Proceedings of the 1984
Winter Simulation Conference
S. Sheppard, U. Pooch, D. Pegden (eds.)

A MULTINOMIAL RANKING AND SELECTION PROCEDURE: SIMULATION AND APPLICATIONS

David Goldsman
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332

Abstract

This paper reports on the role of simulation in the development of a certain statistical ranking and selection procedure.

1. Introduction

In an article elsewhere in these Proceedings, Goldsman (1984), we presented an expository survey of procedures for selecting that multinomial cell which has the largest underlying probability. It is shown in Section 1 of that paper that such procedures are, in a sense, nonparametric. Section 2 of the current paper gives a brief summary of the pertinent notation and terminology. Section 3 focuses on a particular multinomial selection procedure; viz., that of Bechhofer and Goldsman (1984). We discuss the use of simulation in the preparation of the tables necessary for implementation of this selection procedure. In Section 4, it is demonstrated that the selection procedure in question can be used in the simulation environment. Section 5 shows that the procedure can be augmented somewhat by using certain simulation techniques.

2. Preliminaries

The following summarizes the first two sections of Goldsman (1984).

Consider k different competing simulated systems, $\pi_1,\pi_2,\ldots,\pi_k$. Suppose that we take independent vector-observations of the form $(X_1,X_2,\ldots,X_k)$, X_i being from π_i, $i=1,\ldots,k$. For example, X_i could be the (simulated) yearly profit achieved by the i-th of k different inventory policies. Let

$$p_i \equiv P\{X_i \text{ is the 'most desirable' of } X_1,X_2,\ldots,X_k\},$$

where the term 'most desirable' is defined according to some criterion of goodness determined by the experimenter. For instance, 'most desirable' could be taken to mean 'shortest down time' or 'highest expected profit'. Order the p_i's: $p_{[1]} \le p_{[2]} \le \cdots \le p_{[k]}$. (Assume that we have no *a priori* knowledge as to how the $p_{[i]}$'s are paired with the π_i's.) Our goal is to find that π_i associated with $p_{[k]}$, the largest of the p_i's.

In view of the arguments given in Section 1 of Goldsman (1984), we can examine the analogous problem of finding that cell of a k-nomial distribution which has the highest underlying probability (we call this the 'best' cell). Suppose that we take independent observations sequentially from a k-nomial distribution with unknown cell probabilities $p_1,p_2,\ldots,p_k$. We continue to take observations until some stopping criterion is met; one such stopping criterion will be given in Section 3. Denote $x_{i,t}$ as the number of observations from cell i after t multinomial observations (or 'stages') have been taken, $i=1,\ldots,k$; $t=1,2,\ldots$ We refer to the $x_{i,t}$'s as the 'cell counts'. For each t, order the $x_{i,t}$'s: $x_{[1],t} \le x_{[2],t} \le \cdots \le x_{[k],t}$.

Denote T as the stage at which sampling terminates (T is a random variable for the procedure to be presented in the next section.) We choose as best that cell corresponding to the largest count at the termination of sampling, $x_{[k],T}$ (use randomization if necessary). If the cell corresponding to $p_{[k]}$ is chosen, we say that a *correct selection* (CS) has been made.

Now, denote $\Omega_{\theta^*} \equiv \{\underline{p} | \theta^* p_{[k-1]} \le p_{[k]}\}$. In the procedure to be discussed in the next section, the following so-called *indifference-zone probability requirement* must hold:

$$P\{CS | \underline{p} \in \Omega_{\theta^*}\} \ge P^*, \qquad (PR)$$

where $\{P^*,\theta^*\}$ is pre-specified by the experimenter with $1 < \theta^* < \infty$ and $1/k < P^* < 1$.

