Statistical Science
1996, Vol. 11, No. 2, 137–149

New Directions in Adaptive Designs

William F. Rosenberger

Abstract. In any sequential medical experiment on a cohort of human beings, there is an ethical imperative to provide the best possible medical care for the individual patient. This ethical imperative may be compromised if a randomization scheme involving 50–50 allocation is used as accruing evidence begins to favor (albeit not yet conclusively) one experimental therapy over another. Adaptive designs have long been proposed to remedy this situation. An adaptive design seeks to skew assignment probabilities to favor the treatment performing best thus far in the study, proportionately to the magnitude of the treatment effect.

Current researchers in adaptive designs are attempting to provide physicians with a wide choice of design options, and to address practical and ethical concerns within a rigorous mathematical framework. This paper focuses on several broad families of designs, including urn models, random walk rules and other rules. Numerous examples are given along with applications, dose–response studies, clinical trials for efficacy and combined toxicity–efficacy studies.

Key words and phrases: Clinical trials, dose–response studies, ethics, quantile estimation, random walks, randomized play-the-winner rule, urn models.

0. INTRODUCTION

0.1 A Motivating Example

Zelen and Wei (1995) describe a recent clinical trial by Connor et al. (1994) to evaluate the hypothesis that the antiviral therapy AZT reduces the risk of maternal-to-infant HIV transmission. A standard randomization scheme was used to obtain equal allocation to both AZT and placebo, resulting in 239 pregnant women receiving AZT and 238 receiving placebo. The endpoint was whether the newborn infant was HIV-negative or HIV-positive. An HIV-positive newborn could be diagnosed within 12 weeks; a newborn could be safely claimed to be HIV-negative within 24 weeks.

William F. Rosenberger is Assistant Professor, Department of Mathematics and Statistics, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, Maryland 21250, and is Adjunct Assistant Professor, Department of Epidemiology and Preventive Medicine, University of Maryland School of Medicine (e-mail: billr@math.umbc.edu). *This paper won the 1995 David P. Byar Young Investigator Award from the Biometrics Section of the American Statistical Association.*

The results of the trial were compelling: at the conclusion of the trial (which was stopped early), 60 newborns were HIV-positive in the placebo group and only 20 newborns were HIV-positive in the AZT group ($p < 0.01$). However, these numbers mask the harsh reality that three times as many infants on placebo have, for all intents and purposes, received a death sentence by the transmission of HIV from their mothers. Had they been given AZT, one could presume that many would have been saved. Some might say that the high price for medical experimentation on humans is more than made up for by the impact of these studies on our public health (an elegant treatise is Byar et al., 1976). Others might argue that any randomized placebo-controlled clinical trial is unethical. (For a lively interchange, see Royall, 1991, including, in particular, Byar's comments and Royall's rejoinder.) But is there a middle ground?

The study of antiviral therapy in maternal–infant HIV transmission is critically important to neonatal public health—there is no doubt of that. This author will assume that the information desired from such a study was best elicited via a randomized placebo-controlled clinical trial. That being the case, given the seriousness of the outcome of this study, it is reasonable to argue that 50–50 allocation was unethical. As accruing information favoring (albeit, not

conclusively) the AZT arm became available, allocation probabilities should have been shifted from 50–50 allocation proportional to the weight of evidence for AZT. Designs which attempt to do this are called *adaptive designs, response-adaptive designs* or *response-driven designs*.

0.2 Brief History and a Warning

The literature is replete with the concepts and philosophy behind adaptive designs. Herbert Robbins (see, in particular, Robbins, 1952) is perhaps the progenitor, and his work was followed by a flurry of activity in the 1960s, beginning with Anscombe (1963) and Colton (1963). Zelen's (1969) famous play-the-winner paper built on Robbin's ideas. One neglected, but lovely, paper (well ahead of its time), written by Cornfield, Halperin and Greenhouse (1969), synthesized Bayesian, adaptive and multistage designs into an elegant framework. More recent history is discussed by Hardwick (1989) and Rosenberger and Lachin (1993). It is not the intention of this paper to discuss in detail the ethical and logistical aspects of adaptive designs. Such issues have been hotly debated (for a review, see Rosenberger and Lachin, 1993). Rather, this paper focuses on new ideas and variants of old ideas within a rigorous mathematical framework.

Adaptive designs are attractive because they satisfy an ethical imperative of caring for the individual patient in a group experiment, while allowing for the same group inferences. They are attractive to mathematicians and statisticians because the designs impose dependencies which require the full arsenal of martingale techniques and stochastic processes. One important caveat is that in medical experimentation we are dealing with human lives; therefore ethical and logistical considerations must always drive the mathematics. We (the author being no exception) often forget this in our urge to display our mathematical prowess. As statisticians, our job is to develop large families of designs, appropriate for a wide variety of scenarios. Physicians should then be heavily involved in selecting a particular design and its parameters.

