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

RANDOMIZED ADAPTIVE DESIGNS

BY VINCENT F. MELFI

Michigan State University

Abstract

Ethical considerations in the conduct of clinical trials have led to the proposal of adaptive allocation schemes. Because of the deterministic nature of their treatment assignment rules, however, adaptive designs are susceptible to experimenter bias, and are insensitive to time trends in the data. For these and other reasons it is desirable to introduce some randomization into the adaptive allocation. In this paper, a class of randomized adaptive designs is introduced, and some renewal-theoretic tools needed for the analysis of such designs are developed.

1. Introduction. In clinical trials, it is desirable to include some measure of the ethical cost of assigning a patient to an inferior treatment, and then to find a design which minimizes the total cost of the trial. For example, it may be desired to minimize the expected number of patients on the inferior treatment, or to minimize some function of this and the total sample size. Such considerations have led to the proposal of adaptive designs, *i.e.*, designs for which treatment assignments may depend on the responses of previous patients in the trial.

Since adaptive designs have deterministic treatment assignment rules, they are susceptible to experimenter bias, and are insensitive to time trends in the data. (In Section 2, these problems will be illustrated

Received October 1992; revised April 1993.

AMS 1991 subject classification. Primary 60K05, 60K15; secondary 62L05, 60J05.

Key words and phrases. Adaptive allocation, biased-coin allocation, Markov random walk, renewal theory.

in the context of a particular adaptive design.) For this reason, it is desirable to investigate *randomized adaptive designs*, which introduce some randomization into the treatment assignment rules. It is natural to ask what effect randomization has on the properties of the adaptive design; in particular, does the randomized adaptive design have the desirable ethical properties of the adaptive design? These are difficult questions to answer; properties such as the expected number on the inferior treatment and expected sample size require a renewal theory to be developed. In this paper, a class of randomized adaptive designs is introduced, and renewal theory is discussed for such designs.

REMARK 1. Let $\{Z_n : n \geq 1\}$ be a stochastic process, and let $t_a = \inf\{n : Z_n > a\}$. Then $Z_{t_a} - a$ is the overshoot of the process over a. For the purposes of this paper, a renewal theorem is a result which asserts that the distribution of the overshoot converges to a limiting distribution as $a \to \infty$, and which identifies this limiting distribution.

REMARK 2. The role of renewal theory in approximating the characteristics of sequential procedures has a long history. For examples, see Woodroofe (1982).

REMARK 3. Since this paper is expository, the discussion of renewal theory will be somewhat heuristic, and technical details will be relegated to the references.

2. An adaptive design. Let $(X_1^A, X_1^B), (X_2^A, X_2^B), \ldots$ be independent with $X_i^A \sim \mathcal{N}(\theta^A, 1)$ and $X_i^B \sim \mathcal{N}(\theta^B, 1)$ for $i \geq 1$. It is desired to test whether A is superior to B, *i.e.*, whether $\theta := \theta^A - \theta^B$ is greater than 0. [Here X_i^A and X_i^B represent the potential responses of the i^{th} patient to treatments A and B, respectively.]

For each $i \geq 1$, let δ_i equal 1 or 0 depending on whether the i^{th} patient is assigned to treatment A or treatment B. Then $m_k = \delta_1 + \ldots + \delta_k$ and $n_k = k - m_k$ are the numbers of the first k patients assigned to A and B respectively. For $k \geq 1$ and $a, b \in (0, \infty)$, define

$$\begin{aligned} \hat{\theta}_k &= (\frac{1}{m_k}) \sum_{j=1}^{k} \delta_j X_j^A - (\frac{1}{n_k}) \sum_{j=1}^{k} (1 - \delta_j) X_j^B; \\ Z_k &= (\frac{m_k n_k}{k}) \hat{\theta}_k; \\ T_{a,b} &= \inf\{k \geq 1 : Z_k > a \ \ or \ \ Z_k < -b\}. \end{aligned}$$

Then a sequential probability ratio test of $\theta > 0$ takes $T = T_{a,b}$ observations and rejects $\theta > 0$ if and only if $Z_T < -b$.