In the sequel, we will consider the following configurations of $p_{[i]}$'s:

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

and

$$p_{[i]} = 1/k, \quad i=1,\ldots,k. \qquad \text{(EPC)}$$

SC stands for *slippage configuration* and EPC stands for *equal probability configuration*. For many multinomial selection procedures, $P\{CS|\underline{p} \in \Omega_{\theta^*}\}$ is minimized when $\underline{p}$ is in the SC [see Goldsman (1984)]; in these cases, the SC is called the *least-favorable configuration* (LFC) and can be viewed as a 'worst case' configuration (given that $\underline{p} \in \Omega_{\theta^*}$). The EPC is interesting in that we would expect such a configuration to maximize a multinomial procedure's expected sample size, E[T] (i.e., the expected number of multinomial observations needed before the termination criterion is met). A 'good' multinomial procedure should (at least) satisfy the PR and have 'low' $E[T|\underline{p} \in SC]$ and $E[T|\underline{p} \in EPC]$.

3. A sequential multinomial procedure

The following sequential procedure is from Bechhofer and Goldsman (1984).

Procedure P_{BG}:

1. Specify k, P^*, θ^*.

2. Take multinomial observations one at a time until either

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

2-B. the stage $t = N_{BG}$, where N_{BG} is determined by k, P^*, θ^*, and is to be found in B&G's tables for certain values of k, P^*, θ^*. //

Remarks:

1. This procedure is the *closed* version of an *open* sequential procedure from Bechhofer, Kiefer, and Sobel (1968). [By closed, we mean that the number of multinomial observations is bounded.]

2. N_{BG} is chosen as the smallest upper bound on the total number of observations such that the PR is satisfied.

3. It is not known whether the SC is the LFC for this procedure, but we so conjecture.

4. P_{BG} compares favorably with other multinomial procedures in terms of E[T]. See B&G (1984).

Example:

Let k = 3, $P^* = 0.75$, $\theta^* = 3$. It turns out [see Example 3.5.1 of Goldsman (1984)] that $N_{BG} = 5$, with the resulting $P\{CS|\underline{p} = SC\} = 0.757$ and $E[T_{P_{BG}}|\underline{p} = SC] = 3.48$. //

Suppose sampling proceeds as follows:

Stage t	$X_{1,t}$	$X_{2,t}$	$X_{3,t}$
1	0	0	1
2	1	0	1
3	1	0	2
4	1	0	3

Table 1 : An example of P_{BG}

At stage t = 4, stopping criterion 2-A is satisfied, so we terminate sampling and select cell 3 as best. //

Although this procedure is very simple to implement, construction of the necessary tables was a formidable task. For 'small' values of k and P^*, and for 'large' values of θ^*, we used an iterative method (loosely based on a random walk argument) in order to calculate (on a computer) exact values for $P\{CS|\underline{p}\}$ and $E[T|\underline{p}]$. CPU limitations forced us to resort to use of Monte Carlo simulation in order to obtain results for 'large' k and P^*, and for θ^* near 1.

Example:

To illustrate results from our Monte Carlo simulation, a small portion of the k = 10 tables from B&G (1984) is abstracted:

N_{BG}	Estimated $P\{CS\|SC\}$	Estimated $E[T\|SC]$	Estimated $E[T\|EPC]$
31	$.7539^{*}$ (.0014)	18.73^{*} (.08)	$24.46^{\downarrow}$ (.13)
54	$.7535^{*}$ (.0012)	32.03^{*} (.13)	$42.36^{\downarrow}$ (.23)
96	$.7520^{*}$ (.0011)	55.16^{*} (.24)	$72.91^{\downarrow}$ (.41)
244	$.7510^{\#}$ (.0007)	$132.38^{\#}$ (.45)	$175.45^{\downarrow}$ (1.07)

Table 2: Results for P_{BG} for k = 10 based on ($\downarrow$=4000, *=12000, #=20000) independent replications of P_{BG}

We now give an elementary augmentation of P_{BG}; viz., stop sampling when the cell in second place only has a chance to tie.