Short of writing a book, it would be impossible to review the rich assortment of adaptive designs that have been developed over the past few decades. This paper focuses primarily on designs employing several types of stochastic processes, including urn models and random walks, although occasional mention is made of other types of designs. Also, passing reference is made to inference, but the paper's principal theme is design.

Section 1 focuses on the generalized Pólya urn model, with applications to both early-phase and phase III clinical trials discussed in Section 2. Section 3 deals with adaptive designs for clinical trials with continuous responses and survival endpoints. Section 4 focuses on designs for phase I clinical trials, including random walk rules and the continual reassessment method. These designs and a design from Section 2 are briefly compared. Finally, in Section 5, some conclusions are drawn. All results in this paper are based on fixed sample size designs, which are most commonly used in United States clinical trials. There is no reason, in principle, that one cannot incorporate stopping rules with adaptive designs. However, this leads to different inferential problems (which may well be worth the effort to explore).

1. GENERALIZED PÓLYA URN

1.1 Model and Mathematics

One large family of randomized adaptive designs can be developed from the generalized Pólya urn (GPU) model (originally designated by Athreya and Karlin, 1968, as the "generalized Friedman's urn"). The model can be described as follows. Consider an urn containing particles of K types. Initially, the urn contains $\mathbf{Y_0} = (Y_{01}, \ldots, Y_{0K})$ particles, where Y_{0i} denotes the number of particles of type i, $i = 1, \ldots, K$. A particle is drawn or *split* at random from the urn. Its type is observed and the particle is then replaced. When a particle of type i is drawn, we say that a *type i split* occurs. Following a type i split, R_{ij} particles of type j, for $j = 1, \ldots, K$, are added to the urn, or *generated*. In the most general sense, R_{ij} can be random and can be some function of a random process outside the urn process. This is what makes the model so appropriate for adaptive designs (in our case, R_{ij} will be a random function of patient response). A particle must always be generated at each stage (in addition to the replacement), and so $P\{R_{ij} = 0, \text{ for all } j = 1, \ldots, K\}$ is assumed to be 0. Let $\mathbf{E}$ be the matrix comprising elements $\{E_{ij}\} \equiv E\{R_{ij}\}$. We refer to $\mathbf{R}$ as the *rule* and $\mathbf{E}$ as the *design matrix* (although in the branching processes literature it would be more proper to refer to $\mathbf{E}$ as the *generating* matrix), and $\mathbf{E}$ is assumed to be nonnegative and irreducible. After n splits and generations, the *urn composition* is given by the vector $\mathbf{Y_n} = (Y_{n1}, \ldots, Y_{nK})$, where Y_{ni} represents the number of particles in the urn of type i after n splits.

The matrix $\mathbf{E}$ has a maximal eigenvalue ρ with associated *left* eigenvector $\mathbf{v} = (v_1, \ldots, v_K)$ with $\sum_{i=1}^{K} v_i = 1$ (see, e.g., Gantmacher, 1959). The eigenvector $\mathbf{v}$ plays an important role in limiting results for various designs defined by $\mathbf{E}$. Define $X_j \equiv i$

if the jth split was type i, $i = 1, \ldots, K$, and let $I_{ji} = 1$ if $X_j = i$, and $I_{ji} = 0$ otherwise. Then the proportion of type i splits after n splits is $N_i(n)/n \equiv \sum_{j=1}^{n} I_{ji}/n$. Athreya and Karlin (1967, 1968) prove that

$$(1) \qquad \frac{N_i(n)}{n} \to v_i \quad \text{almost surely as } n \to \infty,$$

and

$$(2) \qquad \frac{Y_{ni}}{\sum_{i=1}^{K} Y_{ni}} \to v_i \quad \text{almost surely as } n \to \infty.$$

1.2 Treatment Allocation

If one labels the K particle types with "treatment" identification, then the GPU becomes a randomization machine. Patients are sequentially assigned to a treatment corresponding to the split type, and a rule $\mathbf{R}$ is established for generating new particles based on each possible patient response. The design matrix $\mathbf{E}$ is then the rules averaged over the possible patient responses, and the (i, j) element indicates the expected number of particles added to the urn representing treatment j, given that the previous patient was assigned to treatment i.

