New Developments and Applications in Experimental Design
IMS Lecture Notes – Monograph Series (1998) Volume 34

VARIABLILITY IN ADAPTIVE DESIGNS FOR ESTIMATION OF SUCCESS PROBABILITIES

BY VINCENT MELFI AND CONNIE PAGE

Michigan State University

Three adaptive allocation rules for use when estimating the difference in success probabilities are proposed and studied via simulation. The rules are motivated by the need for randomization and reduction of experimenter bias, and they are adaptive in that the decision about which population to sample at each stage can depend on data collected up to that stage. The empirical mean square errors for these rules and two non-adaptive rules are compared. For moderate total sample sizes and moderate values of the success probabilities, it is shown that "adapting" substantially increases the mean square error over that of the simple totally randomized allocation rule which allocates at each stage to each population with probability 0.5. However, with total sample sizes of 100 and upward, adaptive rules do just as well as the totally randomized rule for moderate success probabilities, and do much better for more extreme values of the success probabilities.

1. Introduction. Suppose we have two Bernoulli populations, A and B, with respective success probabilities, p_A and p_B, and failure probabilities $q_A = 1 - p_A$ and $q_B = 1 - p_B$. The setting could be clinical, with two treatment populations with cure rates equal to the success probabilities, or industrial with two brands of a component with failure rates equal to the success probabilities. In the clinical application, methods of allocating patients to treatments to lower the selection bias, the effect of trends in the data, and the number of patients on the inferior treatment have been proposed. These methods are typically adaptive in that decisions about future allocations depend on past observations. An excellent overview of such methods can be found in Rosenberger (1996).

While many adaptive designs have been suggested and studied in the clinical setting, few have been studied in the industrial setting. The goal of this work is to look at several adaptive allocations where the purpose of experimentation is the estimation of the difference, $p_A - p_B$, and the primary allocation goal is to minimize the variance of the estimator when a fixed total number of observations can be taken. Secondary

---

Received September 1997; revised March 1998.
AMS 1991 subject classifications. Primary: 62L05; secondary 62K99.
Key words and phrases. Adaptive allocation, biased-coin design, randomization.

goals would concern randomization, selection bias, and trends. In the non-adaptive case, where independent random samples of pre-determined sizes n_A and n_B are taken (total number of n fixed) from populations A and B, respectively, the variance of the difference in sample proportions is minimized when $n_A/n_B = (p_A q_A/p_B q_B)^{1/2}$. This optimal allocation is well-known and often called Neyman allocation in survey work. The difficulty with using the allocation is that the population proportions are unknown. Thus, the need for adaptive allocation arises, where the proportions are estimated at each stage and the estimates are then used in allocating the next observation.

Three adaptive rules, derived to do well in terms of the variance of the estimator of $p_A - p_B$, will be studied and compared here. Notation and preliminaries are given in Section 2, the rules are motivated and defined in Section 3, simulation results are presented in Section 4, and conclusions and recommendations are given in Section 5.

2. Motivation and definitions. In this section the reason for considering adaptive allocation rules is given, after which the necessary tools are developed for adaptive allocation in the Bernoulli setting.

2.1. *The need for allocation.* Before developing allocation rules, the actual need for such rules should be assessed by measuring how much is lost if the optimal allocation proportions are not used. Consider independent random samples of sizes n_A and n_B from Bernoulli populations, A and B, where $n = n_A + n_B$ is fixed. Write $n_A = \pi n$ where π is the proportion allocated to A. The variance of the difference in the sample proportions can be written as a function of π, $V(\pi) = n^{-1}(p_A q_A/\pi + p_B q_B/(1-\pi))$. The value of π which minimizes $V(\pi)$ is the optimal proportion, here equal to $\pi_A = \sqrt{p_A q_A}/(\sqrt{p_A q_A} + \sqrt{p_B q_B})$, and the minimum value of $V(\pi)$ is $V(\pi_A) = n^{-1}(\sqrt{p_A q_A} + \sqrt{p_B q_B})^2$. If a proportion π is allocated to A, then the percent increase in variance from not allocating optimally is

$$PIV(\pi, \pi_A) = \frac{V(\pi) - V(\pi_A)}{V(\pi_A)} = \frac{(\pi_A - \pi)^2}{\pi(1-\pi)}.$$

Obviously, as $\pi \to 0$ or 1, this percentage increase tends to ∞, so that great care should be taken when small proportions are used. Substantial increases can also occur in middle ranges. For example, if 20% of the observations are allocated to A when optimal is 40%, then the variance will increase by 25%.

