In *Sequential Analysis* **15** (1996), pp. 71–90. (with black/white diagrams)

Optimal Allocation for Estimating the Mean of a Bivariate Polynomial

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

Keywords: nonlinear estimation, sequential design, robustness, myopic, hyperopic, dynamic programming, forward induction, adaptive allocation, product, Bayesian, fault tolerance

ABSTRACT

Suppose we wish to estimate the mean of some polynomial function of random variables from two independent Bernoulli populations, the parameters of which, themselves, are modeled as independent beta random variables. It is assumed that the total sample size for the experiment is fixed, but that the number of experimental units observed from each population may be random. This problem arises, for example, when estimating the fault tolerance of a system by testing its components individually.

Using a decision theoretic approach, we seek to minimize the Bayes risk that arises from using a squared error loss function. The Bayes estimator can be determined in a straightforward manner, so the problem of optimal *estimation* reduces, therefore, to a problem of optimal *allocation* of the samples between the two populations. This can be solved via dynamic programming. Similar programming techniques are utilized to evaluate properties of a number of ad hoc allocation strategies that might also be considered for use in this problem. Two sample polynomials are analyzed along with a number of examples indicating the effects of different prior parameter settings.

The effects of differences between prior parameters used in the design and analysis stages of the experiment are also examined. For the polynomials considered, the adaptive strategies are found to be especially robust. We discuss computational techniques that facilitate such analyses by permitting rapid re-evaluation of strategies. Capabilities of this sort encourage people to explore designs more fully and to consider them from a number of different viewpoints.

1 INTRODUCTION

Consider a widget whose functionality depends on the error-free operation of two types of independent parts. In order for the widget to function, some specified combination of the two types of parts must be working. For example, the widget might be a communication system consisting of switches and links, which performs properly only if there is a fault-free path from one end to the other. We can model the behavior of the parts using independent Bernoulli random variables. Let Population 1 contain the parts of the first kind and Population 2 contain parts of the second kind. Then define

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

and for Population 2, let

$$Y_i = \begin{cases} 0, & \text{if part } i \text{ from Population 2 is defective;} \\ 1, & \text{if part } i \text{ from Population 2 is ok,} \end{cases} \qquad i = 1, \ldots,$$

where the X's and Y's are mutually independent. If we let p_1 and p_2 denote the respective probabilities of X and Y being defect free, then

$$X_i \ \text{ i.i.d. } B(1, p_1) \quad i = 1, \ldots \qquad \text{and} \qquad Y_j \ \text{ i.i.d. } B(1, p_2) \quad j = 1, \ldots. \tag{1}$$

Our goal is to estimate the chance that the widget will work by doing appropriate sampling from the population of each part type. In the model, it is assumed that the sample size for the experiment is a fixed number, N. However, in several of the sampling designs considered, the number of observations sampled from the different populations, n_1 from Population 1 and n_2 from Population 2, is random. We refer to designs in which n_1 and n_2 are random as *sequential* allocation procedures, and to designs in which n_1 and n_2 are determined prior to the taking of any observations as *fixed* sample size allocation rules or simply as fixed rules.

Let $\delta = (\delta_1, \delta_2, ...\delta_N)$ represent a *sampling procedure* where

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

for $m = 1, ..., N$, and where δ_{m+1} depends only on the information available up until time m. For convenience, we denote information accrued by time m as

$$\mathcal{D}_m = (x_1, \ldots, x_{m_1}; y_1, \ldots, y_{m_2}; \delta(m_1, m_2))$$

where $\delta(m_1, m_2)$ denotes the initial sampling decisions

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

and where $$\sum_{i=1}^{m_1+m_2} \delta_i = m_1 \quad \text{for any} \quad 0 < m_1 + m_2 = m \leq N.$$

We approach this problem using Bayesian decision theory. This approach is taken not only for reasons of mathematical tractability but also because Bayesian methods are gaining acceptance in industrial estimation settings. While one focus of this research is to locate optimal allocation rules, we are equally interested in the evaluation of simpler allocation rules that are easier to implement and perhaps more intuitive.

The problem of determining good sampling strategies for estimating one particular polynomial, the product, has been examined by several authors. Berry (1974) discusses optimal fixed allocation rules for squared error loss plus cost when there are 2 or more independent Bernoulli means. Shapiro (1985) and Page (1995) consider fixed and sequential rules for the same problem, but address more general distributions and allow for generalized squared error loss function. Rekab (1993) examines fixed and sequential rules for the problem using squared error loss. In all of these papers, the behavior of the various rules is explored from an asymptotic perspective. What we add to this series of studies is the optimal sequential allocation rule and exact evaluations of various suboptimal rules, including those proposed in Shapiro (1985), Page (1995) and Rekab (1993). Furthermore, we generalize the problem to include all bivariate polynomials and address issues of design robustness.

In the next section we define the framework used to formulate the problem. In Section 3 we introduce sampling procedures (also called allocation rules), and define the procedures that are considered. In Section 4 we apply the set-up to two specific polynomials and evaluate the efficiency of the different allocation rules.

It is sometimes argued that 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 Sections 4 and 5, therefore, we have given some extra attention to the issues of interpretation of prior parameters and the robustness of the Bayesian approach.

In Section 6 we outline the computational techniques utilized. These techniques are important because they dramatically reduce the time and space required. In particular, for robustness studies one needs to evaluate a large number of alternatives and to re-evaluate a design numerous times. Thus one needs algorithms and implementations which are a couple of orders of magnitude faster than can be tolerated when evaluating a single design on a single criterion.

Finally, in an appendix we give an approximation to the Bayes risk when fixed sample sizes are used. While not used herein, the approximations are very accurate and greatly reduce the computations required.

2 THE BAYES RISK

As mentioned, it is assumed that prior information is available and may be modeled in the form of independent 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:

$$p_1 \sim \text{Be}(a_1, b_1) \qquad \text{and} \qquad p_2 \sim \text{Be}(a_2, b_2). \tag{2}$$

We use the notation $\xi(p_1, p_2) = [\text{Be}(a_1, b_1), \text{Be}(a_2, b_2)]$ to indicate specific prior parameter configurations.

Suppose we wish to estimate the mean of the polynomial $\theta = \theta(p_1, p_2)$, where

$$\theta(p_1, p_2) \;=\; \sum_{r=0}^{u_1}\sum_{s=0}^{u_2} c_{rs} p_1^r p_2^s \quad \text{for some } \; u_1, u_2 \in \mathcal{N} \quad \text{and} \;\; c_{rs} \in \Re.$$

Let $\tilde{\theta}_N$ be any estimate of θ based on N observations. We seek a minimum variance estimate, so our objective is to minimize the integrated risk of $(\theta, \tilde{\theta}_N)$ under quadratic loss $L(\theta, \tilde{\theta}_N) = (\theta - \tilde{\theta}_N)^2$. The *integrated risk*, $\mathcal{R}_N$, based on a sample of size N, is the expected loss of $(\theta, \tilde{\theta}_N)$

$$\mathcal{R}_N(\xi, \tilde{\theta}_N) = \mathbf{E}^{\xi}[L(\theta, \tilde{\theta}_N)] = \mathbf{E}^{\xi}[(\theta - \tilde{\theta}_N)^2], \tag{3}$$

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

Minimizing the integrated risk can 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 solution to the first is well known.

