In *Computing Science and Statistics* **24** (1992), pp. 592-596.

Optimal Allocation for Estimating the Product of Two Means

Janis P. Hardwick[1] Quentin F. Stout[2]
Statistics Department EECS Department
University of Michigan, Ann Arbor, MI 48109

Abstract

Suppose we wish to estimate the product of the means of two independent populations of Bernoulli random variables, the parameters of which, themselves, are modeled as independent beta random variables. Assume that the total sample size for the experiment is fixed, but that the number of experimental units observed from each population may be random. Using a decision theoretic approach, we seek to minimize the Bayes risk that arises from using a squared error loss function.

Although selecting the form of an optimal estimator is critical to solving this problem, the real difficulty lies in determining an optimal strategy for sampling from the two populations. The problem of optimal *estimation* reduces, therefore, to a problem of optimal *allocation* which can be solved exactly using dynamic programming. We utilize similar programming techniques to evaluate exactly some of the other strategies that have been proposed for this problem. (The term *exact* is used repeatedly here to stress the fact that none of the computational results depend on simulation studies.)

1 Introduction

Suppose we have two populations of assembly parts that we intend to use together in the manufacture of a product. Each population has a certain *defect rate*:

$$1 - p_1 \text{ for Population 1}$$
$$1 - p_2 \text{ for Population 2.}$$

In order for the product to work – **both** parts must be defect free. We wish to obtain an estimate of the chance that both parts will work by doing *some* sampling from each population. It is assumed that the sample size for this experiment is a fixed number, N, but that the numbers sampled from each population, n_1 from Population 1 and n_2 from Population 2, may be random.

Now, for Population 1, let

$$X_i = \begin{cases} 0, & \text{if part } i \text{ is defective;} \\ 1, & \text{if part } i \text{ is ok,} \end{cases} \qquad i = 1, \ldots$$

and for Population 2, let

$$Y_j = \begin{cases} 0, & \text{if part } j \text{ is defective;} \\ 1, & \text{if part } j \text{ is ok,} \end{cases} \qquad j = 1, \ldots,$$

where the $X's$ and $Y's$ are mutually independent. Then at each stage, M, of the experiment, one of the random variables X_M or Y_M will be observed, where

$$\begin{aligned} (1) \qquad X_i \quad &\text{i.i.d.} \quad B(1, p_1) \quad i = 1, \ldots \\ Y_i \quad &\text{i.i.d.} \quad B(1, p_2) \quad i = 1, \ldots . \end{aligned}$$

Let θ denote the product of the population means

$$\theta = p_1 * p_2$$

and define the loss due to any estimate $\tilde{\theta}_N$ by the mean squared error of θ and $\tilde{\theta}_N$

$$L(\theta, \tilde{\theta}_N) = (\theta - \tilde{\theta}_N)^2.$$

We take a Bayesian approach to this problem, assuming that prior information of some sort is available and may be modeled in the form of beta priors on p_1 and p_2. Let $\xi(p_1, p_2)$, the prior joint density function on $(p_1, p_2) \in \Omega = (0, 1) \times (0, 1)$, denote the product of the independent beta random variables:

$$(2) \qquad p_1 \sim \text{Be}(a_0, b_0) \quad \text{and} \quad p_2 \sim \text{Be}(c_0, d_0).$$

The Bayes risk, $\mathcal{R}$, for an estimator $\tilde{\theta}_N$ is defined as the expected loss of $(\theta, \tilde{\theta}_N)$

$$\mathcal{R}(\xi, \tilde{\theta}_N) = \mathbf{E}^{\xi}[L(\theta, \tilde{\theta}_N)] = \mathbf{E}^{\xi}[(\theta - \tilde{\theta}_N)^2]$$

where the expectation, $\mathbf{E}^{\xi}$, is taken with respect to the Bayesian model (2) in which (1) holds conditionally given $(p_1, p_2) \in \Omega$.

2 Minimizing the Risk Function

Minimizing the Bayes risk may be tackled in two stages. The first problem is to find the form of the optimal estimator and the second is to determine how many units to sample from each population. Fortunately, these problems are independent and the first is straightforward.