Now suppose that it costs one unit to administer either treatment, and that there is an additional ethical cost of $g(\theta)$ for assigning a patient to the inferior treatment. It is a simple minimization problem to show that if θ were known, allocation in the proportions $p(\theta)$ to A and $1-p(\theta)$ to B minimizes the total cost of the study, where $p(\theta) = 1/(2+\sqrt{g(\theta)})$ if $\theta < 0$ and $p(\theta) = 1 - p(-\theta)$ if $\theta > 0$. This suggests the following adaptive design: Allocate the $(k+1)^{st}$ patient to A if $m_k/k \leq p(\hat{\theta}_k)$; otherwise allocate the patient to B.

REMARK 4. Since anyone with access to the current data can predict, with certainty, the treatment assignment of the next patient, the potential is great for experimenter bias (via screening the patients, for example). To see the problem with possible time trends in the data, consider a study which uses balanced allocation in the beginning (for the first 50 patients, say) in order to obtain a good estimate of θ, and then uses the adaptive scheme until the conclusion of the study. Suppose that the difference between the two treatments changes during this second part of the study. If $p(\theta)$ is 2/3 at the beginning of the study, and then begins decreasing, the study may miss this change, because the design will be allocating all of the patients to A to drive the proportion on A up to 2/3. In short, there is potential for long periods during which no data is gathered on one of the two treatments.

REMARK 5. This allocation scheme is similar in spirit to that given in Robbins and Siegmund (1974); there the goal was to minimize the expected number on the inferior treatment.

3. Randomized adaptive designs. One way to introduce randomization into the allocation scheme of Section 2 is as follows. Let $U_1, U_2, \ldots$ be independent and identically distributed Uniform $(0,1)$ random variables which are independent of the responses $\{(X_i^A, X_i^B) : i \geq 1\}$. Then define δ_k by

$$\delta_k = I\left\{U_k \leq \frac{1+p(\hat{\theta}_k)}{2}\right\} \qquad \text{if } \frac{m_k}{k} < p(\hat{\theta}_k);$$

$$\delta_k = I\left\{U_k \leq \frac{p(\hat{\theta}_k)}{2}\right\} \qquad \text{if } \frac{m_k}{k} > p(\hat{\theta}_k);$$

$$\delta_k = I\left\{U_k \leq \frac{1}{2}\right\} \qquad \text{if} \quad \frac{m_k}{k} = p(\hat{\theta}_k);$$

In words, if the proportion allocated to A is less than $p(\hat{\theta}_k)$, allocate to A with probability $(1 + p(\hat{\theta}_k))/2$ [which is greater than 1/2]; if the proportion allocated to A is greater than $p(\hat{\theta}_k)$, allocate to A with probability $p(\hat{\theta}_k)/2$ [which is less than 1/2]; and if the proportion allocated to A is equal to $p(\hat{\theta}_k)$, randomize.

With this scheme, the potential for bias is greatly reduced, since it is not possible to predict with certainty the treatment assignment of the next patient. In addition, it avoids the problem of runs on one treatment described in Remark 4 above, since at each stage there is positive probability that the next patient will be assigned to each treatment.

This randomization and the adaptive nature of the original allocation scheme introduce a complex dependence structure into the test statistic process $\{Z_k\}$, which complicates the analysis of experiments which use randomized adaptive allocation schemes. The rest of the paper explores the question of developing a renewal theory under such dependence, which is an important step in analyzing such experiments.

REMARK 6. The randomization outlined above is closely related to the biased-coin allocation scheme proposed by Efron (1971), which was intended as a compromise between the competing goals of balance and randomization. Efron's scheme may be described as follows. Let $1/2 < \eta < 1$ be a constant. Let $U_1, U_2, \ldots$ be as above. Define

$$\begin{aligned}
\delta_k &= I\{U_k \leq \eta\} \text{ if } \frac{m_k}{k} < \tfrac{1}{2}; \\
\delta_k &= I\{U_k > \eta\} \text{ if } \frac{m_k}{k} > \tfrac{1}{2}; \\
\delta_k &= I\left\{U_k \leq \tfrac{1}{2}\right\} \text{ if } \frac{m_k}{k} = \tfrac{1}{2}.
\end{aligned}$$

The special case $\eta = 3/4$ corresponds to the above allocation when $p \equiv 1/2$ is known. So randomized adaptive allocation may be thought of as a generalization of Efron's scheme, where the target proportion $p(\theta)$ is unknown and is estimated at each stage. In the next two sections

it will be seen that this difference is what makes randomized adaptive allocation schemes so difficult to analyze.