Suppose that at time $m \le N$, we have obtained m_i observations from Population i, $i \in \{1, 2\}$, with s_i of these being successes and f_i being failures. Thus $m = m_1 + m_2$, $m_i = s_i + f_i$, and the posterior density function, ξ_m, on (p_1, p_2) is the product of the individual posterior densities

$$(p_1 \mid s_1, f_1) \;\sim\; \text{Be}(a_1 + s_1, b_1 + f_1) \qquad\qquad \text{and} \qquad\qquad (p_2 \mid s_2, f_2) \;\sim\; \text{Be}(a_2 + s_2, b_2 + f_2).$$

Note that the vector (s_1, f_1, s_2, f_2) is sufficient for p_1 and p_2.

For any m, the *optimal*, or *Bayes*, *estimate* of θ, denoted $\hat{\theta}_{m_1, m_2}$, is simply its posterior mean given by

$$\hat{\theta}_{m_1, m_2} = E^{\xi}[\theta(p_1, p_2) \mid \mathcal{D}_m] = \sum_{r=0}^{u_1}\sum_{s=0}^{u_2} c_{rs} \widehat{p_1^r}\, \widehat{p_2^s}.$$

The *Bayes risk* of a procedure δ, denoted $\mathcal{R}_N(\xi, \delta)$, is the expected loss when δ is the sampling procedure and $\hat{\theta}_{n_1, n_2}$ is the terminal estimate:

$$\mathcal{R}_N(\xi, \delta) = \mathbf{E}^{\xi}\left[(\theta(p_1, p_2) - \hat{\theta}_{n_1, n_2})^2\right], \tag{4}$$

where $n_1 = \sum_{i=1}^{N} \delta_i$, $n_2 = N - n_1$.

The problem that remains, then, is to find a procedure $\delta^* = \arg\min_{\delta} \mathcal{R}_N(\xi, \delta)$. Such a procedure is said to be *optimal*.

3 SAMPLING PROCEDURES

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 asymptotically. (See, for example, Shapiro (1985).) Here we consider the performance, advantages and disadvantages of 5 types of procedures. These include two fixed allocation schemes: equal allocation (EA) and best fixed allocation (BF); and three sequential schemes: myopic (MP), hyperopic (HP) and optimal (OS). In Section 4, we compute the efficiency of each procedure for various sample sizes and prior parameter settings, for two different example polynomials. The *efficiency* of a procedure δ is $\mathcal{R}_N(\xi, \delta^*)/\mathcal{R}_N(\xi, \delta)$, where δ^* is an optimal procedure.

3.1 FIXED ALLOCATION RULES

The easiest sampling rules to employ are those with predetermined sample sizes. In our examples, we locate the values of n_1 and n_2 that minimize Bayes risk among all fixed allocation rules. We refer to an optimal rule for the predetermined sample size problem as a *best fixed* or BF allocation rule. Since a natural question that arises is how well one fares when sampling from each population equally, we include equal allocation (EA) as a special case of fixed allocation in our examples.

We note that the solution of the minimization problem may indicate that one should use non-integral values of (n_1, n_2). However, this is simply an artifact of the mathematics, and thus our definition of BF procedures is based upon the optimal integral values of (n_1, n_2).

3.2 LOCAL SEQUENTIAL RULES

Rather than fixing n_1 and n_2 at the beginning of the experiment, one can expect to do better by repeatedly updating the estimate of the distribution of (p_1, p_2), and basing upcoming sampling decisions on the most recent information. In this section, "locally" sequential procedures are considered. A local procedure is one that decides which population to sample from next using only a small amount of computation at each stage. The simplicity of such rules makes them attractive, especially because they can incorporate accruing information.

To describe these procedures, it is helpful to utilize various ancillary risk functions. The *conditional expected risk* of procedure δ at Stage m, given an horizon of t, $m \leq t \leq N$, denoted $r_t(\xi_m, \delta)$, is given by

$$r_t(\xi_m, \delta) = \mathbf{E}^{\xi}\left[\left(\theta(p_1, p_2) - \hat{\theta}_t\right)^2 \mid \mathcal{D}_m\right],$$

where $\mathbf{E}^{\xi}(\cdot \mid \mathcal{D}_m)$ denotes expectation with respect to the posterior distribution of p_1 and p_2 at stage t given $\mathcal{D}_m$. The *interim risk* at state (s_1, f_1, s_2, f_2) of Stage m, given an horizon of t, $m \leq t \leq N$, denoted $\mathbf{I}_t^m(s_1, f_1, s_2, f_2)$, is given by

$$\mathbf{I}_t^m(s_1, f_1, s_2, f_2) = \min_{\delta \in \Delta(s_1, f_1, s_2, f_2)} r_t(\xi_m, \delta),$$

where $\Delta(s_1, f_1, s_2, f_2)$ is the set of all procedures that could reach (s_1, f_1, s_2, f_2). Thus, interim risk is the minimum conditional expected risk incurred if one were to start at Stage m, given (s_1, f_1, s_2, f_2), and proceed optimally to Stage t. Note that the interim risk at Stage N is simply the posterior risk computed for N:

$$\mathbf{I}_N^N(s_1, f_1, s_2, f_2) = r_N\left(\xi_N, \delta\right) \quad \text{for any } \delta \in \Delta(s_1, f_1, s_2, f_2).$$

A local procedure that is particularly easy to compute is the *myopic* (MP) or *one-stage look-ahead* rule. To generate the MP rule, we use the posterior risk beginning at Stage $m < N$ and ending at Stage $m + 1$. Let

$$r^i_{m+1}(\xi_m, \delta) = \mathbf{E}^{\xi} \left[\left(\theta(p_1, p_2) - \hat{\theta}_{m+1} \right)^2 \mid \mathcal{D}_m \text{ and } \delta_{m+1} = i \right] ,$$

for $i = 1, 2$. Then the myopic rule is to allocate next from Population i if

$$r^i_{m+1}(\xi_m, \delta) < r^{3-i}_{m+1}(\xi_m, \delta) .$$

In cases where the expected risks are the same, one can either randomize or sample from the population with the fewer number of observations so far. We use the former method in our examples since randomization is desirable in many experimental settings.

The second local rule that we consider is an adaptive version of the best fixed rule. An application of this rule is discussed in Rekab (1993). To generate the rule, at each Stage m one views the problem in the context of Section 3.1 by determining the optimal fixed values to sample from each population, given that you were starting at Stage m. This means locating nonnegative integers n_{1m} and n_{2m} such that $n_{1m}+n_{2m} = N-m$ and such that taking n_{im} more observations from Population i would minimize $r_N(\xi_m, \delta)$ among all fixed-allocation rules. There are two natural ways to formulate the sequential rule in terms of the n_{im}:

1. Sample from Population i if $n_{im} > n_{(3-i)m}$, or
2. Sample from Population i with probability $n_{im}/(n_{1m} + n_{2m})$.

In both cases, we randomize with probability $\frac{1}{2}$ when $n_{1m} = n_{2m}$. We refer to these rules as *hyperopic* because at each stage they look to the end of the experiment and make determinations as to what the "best" sample sizes should be. Because alternative 2 introduces randomization into the sampling process, we prefer it to 1, and thus the hyperopic rule (HP) used in Sections 4 and 5 is alternative 2.

