DOCUMENT RESUME

ED 316 563 TM 014 513

AUTHOR	Flournoy, Nancy
TITLE	Adaptive Sampling Designs.
PUB DATE	89
NOTE	9p.
PUB TYPE	Reports - Evaluative/Feasibility (142)
EDRS PRICE	MF01/PC01 Plus Postage.
DESCRIPTORS	Algorithms; Evaluation Methods; Mathematical Models; *Research Design; *Sampling; Statistical Analysis
IDENTIFIERS	Adaptive Allocation; Clinical Trials; *Sequential Sampling Procedures; *Urn Models

ABSTRACT

Designs for sequential sampling procedures that adapt to cumulative information are discussed. A familiar illustration is the play-the-winner rule in which there are two treatments; after a random start, the same treatment is continued as long as each successive subject registers a success. When a failure occurs, the other treatment is used until failure occurs. Sequential sampling procedures are discussed in relation to clinical trials, but there are many applications of such procedures in statistics. Models are presented for the adaptive allocation of dose levels to subjects arriving sequentially. Natural multivariate design spaces occurring in radiotherapy given at the Fred Hutchinson Cancer Research Center motivate this work. The allocation rules proposed adapt to accumulating failure data, possibly censored, and thus are governed by an underlying, unknown, multivariate failure distribution in the limit. An urn model is used to effect the adaptation of the allocation scheme. A recursive algorithm for updating the urn distribution is derived. The Bernoulli environment and urn models on ordinal sample spaces in a random environment are also considered. (SLD)

ADAPTIVE SAMPLING DESIGNS

Nancy Flournoy
The American University

ADAPTIVE SAMPLING DESIGNS

Nancy Flournoy, The American University
Nancy Flournoy, Clark Hall #211, 4400 Massachusetts Ave. N.W., Washington D.C. 20016-8050

KEY WORDS: dose response surface, failure models, random walks, play-the-winner rules, urn models, random environment

ABSTRACT

Models are presented for the adaptive allocation of dose levels to subjects arriving sequentially. Natural multivariate design spaces occurring in radio-therapy given at the Fred Hutchinson Cancer Research Center motivate this work. The allocation rules proposed adapt to accumulating failure data, possibly censored, and thus are governed by an underlying, unknown, multivariate failure distribution in the limit. An urn model is used to effect the adaptation of the allocation scheme. A recursive algorithm for updating the urn distribution is derived.

1. BACKGROUND

This paper is concerned with designs for sequential sampling procedures that adapt to cumulative information. A familiar illustration is the play-the-winner rule in which there are two treatments; after a random start, the same treatment is continued as long as each successive subject registers a success. When a failure occurs, the other treatment is used until another failure is registered. This allocation strategy is continued until a predetermined stopping criterion is met, at which time the decision as to which treatment is "better" is made by probabilistic criteria. This design, and many variants thereof, have been the source of numerous mathematical analyses such as described in Chow, Robbins and Siegmund (1971). Sequential analysis has become a large subfield in statistics and we briefly sketch the further development of sequential methods as it has motivated this work.

Herein, the term *subject* refers to a generic experimental unit. Exchanging the term subject to *component* or *cells* would connote industrial or laboratory applications, respectively. It is important to note that the designs depend on the spatial and temporal features of the measurements being observed on each experimental unit. The designs considered will be especially useful when the experimental units are expensive.

Sequential sampling procedures were developed during World War II to analyze observations taken in sequence. This was particularly suitable for industrial inspections in which measurements can be taken sequentially from objects arriving on conveyor belts. These early developments assume the outcome measurement is taken on each subject before another subject arrives. Armitage (1960) presents a review of sequential clinical experiments. Hoel, Sobel and Weiss (1975) survey adaptive sampling methods for clinical trials and report that, to the best of their knowledge at that time, no sequential methods were being used in clinical experiments. Two major barriers to the application of sequential methods to clinical studies at that time were that most clinical results are not observed in the same sequence as subjects are treated, and that observations may be truncated (censored).