REMARK 7. Of course, any adaptive allocation scheme may be randomized in the above fashion. Even outside the context of clinical trials, where experimenter bias may not be a concern, there are compelling reasons to include randomization in the allocation scheme. The problem with time trends referred to above is one reason; Efron (1971) gives a good account of some other virtues of randomization.

In addition, there are various ways to implement the randomization. In a recent series of papers, Eisele (1990, 1992) has introduced a randomized adaptive allocation design based on the Wei (1978) biased-coin design (rather than the Efron coin), and has applied it to the Behrens-Fisher problem. Further details may be found in the references.

4. Renewal theory for the Efron coin. This section develops a renewal theory for the special case of the Efron coin allocation scheme. In addition to being of interest in its own right, it will help to unravel the complicated dependence structure in the general randomized adaptive allocation scheme. The basic idea is to decompose the test statistic into $Z_n = S_n + \xi_n$, where $\{S_n\}$ is a process for which a renewal theorem is known, and the "perturbation terms" $\{\xi_n\}$ are negligible in the limit.

It is useful to recast the allocation scheme slightly. Let Y_k be the difference in the numbers allocated to A and B at time k, that is, $Y_k = m_k - n_k$. Notice that $\{Y_k\}$ is a Markov chain. Define the function $\phi(y)$ by

$$\phi(y) = \begin{cases} \eta & \text{if } y < 0; \\ 1/2 & \text{if } y = 0; \\ 1-\eta & \text{if } y > 0. \end{cases}$$

(Here $1/2 < \eta < 1$.) Then δ_k in the Efron scheme is given by

$$\delta_k = I\left\{U_k \leq \phi(Y_{k-1})\right\}.$$

Now rewrite the test statistic.

$$Z_k = \frac{m_k n_k}{k}\left(\frac{1}{m_k}\sum_{j=1}^{k}\delta_j X_j^A - \frac{1}{n_k}\sum_{j=1}^{k}(1-\delta_j)X_j^B\right)$$

$$= \frac{n_k}{k} \sum_{j=1}^{k} \delta_j X_j^A - \frac{m_k}{k} \sum_{j=1}^{k} (1-\delta_j) X_j^B$$

$$= \sum_{j=1}^{k} \frac{1}{2} \left[\delta_j X_j^A - (1-\delta_j) X_j^B \right]$$

$$+ \left(\frac{n_k}{k} - \frac{1}{2} \right) \sum_{j=1}^{k} \left[\delta_j X_j^A + (1-\delta_j) X_j^B \right]$$

$$=: S_k + \xi_k.$$

It remains to show that $\{S_k\}$ and $\{\xi_k\}$ have the properties described above.

5. Digression: Markov random walk; slowly changing sequences. If $Y_0, Y_1, \ldots$ is a Markov chain with stationary probability distribution π, and $S_n = X_1 + \ldots + X_n$, where $\{X_n\}$ has the property that

$$\mathcal{D}\left(X_n \mid \{Y_i : i \geq 0\}, \{X_j : j \neq n\}\right) = F\left(Y_{n-1}, Y_n\right),$$

then $\{S_n\}$ is called a *Markov random walk*. Renewal theory for such processes is well-known; in particular, Kesten (1974) proves a renewal theorem for Markov random walks under quite general conditions on $\{Y_n\}$ and $\{X_n\}$.