3.3 OPTIMAL SEQUENTIAL RULE

The optimal sequential rule (OS) is found via dynamic programming. This method is well known and consists of evaluating and comparing all possible interim risks that can occur beginning at Stage N and moving backwards to Stage 0. Interim risk obeys the recursive optimality principle, and thus, once the interim risk has been determined for all possible sets of outcomes at Stage $m + 1$, it can then be determined for all sets at Stage m, for $m = 0, \ldots, N$. As with the MP and HP procedures, we randomize in the case of ties.

4 EXAMPLES OF ESTIMATING POLYNOMIALS

We consider two example polynomials in this section:

1. $\theta(p_1, p_2) = p_1 p_2$

2. $\theta(p_1, p_2) = (p_1 + p_2 - p_1 p_2) p_1$

The first case is useful because one can easily go through many of the details of the analysis and a number of authors have shown interest in the problem (Berry (1974), Hardwick and Stout (1993), Rekab (1993), Shapiro (1985)). In the second case, the details are more complex, but we include them to indicate that a more general problem is still tractable. An application of this polynomial to a fault-tolerance problem is given in Section 4.2.

4.1 EXAMPLE 1: THE PRODUCT

We begin with some notation. Let $\alpha^{\bar{r}} = \alpha \cdot (\alpha+1) \cdots (\alpha + r - 1)$, and let $\mu_r(v)$ denote the r^{th} moment of a random variable v, i.e., $\mathbf{E}(v^r)$. Recall that for a random variable v from the distribution Be(α, β),

$$\mu_r(v) = \frac{\Gamma(\alpha + r)\,\Gamma(\alpha+\beta)}{\Gamma(\alpha+\beta+r)\,\Gamma(\alpha)} = \frac{\alpha^{\bar{r}}}{(\alpha+\beta)^{\bar{r}}}.$$

Applying some algebra, we get a simple expression for the Bayes risk of $\theta = p_1\, p_2$ when n_1 and n_2 are non-random:

$$\begin{aligned} \mathcal{R}_N\left(\xi, \delta_{n_1,n_2}\right) &= \frac{\mu_2(p_2) a_1 b_1}{(a_1+b_1)^{\bar{2}}(a_1+b_1+n_1)} + \frac{\mu_2(p_1) a_2 b_2}{(a_2+b_2)^{\bar{2}}(a_2+b_2+n_2)} \\ &\quad - \frac{a_1 b_1}{(a_1+b_1)^{\bar{2}}(a_1+b_1+n_1)} \frac{a_2 b_2}{(a_2+b_2)^{\bar{2}}(a_2+b_2+n_2)}. \end{aligned} \tag{5}$$

Equation (5) can be minimized with respect to n_1 by solving a quadratic equation, and it yields the best fixed allocation procedure, BF. Further, by updating the distribution as results accrue, (5) provides the form of the hyperopic (HP) allocation procedure. By updating and setting $n_1 = n_2 = 0$ one obtains the posterior risk needed to determine the optimal sequential procedure (OS).

The third term on the right hand side of (5) is an order of magnitude smaller than either of the first two terms. Thus, a very close approximation to the best fixed allocation rule is attained if one selects n_1 such that

$$\frac{a_1+b_1+n_1}{a_2+b_2+n_2} = \frac{\sqrt{\mathbf{E}^{\xi}[p_2^2\, p_1(1-p_1)]}}{\sqrt{\mathbf{E}^{\xi}[p_1^2\, p_2(1-p_2)]}}. \tag{6}$$

The results reported in this paper do not use this approximation, but it was used in Rekab (1993). A generalization of this first-order approximation scheme appears in the appendix.

The myopic rule (MP) for this example is also quite simple. Suppose we have observed (s_1, f_1, s_2, f_2). If $\alpha_i = a_i + s_i$, $\beta_i = b_i + f_i$, then the posterior distribution of p_i is Be(α_i, β_i), $i = 1, 2$. From (5) one finds that we should sample next from Population 1 if

$$\begin{aligned} &\frac{(\alpha_2+1)\beta_1}{\alpha_1+\beta_1+1} + \frac{(\alpha_1+1)\beta_2}{\alpha_2+\beta_2} + \frac{\beta_1\beta_2}{(\alpha_1+\beta_1+1)(\alpha_2+\beta_2)} \\ &\quad < \frac{(\alpha_2+1)\beta_1}{\alpha_1+\beta_1} + \frac{(\alpha_1+1)\beta_2}{\alpha_2+\beta_2+1} + \frac{\beta_1\beta_2}{(\alpha_1+\beta_1)(\alpha_2+\beta_2+1)}. \end{aligned}$$

If the inequality is reversed, then sample from Population 2, and randomize if there is a tie.

Sample Size N	Efficiency (Risk Ratios) and % on Pop 1				
	OS	MP	HP	BF	EA
20	1.000 50%	1.000 50%	0.997 50%	0.892 50%	0.892 50%
50	1.000 50%	1.000 50%	0.999 50%	0.870 50%	0.870 50%
100	1.000 50%	1.000 50%	1.000 50%	0.861 50%	0.861 50%
∞	1.000 50%	1.000 50%	1.000 50%	0.847 50%	0.847 50%

Table 1: $(p_1, p_2) \sim$ [Be(1,1),Be(1,1)]

4.1.1 ILLUSTRATIVE CONFIGURATIONS AND HEURISTICS

Here we examine five different prior parameter settings for the product of means problem. These examples provide a number of insights. On the one hand, they suggest how prior parameter configurations may be used to characterize prevailing experience and how different prior configurations can affect the sampling schemes discussed in Section 3. They also provide concrete examples of differences among the risks for moderate sample sizes and those obtained in the limiting cases. In Tables 1–3, for various sample sizes, we present the efficiencies of each rule relative to the OS rule, along with the percent sampled from Population 1.

Except for the last row of each table, which contains asymptotic results, all values in the tables are computed using methods described in Section 6. For the asymptotic values, Shapiro (1985) proved that if the true values of the parameters are p_1^* and p_2^*, then the optimal rule will almost surely sample from Population i proportional to $\sqrt{\frac{1-p_i^*}{p_i^*}}$. Further, if $\mathcal{R}_n^{OS}$ is the expected risk of the OS procedure for a sample of size n, then

$$\lim_{n\to\infty} n\mathcal{R}_n^{OS} = \mu_1(p_1)\mu_2(p_2) + \mu_1(p_2)\mu_2(p_1) - 2\mu_2(p_1)\mu_2(p_2) + 2\prod_{i=1}^{2} \frac{\Gamma(a_i+1.5)\Gamma(b_i+0.5)}{\Gamma(a_i)\Gamma(b_i)(a_i+b_i)^2}.$$

Shapiro also showed that MP asymptotically samples in the same proportions as OS, and has asymptotic efficiency 1. Using similar methods, one can show that the asymptotic efficiency of HP is also 1. The asymptotic proportions for BF can be easily derived from Equation (6), which indicates that Population i should be sampled from in proportion to $\sqrt{\frac{b_i}{a_i+1}}$. This result was noted in Shapiro (1985), where it was also shown that the efficiency of BF is at least 0.5.

For our first example, in Table 1, we consider the most commonly used prior parameter configuration: [Be(1,1), Be(1,1)]. In this configuration, BF and EA are the same. Note that as the sample size increases, the efficiency of BF decreases, as is to be expected.