Suppose that, at any time $M \leq N$, we have observed

$$m_1 = \# \text{ parts from Population 1, and}$$
$$m_2 = \# \text{ parts from Population 2.}$$

[1]Research supported in part by National Science Foundation under grants DMS-8914328 and DMS-9157715.

[2]Research supported in part by National Science Foundation/DARPA under grant CCR-9004727.

Then the *posterior density function*, ξ_M on p_1 and p_2, is the product of the individual posterior densities

$$(p_1 \mid i,j) \sim \text{Be}(a,b); \qquad (p_2 \mid k,l) \sim \text{Be}(c,d)$$

where i, j, k and l are sufficient statistics for p_1 and p_2:

$$i = \Sigma_{\alpha=1}^{m_1} X_\alpha; \quad j = m_1 - i; \quad k = \Sigma_{\alpha=1}^{m_2} Y_\alpha; \quad l = m_2 - k,$$
$$\text{and} \qquad a = i + a_0; \quad b = j + b_0; \quad c = k + c_0; \quad d = l + d_0.$$

Due to the independence of the prior distributions, it is easy to show that, for any $M = 0, \ldots, N$, the optimal, or *Bayes*, estimate of θ is simply the product of the posterior means:

$$\hat{\theta}_N = \mathbf{E}^{\xi_M}[p_1] * \mathbf{E}^{\xi_M}[p_2] = \frac{a}{a+b} * \frac{c}{c+d}.$$

What remains to be determined, therefore, is how to carry out sampling from the two populations.

3 Sampling Procedures

A *sampling procedure,* δ, is a sequence of indicator functions, $\delta = (\delta_1, \delta_2, \ldots \delta_N)$, such that

$$\delta_i = \begin{cases} 1, & \text{if Population 1;} \\ 0, & \text{if Population 2.} \end{cases}$$

for $i = 1, \ldots, N$ where δ_{i+1} depends only on the information available up until time i. For convenience, we denote partial sequences as

$$\delta(m_1, m_2) = (\delta_1, \delta_2, \ldots, \delta_{m_1+m_2})$$

$$\text{where} \qquad \sum_{i=1}^{m_1+m_2} \delta_i = m_1 \text{ for any } \ 0 < m_1 + m_2 \le N.$$

The *Bayes risk of a procedure* δ, denoted by $\mathcal{R}_N(\xi, \delta)$, is the expected loss when δ is the sampling procedure and $\hat{\theta}_N$ the terminal estimate. The problem is to find δ^* such that

$$\mathcal{R}_N(\xi, \delta^*) = \min_\delta \mathcal{R}_N(\xi, \delta).$$

One can show that, for any procedure δ,

$$\begin{aligned} \mathcal{R}_N(\xi,\delta) \quad &= \quad \mathbf{E}^\xi \left[\frac{p_1^2\, p_2(1-p_2)}{(c_0+d_0+n_2)} + \frac{p_2^2\, p_1(1-p_1)}{(a_0+b_0+n_1)} \right] \\ &\quad -\mathbf{E}^\xi \left[\frac{p_1(1-p_1)\, p_2(1-p_2)}{(c_0+d_0+n_2)(a_0+b_0+n_1)} \right] \end{aligned} \tag{3}$$

where $n_1 = \sum_{i=1}^N \delta_i$ and $n_2 = N - n_1$. However, locating values of n_1 and n_2 to minimize (3) is non-trivial, since the sample sizes themselves are random variables.

A standard approach to finding good or optimal sampling procedures for allocation problems is to locate asymptotic lower bounds for the Bayes risk of *any* procedure, and then to seek ad hoc sampling rules that either achieve or come close to achieving the lower bounds. Rekab (1989) takes this approach for estimating the product of two normal means using a quadratic loss function. Hardwick (1990) considers a similar approach in estimating the *difference* of two binomial means using a 'loss plus cost' formulation.

Here, we do not study the asymptotic behavior of sampling procedures, but instead determine **exactly** optimal sequential rules (OS). To establish the potential advantages of using such rules, we define and compare three sub-optimal procedures - Equal Allocation (EA), Best Fixed Allocation (BF), and Local Sequential (LS).