A bit of work is required to write down the limiting distribution obtained in Kesten's Theorem. Let P_y and E_y denote probability and expectation given that $Y_0 = y$. Define $\mu = \int E_y(X_1)\pi(dy)$. Also, let $\{(Y_n', X_n') : -\infty < n < \infty\}$ be the two-sided stationary process associated with $\{(Y_n, X_n) : n \geq 0\}$. (For details on its definition and construction, see Kesten (1974).) Define

$$S_n' = \begin{cases} \sum_{i=1}^{n} X_i' & \text{if } n > 0; \\ 0 & \text{if } n = 0; \\ -\sum_{i=n+1}^{0} X_i' & \text{if } n < 0. \end{cases}$$

and define the measure ψ by $\psi(E) = P\{\sup_{n<0} S_n' < 0, Y_0' \in E\}$. Let $\tau_a = \inf\{n : S_n > a\}$; then $S_{\tau_a} - a$ is the overshoot of $\{S_n\}$ over the boundary a. Kesten's Theorem gives conditions under which

$$\lim_{a \to \infty} P_y \left\{ S_{\tau_a} - a > r \right\} = \frac{1}{\mu} \int \psi(dz) \int_r^{\infty} (\lambda - r) P_z \left\{ S_{\tau_0} \in d\lambda \right\}.$$

[The conditions of the Theorem will not be repeated here; it suffices to state that all of the Markov random walks encountered below satisfy these conditions.]

A sequence $\{\xi_n\}$ is said to be slowly changing if the following two conditions hold:

(SC1)

$$P_y\left\{\frac{1}{n}\max_{1\le k\le n}|\xi_k| > \epsilon\right\} \to 0 \text{ as } n \to \infty, \text{ for each } \epsilon > 0 \text{ and } y;$$

(SC2)

$$\limsup_{\delta\to 0}\sup_{n\ge 1} P_y\left\{\max_{0\le k\le n\delta}|\xi_{n+k} - \xi_n| \ge \epsilon\right\} = 0 \forall y \text{ and } \epsilon > 0.$$

Renewal theory also has been investigated for sequences whose terms can be written as the sum of a Markov random walk and a slowly changing perturbation. Define $t_a = \inf\{n : Z_n > a\}$. Then $Z_{t_a} - a$ is the overshoot of $\{Z_n\}$ over the boundary a. The following is proved in Melfi (1992).

THEOREM 1. *If $Z_n = S_n + \xi_n$, where $\{S_n\}$ is a Markov random walk satisfying the conditions of Kesten (1974); if $\{\xi_n\}$ is slowly changing; and if ξ_n is $\sigma(Y_0, \ldots, Y_n, X_1, \ldots, X_n)$- measurable for each n, then $Z_{t_a} - a$ has the same limiting distribution as $S_{\tau_a} - a$, as $a \to \infty$.*

Returning to the problem at hand, it only remains to show that $\{Z_n\}$ satisfies the conditions of the theorem. With S_n and ξ_n defined by the above decomposition, it is easy to see that $\{S_n\}$ is a Markov random walk (the underlying Markov chain is $Y_k = m_k - n_k$). The conditions of Kesten's renewal theorem are easily verified for this process. Also, using results on the rate of convergence of n_k/k to $1/2$, the process $\{\xi_n\}$ can be shown to be slowly changing. Details are in Melfi (1992).

REMARK 8. Lai and Siegmund (1977) proved a result similar to Theorem 1 above in the case where $\{S_n\}$ is an ordinary random walk. This result has found extensive application in sequential analysis. For details, see Woodroofe (1982) and Siegmund (1985).

6. Renewal theory for randomized adaptive designs. The randomized adaptive allocation scheme introduced in Section 3 may be

described as follows. Let $W_k = m_k - kp(\hat{\theta}_k)$ measure the difference between the number allocated to A at time k and the number that should be allocated to A if $\theta = \hat{\theta}_k$. Define the function ϕ by

$$\phi(w,p) = \begin{cases} \frac{1+p}{2}, & \text{if } w < 0; \\ \frac{1}{2}, & \text{if } w = 0; \\ \frac{p}{2}, & \text{if } w > 0. \end{cases}$$

Then

$$\delta_k = I\left\{U_k \leq \phi(W_{k-1}, p(\hat{\theta}_{k-1}))\right\}.$$

From this description, one problem in analyzing such a design is immediately apparent. In contrast to the Efron coin, the design sequence (here called $\{W_n\}$) is not a Markov chain.

