = p$. Then $(p_1, p_2) = (p, \theta^*p)$ or (θ^*p, p) . Assume the former subcase. A similar argument will apply for the latter. Consider P_* and a given 3-vector observation. Then Π_3 is awarded a (pseudo-)success with probability p . Further, Π_1 is awarded a success w. p .

$$\begin{aligned} &P\{\Pi_3 \text{ will not get the success}\} \times P\{\Pi_1 \text{ will get the} \\ &\quad \text{success : only } \Pi_1 \text{ and } \Pi_2 \text{ are under consideration}\}. \\ &= (1-p) \times p/(p+\theta^*p) = p. \quad \text{Similarly,} \end{aligned}$$

$P\{\Pi_2 \text{ will get the success}\} = \theta^*p$. But these success probabilities are exactly the same as those from P_{A1} . Since the termination criteria for both procedures are also identical, we have the result. //

Goldsman and Schruben (1984) consider a more general version of P_* .

Procedure P_{A2} :

1. Specify k, P^*, θ^* .

For $t=1, 2, \dots$

2. Let $I_t = \{i: x_{[k],t} - x_{i,t} \geq N_{A1} - t\}$

(This is the set of Π_i 's that no longer have a chance to win.)

3. For each $i \in I_t$, allocate an interval of length p of $[0, 1]$, where $p=1/(k-1+\theta^*)$.

4. Roll a $U(0,1)$ random number, u .

5. If u falls in an interval allocated for some $j \in I_t$, increment the corresponding $x_{j,t}$ by one (i.e., award

a 'pseudo-success' to Π_j). Otherwise, take actual observations from all Π_i 's such that $i \in \{1, \dots, k\} \setminus I_t$. Increment by one the $x_{i,t}$ corresponding to the 'most desirable' observation.

6. Terminate the procedure (with the usual decision rule) if any of the termination criteria for P_{A1} are satisfied. //

Remarks 4.2:

1. Goldsman and Schruben (1984) prove that

$$P\{CS|P_{A1}, p=LFC\} = P\{CS|P_{A2}, p=LFC\}.$$

2. Clearly, $E(S_{P_{A2}} | p=EPC) \leq E(S_{P_{A1}} | p=EPC)$, where S is

the number of stages (in which actual observations are taken) until termination. It seems likely that this relationship also holds when $p=LFC$, but this has not yet been proven.

3. Tables for P_{A2} are currently being prepared. See Remarks 4.2.1 and 3.6.2 for information concerning $P\{CS\}$.

4. The trick of taking pseudo-observations is particularly suited for the simulation environment.

Example 4.2.1:

Again, let $k=3$, $P^*=0.75$, $\theta^*=3$. Then $P\{CS|p=LFC\} = 0.757$ as before, and $E(S_{P_{A2}} | p=LFC) = 3.12$ //

4.3 Correlation induction

Frequently, it is possible for the simulator to artificially

induce (positive) correlation among the π_i 's. For instance, the simple technique of common random numbers can be used (when applicable). More complicated methods can also be implemented. It stands to reason that as the correlation among the populations increases, it becomes easier for the experimenter to distinguish which of the populations is the 'best.'

Consider the aforementioned selection procedures. Obviously, an increase in θ^* facilitates the distinction of the 'best' multinomial cell. The following crude example illustrates how positive correlation induction can result in increased θ^* .

Example 4.3.1:

Suppose that $k=2$ and that X_i is distributed normally with unknown mean μ_i and known, common variance σ^2 , $i=1,2$. If one observation is larger than another, the first observation is taken to be the more desirable. So, define $p_1 = P(X_1 > X_2)$, $p_2 = 1 - p_1$. Suppose that $\mu_1 > \mu_2$; so we can let $p_1 = \theta p$ and $p_2 = p$, where θ is some number > 1 . Clearly, $\theta = (1-p)/p$. Finally, define $\rho = \text{Corr}(X_1, X_2) \geq 0$.

$$\begin{aligned} \text{Now, } p_1 &= P(X_1 > X_2) = P(X_1 - X_2 > 0) \\ &= P\{[X_1 - X_2 - (\mu_1 - \mu_2)]/\omega > -(\mu_1 - \mu_2)/\omega\}, \\ &\quad \text{where } \omega = \sqrt{2\sigma^2(1-\rho)} \\ &= 1 - \Phi(-(\mu_1 - \mu_2)/\omega) = \Phi((\mu_1 - \mu_2)/\omega), \\ &\quad \text{where } \Phi(.) \text{ is the } N(0,1) \text{ cdf} \\ &= \theta_p p, \text{ say, } = 1-p. \end{aligned}$$

Then $\theta_p = (1-p)/p = \Phi(\eta)/(1-\Phi(\eta))$,

where $\eta = (\mu_1 - \mu_2)/\omega$.