The second example is relevant when we wish to emphasize differing degrees of faith in our prior knowledge. Take the case depicted in Table 2 with $(p_1, p_2) \sim$ [Be(20,10), Be(2,1)]. The mean

	Efficiency (Risk Ratios) and % on Pop 1					Efficiency (Risk Ratios) and % on Pop 1				
N	OS	MP	HP	BF	EA	OS	MP	HP	BF	EA
20	1.000 10%	1.000 10%	0.998 11%	0.979 00%	0.793 50%	1.000 74%	0.999 74%	0.997 72%	0.926 70%	0.862 50%
50	1.000 28%	1.000 27%	0.998 30%	0.937 30%	0.882 50%	1.000 71%	1.000 71%	0.999 70%	0.907 66%	0.846 50%
100	1.000 40%	1.000 40%	1.00 40%	0.930 42%	0.918 50%	1.000 70%	1.000 70%	1.000 70%	0.898 64%	0.838 50%
∞	1.000 53%	1.000 53%	1.000 53%	0.918 54%	0.911 50%	1.000 69%	1.000 69%	1.000 69%	0.885 63%	0.823 50%
	$(p_1, p_2) \sim$ [Be(20,10), Be(2,1)]					$(p_1, p_2) \sim$ [Be(1,2), Be(2,1)]				

Table 2: Nonuniform Distributions

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. The effect of such modeling is that the optimal sampling rule will sample more often from the population with the smaller initial parameters, since that population will have higher variance and hence contribute more risk. Thus, for small N we sample proportionally more from Population 2 than from 1. As the total sample size increases, the proportions sampled from the two populations approach each other since the data eventually take over. The optimal proportion of observations from Population 1 makes a dramatic shift from 10% (when $N = 20$) to 40% (when $N = 100$), with the limiting optimal proportion being 53%. Both MP and HP are virtually fully efficient in this example.

Next, consider the case in which $(p_1, p_2) \sim [\mathrm{Be}(1, 2), \mathrm{Be}(2, 1)]$. With this prior configuration, one might expect EA to be nearly optimal since the distributions are mirror images of one another. As is evident from Table 2, however, this is not the case. More or less independently of N, the OS samples approximately 70% from Population 1, and EA 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, but the prior variances are the same, an optimal rule will tend to sample more from the population with the lower expected mean. This rule is intuitively appealing when one notices that to estimate the product xy of two positive quantities, if $x < y$, then an absolute measurement error of ϵ in x affects the answer more than the same size error in y.

Our last two examples in Table 3, one in which we have $(p_1, p_2) \sim [\mathrm{Be}(0.1, 0.01), \mathrm{Be}(1, 1)]$ and the other in which $(p_1, p_2) \sim [\mathrm{Be}(0.01, 0.1), \mathrm{Be}(1, 1)]$, may appear at first to be pathological since the small values of a_1 and b_1 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 of the distribution that is of most interest. In the case where $p_1 \sim \mathrm{Be}(0.1, 0.01)$, we have a *u*-shaped distribution where the probability that the process is in tolerance is ten times

	Efficiency (Risk Ratios) and % on Pop 1					Efficiency (Risk Ratios) and % on Pop 1				
N	OS	MP	HP	BF	EA	OS	MP	HP	BF	EA
20	1.000 17%	0.998 15%	0.995 14%	0.913 15%	0.626 50%	1.000 92%	0.997 92%	0.993 77%	0.822 35%	0.754 50%
50	1.000 15%	0.997 14%	0.997 13%	0.892 12%	0.583 50%	1.000 92%	0.997 92%	0.997 81%	0.793 32%	0.707 50%
100	1.000 14%	0.997 13%	0.998 12%	0.878 12%	0.565 50%	1.000 91%	0.998 91%	0.998 83%	0.777 31%	0.685 50%
∞	1.000 9%	1.000 9%	1.000 9%	0.818 12%	0.517 50%	1.000 91%	1.000 91%	1.000 91%	0.730 31%	0.636 50%
	$(p_1, p_2) \sim$ [Be$(0.1, 0.01)$, Be$(1, 1)$]					$(p_1, p_2) \sim$ [Be(0.01,0.1), Be(1,1)]				

Table 3: U-Shaped Distributions

greater than the probability that it is out of 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 $p_2 \sim$ Be$(1, 1)$, we know that we should sample more from the population with the smaller expected mean (though the variance also has an effect). 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 one would expect.

An interesting feature of these two examples is illustrated by the behavior of BF. Whereas the two cases under consideration are mirror images of one another, the sampling behavior indicated by BF does not remotely mimic the mirror image set-up. It is also worth noting that the worst efficiency of BF that we observed in any of our examples was 73% which is a good deal higher than the 50% lower bound for efficiency of this rule. For discussion of the situations in which this bound is approached, see Shapiro (1985).

4.2 EXAMPLE 2: A SERIES/PARALLEL POLYNOMIAL

We now consider a polynomial representing a slightly more complex configuration. Suppose we have a circuit involving two parts from Population 1 (Part 1A and 1B) and one part from Population 2 (Part 2). In order for this circuit, illustrated in Figure 1, to work, we must have *either* Part 1A or Part 2 work *and* Part 1B work. Thus, the probability that a randomly selected circuit will work is $\mathbf{P}\{(X \cup Y) \cap X\}$ and the parametric function that we are trying to estimate is

$$p_1^2 + p_1 p_2 - p_1^2 p_2.$$

As before, to develop the allocation rules, we need expressions for the various risk and ancillary functions. To handle this more complex example, we introduce a bit more notation.

Let $$\nu_{k\ell}(p_i) = \mathbf{E}^\xi \left[\mathbf{E}^\xi(p_i^k \mid \mathcal{D}_N)\mathbf{E}^\xi(p_i^\ell \mid \mathcal{D}_N)\right] \qquad \text{for} \quad i = 1, 2.$$

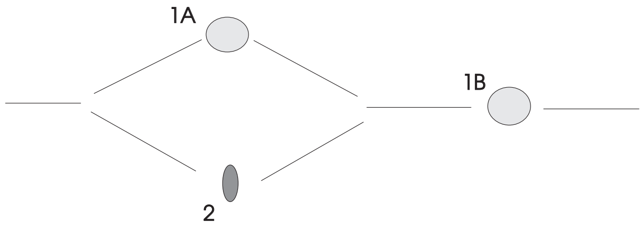

Figure 1: $(X \cup Y) \cap X$

Then $\quad \nu_{11}(p_i) = \mu_2(p_i) + \dfrac{\mu_2(p_i) - \mu_1(p_i)}{(a_i + b_i + n_i)}\,, \qquad \nu_{12}(p_i) = \mu_3(p_i) + 2\,\dfrac{\mu_3(p_i) - \mu_2(p_i)}{(a_1 + b_1 + n_1)}\,,$

and $\quad \nu_{22}(p_i) = \mu_4(p_i) + 4\left[\dfrac{\mu_4(p_i) - \mu_3(p_i)}{(a_1 + b_1 + n_1)}\right] + 2\left[\dfrac{\mu_4(p_i) - 2\mu_3(p_i) + \mu_2(p_i)}{(a_1 + b_1 + n_1)\,(a_1 + b_1 + n_1 + 1)}\right].$