Assume that subject response is dichotomous. Let $T_j = 1$ if subject j's response was a "success" (loosely defined), 0 otherwise. We now impose a simple population model which is homogeneous within treatment. Define $p_i \equiv P\{T_j = 1 \mid X_j = i\}$, for $j = 1, \ldots, n$ and $i = 1, \ldots, K$. Define $\mathbf{p} \equiv \{p_1, \ldots, p_K\}$ and let $q_i \equiv 1 - p_i$.

REMARK 1. The T_j need not be dichotomous, as we shall see in Section 2.5.

REMARK 2. The assumption that the probability of success is homogeneous within a treatment group is not always reasonable over the course of recruitment. There may be a drift in patient characteristics over time. Relaxing this assumption is usually difficult in a rigorous mathematical context. Coad (1991, 1992) proposes prestratification and poststratification techniques to deal with this problem.

For a rule $\mathbf{R}$, the corresponding design matrix $\mathbf{E}$ will be a function of $\mathbf{p}$ and the number of particles generated at each stage. If we assume that the row sums of $\mathbf{E}$ are equal (i.e., $\mathbf{E}$ is a scalar multiple ρ of a stochastic matrix), then ρ, the common row sum, is the maximal eigenvalue of $\mathbf{E}$. This assumption is reasonable in our context and will be assumed throughout. Under this assumption, the computation of $\mathbf{v}$ is much simpler.

REMARK 3. One exception to the constant row sum assumption is if one wanted to define a rule which generates more particles for some treatments than others. This may be applicable in the dose–response context, to be discussed in Section 2.4.

Thus $N_i(n)/n$ is the proportion of patients assigned to level i, $i = 1, \ldots, K$, after n stages, and, by (1), we have that $\mathbf{v}$, the left eigenvector corresponding to ρ, is the asymptotic distribution of treatment assignments. Rosenberger and Sriram (1996) show that

$$(3) \qquad \frac{\sum_{j=1}^{n} T_j I_{ji}}{n} \to \mathbf{p} \cdot \mathbf{v} \quad \text{almost surely,}$$

that is, that the total proportion of successes in the trial has an easily calculable limit, a result of considerable interest from an ethical standpoint.

1.3 Likelihood Results

The maximum likelihood (ML) estimator of p_i is $\hat{p}_i = \sum_j T_j I_{ji} / \sum_j I_{ji}$, the observed proportion of successes observed on treatment i. Rosenberger and Sriram (1996) show that $\hat{p}_i$ is strongly consistent for p_i. Rosenberger, Flournoy and Durham (1996) show that the vector comprising elements $n^{1/2}(\hat{p}_i - p_i)$ is jointly asymptotically normal with mean vector $\mathbf{0}$, variances $p_i q_i / v_i$ and covariances 0. By Slutsky's theorem, we can obtain a similar result with a random norming $[N_i(n)]^{1/2}$ replacing $n^{1/2}$, except that the asymptotic variances are $p_i q_i$. Note that this is the same result one would obtain had independent multinomial sampling been employed, except that the normalization factor is random. Asymptotic inference on the p_i's can be done in the same manner as for the independent case, and the usual contrasts and χ^2 statistics apply.

REMARK 4. One would presume that the rate of convergence for the GPU, with its imposed dependencies, would be considerably slower than for independent multinomial sampling. An Edgeworth-type expansion would allow us to see how the design parameters figure in the rate of convergence. This is an area for future research.

2. APPLICATIONS OF THE GPU

2.1 Randomized Play-the-Winner Rule

The randomized play-the-winner (RPW) rule is an adaptive design introduced by Wei and Durham (1978), motivated as an extension to Zelen's (1969) play-the-winner rule. Wei (1979) first noted that the RPW rule could be formulated as a GPU model. Much of the recent literature on adaptive designs

has focused exclusively on appropriate inferential procedures after a design is implemented (see Wei, 1988; Begg, 1990; Wei, Smythe, Lin and Park, 1990; Rosenberger, 1993; Farewell, Viveros and Sprott, 1993) rather than on the design itself.

REMARK 5. One should mention, but only in passing, that the RPW rule was applied rather disastrously in a very small clinical trial in neonates (see Bartlett et al., 1985). This trial is now part of the clinical trials folklore. The reasons for its failure have been well discussed (see, e.g., Royall, 1991) and argue for considerable care when designing an adaptive clinical trial with a small sample size.

The formulation of the RPW rule as a GPU is as follows:

RPW RULE. Assume there are two treatments (say, A and B), and dichotomous response (success or failure). We start with $\mathbf{Y}_0 = (\alpha, \alpha)$ particles in the urn. If a type A split occurs, the patient is assigned to treatment A; if a type B split occurs, the patient is assigned to treatment B. The particle is replaced and patient response is observed. A success on treatment A or a failure on treatment B generates a type A particle; a success on treatment B or a failure on treatment A generates a type B particle.

