Adaptive Designs
IMS Lecture Notes - Monograph Series (1995) Volume 25

ADAPTIVE ALLOCATION FOR ESTIMATION

BY CONNIE PAGE

Michigan State University

Abstract

Consider I populations and suppose that the goal is estimation of some function of parameters from these populations. Furthermore, a fixed total number of observations can be taken, and the problem is to decide how to allocate this total among the I populations. Some general discussion of this problem is given, then the problem is specialized to dichotomous populations and estimating the product of success probabilities. A Bayesian approach is taken using a family of scaled squared error loss functions for estimation. The best nonrandom allocation and the myopic allocation are derived, and references for the optimal allocation are given. The myopic allocation suggests a simple adaptive rule that is appropriate for all losses in the scaled family, and the limiting properties of this adaptive rule are noted.

1. Introduction. Consider I populations, and suppose that the goal is estimation of some function of parameters from these I populations. A fixed number of observations can be taken, and the problem is to decide how many observations to take from each population or how to allocate the total number of observations available. Since the goal is estimation, the allocation decision is typically made to minimize the variance or Bayes risk of a selected estimator. An early example

Received September 1992; revised July 1993.

AMS 1991 subject classification. Primary 62L05; secondary 62C10.

Key words and phrases. Bayesian sequential allocation, dynamic programming, estimating products, myopic rules, reliability.

of such a problem along with its solution (in the survey sampling area) is the allocation of observations among strata to minimize the variance of the stratified sample mean. The optimal allocation assigns numbers proportional to the subpopulation standard deviations and is called Neyman allocation [see Cochran (1977)]. Of course, the problem with implementing Neyman allocation and many optimal allocation schemes is that they depend on unknown population parameters. Thus, the need for adaptive or sequential allocation is obvious.

Allocation with estimation goals as described above has different properties from allocation for testing, or best treatment assignment goals. In these later cases, the limiting proportion assigned to the best treatment is desired to be 1.0, with other treatments receiving 0.0 in the limit. However, an 'adaptive' Neyman allocation for estimation should produce limiting allocation proportions equal to the population standard deviations, and any allocation scheme giving a limiting proportion of 1.0 or 0.0 would be sub-optimal.

Often estimated functions of population parameters are the difference in two means and linear combinations of means. In both cases, if allocation is done to minimize the variance of the same function of sample means, then Neyman allocation (allocation proportional to population standard deviation) is again optimal. An adaptive approach to estimating the difference is considered by Kelley (1977). Bayesian sequential methods for estimating linear combinations are considered by Louis (1975) and Page Shapiro (1982), among others.

Another function of population means that has broad engineering applications is the product of means. Allocation for estimating a product has been studied by Berry (1977), Page Shapiro (1985), Page (1987) and Noble (1992), among others. If allocation is done to minimize the first order approximation of the variance of the product of sample means, then the 'optimal' allocation is proportional to population coefficients of variation. As with Neyman allocation, this coefficient of variation allocation typically depends on unknown population parameters, and adaptive procedures are needed. Two stage procedures have been studied by Noble (1992), and a Bayesian sequential approach is considered by Page Shapiro (1985).

In this paper, a Bayesian approach to estimation of the product of success probabilities is considered using a family of scaled squared error loss functions. The product of success probabilities can be interpreted as the probability that a series system functions, and this is a major application of estimation of products. For this problem, the best nonrandom allocation is found, myopic allocation rules are derived, and limiting properties of myopic and optimal rules are given. The limiting form of the myopic rule suggests an easily used adaptive rule, and this rule's performance is noted.

2. Bayesian preliminaries. For $i = 1, \ldots, I$, suppose that given $p_i, X_{i1}, \ldots, X_{im}$ are independent identically distributed Bernoulli random variables. Also, assume that given the p_i's, X_{ij}, and X_{rk} are independent when $i \neq r$ (from different populations). The goal is estimation of $\theta = \prod p_i$ under a member of the family of scaled squared error loss functions:

(1)
$$L\left(\theta, \hat{\theta}\right) = \prod_{i=1}^{I} p_i^{t_i} q_i^{s_i} \left(\theta - \hat{\theta}\right)^2,$$

where $q_i \equiv 1 - p_i$. This family includes the usual squared error loss at $t_i = s_i = 0$, and also includes ratio loss, $L\left(\theta, \hat{\theta}\right) = \left(\hat{\theta}/\theta - 1\right)^2$, at $t_i = -2$, $s_i = 0$.

Results will be obtained for squared error loss, then generalized to any member of the family (1). Results for ratio loss will be specially noted.

