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

RESPONSE ADAPTIVE ALLOCATION AND SELECTION BIAS

BY JOHN BATHER

The University of Sussex

Abstract

There have been many papers on biased coin designs and their use in balancing the numbers of patients allocated to different treatments in a clinical trial, without increasing the risk of selection bias. Less attention has been given to the corresponding risk when sequential allocations depend on the previous responses and the aim is to reduce the number of patients on inferior treatments. The ethical requirements may produce a substantial imbalance in the treatment groups. This paper gives a number of examples where selection bias is a serious possibility.

1. Introduction. Selection bias can occur in an experiment designed to compare medical treatments if the experimenter knows, before deciding whether or not to admit a particular patient to the trial, which treatment will be administered next. Blackwell and Hodges (1957) introduced a measure of the bias in a design based on the maximum expected number of correct guesses that an experimenter can achieve when attempting to predict the successive treatment allocations. Their paper and later investigations by Efron (1971), Smith (1984) and many others were concerned with the need to balance the experiment while

Received October 1992; revised April 1993.
AMS 1991 subject classification. Primary 62F03; secondary 62L05.
Key words and phrases. Clinical trials, data dependent allocation, error probabilities.

retaining the principle of randomization. Efron introduced biased coin designs to achieve approximate balance between the numbers of patients allocated to two treatments within a randomized design. More recently, Smith extended this work to a more general class of designs, obtaining further results on the degree of balance achieved and the measure of selection bias.

In his (1984) paper, Smith drew a distinction between biased coin designs and response adaptive designs. In the former, the allocation procedure depends on previous treatment allocations, but not on the patient responses. In what follows, we shall be concerned with adaptive designs where the allocation probabilities at each stage depend on previous responses and the aim is to reduce the number of patients on inferior treatments. This ethical principle is quite different from the statistical one which aims towards equal allocations at the end of the experiment. The danger of selection bias appears to be much greater for response adaptive designs and this will be illustrated by several examples.

We shall restrict attention to experiments for comparing two treatments. Let $\{\psi_t\}$ denote the sequential allocation rule: for $t = 0, 1, \ldots,$ ψ_t is the conditional probability that the next patient will receive treatment 1, given the information available at time t. This information is represented by a σ-algebra $\Im_t$ which is assumed to include the previous patient responses $x_1, x_2, \ldots, x_m$ and $y_1, y_2, \ldots, y_n$ for treatments 1 and 2, respectively, $m + n = t$. Now suppose that the experimenter has some control over the selection of patients. We shall introduce a simple strategy for biasing the results towards a terminal decision in favor of treatment 1, rather than treatment 2. In order to measure the effect of this selection bias, it will be assumed that there is no real difference between the two treatments, but that patients may differ in their mean responses. The experimenter can choose patients with means in the range $\mu \pm \delta$. More precisely, suppose that ψ_t is known at each time and that the experimenter always chooses a "strong patient" with mean $\mu + \delta$ if $\psi_t > \frac{1}{2}$ or a "weak" patient with mean $\mu - \delta$ if $\psi_t < \frac{1}{2}$. When $\psi_t = \frac{1}{2}$, the choice is immaterial, but let us assume that the mean is μ. Thus 2δ represents the range of patient means and we shall be interested in evaluating $g(\delta)$, the probability of a terminal decision that favors treatment 1, for values of $\delta > 0$.

Let $S_t = \sum_{i=1}^{m} x_i - \sum_{j=1}^{n} y_j$ and suppose, for convenience, that $\mu = 0$. Then it is easy to see that, under the above strategy,

(1)
$$E\{S_{t+1} - S_t \mid \Im_t\} = |2\psi_t - 1|\,\delta$$

for all $t \geq 0$. The sequence $\{S_t\}$ is a submartingale and this gives some indication of the bias. In Example 4 we shall consider the special case when $\psi_t = \theta$ or $1 - \theta$ always, for some fixed θ, $\frac{1}{2} < \theta < 1$. Then the expression on the right of (1) reduces to a constant $(2\theta - 1)\,\delta > 0$.