One begins with the Bayes risk for any procedure δ,

$$\mathcal{R}_N(\xi, \delta) \;=\; \mathbf{E}^{\xi}\left[\left(p_1^2 + p_1 p_2 - p_1^2 p_2 - \mathbf{E}^{\xi}[p_1^2 + p_1 p_2 - p_1^2 p_2 \mid \mathcal{D}_N]\right)^2\right]. \tag{7}$$

One can then determine, after some algebra, that

$$\begin{aligned}\mathcal{R}_N(\xi, \delta(n_1, n_2)) \;=\;& \mu_2(p_1)\,\mu_2(p_2) + \mu_4(p_1)\,[1 - 2\mu_1(p_2) + \mu_2(p_2)] \\ &+ 2\mu_3(p_1)\,[\mu_1(p_2) - \mu_2(p_2)] - \Big\{\nu_{11}(p_2)[\nu_{11}(p_1) - 2\nu_{12}(p_1) + \nu_{22}(p_1)] \\ &+ \nu_{22}(p_1) + 2\mu_1(p_2)[\nu_{12}(p_1) - \nu_{22}(p_1)]\Big\},\end{aligned} \tag{8}$$

which is the Bayes risk when n_1 and n_2 are fixed in advance. To determine the BF rule, one substitutes the previously obtained expressions for the μ and ν factors, and then locates the optimal n_1 and n_2. The algebra is much less amenable than it was in the product of means example, but it is computationally easy because the minimization can be carried out in time linear in N regardless of the complexity of the polynomial.

Recall that from the fixed sample size version of the risk, (8), we can also generate the HP procedure. We use the posterior version of (7) to obtain the MP procedure and, by adding dynamic programming, the OS procedure as well.

Table 4 displays the efficiency and percent sampled from Population 1 for the five different allocation procedures applied to the polynomial $p_1 p_2 + p_1^2 - p_1^2 p_2$, using $N = 50$ and the same prior parameter settings that were used for the product of means example. Two properties of these procedures stand out for this polynomial. First, there is the strong bias towards sampling from Population 1, which is to be expected since p_1 affects the polynomial more than p_2. Second, all of the ad hoc rules except for EA perform extremely well.

5 ROBUSTNESS

As seen in Section 4.1.1, the beta family is a rich class of distributions that can be used in a variety of different ways to represent information available before an experiment. Another point to consider,

Sample Size $N = 50$	Efficiency (Risk Ratios)				
Prior Parameter	(% on Pop 1)				
Configuration	OS	MP	HP	BF	EA
$p_1 \sim$ Be$(1,1)$	1.000	1.000	1.000	0.988	0.664
$p_2 \sim$ Be$(1,1)$	(88%)	(88%)	(87%)	(88%)	(50%)
$p_1 \sim$ Be$(20,10)$	1.000	1.000	1.000	0.997	0.857
$p_2 \sim$ Be$(2,1)$	(82%)	(82%)	(82%)	(82%)	(50%)
$p_1 \sim$ Be$(1,2)$	1.000	1.000	1.000	0.993	0.680
$p_2 \sim$ Be$(2,1)$	(89%)	(89%)	(88%)	(88%)	(50%)
$p_1 \sim$ Be$(0.1,0.01)$	1.000	1.000	0.999	0.979	0.616
$p_2 \sim$ Be$(1,1)$	(99%)	(99%)	(98%)	(90%)	(50%)
$p_1 \sim$ Be$(0.01,0.1)$	1.000	1.000	0.999	0.980	0.622
$p_2 \sim$ Be$(1,1)$	(99%)	(99%)	(94%)	(90%)	(50%)

Table 4: Polynomial: $p_1p_2 + p_1^2 - p_1^2p_2$

however, is the extent to which disagreement regarding the parameters of the prior distribution will effect the inference process associated with an experiment. There are a number of ways to examine the sensitivity of a design with respect to its prior and no one method is always best. So, rather than carrying out an exhaustive study of the robustness characteristics of the sequential designs under consideration here, our intention is to show the ease with which such a study can be carried out computationally. We do this in Section 5.1 by considering one measure and contend that many other measures are based on computationally similar approaches. (See Etzioni and Kadane (1993) and the references therein for more in-depth comments on this issue.) In Section 5.2, we consider a different form of robustness, in that we compare the MP procedure in which the "design" and "analysis" priors disagree with the BF procedure in which both the design and analysis prior distributions do agree.

5.1 REANALYSIS

In this discussion, we employ a scenario similar to that utilized in Etzioni and Kadane (1993). Imagine that Marty, a visitor from Mars, arrives on Earth and designs and carries out an experiment using prior parameters Be(a_{M1}, b_{M1}) and Be(a_{M2}, b_{M2}). At a later time, suppose Veni arrives from Venus, and has prior parameters Be(a_{V1}, b_{V1}) and Be(a_{V2}, b_{V2}). Veni could either design and conduct a new experiment, or could acquire and re-analyze Marty's data using Veni's own prior parameters. We wish to assess the loss of efficiency to Veni of re-analyzing Marty's data instead of designing its own experiment. We assess this as follows:

(i) Generate each allocation rule (OS, MP, HP, BF) using Marty's prior.

(ii) Compute the Bayes risk for each method generated in (i), using Veni's prior.

(iii) Generate each allocation rule (OS, MP, HP, BF) using Veni's prior.

(iv) Compute the Bayes risk for each method generated in (iii), using Veni's prior.

(v) Define the *relative efficiency* of each method to be the ratio of risk of that method from (iv) with the risk of that method from (ii).

The idea is to consider the effect of having different design and analysis prior parameters in an experiment where Bayesian design and Bayesian analysis are both being used. Note that the EA rule is unaffected by the choice of prior so that all of the relative risks for EA are one.

We followed the steps just outlined for the problem of estimating the product of two means for a variety of different prior parameter settings for Marty, using a total sample size of $N = 100$. For each of the methods OS, HP, MP, BF, we computed the relative efficiency, (v), for 25 different versions of Marty's prior parameter configurations. We evaluated all combinations of

$$a_{M1} = 0.01,\ 0.1,\ 1,\ 10,\ 100 \quad b_{M1} = 0.01,\ 0.1,\ 1,\ 10,\ 100 \quad a_{M2} = 1 \quad b_{M2} = 1 \tag{9}$$

For all the experiments, Veni's prior parameters are $(p_1, p_2) \sim$ [Be(1,1),Be(1,1)]. In reviewing a variety of tables of relative efficiencies, we observed only two fundamental structures for the five different allocation methods. These are represented in Figures 2 and 3, which display the data for the OS risk in (ii) relative to the OS risk (iii) and the BF risk in (ii) relative to the BF risk in (iii), respectively. These figures are interpolated contour plots based on the relative efficiencies for the grid points in (9). For example, the data point .8653 in Figure 2, corresponding to $a_1 = 0.01$, $b_1 = 0.01$, has $a_{M1} = 0.01$, $b_{M1} = 0.01$, $a_{M2} = 1$, and $b_{M2} = 1$. It represents the relative efficiency of a OS procedure generated using this prior parameter configuration, but where the results were evaluated assuming a uniform distribution on (p_1, p_2).

One interpretation of the data is that, for sample sizes of 100, even if Marty's estimates of (a_1, b_1) differ by an order of magnitude from those of Veni, Marty's design will still be reasonably efficient for Veni, since the efficiency is typically above 90%. However, if Marty's estimates differ by two orders of magnitude then the reduction in efficiency may be unacceptable.