Hence, $\theta_p/\theta_0 = [\Phi(\eta)/\Phi(\eta')] \times [(1-\Phi(\eta'))/(1-\Phi(\eta))]$,

where $\eta' = \eta\sqrt{1-p}$.

This quantity is obviously > 1 . Thus, $\theta_p > \theta_0$ //

Remark 4.3:

More details and examples are given in Goldsman and Schruben (1984).

4.4 Population splitting

We remarked earlier that the LFC represents a 'worst case' configuration for $p \in \Omega_\theta^*$, the preference zone. Of course, such a worst case is rarely encountered in practice. If, in a simulation study, $p \in \Omega_\theta^*$, then it is likely that p is not in the slippage configuration. Perhaps it is possible to take advantage of this likelihood. A generalization of taking pseudo-observations is proposed which avoids taking observations from Π_i 's which seem to be 'inferior' to other populations. Indeed, in the course of sampling, we partition the Π_i 's into a 'good' set G and a 'bad' set B. The populations in G receive real observations while those in B receive only pseudo-observations. If certain populations in G do not garner many successes during sampling, it is possible to exchange them with populations from B. This possibility of exchange gives all of the k populations a chance to win while discouraging actual sampling from 'bad' Π_i 's.

The sampling procedure we consider below is almost purely heuristic. To facilitate the discussion, we consider directly

sampling from a k-nomial distribution.

Procedure P_S :

1. Specify k, P^*, θ^*, n_0 (some initial number of k-nomial observations).
2. Take n_0 observations.
3. Partition cells $1, \dots, k$ into 'good' cells G and 'bad' cells B . Suggestion: Place the cells with \geq the median number of successes (either real or pseudo) into G ; put the others in B .
4. Temporarily be conservative and assume that the 'best' cell is in B . In the LFC, this cell will have probability $\theta^* p$; $P\{B\} = (|B| - 1)p + \theta^* p$. With probability $P\{B\}$, award a pseudo-success to a *random* cell in B . If a pseudo-success is not awarded to a cell in B , take a real multinomial observation from the cells in G (i.e., take observations only from π_i 's corresponding to cells in G).
5. If any of the termination criteria from a previously discussed multinomial procedure (which must be pre-specified) are met, stop sampling and choose as best that cell corresponding to $x_{[k],T}$. Otherwise, go to 3. //

Remarks 4.4:

1. Admittedly heuristic, this procedure is intuitively appealing.
2. Goldsman and Schruben (1984) give more details and some limited computer simulations relevant to this procedure.
Caveat: very little work has been done to date.

3. We feel that this procedure will work well for large k and P^* and for small θ^* .

5. Summary

In this expository paper, we have introduced the reader to the problem of selecting the multinomial cell with the largest underlying probability. A brief review of some of the existing multinomial selection procedures was given. The superior procedures appear to be P_{RA} and P_{A1} . We also argued that these multinomial procedures could actually be viewed as nonparametric procedures; thus, they should be attractive to simulators. Various augmentations for use in the simulation environment were presented. This interesting problem remains an active area of research from the points of view of both statistics and simulation.

Acknowledgement: We thank Prof. Robert E. Bechhofer of Cornell for his many comments and suggestions.

References:

- Bechhofer, R. E., Elmaghraby, S., and Morse, N. (1959). A single-sample multiple-decision procedure for selecting the multinomial event which has the highest probability. Ann. Math. Stat. 30, 102-119.
- Bechhofer, R. E., and Goldsman, D. (1984a). On the Ramey-Alam sequential procedure for selecting the multinomial event which has the largest probability. Cornell SOR&IE technical report.
- Bechhofer, R. E., and Goldsman, D. (1984b). On truncating the Bechhofer-Kiefer-Sobel procedure for selecting the multinomial event which has the largest probability. In preparation.
- Bechhofer, R. E., Kiefer, J., and Sobel, M. (1968). Sequential Identification and Ranking Procedures. The Univ. of Chicago Press, Chicago.
- Bechhofer, R. E., and Kulkarni, R. (1983). Closed sequential procedures for selecting the multinomial events which have the highest probabilities. In preparation.
- Bechhofer, R. E., and Sobel, M. (1958). Nonparametric multiple-decision procedures for selecting that one of k populations which has the highest probability of yielding the largest observation (prelim. report). Abstract. Ann. Math. Stat. 29, p. 325.
- Goldsman, D., and Schruben, L. (1984). On selecting the best of k simulated systems. In preparation.
- Gupta, S. S., and Panchapakesan, S. (1979). Multiple Decision Procedures. John Wiley, New York.
- Kesten, H., and Morse, N. (1959). A property of the multinomial distribution. Ann. Math. Stat. 30, 120-127.
- Ramey, J. T., Jr., and Alam, K. (1979). A sequential procedure for selecting the most probable multinomial event. Biometrika 66, 171-173.