Assume independent natural conjugate beta prior distributions for the success probabilities: p_i are distributed beta (a_{io}, b_{io}). Dropping the subscripts, the density function for p is given by

$$\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1} (1-p)^{b-1}, \quad 0 < p < 1, a > 0, b > 0.$$

Aposteriori, the p_i's have independent beta (a_{im_i}, b_{im_i}) distributions, where $a_{im_i} = a_{io} + \sum_{j=1}^{m_i} X_{ij}, b_{im_i} = b_{io} + m_i - \sum_{j=1}^{m_i} X_{ij}$. Furthermore, assume that $a_{io} + t_i > 0$, and $b_{io} + s_i > 0$, so that there are no problems with integrability. The aposteriori independence is inherited from the apriori independence when the X_i's are conditionally independent. This simplifies loss computations. The case of dependent priors has not been studied, but promises to be an interesting problem.

Loss functions of the form (1) combine easily with the beta prior distributions, making generalization from squared error to the family

(1) very easy. If work is done using scaled squared error loss and prior distributions above, the same results can be obtained using squared error loss and modified parameters, $a_{io} + t_i$, $b_{io} + s_i$, in the prior distributions. Thus, while squared error and ratio loss initially look very different, results under each only differ by the effect of a factor of p_i^{-2} in the ith prior distribution. In general, properties of estimators and allocation rules under different losses in the family (1) will differ 'as much as' the parameters in the beta prior distributions influence the results. The asymptotic results are the same for all losses in (1), given integrability.

There are loss functions of interest which do not have form (1). One example is standardized loss: $\left(\theta - \hat{\theta}\right)^2 / \{\theta(1-\theta)\}$. The techniques used in this paper do not work for this loss.

The Bayes estimator of θ is not sensitive to the allocation, and letting $E_m(\cdot)$ denote expectation given observations up to stage m, the Bayes estimator under squared error loss is given by $\hat{\theta} = \prod_{i=1}^{I} \hat{p}_i$, where $\hat{p}_i = E_m(p_i) = a_{im}/(a_{im} + b_{im})$. For loss functions in the family (1),

$$\hat{p}_i = \frac{E_m\left(p_i^{t_i+1} q_i^{s_i}\right)}{E_m\left(p_i^{t_i} q_i^{s_i}\right)} = \frac{a_{im} + t_i}{a_{im} + b_{im} + t_i + s_i}.$$

For ratio loss, $\hat{p}_i = (a_{im} - 2)/(a_{im} - 2 + b_{im})$. Note that the results for loss in family (1) are obtained from the squared error results by simply changing a_{im} to $a_{im} + t_i$ and b_{im} to $b_{im} + s_i$.

Finally, for use in the next section, an expression for the posterior expected loss and Bayes risk using the Bayes estimator are derived for squared error loss. The posterior expected loss, given information up through stage $m = \sum m_i$ and using the Bayes estimator of θ, is

(2)
$$L^*(m) = \prod_{i=1}^{I} E_m\left(p_i^2\right) - \prod_{i=1}^{I} \left(E_m(p_i)\right)^2.$$

It is convenient to express this posterior expected loss in the form

(3)
$$L^*(m) = \prod_{i=1}^{I} E_m\left(p_i^2\right) - \prod_{i=1}^{I} \left\{E_m\left(p_i^2\right) - \frac{E_m(p_i q_i)}{m_i + a_{io} + b_{io}}\right\}$$

This identity can be confirmed by direct computation. In the case where the m_i are not random, the Bayes risk is equal to $E(L^*(m))$, the expectation of (3) over the data, and is given by

(4)

$$E(L^*(m)) = \prod_{i=1}^{I} E\left(p_i^2\right) - \prod_{i=1}^{I}\left\{E\left(p_i^2\right) - \frac{E(p_i q_i)}{m_i + a_{io} + b_{io}}\right\}.$$

This follows from expression (3), the independence across populations and the nonrandomness of the sample sizes. It is important to note that (4) holds for nonrandom allocation of $m = \sum m_i$ observations, and is not appropriate for sequential or adaptive procedures.

3. Allocation rules. In this section, the optimal nonrandom allocation rule is derived in 3.1, the myopic sequential allocation rule is derived in 3.2 and references for the optimal sequential allocation rule are given in 3.3. Derivation and results are first done for squared error, then the obvious generalization to other loss functions of form (1) is noted.

In all cases, a fixed total number of observations denoted by k is to be allocated. The vector $\mathbf{m} = (m_1, \ldots, m_I)$ gives the number allocated to populations $1, 2, \ldots, I$, respectively, at stage $m \leq k$, and is constrained by $\sum_{i=1}^{I} m_i = m$, $m \leq k$.