Two theoretical thrusts have led to the widespread use of sequential analytic methods when observations are truncated. First, the development of nonsequential methods for truncated observations had to precede the development of sequential designs for such observations. Statistics comparing two treatments were first developed by Gehan (1965) and Mantel (1967), and extended, by Breslow (1970) to accommodate k treatments. By the time the two papers by Peto, Pike, Armitage, Breslow, Cox, Howard, Mantel, McPherson, Peto, and Smith (1976, 1977) codified the now classical design for clinical trials, a fixed two point sampling space already was rapidly becoming a design standard when failure distributions were positive for years after a subjects arrives and is treated. This standard two point sampling space reflected the forefront of the existing statistical tests permitting truncated realizations of the failure process. From 1975 to 1986, we designed and conducted dozens of such experiments at the Fred Hutchinson Cancer Research Center and other medical institutions did likewise.

The value of sequential decision rules for the treatment of patients was apparent, which led to

a major research thrust that centered around the development of sequential tests for two point sampling spaces. Because of logistic difficulties in the retrieval and management of data derived from monitoring subjects to obtain failure times, group sequential approaches became practical in many applications where fully sequential designs were not. Therefore, this research thrust focused on the development of group sequential methods (see Pocock (1977), Harrington, Flemming, and Green (1982), Whitehead (1983) and Fleming, Harrington, and O'Brien (1984)). An adaptive multistage rule was proposed by Tsutakawa (1972) for observations that are logistically distributed. More recently, increased attention is being directed at multistage procedures. For example, see Hall (1981), Cohen and Sackrowitz (1984), Witmer (1986), Clayton and Witmer (1988) and Lorden (1988).

About the time that sequential methods for comparing two truncated failure processes were being developed, attention was focused on studies with more that two treatments. Wampler, Carter, and Williams (1978) and Stablein, Carter, and Wampler (1980) studied experimental designs for estimating parameters from truncated sample failure processes on multidimensional sample spaces. But these have not been extended to models dealing with the sequential arrival of subjects. Flournoy (1986) proposed combining group sequential arrival models with a sequential stopping rule where multiple treatments could be ranked on an ordinal scale, and experimental designs based on this model are now being applied at the Fred Hutchinson Cancer Research Center. These models contain a mix of theoretical and heuristic considerations, including aspects of the sequential allocation rules proposed by Dixon (1965) and Wetherill (1966) with the two-stage Bayesian designs developed by Tsutakawa (1972). The need to formalize the key objectives of these applications motivate this paper.

The sampling designs described are directed at the problem of estimating the relative effects of treatment on the underlying failure distribution. They are based on a modification and extension of Pólya-Eggenberger urn models. This idea arose as an extension of the play-the-winner rule by Zelen (1969), from a design for an outcome that has a Bernoulli distribution, to one in which the outcome is distributed as a general time varying failure process In these models, a treatment decision for an arriving subject must be made before the outcome of the previous trial is observed. Zelen described how urn models could be used to decouple arrival and failure sequences. Wei (1979, 1988) characterized this model mathematically and developed statistical tests based on it.

In Section 2, we describe an extended Pólya-Eggenberger urn model with a generalized sample space in the environment of a failure process. In Section 3, we exemplify this generalized urn model by assuming the failure process has a conditional Bernoulli distribution and, in Section 4, a specific replacement strategy is incorporated into the model. Finally, in Section 5, a recursive algorithm for the evolving urn distribution is developed in terms of the observed failures and the replacement algorithm.

2. URN MODELS ON ORDINAL SAMPLE SPACES IN A RANDOM ENVIRONMENT

The class of Pólya-Eggenberger urn models was designed to model potential contagion (for a general description, see Johnson and Kotz (1977)). A ball is taken from an urn with b black and r red balls, its color is noted, and that ball plus s b s of the same color are returned to the urn. e that Pólya-Eggenberger urn models form a general class in which balls are distinguished by a categorical feature, such as color. Hill, Lane and Sudderth (1980) generalized the Pólya-Eggenberger urn models, retaining the categorical labeling of balls, while studying the asymptotic properties of several stochastic replacement strategies. Chen and Starr (1980) studied optimal stopping when drawing balls with ordinal labels from an urn, but their definition of 'optimal' reflects a different goal from the one described here.