If A is "doing better," the urn composition is skewed to favor treatment A. Under a simple population model $p_A = P\{\text{success}\,|\,\text{A}\}$, $p_B = P\{\text{success}\,|\,\text{B}\}$, $q_A = 1 - p_A$, $q_B = 1 - p_B$, it is easy to see we have a GPU with

$$\mathbf{E} = \begin{bmatrix} p_A & q_A \\ q_B & p_B \end{bmatrix}.$$

From (1), we can obtain results on the proportion of patients assigned to each treatment as

$$\frac{N_A(n)}{n} \to \frac{q_B}{q_B + q_A} \quad \text{almost surely} \tag{4}$$

and

$$\frac{Y_{nA}}{Y_{nA} + Y_{nB}} \to \frac{q_B}{q_B + q_A} \quad \text{almost surely.} \tag{5}$$

REMARK 6. When $p_A + p_B \leq 3/2$, the joint limiting distribution of the two statistics in (4) and (5), suitably normalized, is Gaussian, and the asymptotic variance–covariance matrix is given in Rosenberger (1992). When $p_A + p_B > 3/2$, the limiting distributions of the urn composition and the proportion of patients assigned to A are unknown (see the Appendix for details).

From (3), we obtain that the total number of successes converges almost surely to

$$\mathbf{p} \cdot \mathbf{v} = \frac{p_A q_B + p_B q_A}{q_A + q_B}. \tag{6}$$

2.2 Example

To illustrate the utility of the results in Section 2.1, we will apply them to the observed data from the maternal–infant HIV transmission data from Section 0.1. Suppose we assume the observed proportions of success (i.e., the newborn is HIV-negative) are the true proportions of success, that is, that $p_A = 0.75$ and $p_B = 0.92$, where A is placebo and B is AZT. Equation (4) tells us that the 50–50 allocation would be skewed to approximately 25–75 in favor of the AZT arm and, by (6), 12% of the infants would have been HIV-positive (compared to approximately 17% in the actual trial).

As stated earlier, standard inference methods could be used on $\widehat{p}_A$ and $\widehat{p}_B$. Because of the imbalance in allocation with the RPW rule, modest loss of power will result, leading to increased sample size requirements (therefore dampening some of the beneficial effects of the adaptive design). For the $3:1$ resulting allocation in this example, standard sample size formulas for the difference of two proportions indicate an 18% increase in requisite sample size for the RPW rule. Even so, there would have been seven fewer treatment failures had the RPW rule been employed.

REMARK 7. The astute reader of Connor et al. (1994) will note that this discussion has oversimplified their study, which involved an early stopping rule and an analysis based on time to transmission. However, this example was simply an attempt to show that an adaptive design can be used to obtain more agreeable allocation proportions and to increase the success rate. A simulation study incorporating the delayed response is presented in Yao and Wei (1996). They conclude that the RPW rule is better than equal allocation for this example, and, in fact, there is minimal loss of power.

REMARK 8. This illustration should not be confused with a blanket endorsement of the RPW rule. It is, perhaps, the most simplistic randomized GPU design for two treatments. However, the rule is quite arbitrary. In addition, early stopping may have been a possibility in this illustration, had an appropriate stopping rule been determined a priori (see Rosenberger and Sriram, 1996, for one such rule). Also, p_A and p_B are not known in advance. A design should

be chosen which protects patients well under a variety of models.

REMARK 9. Care must be taken in choosing the initial urn composition, as a very successful experimental therapy may induce a run that results in little control data—data that may be critical for long-term follow-up and evaluation. In practice, simulation may be an essential tool in visualizing the role of the initial urn composition.

REMARK 10. The RPW rule was recently used in a trial of fluoxetine in depression (Tamura, Faries, Andersen and Heiligenstein, 1994).

2.3 Extension to *K* Treatments

Wei (1979) discusses using the GPU for clinical trials of K treatments. One example of a GPU would be that a success on treatment i generates $K-1$ type i particles, and a failure on treatment i generates one particle for each of the $K-1$ other types. However, it may seem counterintuitive to add particles to an urn of other types when a failure on i gives you no information about efficacy for the other $K-1$ treatments. In fact, if one treatment is doing particularly badly, some might argue that it would be unethical to add particles of that type to the urn as a result of another treatment's failure.

Several other methods seem more reasonable in this scenario. One idea, due to Li (1995), is to generate particles only of the same type if there is a success, and to do nothing if there is a failure. This leads to a diagonal design matrix. This design is no longer a GPU, but it has other nice properties. In fact, the urn composition will converge to a single particle type, representing the best treatment. The rate of convergence is obviously important here and should be explored.