3.1 The optimal nonrandom allocation. Nonrandom allocation rules are simply rules that give the number of observations to take from each population before any data is obtained. The rules are not sequential and are not data dependent. Outside the Bayesian framework, optimal nonrandom allocation rules typically depend on unknown parameters, and thus are often used to suggest an adaptive version of the 'optimal' allocation. Within the Bayesian framework, the optimal nonrandom allocation rule is derived by finding the allocation that minimizes the Bayes risk (over all nonrandom allocations). This is a straightforward minimization problem, and the answer (optimal allocation) depends on the prior distribution and involves prior expectations of success probabilities and their squares.

Recall that a *total* of k observations will be allocated, and let k_i be the number allocated to population i, $i = 1, \ldots, I$ and where $k = \sum k_i$. Now express (4), the Bayes risk, at the kth (final) stage: $E(L^*(k))$. Treating the k_i's as continuous variables and minimizing $E(L^*(k))$ under the constraint that $\sum_{i=1}^{I} k_i = k$ yields the following allocation:

(5)
$$k_i + a_{io} + b_{io} \text{ is proportional to } \left\{E(p_i q_i) / E\left(p_i^2\right)\right\}^{\frac{1}{2}}.$$

Evaluating these expectations for squared error loss yields the allocation

(6)
$$k_i + a_{io} + b_{io} \text{ proportional to } \{b_{io} / (a_{io} + 1)\}^{\frac{1}{2}}.$$

From (6), the optimal nonrandom allocation under the general scaled squared error loss is

$$k_i + a_{io} + t_i + b_{io} + s_i \text{ proportional to } \{(b_{io} + s_i) / (a_{io} + t_i + 1)\}^{\frac{1}{2}}.$$

The optimal nonrandom allocation for ratio loss is $k_i + a_{io} - 2 + b_{io}$ proportional to $\{(b_{io} / (a_{io} - 2 + b_{io}))\}^{\frac{1}{2}}$.

Note that the best nonrandom allocation is totally dependent on the choice of prior parameters. The k_i notation will always refer to the nonrandom allocation of k, the total number of available observations.

3.2 Myopic allocation rules. The myopic or one-step-ahead allocation rules are purely sequential in that at each stage (beginning with stage 1), a decision is made about where to take the next observation, and that decision will be to take the action (observe from the population) that minimizes the posterior expected loss for the next stage given present information.

The tractable form of the posterior expected loss (3) allows the actual computation of myopic rules for this case. The derivation goes as follows. Recall that m represents an "intermediate" stage and that a total number of k observations are to be allocated. Let e_i be an I-dimensional vector with 1 in the ith component, and 0 in all other components. For example, $e_2 = (0, 1, 0, \ldots, 0)$. Consider $L^*(\mathbf{m} + e_i)$ as the posterior expected loss using the Bayes estimator when the next observation (at stage $m+1$) is taken from population i. Then the expectation of squared error loss given information up to stage m is

(7)
$$E_m\left(L^*\left(\mathbf{m}+e_i\right)\right)=L^*\left(\mathbf{m}\right)-\prod_{i=1}^{I}\left[E_m\left(p_i\right)\right]^2 B_m\left(i,m_i\right),$$

where the score functions $B_m\left(\cdot,\cdot\right)$ are defined by

(8)
$$\begin{aligned} B_m\left(i,m_i\right) &= \frac{E_m\left(p_i,q_i\right)/\left[E_m\left(p_i\right)\right]^2}{\left(m_i+a_{io}+b_{io}\right)\left(m_i+a_{io}+b_{io}+1\right)} \\ &= \frac{\left(b_{im}/a_{im}\right)}{\left(a_{io}+b_{io}+m_i+1\right)^2}. \end{aligned}$$

The myopic procedure takes an observation from population i at stage $m+1$ if $E_m\left(L^*\left(\mathbf{m}+e_i\right)\right)\leq E_m\left(L^*\left(\mathbf{m}+e_j\right)\right)$ for all j not equal to i, with ties broken arbitrarily. That is, observe from the population where the loss using this additional observation is expected to be lowest at the next stage given present information. Using the expression in (7), the myopic allocation will take an observation from the population with the highest score function $B_m\left(j,m_j\right)$.

Note that if an observation is taken from population i at stage m, then only its score function, $B_{m+1}\left(i,m_i\right)$ is affected; the other score functions are the same for the next stage, $B_{m+1}\left(j,m_j\right)=B_m\left(j,m_j\right)$. Furthermore, the numerators of these score functions converge almost surely as the number of stages tends to infinity. These properties give the limiting behavior (as the number of stages tends to infinity) of myopic allocation as described in Section 4.