We extend the Pólya-Eggenberger models in two central ways: we assume (1) a random process governs not only the drawing of balls, but their replacement as well, and (2) balls are categorized according to an ordinal attribute that maps to a parameter in the distribution governing the replacement of balls. The first feature places the urn model in a 'random environment'; the second links the drawing and replacement distributions by means of a bivariate time series so that the drawing distribution may adapt to the characteristics of the replacement distribution. The replacement strategy is specified as a mapping that is a function of the experimental results. This map

links the urn distribution to the failure distribution.

Consider a sequence of trials, in each of which a subject receives a treatment that involves a particular level or dose. The level that will be most effective is unknown, and it may be that a single optimal treatment level does not exist. Instead, it is expected that efficacy follows a probability distribution over the treatment space. Indeed, the failure density may well be composed of a mixture in the sense that one failure process may predominate if the treatment level is too low and another may predominate if the treatment level is too high. If a treatment level is too low, then we want to increase the level for subsequent subjects. Also, the treatment may be toxic or expensive, so we want to decrease the level for subsequent subjects if it is higher than necessary.

3. SAMPLING FROM AN ORDINAL URN PROCESS

For each subject, a ball is drawn that determines the treatment level. We construct a sampling protocol that adapts to the distribution of the *minimum* effective treatment level. This is accomplished by letting the unknown random failure process(es) that governs the outcome of treatment also govern the replacement of balls in the urn.

Consider an urn with balls labeled j=1, ..., J (J large) where the label determines the treatment level. The number of balls in the urn at each level is a random variable that changes as outcomes are observed. Let $\underset{\sim}{\beta}_k=(\beta_{k1}, \ldots, \beta_{kJ})$ be the random vector whose j-th element, β_{kj}, is the number of balls at level j just prior to the time that the k-th ball is drawn. Let $\underset{\sim}{\beta}_k$ assume vector values $\underset{\sim}{b}_k$ with elements

$$b_{kj} \in \left\{0, 1, 2, \ldots, B_k : \sum_j b_{kj} = B_k\right\}.$$

Then $P\{\underset{\sim}{\beta}_k = \underset{\sim}{b}_k\}$ denotes the urn distribution upon the arrival of the k-th subject. The initial urn distribution and the total starting number of balls in the urn are fixed, that is

$$P\{\underset{\sim}{\beta}_0 = \underset{\sim}{b}_0 : \sum_{j=1}^{J} b_{0j} = B_0\} = 1.$$

We restrict this analysis to the case when $\Sigma_j b_{kj} = B$ for all k which implies that the k-th ball drawn must be replaced before the (k+1)st draw. This means that failures are observed in the same sequence as subjects arrive for treatment assignments. By permitting B_k to vary we would accommodate more general failure distributions than the one being studied.

The urn distribution depends on the density governing the draw of the balls and the replacement. The drawing of balls conditional on the urn distribution as each subject arrives has a multinomial distribution. This portion of the model remains constant across the entire general class of models we discuss while the algorithms for replacing balls and the failure distributions will vary.

Let the balls be drawn at random and consider a specific urn distribution $\underset{\sim}{b}_k$ upon the arrival of the k-th subject. Let $\underset{\sim}{X}_k \equiv (X_{k1}, \ldots, X_{k(J-1)})$ be a random vector with elements X_{kj}. Also let $\sum_{i=1}^{J} X_{ki} = 1$, where

$$X_{kJ} \equiv 1 - \sum_{i=1}^{J-1} X_{ki}.$$

Let $X_{kj}=1$ if a j-level ball is drawn upon the arrival of the k-th subject and $X_{ki} = 0$ for all $i \neq j$ for j=1, ..., J. Let $\underset{\sim}{p}_k \equiv (p_{k1}, \ldots, p_{kJ})$ be a vector with elements

$$p_{kj} \equiv b_{kj}/B, \qquad j=1, \ldots, J$$

where $\sum_{i=1}^{J} p_{ki} = 1$ for all k.