The next section is concerned with two-stage designs where it is feasible to evaluate $g(\delta)$ explicitly. It starts with a very simple illustration using Bernoulli trials. Example 2 involves normal data and unequal samples in the second stage. In Section 3, we shall investigate two sequential designs, one of which is based on the ECMO studies, that have been widely discussed from an ethical point of view. Example 4 relies on a remarkable result of Robbins and Siegmund (1974) for sequential probability ratio tests with normal observations. Simulations show that selection bias can have a serious effect on terminal decisions. The final section contains a brief discussion of the advantages of early stopping clinical trials, with or without data-dependent allocations.

2. Two-stage procedures.

EXAMPLE 1. The response to treatment i is success (s_i) or failure (f_i), and we denote the probability of success by p_i, $i = 1, 2$. Suppose that $p_1 = p_2 = p$ and consider just three possible values: $p = \frac{1}{2}$, $\frac{1}{2} + \delta$ or $\frac{1}{2} - \delta$, where $0 \leq \delta \leq \frac{1}{2}$. The first stage of the experiment uses two patients, one allocated to each treatment. If the results are (s_1, f_2), indicating a preference for treatment 1, the second stage consists of 3 patients on treatment 1 and 1 patient on treatment 2. Similarly, (f_1, s_2) in the first stage leads to allocating 1 patient to treatment 1 and 3 patients to treatment in the second stage. However, if the initial results are (s_1, s_2) or (f_1, f_2), the second stage has two patients on each treatment, allocated with appropriate randomization.

The experimenter uses "average" patients with $p = \frac{1}{2}$ in the first stage and also in the second, if the results are equal. Patients with $p = \frac{1}{2} + \delta$ or $p = \frac{1}{2} - \delta$ are used throughout the second stage when the allocations are 3 + 1 or 1 + 3, respectively.

At the end of stage 2, when six patients have been treated, the "winner" is the treatment with the higher proportion of successes. The possible numbers on treatments 1 and 2 at the end of stage 2 are (3,3), (4,2) and (2,4) depending on whether the results from the first stage

are the same, (s_1, f_2) and (f_1, s_2), respectively. Let us make the convention the 2 successes out of 4 beats 1 success out of 2 and determine the winner by tossing a fair coin if, for example, 2 successes out of 3 are obtained from both treatments. It is a straightforward matter to evaluate the probability that treatment 1 emerges as the winner:

(2)
$$g(\delta) = \frac{1}{2} + \frac{1}{8}\delta\left(1 - 4\delta^2\right),$$

for $0 \leq \delta \leq \frac{1}{2}$. The value of $g(\delta)$ cannot exceed 0.5241 in this simple example and the maximum occurs when $\delta = 0.2887$.

EXAMPLE 2. Let $w_i(t)$ be independent random variables with distributions $N(\mu_i t,\ t)$, $i = 1, 2$. We can think of diffusion processes in continuous time or regard each $w_i(t)$ as the sum of t independent observations from $N(\mu_i,\ 1)$. Consider a first stage with two equal samples of size a and suppose that the second stage has unequal samples with sizes b and c, $b > c$. If $w_1(a) = x$, $w_2(a) = y$ and $x \geq y$, then treatment 1 receives the larger sample in the second stage. In the second stage, we observe $u = w_1(a+b) - w_1(a)$ and $v = w_2(a+c) - w_2(a)$. Similarly, if $x < y$, we observe $u = w_1(a+c) - w_1(a)$ and $v = w_2(a+b) - w_2(a)$ in the second stage.

As before, we assume that the two treatments are equivalent, with mean scores $\mu_1 = \mu_2 = \mu$. For convenience, let $\mu = 0$ in the first stage and suppose that selection bias produces a mean score δ or $-\delta$ for both treatments in the second stage, according as $x \geq y$ or $x < y$.

The winner is determined by comparing average scores at the end of stage 2. Thus, the event that treatment 1 wins is

(3)
$$A = \left\{[x \geq y] \cap \left[\frac{x+u}{a+b} \geq \frac{y+v}{a+c}\right]\right\} \cup \left\{[x < y] \cap \left[\frac{x+u}{a+c} \geq \frac{y+v}{a+b}\right]\right\}.$$

In fact, $g(\delta) = P(A)$ has a simple expression in terms of the standard normal distribution function Φ:

(4)
$$g(\delta) = \Phi(\delta/\lambda),$$

where

$$\lambda^2 = \frac{(a+b)(a+c)(2a+b+c)}{a^2(b-c)^2}.$$

A proof of this formula is outlined in the Appendix.