An extension of this design is to add particles representing the same treatment for a success and to remove particles from the urn representing the same treatment for a failure (I am indebted to Professor Stephen Durham for suggesting this design). This leads to a class of diagonal design matrices with potentially negative entries. There is a potential in the latter design for a particle type to die out. That may be ethically desirable. If not, one could impose some outside random immigration process to replenish the urn periodically. Such designs have a lot of potential theoretically, as they can be embedded into continuous-time birth-and-death processes, for which exact and asymptotic results are already established.

Andersen, Faries and Tamura (1994) describe an urn scheme where a success on treatment i generates a type i particle, and a failure generates fractional particles for the other $K-1$ types, allocated in the same proportion as the urn composition at the previous stage. Such a design may provide a more logical and attractive allocation rule than a GPU, but theoretical results would be difficult to obtain because the generations are random functions dependent on all previous splits and generations.

2.4 Dose–Response Studies

Rosenberger, Flournoy and Durham (1996) suggest using the GPU model for a dose–response study. Consider now the sequential random allocation of K dose-levels, $x_1 < \cdots < x_K$, of a single therapy. The goal of the study is to find the optimal dose level, by estimating a quantile μ corresponding to a target percentile Γ of the dose–response curve (see Figure 1). Here, instead of $\mathbf{p}$ being a vector of success probabilities, it represents the vector of toxicity probabilities.

In contrast to the K treatment problem, we can assume a monotonically nondecreasing response function, $p_i = F(x_i)$, which need not be of any particular parametric form, and thus we can establish ethical and logical rules for the generation of new particles. For example, a toxicity at level x_i generates particles at level x_{i-1}; a nontoxicity at level x_i generates particles at level x_i and x_{i+1}, with the specific rule dependent on ethical considerations (as is the choice of Γ). This gives rise to a class of tridiagonal design matrices. In Section 5, we will compare the following specific rule with two other adaptive designs for dose–response studies:

GPU RULE 1. If patient j experiences a toxicity at level x_i add ρ particles at level x_{i-1}; if patient j does not experience a toxicity, r particles are added at level x_{i+1} and $\rho - r$ particles are added at level x_i.

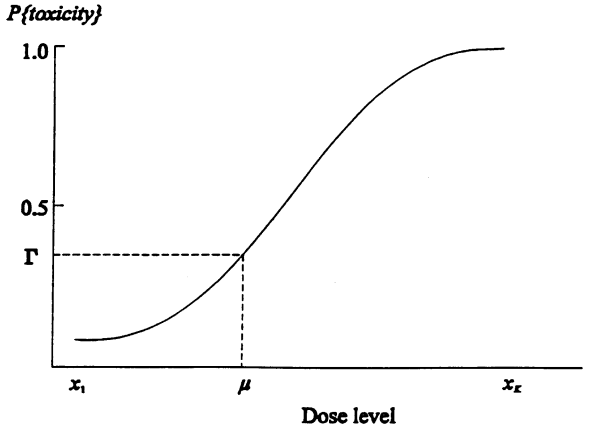

FIG. 1. *Typical dose–response curve.*

$$\mathbf{E} = \begin{bmatrix} \rho - rq_1 & rq_1 & 0 & 0 & \cdots & 0 & 0 \\ \rho p_2 & (\rho - r)q_2 & rq_2 & 0 & \cdots & 0 & 0 \\ 0 & \rho p_3 & (\rho - r)q_3 & rq_3 & \cdots & 0 & 0 \\ 0 & 0 & \rho p_4 & (\rho - r)q_4 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & (\rho - r)q_{K-1} & rq_{K-1} \\ 0 & 0 & 0 & 0 & \cdots & \rho p_K & \rho q_K \end{bmatrix}.$$

Suitable boundary conditions give the design matrix shown above.

REMARK 11. Occasionally in toxicity studies, it will be discovered that $\mathbf{p}$ is not monotonic. There may be a tail-off at the upper end of the dose–response curve. The GPU designs are inappropriate in this setting. Schmoor and Schumacher (1992) discuss adaptive designs for this scenario.

REMARK 12. Prior information on toxicity can be incorporated into the design via the initial urn composition, and $\mathbf{v}$ can be computed under a wide variety of choices of $\mathbf{p}$ and different rules, so that an appropriate design can be established.