Then conditional on the urn distribution $\underset{\sim}{\beta}_k=\underset{\sim}{b}_k$, $\underset{\sim}{X}_k$ has a multinomial distribution with parameter $\underset{\sim}{p}_k$. The event that a ball labeled with the j-th treatment level is chosen is $\{X_j = 1\}$. This event is equivalent to the vector valued event

$$\{\underset{\sim}{X}_k = \ldots, x_{kj} = 1 \;\&\; x_{ki} = 0, \forall\, i \neq j, i = 1, \ldots, J\}.$$

It follows that X_j is conditionally distributed as a Bernoulli with parameter p_{kj} and the conditional probability of the k-th ball being drawn having label j is written as:

$$P\{X_{kj} = 1 \mid \underset{\sim}{p}_k\} = p_{kj}. \qquad (1)$$

If $\underset{\sim}{b}_k^+$ is the (1xJ) vector tabulating the number of balls at each level in the urn immediately after the k-th draw, and before any additional balls are placed in the urn, then

$$\underset{\sim}{b}_k^+ \equiv \underset{\sim}{b}_k - \underset{\sim}{x}_k. \qquad (2)$$

4. THE BERNOULLI ENVIRONMENT

The law governing the drawing of balls from the urn upon arrival of the (k+1)st subject not only depends on the history of the experiment through the urn's distribution at the k-th arrival, but also on realizations of the failure process during this inter-arrival time. Let U_k denote the random failure process for the k-th subject. If the failure process U_k is a function of time since arrival (time on trial), then naturally the replacement operator also would be a function of time since arrival. To define a specific urn process in the environment of a random failure process, select the failure distribution and the replacement operator.

At this stage of our development, we consider a simple failure process:

$$U_k = \begin{cases} 1 & \text{if the k-th outcome is a failure} \\ 0 & \text{otherwise.} \end{cases}$$

Also assume that the subjects k=1, 2, ... are homogeneous; and that the processes determining the subjects' arrival times are independent of each other and independent of the failure process. The outcome U_k is not observed unconditionally, but only conditional given a treatment level j. Our concern is with the relationship between the conditional failure probabilities as the treatment level varies. Let

$$P\{U_k = 1 \mid \underset{\sim}{X}_k:\ X_{kj}=1\} \equiv \pi_j$$

denote the probabilities of failure given treatment level j and let

$$P\{U_k = 0 \mid \underset{\sim}{X}_k:\ X_{kj}=1\} \equiv 1-\pi_j \equiv \bar{\pi}_j$$

$$j = 1, \ldots, J \text{ and } k=1, 2, \ldots .$$

Conditional on the k-th draw, the failure process U_k has a Bernoulli distribution:

$$P\{U_k = u_k \mid \underset{\sim}{X}_k:\ X_{kj}=1\} = \pi_j^{u_k} \bar{\pi}_j^{\bar{u}_k}, \quad (3)$$

$$j = 1, \ldots, J \text{ and } k=1, 2, \ldots ,$$

where we write $\bar{z} = 1 - z$.

From equation (3), it is clear that once the treatment level is determined by the draw $\underset{\sim}{x}_k$, the failure process U_k has no further dependence on the urn distribution or on its prior evolution. When equations (1) and (3) are applied to the joint probability mass function of the treatment and outcome for subject k conditional on the urn distribution $\underset{\sim}{\beta}_k$, the joint probability factors into two terms, only one of which depends on the urn distribution: That is, when $\{X_{kj}=1\}$, we have

$$P\{U_k = u_k, \underset{\sim}{X}_k = \underset{\sim}{x}_k \mid \underset{\sim}{p}_k\} \quad (4)$$