2.2. *Notation and definitions.* As before, let the optimal proportion to be allocated to population A be denoted by $\pi_A = \sqrt{p_A q_A}/(\sqrt{p_A q_A} + \sqrt{p_B q_B})$. Assume that $0 < \pi_A < 1$.

Some notation is required. Again, fix n as the total number of observations to be allocated. Consider $X_1, \ldots, X_n$ independent and identically distributed Bernoulli(p_A), independent of $Y_1, \ldots, Y_n$ independent and identically distributed Bernoulli(p_B). Allocation is specified by the sequence $\delta_1, \ldots, \delta_n$ where $\delta_i = 1$ or 0 if the i^{th} observation is from population A or B, respectively. For each stage, $i = 1, \ldots, n$, define $S_A(i) = \sum_{j=1}^{i} \delta_j X_j$, the number of successes observed from population A,

and define $S_B(i) = \sum_{j=1}^{i}(1-\delta_j)Y_j$, the number of successes observed from population B, by the ith stage. Also, define $N_A(i) = \sum_{j=1}^{i}\delta_j$, the total number of observations from population A and $N_B(i) = \sum_{j=1}^{i}(1-\delta_j)$, the total number of observations from population B, at the ith stage. Note that $N_A(i) + N_B(i) = i$ for $i = 1, \ldots, n$. Obviously, δ_i can depend only on the previous stages; that is, δ_i is measurable $F_{i-1} = \sigma\{\, \delta_j X_j, (1-\delta_j)Y_j, \delta_j, j = 1, \ldots, i-1 \,\}$.

Finally, defining adaptive rules to get proportion π_A from population A requires estimating an unknown π_A at each stage. Here, $\pi_A = \sigma_A/(\sigma_A + \sigma_B)$ is a function of the population standard deviations $\sigma_A = \sqrt{p_A q_A}$ and $\sigma_B = \sqrt{p_B q_B}$, and its estimator will involve estimators of standard deviations of Bernoulli populations.

2.3. *Estimation of σ^2 for Bernoulli populations.* When used in an adaptive setting, estimators of a Bernoulli population variance such as the product of the sample proportions of successes and failures cause difficulties and often lead to eventually observing from only one population. Bayesian methods [Chew (1971)] suggest adding on a positive constant to the number of successes and to the number of failures before taking the proportions. That is, if S successes are observed in m trials, then estimate $\sigma^2 = pq$ by $\hat{\sigma}^2 = (S+b)(m-S+b)/(m+2b)^2$. The b comes from a symmetric beta(b,b) prior on the success probability. If $b = 1$, this is derived from a uniform prior; as b increases, more weight is concentrated at 0.5; for b small, more weight is placed on the extremes.

This estimator of σ^2 will be used for estimating the variance in each of the Bernoulli populations A and B.

3. Adaptive allocation rules. Three adaptive rules will be defined in this section, all aiming for allocation proportion π_A from population A, but having different randomization properties. The estimator of π_A at stage i will be

$$\hat{\pi}_A(i) = \frac{\hat{\sigma}_A(i)}{\hat{\sigma}_A(i) + \hat{\sigma}_B(i)}$$

where

$$\hat{\sigma}_A^2(i) = \frac{(S_A(i)+b)(N_A(i)-S_A(i)+b)}{(N_A(i)+2b)^2}$$

and

$$\hat{\sigma}_B^2(i) = \frac{(S_B(i)+b)(N_B(i)-S_B(i)+b)}{(N_B(i)+2b)^2}$$

for $b > 0$. We will typically use $b = 0.5$, somewhat motivated by correcting for continuity. However, b can be considered a design parameter and as is demonstrated in the simulations of Section 4, the choice of $b = 0.1$ can substantially increase the mean square error of the estimator of $p_A - p_B$. Values of b ranging from $b = 0.5$ to $b = 4$ gave simulated mean square errors similar to $b = 0.5$. As noted earlier, $b = 1.0$ represents a uniform prior on the success probability. Other work involving a Bayesian approach and uniform priors is contained in Hardwick and Stout (1996) and Hardwick (1991).

All of the adaptive rules will take an initial sample of size $n_o \geq 1$ from each population to allow for easy beginning of the adaptive process.