The expression for $g(\delta)$ can be compared with the corresponding probability, given that there is a real difference between the treatments but no selection bias. Suppose that $\mu_1 = \nu$ and $\mu_2 = -\nu$ throughout the experiment. In this case, it can be shown that

(5)
$$P(A) = \Phi(2\nu/\rho),$$

where

$$\rho^2 = \frac{(2a+b+c)}{(a+b)(a+c)}.$$

This result is an extension of one obtained by Coad (1992) for the case $c = 0$. Note that (5) is exactly what would be obtained in the case that samples of sizes $a + b$ and $a + c$ are assigned to the two treatments in advance, without regard to the results of the first stage. The earlier result (4) is also equivalent to one for prescribed sample sizes and allocations. For example, the same formula holds if $\mu_1 = \mu_2 = 0$ in the first stage and $\mu_1 = \mu_2 = \delta$ in the second, with sample sizes b and c, respectively. It is also worth remarking that $g(\delta)$ can be close to 1 even when δ is small. In particular, this holds if $c = a$ and both b/a and $a\delta^2$ are large.

3. Sequential allocation. The next illustration is based on an adaptive design used by Bartlett, Roloff, Cornell, Andrews, Dillon and Zwischenberger (1985) to study a treatment called extracorporeal membrane oxygenation (ECMO) for babies with respiratory failure. A randomized play-the-winner rule was used to allocate patients between this and a control treatment: see Wei (1979 and 1988) and Wei and Durham (1978). In fact, only one out of ten patients was allocated to the control and this led to a further investigation and much discussion in the literature: see Ware (1989) and Royall (1991). The original clinical trial prescribed a sample of size 10 and a non-Markovian decision rule,

so it will be easier to evaluate the effect of selection bias for a slightly different procedure.

EXAMPLE 3. Consider an urn containing balls of two types, representing treatments 1 and 2. The state (j,k) denotes j balls of type 1 and k of type 2. In general, a ball is drawn at random to give the treatment for the next patient and then it is replaced, together with an additional ball of the same type if the result is a success. In the case of failure, an extra ball of the other type is added. The initial state is (1,1) and we suppose that the experiment ends as soon as 10 further balls of type 1 or 10 balls of type 2 have been added. The terminal states (11,1), (11,2), ..., (11,10) produce a decision in favor of treatment 1. A slightly different decision rule would be to prefer the treatment that achieves the higher proportion of successes, but this is more awkward to evaluate.

We suppose that the two treatments are equivalent, with a common probability of success depending on the state (j,k) because of selection bias. Let the probability of success be $p+\delta$ when $j>k$, $p-\delta$ when $j<k$, and p if $j=k$, where $0 \le p-\delta \le p+\delta \le 1$. The result from the next patient leads to a transition $(j,k) \to (j+1,k)$ or $(j,k+1)$. Let Π_{jk} be the probability of reaching $(j+1,k)$ next. Then

(6)
$$\prod_{jk} = \frac{(j-k)(p \pm \delta)+k}{(j+k)},$$

where the term $+\delta$ applies if $j>k$ and it is replaced by $-\delta$ if $j<k$. The expression (6) reduces to $\frac{1}{2}$ if $j=k$. Now consider the probability of a terminal decision in favor of treatment 1. This can be evaluated as $g(\delta)=\gamma_{11}$, where γ_{jk} is the probability of eventually reaching one of the appropriate terminal states from (j,k). We have

(7)
$$\gamma_{jk} = \prod_{jk} \gamma_{j+1,k} + \left(1 - \prod_{jk}\right) \gamma_{j,k+1}$$

for $j,k = 1, 2, \ldots, 10$ with the boundary conditions that $\gamma_{jk}=1$ for $j=11$, $k=1, 2, \ldots, 10$, and $\gamma_{jk}=0$ for $j=1, 2, \ldots, 10$, $k=11$. Table 1 gives values of $g(\delta)=\gamma_{11}$ obtained from relation (7).

Table 1.
Example 3: probability of deciding in favor of treatment 1.