The relative efficiencies for the MP and HP rules mirror those for OS quite closely, so the plots for these rules look similar to Figure 2 and are omitted. For all three adaptive rules, (OS, HP, MP), the worst case scenarios occur when Marty's prior parameters are [Be(100,100), Be(1,1)] and Veni analyzed the data using [Be(1,1), Be(1,1)]. For the BF rule, however, we see a different picture. Throughout, the relative efficiency is much lower, and the worst case scenario occurs when Marty has the very extreme setup of [Be(.01,100), Be(1,1)].

5.2 ADAPTIVE vs. FIXED SAMPLE SIZE RULES

Another important form of robustness concerns adaptive versus fixed allocation rules. Here we show the efficiency of the myopic rule MP versus that of the best fixed rule BF. For the product polynomial, we create the MP rule using uniform priors, denoted as MP[Be(1,1), Be(1,1)]. We

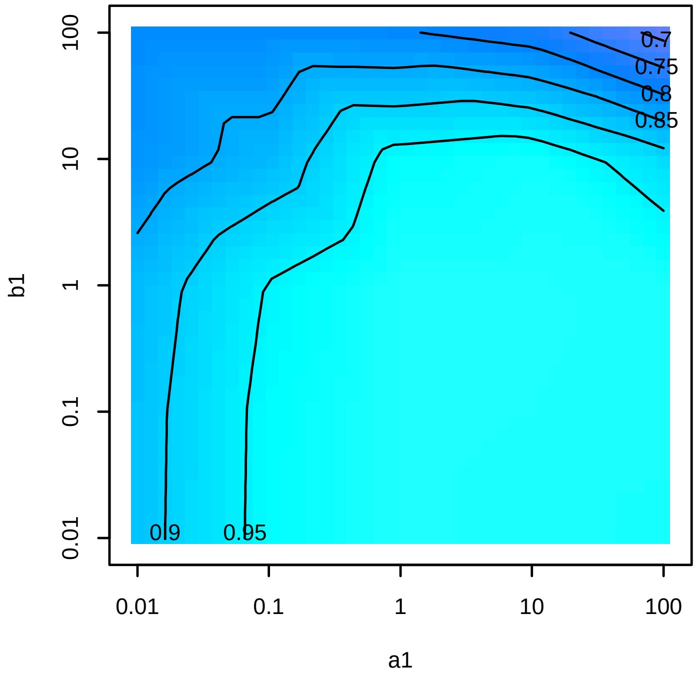

Figure 2: Efficiency of OS[Be(a_1, b_1),Be(1,1)] relative to OS[Be(1,1),Be(1,1)], analyzed using [Be(1,1), Be(1,1)]

design BF assuming $(p_1, p_2) \sim [\mathrm{Be}(a_1, b_1), \mathrm{Be}(1, 1)]$, for $a_1, b_1 \in \{0.01, 0.1, 1, 10, 100\}$, and then compare the efficiency of MP relative to BF, analyzed using [Be(a_1, b_1), Be(1,1)].

Figure 4 shows that the efficiency of MP is very high throughout most of the region, and often exceeds that of BF even when MP has used design and analysis priors that differ by an order of magnitude. It is only when MP's design and analysis priors differ by two orders of magnitude that the efficiency becomes unacceptable. Thus MP could be viewed as being a more conservative, robust choice than BF. Also note that MP's relative efficiency would have been even greater had it been compared to EA.

6 COMPUTATIONS

All results presented herein are exact and have been obtained by a variety of techniques. As mentioned, dynamic programming was used to determine the OS rule. For the other rules the material in the previous sections shows how to decide each rule's allocation directly.

Still, for all of the rules considered, it was necessary to compute several quantities, such as the risk incurred and the percent sampled from Population 1. For EA and BF these values can be determined analytically. However, for the adaptive procedures, these calculations were performed using backward induction, which, like dynamic programming, starts with terminal states and works towards the initial state. While backward induction is well-known in theory, in practice one rarely

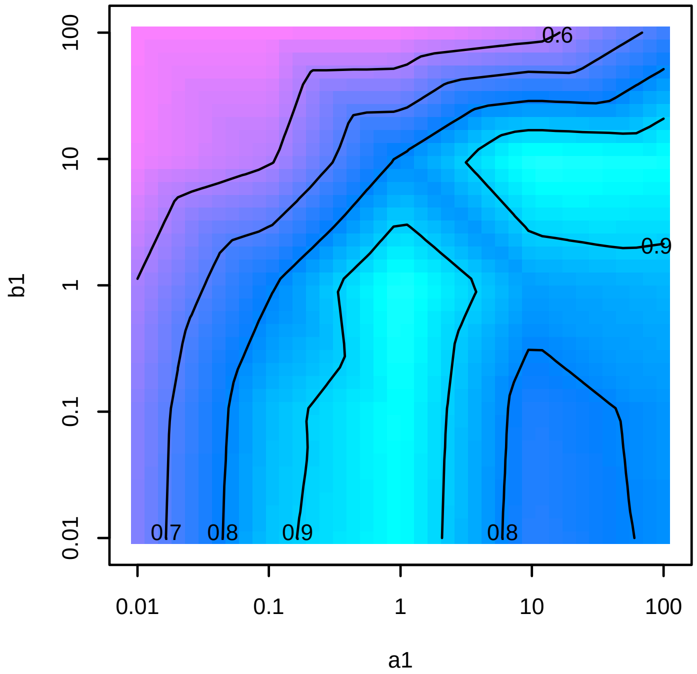

Figure 3: Efficiency of BF[Be(a_1, b_1),Be(1,1)] relative to BF[Be(1,1),Be(1,1)], analyzed using [Be(1,1), Be(1,1)]

sees it used to evaluate allocation rules (see, however, Berry and Eick (1995) and Hardwick and Stout (1992)). The time for dynamic programming and backward induction is $\Theta(N^4)$ and the space needed is $\Theta(N^3)$ (see Hardwick and Stout (1995)).

To compute the figures in Section 5 efficiently, an additional technique is useful. There, a single allocation rule needed to be evaluated several times, once for each different prior. If one needed to make k evaluations, then repeated use of backward induction would take $\Theta(kN^4)$ time. However, by using *forward induction*, the time required can be reduced to $\Theta(N^4 + kN^3)$, and the space required can be kept at $\Theta(N^3)$. The forward induction method consists of first going through the state space from the initial state towards the terminal states. Then, for each state, the number of different paths that could reach that state is computed. In this way, for each evaluation, only the terminal states need be visited. The path counts are combined with the appropriate probabilities to determine the probability of reaching that terminal state. This can then be used to determine whatever average quantities one desires. Further details concerning solving allocation problems using forward induction, backward induction, and dynamic programming appear in Hardwick and Stout (1995).

One minor point to note is that, while OS, MP, and HP are defined by comparing various quantities and randomizing only when these values are equal, in practice we randomize whenever the difference between values being compared is less than 10^{-5}. This slightly improves the numeric stability of the results.