In Section 4, we compute the efficiency of each procedure for various sample sizes and parameter configurations. The *efficiency* of a procedure δ is defined as

$$\frac{\mathcal{R}_N(\xi,\delta)}{\mathcal{R}_N(\xi,\delta^*)} \tag{4}$$

where, for a *fully* efficient sampling procedure, the ratio in (4) is one.

3.1 Best Fixed Allocation Rules

When examining this problem, an obvious question that arises is how well one fares when sampling from each population equally. However, since equal allocation is a special case of fixed allocation (in which n_1 and n_2 must be specified in advance) we consider the more general problem here, and the former in Sections 4 and 5.

For fixed allocation rules, $\delta(n_1, n_2)$, the risk function takes on a simplified form. With n_1 and n_2 no longer random, the Bayes risk reduces to

$$\begin{aligned} \mathcal{R}_N(\xi, \delta(n_1,n_2)) &= \frac{\mathbf{E}^\xi[p_1^2\, p_2(1-p_2)]}{(c_0+d_0+n_2)} \\ &+ \frac{\mathbf{E}^\xi[p_2^2\, p_1(1-p_1)]}{(a_0+b_0+n_1)} - \frac{\mathbf{E}^\xi[p_1(1-p_1)\, p_2(1-p_2)]}{(c_0+d_0+n_2)(a_0+b_0+n_1)} \end{aligned} \tag{5}$$

which can be minimized directly since the moments of the prior distributions are constants. In this case, the optimal sample sizes n_1 and n_2 are obtained by solving a quadratic equation.

Note that the third term on the right hand side of (5) is an order of magnitude smaller than either of the first two terms. If we ignore this term and choose n_1 to minimize only the first order approximation of the risk function, the difference in the risks obtained is barely visible (on the order of 10^{-4}). A very close approximation to the best fixed allocation rule, then, is attained if we select n_1 such that

$$\frac{a_0+b_0+n_1}{c_0+d_0+n_2} = \frac{\sqrt{\mathbf{E}^{\xi}[p_2^2 p_1(1-p_1)]}}{\sqrt{\mathbf{E}^{\xi}[p_1^2 p_2(1-p_2)]}}.$$

3.2 Local Sequential Rules

Because optimal sample sizes actually depend on the true parameter values, we can diminish the risk by using sequential sampling procedures that utilize information in the data as it comes in. Thus, instead of selecting n_1 and n_2 at the beginning of the trial, one can expect to do better by continually re-estimating parameters and basing sampling decisions on the most recent estimates.

One ad hoc approach to sequential sampling is to use a *local* or *myopic* rule (LS). Here, we consider a local rule in which we compare the expected risk associated with each sampling option given the data observed so far. For example, suppose that at time M we have sampled m_1 and m_2 from Populations 1 and 2 respectively; then the *posterior Bayes risk* at M is

$$\begin{aligned}\mathcal{R}_M\left(\xi_M, \delta(m_1,m_2)\right) &= \frac{\mathbf{E}^{\xi_M}[p_1^2]\,\mathbf{E}^{\xi_M}[p_2(1-p_2)]}{(c_0+d_0+m_2)} \\ &+ \frac{\mathbf{E}^{\xi_M}[p_2^2]\,\mathbf{E}^{\xi_M}[p_1(1-p_1)]}{(a_0+b_0+m_1)} \\ &- \frac{\mathbf{E}^{\xi_M}[p_1(1-p_1)]\,\mathbf{E}^{\xi_M}[p_2(1-p_2)]}{(c_0+d_0+m_2)(a_0+b_0+m_1)}.\end{aligned}$$

The expected risk if we sample next from Population 1 and then stop is

$$\mathcal{R}_{M+1}\left(\xi_M, \delta(m_1+1,m_2)\right),$$

and if we sample next from Population 2 and then stop, we expect

$$\mathcal{R}_{M+1}\left(\xi_M, \delta(m_1,m_2+1)\right).$$

The myopic rule indicates that we sample from Population 1 if

$$\mathcal{R}_{M+1}\left(\xi_M, \delta(m_1+1,m_2)\right) < \mathcal{R}_{M+1}\left(\xi_M, \delta(m_1,m_2+1)\right).$$