δ	0.025	0.050	0.100	0.125	0.167	0.250
$p = 0.3$	0.511	0.522	0.543	0.554	0.572	0.608
$p = 0.4$	0.512	0.524	0.548	0.559	0.579	0.618
$p = 0.5$	0.513	0.526	0.551	0.564	0.585	0.626
$p = 0.6$	0.514	0.527	0.554	0.568	0.590	0.633
$p = 0.7$	0.514	0.528	0.556	0.570	0.593	0.638

EXAMPLE 4. We return to normal responses for our final illustration. Let $x_1, x_2, \ldots, x_m$ and $y_1, y_2, \ldots, y_n$ be independent observations from the distributions $N(\mu_1, 1)$ and $N(\mu_2, 1)$, respectively. Let $\bar{x}_m$and $\bar{y}_n$ be the sample means at time $m + n = t$ and define

(8)
$$z_{m,n} = \frac{mn}{m+n}(\bar{x}_m - \bar{y}_n).$$

Consider sequential allocation rules $\{\psi_t\}$, where the conditional probability ψ_t depends only on the differences between responses up to that time: in other words, the σ-algebra $\Im_t$ is generated by $x_2 - x_1$, $x_3 - x_1$, $\ldots$, $x_m - x_1$, $y_1 - x_1$, $y_2 - x_1$, $\ldots$, $y_n - x_1$.

Robbins and Siegmund (1974) investigated sequential probability ratio tests defined by stopping the sequence $\{z_{m,n}\}$ at the point (M, N), which is the first time that $z_{m,n} \notin (-b, a)$ for some fixed a and b. They showed that, for a large class of allocation rules, the error probability of the test is approximately independent of the rule used. Their result involves neglecting any overshoot at the stopping barriers, but it yields good approximations when a and b are large. We shall confine our attention to the symmetric case $a = b$. Let $\mu_1 = \nu$, $\mu_2 = -\nu$ and suppose that $\nu \geq 0$. The error probability for the sequential probability ratio test based on (8) with barriers at $\pm b$ is given by Wald's approximation as $1 - h(\nu)$, where

(9)
$$h(\nu) \simeq \{1 + \exp(-4b\nu)\}^{-1}.$$

Thus, $h(\nu)$ corresponds to stopping at the upper barrier. In particular (9) holds for the following allocation rule:

$$\begin{aligned} \psi_0 &= \tfrac{1}{2}, \\ \psi_t &= \theta \text{ if } z_{m,n} \geq 0, \\ \psi_t &= 1-\theta \text{ if } z_{m,n} < 0. \end{aligned}$$

The constant θ must be chosen in the range $\frac{1}{2} \leq \theta < 1$. We shall examine the effect of selection bias when this rule is applied.

Now consider the situation when $\mu_1 = \mu_2$ always and suppose the experimenter favors treatment 1 by choosing patients with mean response $\mu_1 = \mu_2 = \delta$ whenever $\psi_t = \theta$ and patients with mean $-\delta$ when $\psi_t = 1-\theta$. Under this strategy, the behavior of the random process $\{z_{m,n}\}$ is more complicated and it is difficult to extend the results of Robbins & Siegmund. However, we can obtain a rough bound on the probability of selecting treatment 1 as the winner at the end. The heuristic argument given below indicates that

(10)
$$g(\delta) \leq \{1 + \exp(-4b\delta(2\theta - 1))\}^{-1}.$$

Comparison with (9) suggests that, at worst, selection bias may be equivalent to a true mean difference $\mu_1 - \mu_2 = 2\nu$, where $\nu = \delta(2\theta - 1)$.

For the moment, let us imagine that there are no barriers ($b = \infty$) and consider the situation when t is large. Write $m = m^+ + m^-$ and $n = n^+ + n^-$, where for example, n^- is the number of patients with mean response $-\delta$ receiving treatment 2 before time t. It is reasonably clear that all the numbers m^+, m^-, n^+, n^- will grow without limit as $t \to \infty$. Then the law of large numbers will ensure that

$$\frac{m^+}{m^+ + n^+} \simeq \theta, \qquad \frac{m^-}{m^- + n^-} \simeq 1 - \theta,$$

when t is large. However, not much can be said about the limiting behavior of the ratio $\varphi_t = (m^+ + n^+)/(m+n)$. We have