To attempt to circumvent this problem, reason as follows. The reason that the design sequence is non-Markovian is that it involves the estimate $p(\hat{\theta}_k)$, which depends on the whole history of the experiment. Since $p(\hat{\theta}_k)$ is converging to the true optimal proportion $p(\theta)$ as k gets large, however, the process $\{W_k\}$ should "behave" like a Markov chain for large k. This may be made precise via a Taylor expansion of $p(\hat{\theta}_k)$.

Call the Markov chain obtained in the above manner $\{Y_k\}$. The next step is to decompose the test statistic Z_k as in the previous section. Unfortunately, a more serious problem arises here: The perturbation term obtained in this decomposition does not satisfy the slowly changing conditions. (Again the culprit is the fact that the optimal allocation proportion must be estimated. This introduces additional variability into the process, which makes the perturbation terms too large.)

The way out of this problem leads to a significant reformulation of Theorem 1. Instead of requiring closeness of Z_n to a Markov random walk in the sense that the difference between the two be slowly changing, the new theorem will just require that the (conditional) distributions of the two processes become close. The details follow.

Some additional definitions and notation are needed in order to make the notion of distributional closeness precise. The Prokhorov metric will be used to measure distributional closeness. For probability measures P and Q defined on the Borel sets $\mathcal{A}$ of a metric space $(\mathcal{X}, d)$, the Prokhorov metric ρ is defined by

$$\rho(P,Q) = \inf\{\epsilon > 0 : P(A) \leq Q(A^\epsilon) + \epsilon \ \ \forall A \in \mathcal{A}\},$$

where $A^\epsilon = \{z \in \mathcal{X} : d(z,A) < \epsilon\}$. As long as $(\mathcal{X}, d)$ is separable, convergence in the Prokhorov metric is equivalent to weak convergence, so it is a natural measure of distributional closeness. For this and other properties, see Dudley (1989).

Let $\mathcal{F}_n = \sigma(W_0, \ldots, W_n, Z_1, \ldots, Z_n)$. Recall that $t_a = \inf\{n : Z_n > a\}$ is the first passage time of the process $\{Z_n\}$ over a. The post-t_a delayed process and the prior σ-algebra are defined by

$$Z_{a,k} = Z_{t_a+k} - Z_{t_a}, \quad k \geq 1;$$

$$\mathcal{F}_a = \mathcal{F}_{t_a} = \{E : E \cap \{t_a = n\} \in \mathcal{F}_n \quad \text{for all} \quad n \geq 1\}.$$

Denote the m-step distribution of a Markov random walk by Q_m^* and the m-step conditional distribution of the post-t_a delayed process by $Q_{a,m}$, *i.e.*

$$Q_m^*(y; B) = P_y\{(S_1, \ldots, S_m) \in B\}; \text{ and}$$

$$Q_{a,m}(\omega, B) = P\{(Z_{a,1}, \ldots, Z_{a,m}) \in B | \mathcal{F}_a\}(\omega).$$

The following condition replaces the requirement that Z_n differ from S_n by a slowly changing sequence.

CONDITION I. *There exists a Markov random walk* $\{S_n\}$ *satisfying the conditions of Kesten's Renewal Theorem for which*

$$\rho[Q_{a,m}, Q_m^*(W_{t_a}; \cdot)] \to 0 \text{ in probability as } a \to \infty, \text{ for every } m \geq 1.$$

Basically, the condition says that, in the limit, the m-step conditional distribution of the process $\{Z_n\}$ must be close to the m-step conditional distribution of a Markov random walk, started at the same place (W_{t_a}). Notice that the closeness required is weak: First, it is only distributional closeness; second, the convergence is only in probability. The weakness of this condition requires the addition of the following two conditions:

CONDITION II. $\{Z_{t_a} - a : a > 0\}$ *is tight.*
CONDITION III. $\{W_{t_a} : a > 0\}$ *is tight.*

THEOREM 2. *Assume Conditions I-III. Then* $Z_{t_a} - a$ *has the same limiting distribution as* $S_{\tau_a} - a$, *as* $a \to \infty$.