$$= P\{U_k = u_k \mid \underset{\sim}{X}_k:\ X_{kj}=1\}\ P\{X_{kj} = 1 \mid \underset{\sim}{p}_k\}$$

$$= \pi_j^{u_k} \bar{\pi}_j^{\bar{u}_k} p_{kj} .$$

Let y_{kj} be the number of failures among the first k subjects treated at the j-th dose level and n_{kj} the number of trials at the j-th dose level. Then

$$y_{kj} = \sum_{i=1}^{k} u_i x_{ij} \text{ and } n_{kj} = \sum_{i=1}^{k} x_{ij}, \quad (5)$$

$$j = 1, \ldots, J.$$

The difference $n_{kj} - y_{kj}$ is the number of successes among the first k trials. The conditional distribution of y_{kj} given n_{kj} is Binomial (n_{kj}, π_j). Furthermore, the k-th subject's marginal failure density, conditional on $\underset{\sim}{\beta}_k$ is

$$P\{U_k = u_k \mid \underset{\sim}{\beta}_k = \underset{\sim}{p}_k B\}$$

$$= \sum_j \pi_j^{u_k} \bar{\pi}_j^{\bar{u}_k} p_{kj}.$$

Let $(\mathcal{U}_k, \mathfrak{X}_k)$ be a $(k \times (J+1))$ matrix in which the m-th row is $(U_m, \underset{\sim}{X}_m)$. Then the total accumulation of random variables through the k-th trial is contained in $(\mathcal{U}_k, \mathfrak{X}_k)$. Analogously, let be the $(k \times (J+1))$ matrix of observations accumulated through the k-th trial: That is, the m-th row of $(u, \underset{\sim}{x}_k)$ is $(u_m, \underset{\sim}{x}_m)$ for $m = 1, \ldots, k$.

Then using Equation (4), the density of $(\mathcal{U}_k, \mathfrak{X}_k)$ conditional on the evolving urn distribution parameters $\{\underset{\sim}{p}_1, \ldots, \underset{\sim}{p}_k\}$ factors into a product of densities:

$$P\left\{(\mathcal{U}_k, \mathfrak{X}_k) = (\underset{\sim}{u}_k, \underset{\sim}{x}_k) \mid \underset{\sim}{p}_1, \ldots, \underset{\sim}{p}_k\right\}$$

$$= \prod_{m=1}^{k} P\{U_m = u_m \mid \underset{\sim}{X}_m\}\ P\{\underset{\sim}{X}_m = \underset{\sim}{x}_m \mid \underset{\sim}{p}_m\}.$$

Note that only the second term in the product depends on the history of the urn distribution.

5. AN EXEMPLARY REPLACEMENT STRATEGY

For an initial replacement strategy, let $B_k=B$ for all k and define the replacement operator to restrict the level on the replacement ball to be adjacent to the level on the ball drawn. If the j-th ball is drawn upon the arrival of the k-th subject, then one such strategy is defined by the following recursive formula:

$$b_{(k+1)j} = b_{kj}^{+} + (x_{k(j-1)}u_k + x_{k(j+1)}\bar{u}_k), \quad (6)$$

$$j=1, \ldots, J,$$

where the subscripts are taken mod(J) so that

$$b_{(k+1)1} = b_{k1}^{+} + (x_J u_k + x_2 \bar{u}_k)$$

and

$$b_{(k+1)J} = b_{kJ}^{+} + (x_{J-1}u_k + x_1 \bar{u}_k).$$

The cyclic behavior is adopted in order to make the transformation one to one. In the applications motivating this work, the number and range of treatment levels is sufficiently large that the replacement process should reach the boundaries with null probability. In other words, the treatment will be 'successful' at the highest levels, J−1 and J, and may well be far above the minimally successful level with severe toxicities. Similarly, at the lowest levels the treatment is known to be ineffective (the lowest level may be no treatment at all). If these conditions hold, the replacement process will never actually cycle.