$$E\left\{\bar{x}_m - \bar{y}_n \mid m^+, m^-, n^+, n^-\right\} = \delta\left\{\frac{(m^+ - m^-)}{m} - \frac{(n^+ - n^-)}{n}\right\},$$

and this is not difficult to express in terms of θ and $\varphi = \varphi_t$. Some rearrangement leads to

$$E\left\{\bar{x}_m - \bar{y}_n \mid m^+, m^-, n^+, n^-\right\}$$

$$\simeq \frac{2\delta(2\theta - 1)\varphi(1-\varphi)}{\{\theta\varphi + (1-\theta)(1-\varphi)\}\{\theta(1-\varphi) + (1-\theta)\varphi\}}.$$

This quantity is non-negative, for all possible values $0 \leq \varphi \leq 1$. We can obtain a simple upper bound by noting that the denominator on the right is decreasing in θ, attaining its minimum value $\varphi(1-\varphi)$ when θ is replaced by 1. Hence, the unconditional expectation has approximate bounds

$$0 \leq E\left\{\bar{x}_m - \bar{y}_n\right\} \leq 2\delta(2\theta - 1).$$

The corresponding expectation under the original conditions of (9) is $E\left\{\bar{x}_m - \bar{y}_n\right\} = 2\nu$ and the inequality (10) is obtained on replacing the process $\{z_{m,n}\}$ by a Brownian motion with constant drift $2\delta(2\theta - 1)$. Of course, this argument is plausible for finite barriers only when b is large.

Table 2 gives estimates of $g(\delta)$ based on 10,000 simulations for a stopping rule determined by $b = 6$. In each case, the upper figure is the proportion of decisions in favor of treatment 1 and the lower figure is the bound obtained from (10). In general, this is a useful guide. It is also clear from the table that $g(\delta)$ can be close to 1.

Table 2.

Example 4: proportion of terminal decisions in favor of treatment 1 and upper bound (lower figure).

δ	0.010	0.025	0.050	0.100	0.250
$\theta = 0.5$	0.494	0.499	0.496	0.503	0.517
	0.500	0.500	0.500	0.500	0.500
$\theta = 0.6$	0.511	0.528	0.553	0.593	0.721
	0.512	0.530	0.560	0.618	0.769
$\theta = 0.7$	0.518	0.544	0.600	0.674	0.843
	0.524	0.560	0.618	0.723	0.917
$\theta = 0.8$	0.527	0.565	0.629	0.742	0.904
	0.536	0.589	0.673	0.809	0.973
$\theta = 0.9$	0.532	0.571	0.643	0.763	0.921
	0.548	0.618	0.723	0.872	0.992
$\theta = 0.95$	0.524	0.585	0.666	0.789	0.935
	0.554	0.632	0.747	0.897	0.996

4. Discussion. The examples we have considered show that selection bias can have a substantial effect in distorting the results of comparative experiments. The assumptions made here about the strategy of the experimenter are rather artificial: bias is not easy to model. However, the difficulty of detecting it is much more important and this is a strong argument for avoiding the possibility. It seems that this is best done by avoiding data-dependent allocation of treatments altogether. In the present context, this means restricting to allocation rules with $\psi_t = \frac{1}{2}$ always. Alternatively, we can rely on group sequential designs where the patients are arranged in small groups, before randomization.

There are disadvantages in restricting allocation to randomization with unbiased coins. As we mentioned earlier, one consequence is loss

of balance in the treatment allocations, but this seems minor in comparison with the ethical cost associated with patients receiving inferior treatments. The ethical cost is worth examining further.

Sequential methods can be based on random stopping, without using data-dependent allocation, and it is arguable that most of the possible gains in efficiency can be achieved in this way. Procedures for comparing several medical treatments have been studied recently from this point of view: see Bather and Coad (1992). The procedures there rely on eliminating treatments at suitable stopping times until only the winner is left. Allocations at intermediate times are always based on equal randomization between the surviving treatments. The paper shows that such procedures can achieve a pattern of error probabilities equivalent to the obvious fixed-sample procedure for a much lower level of expected successes lost. Coad has also shown (in a manuscript found in this volume) that the results obtained by Robbins and Siegmund for sequential probability ratio tests can also be extended to experiments with more than two treatments. His simulations of several data-dependent allocation rules and procedures with equal randomization indicate that the extra reductions in expected successes lost can be as high as 20 – 25%. This gives some idea of the advantages to be set against the unknown risk of selection bias.