Since the major goal of allocation is to achieve proportion π_A from population A, a natural adaptive rule is to estimate π_A at each stage, then allocate at the next stage to "correct" the observed proportion according to the estimated π_A. This type of rule was proposed by Thompson (1933). Specifically, define D-Allocation (Adaptive-Deterministic) as take an initial sample of size $n_o \geq 1$ from each population. Then for $i = 2n_o + 1, \ldots, n-1$, at the ith stage, allocate the $(i+1)$th observation to A (to B) if $N_A(i)/i < (\geq)\hat{\pi}_A(i)$. Such allocations have been studied by Robbins, Simons, and Starr (1967) for normal populations.

The D-rule lacks randomization, and a simple randomized rule allocates to population A at each stage with probability $\hat{\pi}_A$. Specifically, define R-Allocation (Adaptive-Randomized) as take an initial sample of size $n_o \geq 1$ from each population. Then for $i = 2n_o + 1, \ldots, n-1$, at the ith stage, allocate the $(i+1)$th observation to A with probability $\hat{\pi}^*_A(i) = (\frac{\hat{\pi}_A(i)n - n_o}{n-2n_o})^+$.

Note that the $\hat{\pi}_A$ is modified to $\hat{\pi}^*_A$ to adjust for initial samples sizes of n_o from each population. The R-rule is a randomized rule, and thus, is protected from selection bias and some accidental bias.

We will study a second more complicated randomized rule which is a special doubly biased coin design considered by Eisele (1990, 1994). In general, these doubly biased coin designs are defined by an allocation probability function, $\phi : [0,1]^2 \to [0,1]$. The allocation rule then allocates the $(i+1)$th observation to population A with probability $\phi(N_A(i)/i, \hat{\pi}_A(i))$. We will consider the allocation probability function $\phi(r,\pi) = (1 - (1-\pi)r/\pi)^+$ which was used by Eisele (1990, 1994) in the case of normal populations.

Specifically, define B-Allocation (Adaptive Biased Coin) as take an initial sample of size $n_o \geq 1$ from each population. Then for $i = 2n_o + 1, \ldots, n-1$, at the ith stage, allocate the $(i+1)$th observation to A with probability

$$\phi\left(\frac{N_A(i)}{i}, \hat{\pi}_A(i)\right) = \left(1 - \frac{1-\hat{\pi}_A(i)}{\hat{\pi}_A(i)}\frac{N_A(i)}{i}\right)^+.$$

This rule (choice of ϕ) can be derived from Wei's biased coin designs (1978). Wei's chosen allocation probability function to achieve balance is $1 - N_A(i)/i$, which, when generalized to achieve proportion π_A, leads to $\phi(r,\pi) = (1 - 1 - \pi r/\pi)1[r < \pi] + \pi(1 - r)/1 - \pi 1[r \geq \pi]$. Here 1[.] represents the indicator function. As a function of r with π fixed, these are two line segments. Eisele's and our choice of ϕ is the positive part of the first line segment.

4. Simulations. Two other allocation rules will be studied here along with the adaptive rules of Section 3. The first rule (T) is total randomization for balance at each stage. Define T-Allocation (Totally Randomized) as take an initial sample of size $n_o \geq 1$ from each population. Then for $i = 2n_o + 1, \ldots, n-1$, at the ith stage, allocate the $(i+1)$th observation to A with probability 0.5.

The second rule is O-Allocation (Totally Randomized Optimal) which is defined as follows: Take an initial sample of size $n_o \geq 1$ from each population. Then for

$i = 2n_o + 1, \ldots, n-1$, at the ith stage, allocate the $(i+1)$th observation to A with probability π_A.

Obviously, the O-rule cannot be implemented in practice since π_A is unknown. However, the R-rule from Section 3 can be considered the adaptive version of the O-rule, and comparison of the two rules will give insights into the effect of "adapting" the allocation probability at each stage.

The simulation results reported below are for 5000 samples of size n, where $n = 30$ and $n = 100$ are used. Throughout, the initial sample size from each population is $n_o = 1$. For each sample, the estimate, $S_A(n)/N_A(n) - S_B(n)/N_B(n)$ of $p_A - p_B$ is computed, and its empirical mean square error (MSE) then found as an average of the squared difference of the estimate and the parameter values over the 5000 samples. The scaled mean square errors, $n \times MSE$, are reported on the vertical scale, and labeled consistently with the rule labels: D, R, B, T, and O. They are also compared with the target, $nV(\pi_A) = (\sqrt{p_A q_A} + \sqrt{p_B q_B})^2$, given as the solid line on the graphs.