For loss (1) the myopic allocation maximizes score function

$$\frac{\left(b_{im}+s_i\right)/\left(a_{im}+t_i\right)}{\left(a_{io}+b_{io}+m_i+t_i+s_i\right)^2},$$

whereas ratio loss maximizes score function

$$\frac{b_{im}/a_{im} - 2}{\left(a_{io} + b_{io} + m_i - 2\right)^2}.$$

3.3. Optimal allocation rules. Since the total number of observations to be allocated is fixed at k, and Bayesian procedures have a Markov property, the technique of dynamic programming or backward induction can be used to derive the optimal allocation rule. The recursive equations needed to define the optimal allocation will not be presented here, but the general set-up can be found in Page Shapiro (1985), and specific equations for $I = 2$ can be found in Hardwick and Stout (1992). In the past, even computer implementation of such solutions has been difficult to impossible. However, that situation is changing with improved computers, and Hardwick and Stout (1992) are presently implementing these dynamic programming solutions for the case of square error loss. Still, even with implementation, the derivation of properties of rules obtained in this way is difficult. However, when the optimal rule is known and its Bayes risk can be computed for small to moderate sample sizes, then its risk can be compared with risks of simpler adaptive rules, and the moderate sample size performance of such rules can then be evaluated. Hence, exact computation of the optimal Bayes risk is very important, and there is a wealth of problems needing this computation.

4. Limiting behavior and comparisons.

4.1. Limiting behavior. Myopic allocation is defined in terms of score functions $B_m\left(i, m_i\right)$, and the numerator of these score functions tends almost surely to q_i/p_i, which implies that the limiting score functions tend almost surely to the square root of the odds of failure, $\left(q_i/p_i\right)^{1/2}$. The limiting form of the Bayes risk can also be derived for myopic allocation:

(9)
$$\lim_{k \to \infty} kE\left(L^*\left(k\right)\right) = \theta^2 \left\{ \sum_{i=1}^{I} \left(\frac{q_i}{p_i}\right)^{\frac{1}{2}} \right\}^2 .$$

Furthermore, the limiting form of the Bayes risk for the optimal allocation can be shown to satisfy (9), so that the myopic rule and other rules with this same limit (9) are asymptotically optimal.

The results suggest a natural adaptive rule: allocate to get the number from each population proportional, in the limit, to the square root of the odds failure. The odds ratio would be estimated at each stage, and allocation carried out using the estimates. It can be shown that such allocations have the same limiting risk as the optimal allocation, and are thus, asymptotically optimal. For details in a general setting, see Page Shapiro (1985).

4.2. Comparisons. The easiest rule to implement is the nonrandom allocation rule described in Section 3.1. However, this rule is totally prior dependent and nonadaptive, and thus, the most sensitive to the prior distributions. The optimal rule requires dynamic programming, and is just beginning to seem tractable for a small number of populations. The most attractive of the three rules on the basis of implementation is the myopic rule, or its suggested adaptive form which allocates proportional to updated estimates of the square root of the odds of failure.

The myopic rule and its adaptive version are asymptotically optimal, but the question of 'how optimal' they are for small and moderate sample sizes remains. Of course, any answer will depend on the prior parameters to some extend, and thus will also depend on the loss chose from the family (1). Recent work of Hardwick and Stout (1992) for squared error loss, and $I = 2$, suggests that the procedures are indeed very good. For sample sizes as low as 20, they have found Bayes risks for optimal and myopic are within 1% of each other. This is very encouraging because the myopic rule (and its adaptive version) is easy to define and to use, and its properties are easy to derive.

References

BERRY, D. (1977). Optimal sampling schemes for estimating system reliability by testing components-I: fixed sample size. *Journal of the American Statistical Association* **69** 485-491.

COCHRAN, W. (1977). *Sampling Techniques.* New York: John Wiley.

HARDWICK, J. and STOUT, Q. (1992). Optimal allocation for estimating the product of two means. *Proceedings of the 23rd symposium on the Interface: Computing Science and Statistics.* Fairfax, VA: Interface Foundation of North America, 592-596.

KELLEY, T. (1977). Sequential Bayes estimation of the difference between means. *The Annals of Statistics* **5** 379-384.

LOUIS, T. (1975). Optimal allocation in sequential tests comparing the means of two Gaussian populations. *Biometrika* **62** 359-369.

NOBLE, W. (1992). *First Order Allocation,* Department of Statistics and Probability, Michigan State University (doctoral dissertation).

PAGE, C. (1987). Sequential designs for estimating products of parameters. *Sequential Analysis* **6** 351-371.

PAGE SHAPIRO, C. (1982). Sequential allocation with applications to Bayesian simultaneous estimation. *Sequential Analysis* **1** 145-158.

PAGE SHAPIRO, C. (1985). Allocation schemes for estimating the product of positive parameters. *Journal of the American Statistical Association* **80** 449-454.

DEPARTMENT OF STATISTICS AND PROBABILITY
MICHIGAN STATE UNIVERSITY
A415 WELLS HALL
EAST LANSING, MI 48824