The computations necessary to define the LS rule are straightforward, but determining the properties of the rule is more complicated. To obtain the Bayes risk (or other summary characteristics) of the LS procedure, it is necessary to keep track of what can happen at each stage. However, instead of following the sampling process forward in time, as one does in a simulation study, we record what happens beginning from the last stage and work backwards toward the first. This technique is known as *backward induction;* and, while one rarely sees it used to evaluate allocation rules, similar approaches have been taken in Berry and Eick (1987) and Hardwick and Stout (1990).

3.3 Optimal Rules

One expects the LS rule to improve on fixed allocation rules because information about study parameters accumulates and is put to use as the experiment progresses. To improve still further, however, one needs not only to utilize incoming data, but also to *anticipate* future losses. Suppose, for example, that we are at stage M, having observed (i,j,k,l). Let $\Delta(i,j,k,l)$ be the set of procedures, $\delta(i,j,k,l) = (\delta \mid i,j,k,l)$, that are consistent with the data at stage M. It is useful to work with the *optimal interim risk* function

$$\mathbf{I}_N^M(i,j,k,l) = \min_{\delta \in \Delta(i,j,k,l)} \mathbf{E}^{\xi_M}[\mathcal{R}_N(\xi,\delta)],$$

which is the minimum expected risk incurred if one were to start at stage M, given (i,j,k,l), and proceed optimally to stage N. Note that the optimal interim risk at stage N is simply the posterior risk computed for $M=N$:

$$\mathbf{I}_N^N(i,j,k,l) = \mathcal{R}_N\left(\xi_N, \delta(i,j,k,l)\right).$$

If we take the next observation from Population 1 and then proceed optimally to the end of the experiment, at stage $M+1$ we will incur a conditional expected interim risk of

$$\begin{aligned}\mathbf{E}^{\xi_M}[\mathbf{I}_N^{M+1}(i,j,k,l) \mid \delta_{M+1}=0] &= \\ &\mathbf{E}^{\xi_M}[p_1]\,\mathbf{I}_N^{M+1}(i+1,j,k,l) + \\ &\mathbf{E}^{\xi_M}[1-p_1]\,\mathbf{I}_N^{M+1}(i,j+1,k,l).\end{aligned}$$

Similarly the conditional expected interim risk if the next observation is from Population 2 is

$$\begin{aligned}\mathbf{E}^{\xi_M}[\mathbf{I}_N^{M+1}(i,j,k,l) \mid \delta_{M+1}=1] &= \\ &\mathbf{E}^{\xi_M}[p_2]\,\mathbf{I}_N^{M+1}(i,j,k+1,l) + \\ &\mathbf{E}^{\xi_M}[1-p_2]\,\mathbf{I}_N^{M+1}(i,j,k,l+1).\end{aligned}$$

Interim risk obeys the recursive optimality principle

$$\begin{aligned}\mathbf{I}_N^M(i,j,k,l) &= \min\{\mathbf{E}^{\xi_M}[\mathbf{I}_N^{M+1}(i,j,k,l) \mid \delta_{M+1}=0], \\ &\qquad \mathbf{E}^{\xi_M}[\mathbf{I}_N^{M+1}(i,j,k,l) \mid \delta_{M+1}=1]\}. \end{aligned} \tag{6}$$

This optimality principle is the heart of the dynamic programming approach for determining the optimal rule, working from the end of the experiment towards the beginning. Once the optimal interim risk has been determined for all possible observations at stage $M+1$, it can then be determined for observations at stage M by using (6). If at stage M we have observed (i,j,k,l), then the optimal population choice for the next observation is the one which attains the minimum in (6), randomizing in the case of ties.

4 Examples and Heuristics

Often, prior distributions utilized in Bayesian decision problems tend to be chosen more for their mathematical tractability than for their inherent reflection of beliefs about design parameters. In this problem, however, one finds that a wide variety of perspectives can be represented by judicious selection among the beta priors at our disposal.