APPENDIX

It remains to sketch the proof of (4). We have to show that $P(A) = \Phi(\delta/\lambda)$, where A is the event defined in (3), Φ is the standard normal distribution function and λ is given by (4). Let f be the density of the distribution $N(0, a)$. In the notation of Example 2, let $\Pi(x, y)$ be the conditional probability of A given x and y. It is easily verified, for $x \geq y$,

$$\prod(x, y) = \Phi\left\{(x(a+c) - y(a+b) + \xi)\sigma^{-1}\right\},$$

where $\xi = a(b-c)\delta$ and $\sigma^2 = b(a+c)^2 + c(a+b)^2$. Given any $x \geq y$, we can evaluate $\Pi(-x, -y)$ in a similar way and this leads to

$$\prod(-x, -y) = \Phi\left\{(y(a+c) - x(a+b) + \xi)\sigma^{-1}\right\}.$$

Hence

$$g(\delta) = P(A) = \iint_{x \geq y} \left\{ \prod(x, y) + \prod(-x, -y) \right\} f(x) f(y) \, dxdy.$$

We now transform the integral of the second term $\Pi(-x, -y)$ by writing $x = y'$, $y = x'$. It becomes

$$\iint_{x' \leq y'} \Phi \left\{ (x'(a+c) - y'(a+b) + \xi) \sigma^{-1} \right\} f(x') f(y') \, dx'dy'$$

and we can now express $g(\delta)$ as an integral over $\Re^2$:

$$g(\delta) = \iint \Phi \left\{ (x(a+c) - y(a+b) + \xi) \sigma^{-1} \right\} f(x) f(y) \, dxdy.$$

Finally, we can evaluate this integral as the probability that

$$y(a+b) - x(a+c) + \sigma z \leq \xi$$

when x, y and z are independent normal random variables with common mean zero and variances a, a and 1 respectively. It is a straightforward matter to verify that $g(\delta) = \Phi(\delta/\lambda)$, as required.

Acknowledgment. I am grateful to Steve Coad for the computations leading to Tables 1 and 2 and also for helpful discussions.

References

BARTLETT, R.H., ROLOFF, D.W., CORNELL, R.G., ANDREWS, A.F., DILLON, P.W. and ZWISCHENBERGER, J.B. (1985). Extracorporeal circulation in neonatal respiratory failure: a prospective randomized study. *Pediatrics* **76** 479-487.

BATHER J.A. and COAD, D.S. (1992). Sequential procedures for comparing several medical treatments. *Sequential Analysis* **11** 339-376.

BLACKWELL, D. and HODGES, J.L. (1957). Design for the control of selection bias. *The Annals of Mathematical Statistics* **28** 449-460.

COAD, D.S. (1992). Some results on estimation for two-stage clinical trials. *Sequential Analysis* **11** 299-312.

EFRON, B. (1971). Forcing a sequential experiment to be balanced. *Biometrika* **58** 403-17.

ROBBINS, H. and SIEGMUND, D.O. (1974). Sequential tests involving two populations. *Journal of the American Statistical Association* **69** 132-139.

ROYALL, R.M.(1991). Ethics and statistics in randomized clinical trials (with discussion). *Statistical Science* **6** 52-88.

SMITH, R.L. (1984). Sequential treatment allocation using biased coin designs. *Journal of the Royal Statistical Society B* **46** 519-543.

WARE, J.H. (1989). Investigating therapies of potentially great benefit: ECMO (with discussion). *Statistical Science* **4** 298-340.

WEI, L.J. (1979). The generalized Polya's urn for sequential medical trials. *The Annals of Statistics* **7** 291-296.

WEI, L.J. (1988). Exact two-sample permutation tests based on the randomized play-the-winner rule. *Biometrika* **75** 603-606.

WEI, L.J. and DURHAM, S. (1978). The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association* **73** 830-843.

MATHEMATICS DIVISION
UNIVERSITY OF SUSSEX
FALMER, BRIGHTON BN1 9QH
UNITED KINGDOM