Maximum likelihood estimation of μ can be accomplished if we assume that p_i follows a location-scale family, that is,

$$p_i = F\left(\frac{x_i - \alpha}{\beta}\right),$$

for location-scale parameters α and β. The ML estimators $\widehat{\alpha}$ and $\widehat{\beta}$ can be computed using the chain rule, and the equation $\mu = \alpha + \beta F^{-1}(\Gamma)$ can be used to estimate μ. The asymptotic normality of $\widehat{\mu}$ follows from the continuous mapping theorem.

REMARK 13. If F is, for instance, a logistic distribution, then $\widehat{\alpha}$ and $\widehat{\beta}$ are obtainable using a standard logistic regression procedure, and $F^{-1}(\Gamma)$ is simply $\ln\{\Gamma/(1-\Gamma)\}$, the logit. The correct variances can be obtained by computing the Fisher's information matrix directly using the GPU likelihood (see Rosenberger and Grill, 1996).

2.5 Combination Toxicity–Efficacy Studies

Some dose–response studies are not only interested in the toxicity of a dose level of a drug, but also its efficacy in treating the disease. Higher dose levels may be more efficacious, but there may be a threshold at which the drug becomes toxic. Likewise, nontoxic dose levels that are not efficacious are of little use. In this section, we consider a GPU model with trichotomous response. Let p_{i2} be the underlying probability that dose level i is nontoxic, but also not efficacious. Let p_{i1} be the probability that dose level i is not toxic and efficacious (what we want). Let p_{i0} be the probability of toxicity at level i (in which case we are not interested in efficacy considerations). Then appropriate designs can be developed which generate particles at higher levels if there is no toxicity and no efficacy response, at the same level if there is no toxicity and a response and at lower levels if there is toxicity.

As an example, consider the following rules.

GPU RULE 2. If the patient is assigned to level x_i and there is no response and no toxicity, add two particles at level x_{i+1}. If there is response and no toxicity, add one particle at level x_i and one at level x_{i-1}. If there is toxicity, add one particle at level x_{i-1} and one at level x_{i-2}. (Of course, we also need appropriate boundary conditions at $i = 1$ and $i = K$.)

GPU RULE 3. Same as rule 2, except if there is no response and no toxicity, add one particle to level x_{i+1} and one particle to level x_i. This rule is slightly more conservative at the "high end."

If we have five dose levels, rule 2 gives the design matrix shown below.

$$\mathbf{E} = \begin{bmatrix} 2p_{10} + 2p_{11} & 2p_{12} & 0 & 0 & 0 \\ 2p_{20} + p_{21} & p_{21} & 2p_{22} & 0 & 0 \\ p_{30} & p_{30} + p_{31} & p_{31} & 2p_{32} & 0 \\ 0 & p_{40} & p_{40} + p_{41} & p_{41} & 2p_{42} \\ 0 & 0 & p_{50} & p_{50} + p_{51} & p_{51} + 2p_{52} \end{bmatrix}.$$

TABLE 1
Asymptotic proportions of assignment for a reasonable response curve; R = *efficacy response,* NR = *no efficacy response,* T = *toxicity,* NT = *no toxicity*

i	p_{i2}	p_{i1}	p_{i0}	v_i	
Dose level	NR, NT	R, NT	T	Rule 2	Rule 3
1	0.882	0.098	0.020	0.080	0.250
2	0.475	0.475	0.050	0.209	0.351
3	0.414	0.506	0.080	0.258	0.231
4	0.356	0.534	0.110	0.245	0.118
5	0.301	0.559	0.140	0.208	0.050

The design matrix for rule 3 can be similarly derived. Table 1 shows the asymptotic proportions assigned to each dose level, **v**, under rules 2 and 3 when a typical dose-response relationship exists. Table 2 gives **v** when higher dose levels are extremely toxic. GPU rule 3 appears to be better because it protects patients particularly well when there is extreme toxicity at the higher dose levels. For the probabilities in Table 1, (3) gives the expected asymptotic proportion of toxicity as 8.9% for rule 2 and 6.1% for rule 3. In Table 2, the asymptotic proportion of toxicity is 33% for rule 2 and only 22% for rule 3. GPU rule 3 again appears better. The two designs give surprisingly different results, given the slight design variations. This emphasizes the need for careful design selection.

2.6 Preliminary Consideration of Rates of Convergence

The rate of convergence of $N_i(n)/n$ to v_i is of considerable interest in practice. Some very preliminary simulations have shown that, for $K = 2$, the RPW has a very fast rate of convergence. The asymptotic distribution was approximated in some cases for trials as small as 7–12. Similarly, GPU rule 1 in Section 2.4 converged very rapidly. However, tridiagonal designs in Section 2.4 were very slow to converge and were quite unstable when $K = 6$.