Here we examine four different parameter configurations. These examples are intended to illustrate not only how prior configurations may be used to characterize prevailing experience, but also how different prior configurations can affect the sampling schemes discussed in 3.

The first example is relevant when we wish to emphasize differing degrees of faith in our prior knowledge. In such cases, for example, we might have $\mathbf{E}^{\xi}[p_1] = \mathbf{E}^{\xi}[p_2]$, but increase proportionately the values of the parameters for the distribution we believe we know more about. The effect of such modeling is that the optimal sampling rule will sample more often from the population with the smaller initial parameters.

Take the case depicted in Table 1, where

$$p_1 \sim \mathrm{Be}(20, 10) \text{ and } p_2 \sim \mathrm{Be}(2, 1).$$

The mean for each distribution is $\frac{2}{3}$, but we have less faith in the accuracy of the prior information for p_2 than for p_1. While the sampling scheme reflects this, one can see that, as the total sample size increases, the proportions sampled from the two populations approach each other.

Next, consider the case in which $p_1 \sim \mathrm{Be}(1, 2)$ and $p_2 \sim \mathrm{Be}(2, 1)$. In this parameter configuration, one might expect the EA rule to be nearly optimal, since with $a_0 + b_0 = c_0 + d_0$ and $\mathbf{E}^{\xi}[p_1] = 1 - \mathbf{E}^{\xi}[p_2]$, the distributions are mirror images of one another. As is evident from the Table 2, however, this is not the case. On average, the optimal rule samples approximately 70% from Population 1, and the EA rule turns out to be only about 85% efficient. This observation leads to a second heuristic regarding sample size selection: if the prior information indicates that one mean is less than the other (all else being equal), an optimal rule will tend to sample more from the population with the lower expected mean.

Our last two examples, Table 3 (in which $a_0 = 0.1, b_0 = 0.01, c_0 = 1, d_0 = 1$) and Table 4 (in which $a_0 = 0.01, b_0 = 0.1, c_0 = 1, d_0 = 1$) may appear at first to be pathological since the small values of a_0 and b_0 indicate that very little information is known in advance. Suppose, however, that the first of the two batches of parts under consideration comes from a manufacturing process that is either *in tolerance* (working) or *out of tolerance* (not working). In such a case, it is the third moment (shape) of the distribution that is of interest rather than the first two moments (center and spread). In the case where $p_1 \sim \mathrm{Be}(0.1, 0.01)$, we face a u-shaped distribution where the probability that the process is out of tolerance is ten times greater than the probability that it is in tolerance. After only a couple of observations we should be able to determine, with high probability, which state we have encountered. From that point on, since $p_2 \sim \mathrm{Be}(1, 1)$, we know that we should sample more from the population with the smaller expected mean. Given the circumstances, the fact that the information in the prior with the small parameters is dominated so quickly by the data is precisely what we expect.

Sample Size	Efficiency (Risk Ratios) (% on Pop 1)			
	OS	LS	BF	EA
$N = 20$	1.000 (10%)	1.000 (10%)	0.927 (00%)	0.792 (50%)
$N = 50$	1.000 (28%)	1.000 (29%)	0.937 (30%)	0.882 (50%)
$N = 100$	1.000 (40%)	1.00 (40%)	0.930 (42%)	0.917 (50%)

Table 1: $p_1 \sim \mathrm{Be}(20, 10)$ and $p_2 \sim \mathrm{Be}(2, 1)$

Sample Size	Efficiency (Risk Ratios) (% on Pop 1)			
	OS	LS	BF	EA
$N = 20$	1.000 (74%)	0.999 (74%)	0.927 (67%)	0.862 (50%)
$N = 50$	1.000 (70%)	1.000 (71%)	0.906 (65%)	0.846 (50%)
$N = 100$	1.000 (70%)	1.000 (70%)	0.898 (64%)	0.837 (50%)

Table 2: $p_1 \sim \mathrm{Be}(1, 2)$ and $p_2 \sim \mathrm{Be}(2, 1)$

5 Conclusions

In writing this paper we were working towards two goals. The first was to solve a fairly straightforward nonlinear estimation problem, and the second was to acquaint the reader with certain useful, but relatively underutilized, computational techniques.