Procedure P_{BG2}:

1. Specify k, P^*, θ^*.

2. Take observations until

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

2-B. $t = N_{BG2} = N_{BG}$, where N_{BG} is from P_{BG} or

2-C. $x_{[k],t}-x_{[k-1],t} = N_{BG2}-t$. //

Remarks:

1. Clearly, $E[T_{P_{BG2}}] \leq E[T_{P_{BG}}]$.

2. It can be shown that $P\{CS|P_{BG2}\} = P\{CS|P_{BG}\}$. I.e., no P{CS} is lost between the two procedures.

3. Tables for P_{BG2} are currently being prepared. Concerning the P{CS}, see the above remark.

Example:

Again, let k = 3, $P^* = 0.75$, $\theta^* = 3$. Then $N_{BG2} = 5$ and $P\{CS|\underline{p} = SC\} = 0.757$ as before. Now, $E[T_{P_{BG2}}|\underline{p} = SC] = 3.24 < 3.48 = E[T_{P_{BG}}|\underline{p} = SC]$. //

4. An example in the simulation environment

We are now interested in the more general problem of determining which of k *arbitrary* populations $\pi_1,\ldots,\pi_k$, is the 'best.' Suppose X_i is an independent observation from π_i, i=1,...,k. Recall that we can correspond each of the k π_i's with a cell of a k-nomial distribution with cell probabilities $p_1,\ldots,p_k$, where $p_i = P\{X_i$ is the 'most desirable' of $X_1,\ldots,X_k\}$; so the multinomial procedures P_{BG} and P_{BG2} 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.

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; we therefore resort to the use of simulation and multinomial selection techniques.

For the sake of simplicity, suppose that k=3, P^*=0.75, and θ^*=3; so we must choose among three (s,S) policies. Further, it is desired that the usual indifference-zone PR hold: $P\{CS|\theta^* p_{[k-1]} \leq p_{[k]}\} \geq P^*$, 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_{BG2}.

We simulate each of the three (s,S) policies (with different pseudo-random number sequences) to obtain vector observations $\underline{Y}_1,\underline{Y}_2,\ldots$ Let $y_{j,t} \equiv$ the profit from policy j on the t-th simulation run. $\underline{Y}_t \equiv (y_{1,t},y_{2,t},y_{3,t})$, t=1,...,T, where T is the stage of sampling at which P_{BG2} terminates. After the t-th stage of sampling is completed, we identify the policy which yields the highest profit among $(y_{1,t},y_{2,t},y_{3,t})$. Randomization is used to break ties. We 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 $\underline{X}_1 = (x_{1,1},x_{2,1},x_{3,1}) = (0,1,0)$. //

We take 3-vector simulated observations until P_{BG2} calls for the termination of sampling. Recall from Section 3 that P_{BG2} terminates when:

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

2. $t = N_{A1}$ (= 5) or

3. $x_{[k],t} - x_{[k-1],t} = N_{A1}-t$ (= 5-t)

In the table below, we continue the example. The first column gives the sampling stage t - 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.

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

Table 3: Vector-observations and corresponding cell counts

At stage t = 4, P_{BG2} calls for procedure termination since $x_{[k],t}-x_{[k-1],t} = N_{BG2}-t$. We choose policy two as 'best,' since that is the policy corresponding to $x_{[3],T}$.//

5. Augmentations using simulation

5.1 Pseudo-observations

We discuss an augmentation of P_{BG2} that eliminates populations which seem to be 'inferior.' This 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, $\underline{x}_3 = (x_{1,3},x_{2,3},x_{3,3}) = (1,2,0)$.

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 $N_{BG2} = 5$).

Case 2: If $\underline{x}_4 = (1,3,0)$, then sampling terminates and π_2 wins (since stopping criterion 2-A of P_{BG2} holds).

Case 3: If $\underline{x}_4 = (1,2,1)$, then π_2 wins (since $x_{[k],t}-x_{[k-1],t} = N_{BG2}-t$).//

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 $0 \le u \le 1$. 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 $u \le p < 1$, 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* $\underline{y}_4$. [We have generously given π_3 a 'free success'. This non-observation is called a *pseudo-observation* (or *pseudo-success*).] 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_{BG2} 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_{BG2}, the underlying configuration of the p_i's is the SC, and $\underline{x}_3=(1,2,0)$}. Define C similarly except that P_{BG2} 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 $P\{CS|p_3=\theta^* p, B\} = P\{CS|p_3=\theta^* p, C\} = 0$. /