TABLE 2
Asymptotic proportions of assignment when high dose levels are very toxic; R = *efficacy response,* NR = *no efficacy response,* T = *toxicity,* NT = *no toxicity*

i	p_{i2}	p_{i1}	p_{i0}	v_i	
Dose level	NR, NT	R, NT	T	Rule 2	Rule 3
1	0.855	0.045	0.100	0.195	0.403
2	0.400	0.400	0.200	0.320	0.366
3	0.610	0.030	0.360	0.216	0.146
4	0.430	0.010	0.560	0.168	0.066
5	0.260	0.000	0.740	0.100	0.020

Some conclusions can be drawn from these preliminary simulations and also some theoretical results on rates of convergence in the GPU (see the Appendix):

1. The rate of convergence is highly dependent upon how "spread out" the mass in the matrix **E** is. In particular, convergence is slower when most of the mass is near the diagonal.
2. The larger K is, the slower convergence is (unless the mass is well distributed throughout **E**).
3. The initial urn composition is critical in the rate of convergence.

Point 3 is obvious, as one would expect more rapid convergence when the initial urn distribution is close to the asymptotic distribution. Points 1 and 2 are more subtle, and although the theoretical justification is available (see the Appendix), one can offer the following heuristic justification. Putting most of the mass of **E** on or close to the diagonal allows generations only around the same levels as the previous splits. So reaching the stationary distribution in the tails may take longer than if generations occurred over a wider spread. There is a potential, when most of the mass is on the diagonal, to get "hung up" at a specific treatment or level for a long period of time. The three points listed above merit further theoretical study and simulation.

3. MORE GENERAL RESPONSES

3.1 Continuous Outcomes

The urn models we have discussed thus far are applicable for binary or polychotomous response. Some clinical trials have continuous outcomes, such as blood pressure, for instance. Rosenberger (1993) develops a biased coin randomization scheme for continuous outcomes based on a linear rank statistic. Assume there are two treatments, say A and B. Define T_n to be 1 if patient n was assigned to A, 0 otherwise. Each patient is randomized with a probability that is a function of the current value of the rank statistic. Let r_{ij} be the rank of the ith patient based on some outcome variable after j outcomes are available ($i < j$), and let a_{ij} be some score function of the r_{ij}'s. Let the scores be centered so that $\sum_{i=1}^{j} a_{ij} = 0$. Then patient n is randomized to treatment A with probability

$$p_n = \frac{1}{2}\left\{1 + \frac{\sum_{i=1}^{n-1} a_{i,n-1}(T_j - \frac{1}{2})}{\sum_{i=1}^{n-1} a_{i,n-1}^{+}}\right\},$$

where the denominator represents the sum of the positive scores.

One can use a permutation test based on the rank scores to test hypotheses of the treatment effect. The form of the statistic is given in Rosenberger (1993), and simulations indicate that it is asymptotically standard normal.

3.2 Survival Trials

Many clinical trials are survival trials with staggered entry and censoring. The adaptive designs discussed so far are inappropriate for this type of trial, where events may occur years after randomization. An idea similar to that in the previous section was proposed, but never explored, by Rosenberger and Lachin (1993). The proposal is randomization according to an adaptive biased coin, the bias being some function of the difference in the survival curves for two competing treatments. The Kaplan–Meier estimate is not appropriate, due to the instability in the tails. One possibility is using the standard Mantel (1966) formulation of the logrank statistic. The idea is similar in principle to that proposed by Flehinger and Louis (1971), who assumed exponential survival, established a stopping rule and studied adaptive allocation schemes (primarily involving the absolute difference in numbers of events between two treatments). Recently, Yao and Wei (1996) proposed a similar technique using the Gehan–Wilcoxon statistic.

Let $\{0 < t_1 < \cdots < t_L\}$ be ordered event times for all patients in the trial (where t_i is the time from randomization to event; i.e., allowing for staggered entry). In the Mantel formulation, at each event time t_i, a 2×2 table is constructed with $\delta_i = 1$ if the event occurred on treatment 1, $\delta_i = 0$ otherwise (in the continuous time model, ties are assumed not to occur). Let n_{ji} be the number of patients at risk on treatment j, $j = 1, 2$, immediately prior to the event time, and let $N_i = n_{1i} + n_{2i}$. Then under the hypergeometric model, the numerator of the logrank statistic computed at time t is given by $\sum_{i=1}^{L(t)} \{\delta_i - n_{1i}/N_i\}$, where, for convenience, we make L a function of t (i.e., the total number of events occurring up to and including the time of computation, t).