After reducing the statistical decision problem to one of optimal allocation, we considered four classes of rules that could be used to address the problem.

Sample Size	Efficiency (Risk Ratios) (% on Pop 1)			
	OS	LS	BF	EA
$N = 20$	1.000 (16%)	0.998 (15%)	0.917 (12%)	0.626 (50%)
$N = 50$	1.000 (15%)	0.997 (14%)	0.892 (12%)	0.583 (50%)
$N = 100$	1.000 (14%)	0.997 (13%)	0.878 (12%)	0.565 (50%)

Table 3: $p_1 \sim \text{Be}(0.1, 0.01)$ and $p_2 \sim \text{Be}(1, 1)$

Sample Size	Efficiency (Risk Ratios) (% on Pop 1)			
	OS	LS	BF	EA
$N = 20$	1.000 (92%)	0.997 (92%)	0.823 (33%)	0.754 (50%)
$N = 50$	1.000 (90%)	0.998 (92%)	0.793 (32%)	0.707 (50%)
$N = 100$	1.000 (90%)	0.998 (91%)	0.778 (31%)	0.686 (50%)

Table 4: $p_1 \sim \text{Be}(0.01, 0.1)$ and $p_2 \sim \text{Be}(1, 1)$

- OS = Optimal sequential rule - defined recursively; located and evaluated via dynamic programming.
- LS = Local sequential rule - an ad hoc myopic strategy which can be located directly (at each stage), but which is evaluated using backward induction.
- BF = Best fixed allocation rule - defined to be the optimal rule when n_1 and n_2 are fixed in advance. Sample sizes and risk may be obtained directly.
- EA = Equal allocation - defined as any allocation rule such that $\frac{N}{2}$ observations are sampled from each population.

Our statistical conclusions are typified by the results in Tables 1-4. By definition, the OS rule is uniformly best, and subject to heuristic arguments of the sort detailed in previous section, the performance of the other three allocation rules depends on the configuration of the prior parameters and the sample size for the experiment.

In particular, we found that, without exception, the LS rule is so close to being fully efficient that it appears to offer the best overall combination of practicality and efficiency. Furthermore, for a variety of parameter configurations, both the BF and EA rules performed surprisingly well. Still, it is interesting to note how poorly the EA rule performs in settings such as those described in Tables 3 and 4. On the whole, however, given the performance of the LS rule, it seems reasonable to acknowledge that locating the *optimal* rule here is more an academic question than a practical one. Nevertheless, as mentioned, we had another reason for pursuing an entire solution. We believe that this problem serves as an excellent vehicle for illustrating the diverse applications of the computational technique of backward induction.

For many years, statisticians have known that, in theory, backward induction serves as a useful means for computing solutions to sequential allocation problems. For example, in discrete state space allocation problems, one can often express the optimal solution via a set of recurrence equations that can be solved dynamically. Historically, such programming has been too computationally intensive to be very practical, but increasingly, realistic problems *can* be solved in this manner. A less well understood ability of backward induction is its capacity to evaluate, *exactly,* properties of all types of allocation rules. Even if an allocation strategy is not defined through a set of recursion equations, one can still determine its attributes through backward induction. This exact approach is highly preferable to the usual method of seeking approximate characteristics of an allocation rule asymptotically or through simulation studies.

References

[1] Rekab, K (1989), Asymptotic efficiency in sequential designs for estimation. *Sequential Analysis* **8**, 269-280.

[2] Hardwick, J. P. (1989), Computational problems associated with minimizing the risk of a simple clinical trial, **Contemporary Mathematics 115**, 239-256.

[3] Hardwick, J. P. and Stout, Q. F. (1990), Bandit strategies for ethical sequential allocation. *Symposium on the Interface: Computing Science and Statistics.*

[4] Berry, D. A. and Eick, S. G. (1987), Decision analysis of randomized clinical trials: comparison with adaptive procedures. *Unpublished manuscript.*