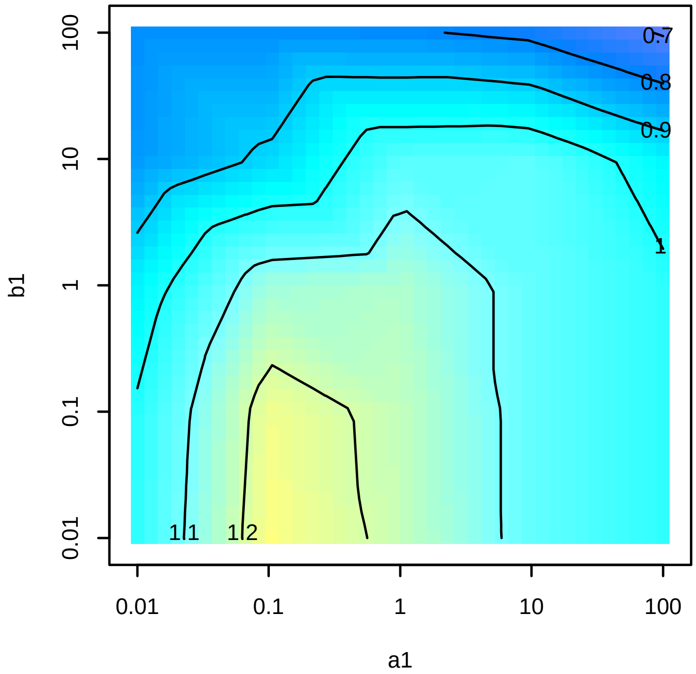

Figure 4: Efficiency of MP[Be(1,1),Be(1,1)] relative to BF[Be(a_1, b_1),Be(1,1)], analyzed using [Be(a_1, b_1), Be(1,1)]

7 CONCLUSIONS

In writing this paper we had a couple of different goals. The first was to solve a practical set of estimation problems. After reducing the estimation problem to one of allocation, we considered a variety of sampling procedures that could be used in actual experiments.

Our strategy for generating the different rules can be summarized as follows. Find an expression for the Bayes risk as a function of n_1 so that the fixed sample problem can be solved. From this locate and evaluate the BF, EA and HP rules. To generate the remaining two rules, MP and OS, one need work only with the interim risk functions.

Our statistical conclusions are typified by the results in Tables 1–4. By definition, the OS rule is best. Furthermore, subject to heuristic arguments of the sort detailed in Section 4.1.1, the performance of the other four allocation rules depends on the polynomial, the configuration of the prior parameters, and the sample size for the experiment.

In particular, we found that, without exception, the local sequential rules are so close to being optimal that the myopic rule appears to offer the best overall combination of practicality and efficiency. Furthermore, for a variety of prior parameter configurations, both the BF and EA rules performed reasonably well. Still, it is interesting to note how poorly the EA rule performs in settings such as those described in Table 3.

Hyperopic rules have received significantly less attention than the others considered here. However, they have several appealing features. Because they are adaptive versions of the best fixed rule, they can be used whenever BF can be determined. Further, at any point in the experiment, they decide the optimal allocation to use for the rest of the experiment if no new information is collected. This makes them natural candidates for situations in which responses are delayed. In such situations one would determine the optimal allocation for the remainder of the experiment, assuming that the only allocated samples are the ones that have been observed, and then subtract the unfinished observations to decide the remaining allocations. Thus, depending on the delays encountered in the responses, hyperopic rules could range from the fully sequentially HP considered here to the fixed allocation BF rule.

On the whole, given the performance of the local rules, MP and HP, (including the robustness exhibited by MP in Figure 4), it seems clear that locating the *optimal* rule for these problems is important primarily because it provides a precise basis for comparison for the ad hoc rules.

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 techniques of dynamic programming, backward induction, and forward induction.

For many years, statisticians have known that, in theory, dynamic programming serves as a useful means for computing solutions to sequential allocation problems. Historically, such programming has been too computationally intensive to be very practical, but increasingly, realistic problems can be solved in this manner. Less well understood is the ability of backward induction to evaluate properties of all types of allocation rules exactly. Even if an allocation rule is not defined through a set of recursion equations, one can still determine its attributes through backward induction. If an allocation rule needs to be evaluated several times, as in Section 5, then forward induction can be used to greatly reduce the computational time while still yielding exact results.

This approach has some advantages over the usual method of seeking approximate characteristics of an allocation rule asymptotically or through simulation studies. Asymptotic analyses are very important, but in practice it is hard to know how good an approximation they are for a given situation (e.g., see Table 2). In particular, if one cannot determine the optimal rule, then it is difficult to know just how efficient an ad hoc rule will be for a given sample size.

The computational approach also allows one to explore designs more thoroughly, considering criteria beyond the design criterion that may not be amenable to analytic evaluation. To do so, however, one needs algorithms that enable rapid evaluations, since a great many evaluations are needed to produce views such as were shown in Section 5.

Computationally intensive examination of nontrivial adaptive designs is a new possibility which we expect will rapidly evolve and expand over time. By making computations very fast, and utilizing graphical output to make the results more comprehensible, researchers can engage in an interactive design and optimization approach heretofore impossible. Tentative designs can be explored and adjusted multiple times to better meet a mix of criteria. Researchers are no longer limited to those designs that have simple analytic behavior, nor are they limited to evaluating only certain criteria. It is this flexibility that we view as being the most important aspect of this work. The specific examples shown here are merely indications of the types of problems and approaches one might try.

8 APPENDIX

We are interested in minimizing the Bayes risk for the case in which n_1 and n_2 are fixed integers. While one can write out the fixed sample size Bayes risk with respect to squared error loss directly for any polynomial, in general the expression is quite complicated. In order to help understand the general behavior of the function, we provide a first order approximation to the Bayes risk that applies to any bivariate polynomial $\theta(p_1, p_2)$.

As was shown in Section 4.1.1, this approximation can be used to obtain the asymptotic properties of the BF rule. The approximation is also useful with the HP rule, since that rule must make multiple risk evaluations. Since the approximation is quite accurate, and decisions are constantly being updated, the effect on HP of using the approximation instead of the true Bayes risk is extremely small.

As seen in the Section 2, the Bayes risk is given by

$$\mathbf{E}^{\xi}[(\theta - \hat{\theta}_{n_1,n_2})^2] = \mathbf{E}^{\xi}\left[\left(\sum_{r=0}^{u_2}\sum_{s=0}^{u_1} c_{rs}p_1^r p_2^s - \sum_{r=0}^{u_2}\sum_{s=0}^{u_1} c_{rs}\mathbf{E}^{\xi}(p_1^r \mid \mathcal{D}_N)\mathbf{E}^{\xi}(p_2^s \mid \mathcal{D}_N)\right)^2\right]. \quad (10)$$

If we expand this expression, we will be left with a series of terms of the form

$$\begin{aligned} &\mathbf{E}^{\xi}\left[c_{ij}\, p_1^i\, p_2^j \cdot c_{k\ell}\, p_1^k\, p_2^\ell - c_{ij}\,\mathbf{E}^{\xi}(p_1^i \mid \mathcal{D}_N)\,\mathbf{E}^{\xi}(p_2^j \mid \mathcal{D}_N) \cdot c_{k\ell}\,\mathbf{E}^{\xi}(p_1^k \mid \mathcal{D}_N)\,\mathbf{E}^{\xi}(p_2^\ell \mid \mathcal{D}_N)\right] \\ &= \quad c_{ij}c_{k\ell}\Big\{\mathbf{E}^{\xi}\left[p_1^{i+k}\right]\,\mathbf{E}^{\xi}\left[p_2^{j+\ell}\right] \qquad (11) \\ &\qquad - \mathbf{E}^{\xi}\left[\mathbf{E}^{\xi}(p_1^i \mid \mathcal{D}_N)\mathbf{E}^{\xi}(p_1^k \mid \mathcal{D}_N)\right]\,\mathbf{E}^{\xi}\left[\mathbf{E}^{\xi}(p_2^j \mid \mathcal{D}_N)\mathbf{E}^{\xi}(p_2^\ell \mid \mathcal{D}_N)\right]\Big\} \end{aligned}$$