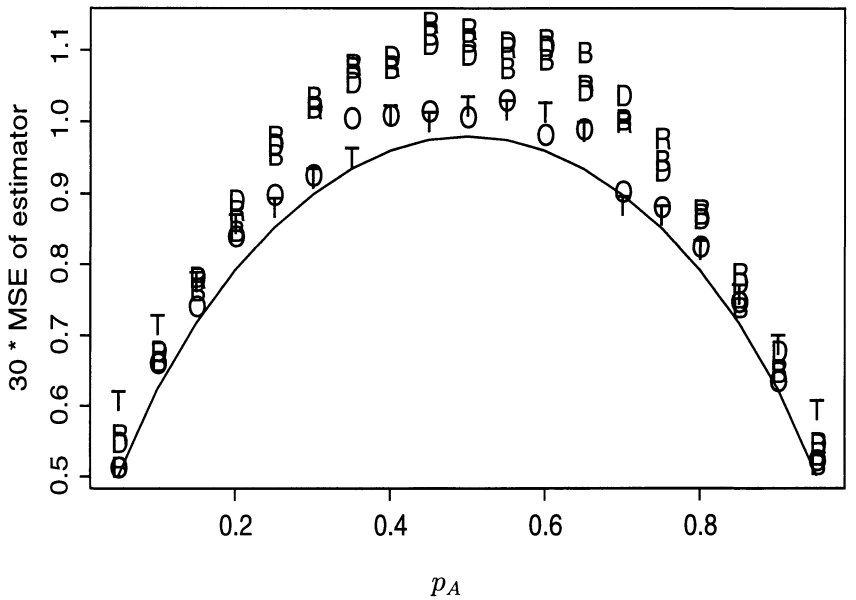

FIG. 1. *Scaled mean squared errors versus p_A for $n = 30, p_B = 0.4, b = 0.5$. The solid line represents the target $nV(\pi_A) = (\sqrt{p_A q_A} + \sqrt{p_B q_B})^2$. The symbols used correspond to the definitions in Section 3 and the beginning of Section 4. For example, the points labelled "D" give the scaled mean squared error for the adaptive-deterministic rule.*

Throughout, p_B is fixed at 0.4, and p_A ranges over values, $0.05, 0.10, \ldots, 0.95$. on the horizontal axes. Figures 1 and 2 use $b = 0.5$ in the $\hat{\pi}_A$ estimator; Figure 3 uses $b = 0.1$. Total sample sizes are $n = 30$ for Figures 1 and 3, and $n = 100$ for Figure 2.

In comparing all five rules, note how well the totally randomized rule (T) does compared to both the optimal randomized (O) and to the adaptive rules in the $n = 30$

case and for moderate success probabilities. For $p_B = 0.4$, the T-rule clearly beats all the adaptive rules for p_A between 0.25 and 0.80. However, the T-rule becomes clearly worse when $p_A < 0.10$. Thus, in moderate samples, adaptive rules work the best for small or large success probabilities. For $n = 100$, the T-rule is not always the best choice even for middle range success probabilities, and thus, T loses its advantage over the adaptive rules when sample sizes increase.

For both sample sizes and all ranges of p_A, there seems to be little difference in the mean square errors for the adaptive rules, D, R, and B.