Case 2: Suppose $p_3 = p$. Then $(p_1,p_2) = (p,\theta^* p)$ or $(\theta^* p,p)$. 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 with probability:

$[P\{\pi_3$ will not get the success$\}$ x $P\{\pi_1$ will get the success | only π_1 and π_2 are under consideration$\}] =$

$(1-p) \times p/(p+\theta^* p) = p.$

Similarly, $P\{\pi_2$ will get the success$\} = \theta^* p$. But these success probabilities are exactly the same as those from P_{BG2}. 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_{SIM}:

1. Specify k, P^*, θ^*. For t=1,2,...

2. Let $I_t = \{i | x_{[k],t} - x_{i,t} \geq N_{BG2} - t\}$ (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 $i \in \{1,\ldots,k\} \setminus I_t$. 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_{BG2} are satisfied. //

Remarks:

1. G&S prove that $P\{CS|P_{BG2}, \underline{p} = SC\} = P\{CS|P_{SIM}, \underline{p} = SC\}$.

2. Clearly, $E[T_{P_{SIM}}|\underline{p} = EPC] \leq E[T_{P_{BG2}}|\underline{p} = EPC]$, where T_P is the number of stages (in which *actual* observations are taken) until termination. It seems likely that this relationship also holds when $\underline{p} = SC$, but this has not yet been proven.

3. Tables for P_{SIM} are currently being prepared.

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

Example:

Again, let k=3, P^*=0.75, and θ^*=3. Then $P\{CS|\underline{p} = SC\} = 0.757$ as before, and $E[T_{P_{SIM}}|\underline{p} = SC] = 3.12$. //

5.2 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:

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. Define $p_1 \equiv P(X_1 > X_2)$ and $p_2 \equiv 1 - p_1$. Suppose that $\mu_1 > \mu_2$; so we can let $p_1 \equiv \theta p$ and $p_2 \equiv p$, where $\theta = (1-p)/p > 1$. Finally, define $\rho \equiv \mathrm{Corr}(X_1, X_2) \geq 0$.

Then

$$p_1 = P(X_1 > X_2) = P(X_1 - X_2 > 0)$$

$$= P\{[X_1 - X_2 - (\mu_1 - \mu_2)]/\omega > -(\mu_1 - \mu_2)/\omega\}, \quad \text{where } \omega = \sqrt{(2\sigma^2(1-\rho))}$$

$$= 1 - \Phi(-(\mu_1 - \mu_2)/\omega) = \Phi((\mu_1 - \mu_2)/\omega), \quad \text{where } \Phi(.) \text{ is the N(0,1) cdf}$$

$$= \theta_\rho p, \text{ say, } = 1-p.$$

So $\theta_\rho = (1-p)/p = \Phi(\eta)/(1-\Phi(\eta))$, where $\eta = (\mu_1 - \mu_2)/\omega$.

Hence,

$$\theta_\rho/\theta_0 = [\Phi(\eta)/\Phi(\eta')] \times [(1-\Phi(\eta'))/(1-\Phi(\eta))], \quad \text{where } \eta' = \eta\sqrt{(1-\rho)}.$$

This quantity is obviously > 1; $\theta_\rho > \theta_0$. //

Summary

In this article, we have shown that multinomial selection procedures can be adapted for use in the simulation environment as nonparametric selection procedures. It was also demonstrated that simulation plays an important role in the development and implementation of such procedures. It is hoped that the experimenter will make use of these procedures when they are applicable to the problem at hand.

References

Bechhofer, R. E., and Goldsman, D. (1984). Truncation of the Bechhofer-Kiefer-Sobel sequential procedure for selecting the multinomial event which has the largest probability. School of OR&IE technical report, Cornell University, Ithaca, NY.

Bechhofer, R. E., Kiefer, J., and Sobel, M. (1968). Sequential Identification and Ranking Procedures. The Univ. of Chicago Press, Chicago.

Goldsman, D. (1984). On selecting the best of k systems: an expository survey of indifference-zone multinomial procedures. Proceedings of the 1984 Winter Simulation Conference, Dallas.

Goldsman, D., and Schruben, L. (1984). On selecting the best of k simulated systems. In preparation.