REMARK 9. The convergence of the overshoot of the Markov random walk and the requirement that Z_n differ from S_n by a slowly changing perturbation term in Theorem 1 imply that Condition II holds. Under the weaker Condition I, the tightness of the overshoot of $\{Z_n\}$ must be assumed. The reason that Condition III is needed has to do with the proof of the Theorem. In the proof of Theorem 1, it is necessary to condition on $\mathcal{F}_N$ for fixed (nonrandom) times N, and then use the tightness of the Markov chain $\{Y_n\}$ and a result on uniform convergence on compacts in Kesten's Theorem. In Theorem 2, it is necessary to condition on $\mathcal{F}_{t_a}$, and so the tightness of $\{W_n\}$ sampled at the random times t_a is needed. A proof of Theorem 2, and more on the relation of the conditions of Theorem 2 to those of Theorem 1, may be found in Melfi (1994).

6. Further research. As mentioned in the introduction, the purpose of developing a renewal theory for randomized adaptive designs is to obtain information about the properties of studies which use these designs. There is more work to be done before such questions can be answered for general randomized adaptive designs. Below are a few of the more challenging questions that remain.

A basic problem is that Condition III of Theorem 2, which requires that $\{W_{t_a} : a > 0\}$ be tight, is difficult to verify. The difficulty is due to two factors. First, the process $\{W_n\}$ itself has a complicated dependence structure. Second, tightness is needed for the process sampled at the random times t_a. Two approaches to this problem are being considered. On is to try to weaken Condition III by modifying the proof of Theorem 2. The other is to utilize the fact that $\{W_n\}$ behaves like a Markov chain in the limit to verify Condition III.

The limiting distribution of the overshoot in Theorems 1 and 2 is related to the measure ψ and to the distribution of S_{τ_0}. For the Efron biased-coin allocation scheme treated in Section 4, these quantities can be related to the distributions of random walks, and thereby can be simplified sufficiently to become computationally tractable [the details of this simplification are given in Melfi (1992)]. In the general case, this is not possible, and some other method must be discovered.

In the case where $\{S_n\}$ is an ordinary random walk, (*i.e.*, the sum of independent random variables with common distribution F), this sort of simplification has been carried out. The main tool used is Spitzer's Identity, which relates the distribution of S_{τ_0} to the distribution F of the summands [see Woodroofe (1982) for details]. It remains to be seen whether some analogue of Spitzer's Identity holds in the Markov random walk setting.

References

DUDLEY, R.M. (1989). *Real Analysis and Probability.* Pacific Grove, CA: Wadsworth.

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. *Sequential Analysis* **9** 343-359.

EISELE, J. (1994). The doubly adaptive biased coin design for sequential clinical trials. *Journal of Statistical Planning and Inference* **38** 249-262.

KESTEN, H. (1974). Renewal theory for functionals of a Markov chain with general state space. *The Annals of Probability* **2** 355-386.

LAI, T.L. and SIEGMUND, D. (1977). A nonlinear renewal theory with applications to sequential analysis I. *The Annals of Statistics* **5** 946-954.

MELFI, V. (1992). Nonlinear Markov renewal theory with statistical applications. *The Annals of Probability* **20** 753-771.

MELFI, V. (1994). Nonlinear renewal theory for Markov random walks. *Stochastic Processes and Their Applications*, in press.

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

SIEGMUND, D. (1985). *Sequential Analysis.* New York: Springer.

WEI, L.J. (1978). The adaptive biased coin design for sequential experiments. *The Annals of Statistics* **6** 92-100.

WOODROOFE, M. (1982). *Nonlinear Renewal Theory in Sequential Analysis.* Philadelphia: Society for Industrial and Applied Mathematics.

DEPARTMENT OF STATISTICS AND PROBABILITY
MICHIGAN STATE UNIVERSITY
A-421 WELLS HALL
EAST LANSING, MI 48824