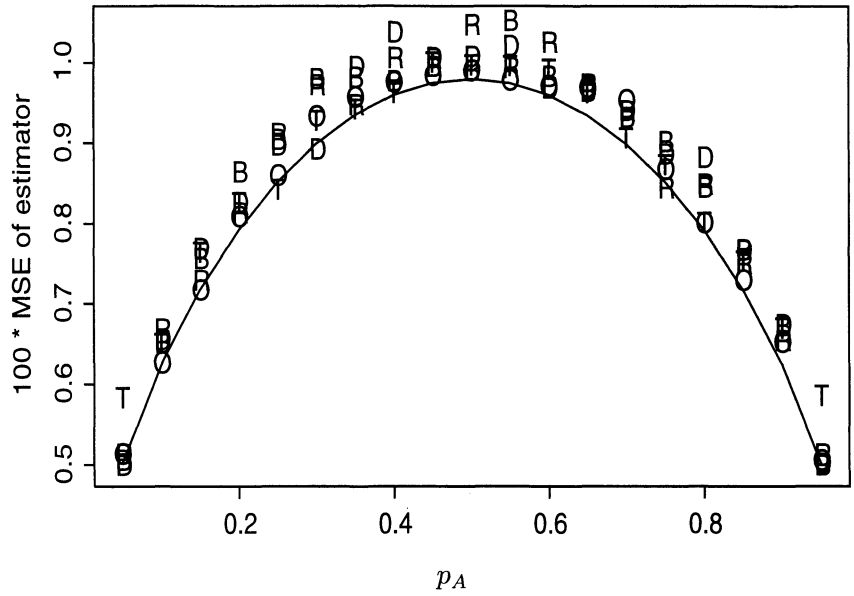

FIG. 2. *Scaled mean squared errors versus p_A for $n = 100, p_B = 0.4, b = 0.5$. The solid line represents the target $nV(\pi_A) = (\sqrt{p_A q_A} + \sqrt{p_B q_B})^2$. The symbols used correspond to the definitions is Section 3 and the beginning of Section 4. For example, the points labelled "D" give the scaled mean squared error for the adaptive-deterministic rule.*

There is also a lesson to be learned from comparing the O-rule to the R-rule. The R-rule is the adaptive version of O, but the MSE for R is much worse than the MSE for O in the $n = 30$ case and for moderate success probabilities. This indicates that moving the probability of allocation to A at each stage has substantially increased the MSE.

Finally, the effect of the choice of b on the mean square errors for the adaptive rules can be substantial if b is too small, as seen in Figure 3. When b is large, the resulting rule acts more like the T-rule, since increased weight is put at 0.5.

5. Conclusions and recommendations. Recall that the O-rule cannot be implemented since π_A is unknown. However, the T-rule is easily implemented, and for

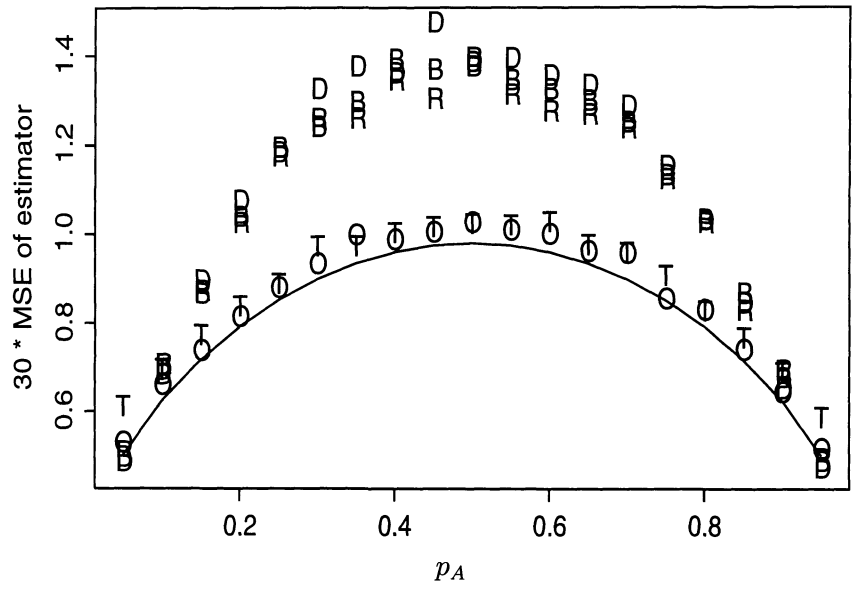

FIG. 3. *Scaled mean squared errors versus p_A for $n = 30, p_B = 0.4, b = 0.1$. The solid line represents the target $nV(\pi_A) = (\sqrt{p_A q_A} + \sqrt{p_B q_B})^2$. The symbols used correspond to the definitions in Section 3 and the beginning of Section 4. For example, the points labelled "D" give the scaled mean squared error for the adaptive-deterministic rule.*

$n = 30$, the T-rule is clearly better than any adaptive rule for p_A between 0.25 and 0.75 and $p_B = 0.4$. There are several reasons for this behavior. First, the MSE function is relatively flat over a wide middle range of p_A and p_B, so that non-optimal designs do almost as well as optimal designs in that range. Therefore, any design component which even slightly increases the variance would inflate the MSEs. And all of the adaptive rules have such a component since the probability of allocation to A depends on either $\hat{\pi}_A(i)$ or $N_A(i)$ or on both, and these components change at each stage. The comparison of the O-rule and the R-rule (its adaptive counterpart) emphasizes how much these changing components can affect the MSE for moderate sample sizes.