The strategy expressed in (6) is just one of a wide class of interesting alternative replacement strategies, but it's simplicity is useful to illustrate the general class of adaptive models while permitting a number of their characteristics to be developed. Other strategies will be needed for varying underlying shapes of the dose-response surface. This strategy has the following interpretation: If U=1 there has been a failure and consequently a ball labeled with the next higher level is added to the urn, whereas if U=0, a ball marked with the next lower level is added to the urn.

We now extend (6) to a matrix representation. Let $R_k \equiv R(u_k)$ be the JxJ matrix operating the replacement strategy. If the failure process was more general than the Bernoulli, R_k would be a function the multiple outcomes under observation at the time of the k-th arrival, and thus a function of both the failure times and the interarrival times.

Let $\mathrm{Circ}(c_1, \ldots, c_J)$ denote a circulant matrix:

$$\mathrm{Circ}(c_1, \ldots, c_J) \equiv \begin{bmatrix} c_1 & c_2 & \cdots & c_J \\ c_J & c_1 & \cdots & c_{J-1} \\ \vdots & \ddots & \ddots & \vdots \\ c_2 & \cdots & c_J & c_1 \end{bmatrix}$$

In the example described by (6), the replacement operator is the circulant matrix

$$R_k \equiv \mathrm{Circ}(0, u, 0, \ldots, 0, \bar{u}).$$

Therefore the change in the distribution of balls in the urn induced by the k-th subject can be written as

$$\underset{\sim}{b}_{k+1} - \underset{\sim}{b}_k = \underset{\sim}{x}_k(R_k - I), \quad (7)$$

where subtracting the identity matrix accounts for the influence of the draw from the urn as given in (2).

6. ADAPTING THE URN PROCESS

To find the distribution of the urn as it evolves, we first decompose the replacement operator. Then we transform the joint density of the accumulating draws and the experimental results to a density of partial sums. Conditioning on one of the partial sums, we derive the urn distribution.

Define JxJ dimensional circulant matrix $P \equiv \mathrm{Circ}(0, \ldots, 0, 1)$, for which

$$P^{-1} = P^T = \mathrm{Circ}(0, 1, 0, \ldots, 0),$$

where P^{-1} is the inverse and P^T is the transpose of P. Note that P operating on $\underset{\sim}{x}$ shifts the elements in $\underset{\sim}{x}$ to the next higher dose level, whereas P^T shifts the elements in $\underset{\sim}{x}$ to the next lower dose level. Let $R_k = \bar{u}_k P + u_k P^T$. So using the difference equation (7), the urn distribution in (6) can be expressed in terms of the partial sums defined in (5):

$$\underset{\sim}{b}_{k+1} = \underset{\sim}{b}_1 + \sum_{i=1}^{k} \underset{\sim}{x}_i (R_i - I) \quad (8)$$

$$= \underset{\sim}{b}_1 + \underset{\sim}{n}_k(P - I) + \underset{\sim}{y}_k(P^T - P).$$

Thus we have derived an general algorithm by which the urn distribution can be updated. When the urn distribution reaches a steady state, the influence of the initial urn distribution will be negligible and it will reflect the failure distribution. A more precise characterization of this evolving reflection is currently under study.

REFERENCES

Armitage, P. (1960). *Sequential Medical Trials.* Blackwell, Oxford.

Breslow, N. (1970). A generalized Kruskal-Wallis test for comparing k samples subject to unequal patterns of censorship. *Biometrika 57*, 579-594.

Carter, Jr., W. H., Wampler, G. L., and Stablein, D. M. (1983). *Regression Analysis of Survival Data in Cancer Chemotherapy*, Marcel Dekker, Inc., New York.

Chen, W. and Starr, N. (1980). Optimal stopping in an urn. *Annals of Probability 8*, 451-464

Chow, Y. S., Robbins, H., Siegmund, D. (1971). Great Expectations, *The Theory of Stopping.* Houghton-Mifflin, Boston.

Clayton, M. and Witmer, J. (1987). Two-stage bandits. *Annals of Statistics 16*, 887-884.