Now, note that (11) can be expressed as

$$c_{ij}c_{k\ell}\Big\{\mu_{i+k}(p_1)\,\mu_{j+\ell}(p_2) - \mathbf{E}^{g}\left[\mu_i(p_1 \mid \mathcal{D}_N)\mu_k(p_1 \mid \mathcal{D}_N)\right]\mathbf{E}^{h}\left[\mu_j(p_2 \mid \mathcal{D}_N)\mu_\ell(p_2 \mid \mathcal{D}_N)\right]\Big\},$$

where the expectations $\mathbf{E}^g$ and $\mathbf{E}^h$ are with respect to the marginal distributions g of x and h of y respectively:

$$g(x) \sim \text{Beta-Binomial}(n_1, a_1, b_1) \;\text{ and }\; h(y) \sim \text{Beta-Binomial}(n_2, a_2, b_2).$$

Given the moments for g and h, it is clear that each term of the form (11) can be evaluated, and thus, with a little algebra, one can compute the Bayes risk whenever n_1 and n_2 are fixed. Still, while in Section 4, we carried out these calculations exactly for two examples, in general, such manipulations are tedious at best. It is useful, therefore, to have approximations for (10) that can be used in its place. The following proposition is useful in this respect in that it provides first order approximations for all of the terms that comprise (10).

Proposition 1

$$\mathbf{E}^{\xi}\left[p_1^i\,p_2^j\,\cdot\,p_1^k\,p_2^\ell\;-\;\mathbf{E}^{\xi}(p_1^i\mid\mathcal{D}_N)\,\mathbf{E}^{\xi}(p_2^j\mid\mathcal{D}_N)\;\cdot\;\mathbf{E}^{\xi}(p_1^k\mid\mathcal{D}_N)\,\mathbf{E}^{\xi}(p_2^\ell\mid\mathcal{D}_N)\right] \tag{12}$$
$$= \mu_{j+\ell}(p_2)\,ik\left[\frac{a_1^{\overline{i+k-1}}b_1}{(a_1+b_1)^{\overline{i+k}}(a_1+b_1+n_1)}\right]$$
$$+\,\mu_{i+k}(p_1)\,j\ell\left[\frac{a_2^{\overline{j+\ell-1}}b_2}{(a_2+b_2)^{\overline{j+\ell}}(a_2+b_2+n_2)}\right]+O\left(\frac{1}{n^2}\right)$$

where $n=\min\{n_1,n_2\}$.

Proof. Note that if $p\sim$ Beta(a,b) and $g\sim$ Beta-Binomial(n,a,b), then for natural numbers i, k,

$$\begin{aligned}
\mathbf{E}^{g}\left[\mu_i(p\mid\mathcal{D}_n)\mu_k(p\mid\mathcal{D}_n)\right] &= \frac{\sum_{s=0}^{n}\binom{n}{s}a^{\overline{s}}b^{\overline{n-s}}(a+s)^{\overline{i}}(a+s)^{\overline{k}}}{(a+b)^{\overline{n}}(a+b+n)^{\overline{i}}(a+b+n)^{\overline{k}}}\\
&= \frac{\sum_{x=0}^{k}\binom{k}{x}(-i)^{\overline{k-x}}a^{\overline{i+x}}(a+i+x+b)^{\overline{n}}}{\sum_{x=0}^{k}\binom{k}{x}(-i)^{\overline{k-x}}(a+b)^{\overline{n+i+x}}}\\
&= \mu_{i+k}(p)-ik\left[\frac{a^{\overline{i+k-1}}b}{(a+b)^{\overline{i+k}}(a+b+n)}\right]+O\left(\frac{1}{n^2}\right)\,.
\end{aligned}$$

Then, as $n_1,\,n_2\to\infty$, the entire expression (12) can be written as

$$\mu_{i+k}(p_1)\,\mu_{j+\ell}(p_2)-\left(\mu_{i+k}(p_1)-ik\left[\frac{a_1^{\overline{i+k-1}}b_1}{(a_1+b_1)^{\overline{i+k}}(a_1+b_1+n_1)}\right]\right)\cdot$$
$$\left(\mu_{j+\ell}(p_2)-j\ell\left[\frac{a_2^{\overline{j+\ell-1}}b_2}{(a_2+b_2)^{\overline{j+\ell}}(a_2+b_2+n_2)}\right]\right)+O\left(\frac{1}{n^2}\right)$$
$$= \mu_{j+\ell}(p_2)\,ik\left[\frac{a_1^{\overline{i+k-1}}b_1}{(a_1+b_1)^{\overline{i+k}}(a_1+b_1+n_1)}\right]$$
$$+\,\mu_{i+k}(p_1)\,j\ell\left[\frac{a_2^{\overline{j+\ell-1}}b_2}{(a_2+b_2)^{\overline{j+\ell}}(a_2+b_2+n_2)}\right]+O\left(\frac{1}{n^2}\right)$$

where $n=\min\{n_1,n_2\}$. $\square$

ACKNOWLEDGEMENTS

We thank C. Page for her helpful comments. This research was supported in part by National Science Foundation grants DMS-9157715 and DMS-9504980.

BIBLIOGRAPHY

Berry, D. A. (1974), "Optimal sampling schemes for estimating system reliability by testing components – I: fixed sample sizes", *J. Amer. Statist. Assoc.* **69**: 485–491.

Berry, D. A. and Eick, S. G. (1995), "Adaptive assignment versus balanced randomization in clinical trials— a decision-analysis", *Stat. in Medicine* **14**, pp. 231–246.

Etzioni, R. and Kadane, J. (1993), "Optimal experimental design for another's analysis", *J. Amer. Statist. Assoc.* **88**: 1404–1411.

Hardwick, J. and Stout, Q. F. (1992), "Bandit strategies for ethical sequential allocation", *Computing Science and Stat.* **23**: 421–424.

Hardwick, J. and Stout, Q. F. (1993), "Optimal allocation for estimating the product of two means", *Computing Science and Stat.* **24**: 592–596.

Hardwick, J. and Stout, Q. F. (1995), "Exact computational analyses for adaptive designs", in *Adaptive Designs*, N. Flournoy and W. F. Rosenberger, eds., IMS Lecture Notes–Monograph Series **25**: 223–237.

Page, C. (1995), "Adaptive allocation for estimation", in *Adaptive Designs*, N. Flournoy and W. F. Rosenberger, eds., IMS Lecture Notes–Monograph Series **25**: 213–222.

Rekab, K. (1993), "A sampling scheme for estimating the reliability of a series system", *IEEE Trans. Reliability* **42**: 287-290.

Shapiro, C. Page (1985), "Allocation schemes for estimating the product of positive parameters", *J. Amer. Statist. Assoc.* **80**: 449–454.