Thus the flatness of the MSE curves in the middle range of p_A and p_B and the increase in MSE caused by changing probabilities of allocation in the adaptive rules causes the T-rule to outperform the adaptive rules in the middle ranges and for small to moderate sample sizes.

An exact formula for the variance of allocation proportions in the setting of randomized play-the-winner allocation is given in Matthews and Rosenberger (1997). Such allocations are not designed to achieve the optimal allocation proportion π_A as in the current case, but are interesting for comparative purposes.

The choice of allocation rule obviously depends on experimenter goals. If the only

goal is the minimization of the variance of the resulting estimator, and randomization is not important, then a simple balanced sample ($n/2$ from each population) should be considered. The balanced sample gives a scaled MSE equal to $2(\sigma_A^2 + \sigma_B^2)$, and the difference in this scaled MSE and the target curve is $(\sigma_A - \sigma_B)^2$. For success probabilities between 0.2 and 0.8, this difference is less than 0.01. If randomization is important, then for moderate sample sizes, the completely randomized rule (T) is very effective for p_A and p_B both between 0.10 and 0.90. If either p_A or p_B is possibly extreme (below 0.10 or above 0.90), then adaptive rules give lower MSE for total sample size of at least 30.

As for choosing between adaptive rules (their MSE's are all comparable), both the R-rule and the D-rule are easily described and motivated for an experimenter. The B-rule looks very arbitrary in its choice of allocation probability, and although its mathematical properties lead to its good performance, it may not appeal to many experimenters. Choosing between D and R depends on how important randomization is to the experimenter.

One of the referees brought to our attention the existence of methods to compute mean squared errors exactly for some allocation procedures. See Hardwick and Stout (1998). Such computations would not change the conclusions of our work, of course, but should prove very useful in practice.

Acknowledgments. It is a pleasure to thank both referees and the editor, all of whom provided useful input on presentation. One of the referees also brought several of the references to our attention. Thanks also to Margarida Geraldes for helpful conversations.

REFERENCES

CHEW, V. (1971). Point estimation of the parameter of the binomial distribution. *Am. Statist.* **25** 47–50.

EFRON, B. (1971). Forcing a sequential experiment to be balanced. *Biometrika* **58** 404–418.

EISELE, J. (1990). An adaptive biased coin design for the Behrens-Fisher problem. *Seq. Anal.* **9** 343–359.

EISELE, J. (1992). The doubly adaptive biased coin design for sequential clinical trials. *J. Statist. Plann. Inf.* **36** 249–261.

HARDWICK, J. and STOUT, Q. (1996). Optimal allocation for estimating the mean of a bivariate polynomial. *Seq. Anal.* **15** 71–90.

HARDWICK, J. and STOUT, Q. F. (1998). Forward induction for evaluating sequential allocation procedures. *SIAM J. Sci. and Statist. Comp.*, in press.

HARDWICK, J. P. (1991). Computational problems associated with minimizing the risk in a simple clinical trial. In *Statistical Multiple Integration* (N. Flournoy and R. Tsutakawa, eds.), 239–256 American Mathematical Society, Providence.

MATTHEWS, P. and ROSENBERGER, W. (1997). Variance in randomized play-the-winner clinical trials. *Stat. Prob. Letters* **35** 233–240.

ROBBINS, H., SIMONS, G. and STARR, N. (1967). A sequential analogue of the Behrens-Fisher problem. *Ann. Math. Statist.* **38** 1384–1391.

ROSENBERGER, W. (1996). New Directions in Adaptive Designs. *Statist. Sci.* **11** 137–149.

THOMPSON, W. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25** 275–294.

WEI, L. J. (1978). The adaptive biased coin design for sequential experiments. *Ann. Statist.* **6** 92–100.

WEI, L. J., SMYTHE, R. T., LIN, D. Y. and PARK, T. S. (1990). Statistical Inference with Data-Dependent Treatment Allocation Rules. *J. Am. Statist. Assoc.* **85** 156–162.

DEPARTMENT OF STATISITICS AND PROBABILITY
MICHIGAN STATE UNIVERSITY
EAST LANSING, MI 48824 U.S.A.