Cohen, A. and Sackrowitz, H. B. (1984). Results on double sample estimation for the binomial distribution. *Annals of Statistics 12*, 1109-1116.

Cox, D.R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society, Series B 34*, 187-200.

Dixon, W. (1965). The up-and-down method for small samples. *Journal of the American Statistical Association 60*, 967-978.

Johnson, N.L. and Kotz, S. (1977). *Urn Models and Their Applications*, John Wiley and Sons, Inc., New York.

Fleming, T.R., Harrington, D.P., and O'Brien, P.C. (1984). Designs for group sequential tests. *Controlled Clinical Trials 5*, 348-361.

Flournoy, N. (1986). New designs for clinical experiments, in *The Adult Leukemia Center Grant Proposal CA18029* to the National Cancer Institute (E. D. Thomas, M. D. Principal Investigator).

Gehan, E. (1965). A generalized Wilcoxon test for comparing arbitrarily singly censored samples. *Biometrika 52*, 203-223.

Hall, P. (1981). Asymptotic theory of triple sampling for sequential estimation of the mean. *Annals of Statistics 9*, 1229-1238.

Harrington, D.P., Fleming T.R., Green, S.J. (1982). Procedures for serial testing in censored survival data. Pages 287-301 in: *Lecture Notes-Monograph Series: Survival Analysis*, (Crowley, J. and Johnson, R.A., Editors) The Institute of Mathematical Statistics, Hayward, CA.

Hill, B., Lane, D. and Sudderth, W. (1980) Strong law for some generalized urn processes. *Annals of Probability* 8, 214-226.

Hoel, D.G., Sobel, M. and Weiss, G.H. (1975) A survey of adaptive sampling for clinical trials. P. 29-61 in: *Perspectives in Biometrics* 1 (Elashoff, R.M., Editor), Academic Press, New York.

Lorden, G. (1988). Testing in Stages. Abstract, P. 206-66 in: *IMS Bulletin* 17, 234.

Mantel, N. (1967). Ranking procedures for arbitrarily restricted observations. *Biometrics 23*, 65-78.

Marshall, A.W. and Olkin, I. (1989). Bivariate distributions generated from Pólya-Eggenberger urn models. *Unpublished paper.*

Peto, R., Pike, M.C., Armitage, P., Breslow, N.E., Cox, D.R., Howard, S.V., Mantel, N., McPherson, K., Peto, J., and Smith, P.G. (1976). Design and analysis of randomized clinical trials requiring prolonged observation of each patient. Part I. Introduction and Design. *British Journal of Cancer* 34, 585-612.

Peto, R., Pike, M.C., Armitage, P., Breslow, N.E., Cox, D.R., Howard, S.V., Mantel, N., McPherson, K., Peto, J., and Smith, P.G. (1977). Design and analysis of randomized clinical trials

requiring prolonged observation of each patient. Part II. Analysis and Examples. *British Journal of Cancer* 35, 1-39.

Pocock., S.J. (1977) Group sequential methods in the design and analysis of clinical trials. *Biometrika 64*, 1-9.

Simon, R. (1977). Adaptive treatment assignment methods and clinical trials. *Biometrics 33*, 743-9.

Stablein, D.M., Carter, Jr., W.H., and Wampler, G.L. Survival analysis of drug combinations using a hazards model with time-dependent covariates. *Biometrics 36*, 537-546.

Tsutakawa, R.K. (1980). Selection of dose levels for estimating a percent point of a quantal response curve. *Applied Statistics 29*, 25-33.

Wei, L.J. (1979). The generalized Polya's urn for sequential medical trials. *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.

Wetherill, G.B. (1966). *Sequential Methods in Statistics*. Methuen, London.

Whitehead, J. (1983). *The Design and Analysis of Sequential Clinical Trials*. Ellis Horwood, Chichester.

Witmer, J. (1986). Bayesian multistage decision problems. *Annals of Statistics 14*, 283-297.

Zelen, M (1969). Play-the-winner rule and the controlled clinical trial. *Journal of the American Statistical Association 64*, 131-146.