Now the idea is to find a mapping of the logrank statistic to a function on $(0, 1)$ which is symmetric (to the extent possible) about $1/2$. Note that when there is no censoring

$$-n_1 \sum_{i=1}^{L(t)} \frac{1}{N-i} \le \sum_{i=1}^{L(t)} \left\{ \delta_i - \frac{n_{1i}}{N_i} \right\} \le n_2 \sum_{i=1}^{L(t)} \frac{1}{N-i},$$

where n_j is the total number randomized to treatment j up to time t and $N = n_1 + n_2$. (This bound is conservative if there is censoring.) Let $\Im_{t^-}$ be the history of the events and censorings to t. Let $Y(t) = 1$ if a patient is randomized to treatment 1 at time t, and 0 if the patient is randomized to treatment 2 at time t. It seems that the most natural mapping for an adaptive biased coin parameter would be

$$p_{t^-} \equiv E\{Y(t) | \Im_{t^-}\} = \frac{1}{2}\left(1 - \frac{\sum_{i=1}^{L(t^-)} \{\delta_i - (n_{1i}/N_i)\}}{\max\{n_1, n_2\} \sum_{i=1}^{L(t^-)} [1/(N-i)]}\right).$$

Properties of the allocation scheme, in particular the proportion assigned to the inferior treatment, should be explored by simulation. Of particular concern with staggered entry is that the allocation rule may put heavy weight on the early events. Note that the logrank statistic is simply the coefficient in the Cox proportional hazard model if there are no covariates (see, e.g., Kalbfleisch and Prentice, 1980). Hence, the treatment effect coefficient in the Cox model can be used in the mapping if there are covariates. This is an important point, because most survival trials have numerous covariates which are being constantly monitored besides end-point status.

REMARK 14. Some clinical trials are long-term survival trials, where outcomes are sometimes not available until years after the accrual period. Obviously, adaptive designs are not applicable for these trials.

4. MORE DESIGNS FOR DOSE–RESPONSE

4.1 Random Walk Rules

Let us return to the dose–response scenario of Section 2.4. While randomization may be desirable to protect from selection bias, and has been promoted by many (see, e.g., Temple, 1981, and Storer, 1989), the GPU allows positive probabilities of large jumps in dose levels between successive patients. Clinicians may be wary about this. An alternative to the GPU, with similar goals, is the random walk rules of Durham and Flournoy (1994). These random walk rules are a variant on the familiar up-and-down rules (see Anderson, McCarthy and Tukey, 1946; Dixon and Mood, 1948; Derman, 1957). Here, if the previous patient was assigned to level x_i, $i = 2, \ldots, K-1$, the next patient will be assigned to level x_{i+1} with probability p_i, to level x_i with probability r_i and to level x_{i-1} with probability q_i, such that $p_i + r_i + q_i = 1$ (with suitable boundary conditions for $i = 1$ and $i = K$). The parameters p_i, r_i and q_i depend on the previous patient's response and some random event, such as the result of a biased coin flip, where the bias is b.

From random walk theory, we have $N_i(n)/n$ converges to a constant π_i almost surely, where π_i is the ith element of the solution to the equation $\boldsymbol{\pi}\mathbf{P} = \boldsymbol{\pi}$, where $\mathbf{P}$ is the random walk transition matrix (see Karlin and Taylor, 1975). Exact results are derived both for the treatment assignment distribution (Durham, Flournoy and Montazer-Haghighi, 1993) and for the total number of toxicities (Flournoy, Durham and Rosenberger, 1995).

One example of a random walk rule is as follows:

RWR RULE 1. If patient j experiences a toxicity at level x_i, assign patient j to level x_{i-1}. If patient j has no toxicity at level x_i, flip a biased coin. If the coin lands heads up, assign patient $j+1$ to level x_{i+1}. If the coin lands heads down, assign patient $j+1$ to level x_i. (Again, we require suitable boundary conditions.)

Obviously, the bias b will depend on ethical requirements. In our case, referring to Figure 1, b will depend on the choice of Γ. Heuristically, in targeting low values of Γ, it would make sense not to go up too quickly. A small value of b will allow for this. In fact, Durham and Flournoy (1994) prove the elegant result that if $b = \Gamma/(1-\Gamma)$, $\Gamma \leq 0.5$, for RWR rule 1, the asymptotic distribution of assignment π_i will be unimodal around the target quantile μ. Similar results are obtainable for any random walk design, assuming certain monotonicity conditions.

The maximum likelihood theory for the GPU applies analogously for random walk rules (Rosenberger, Durham and Flournoy, 1996). Again, there is no requirement that responses be dichotomous, and designs similar to those in Section 2.5 are possible. Finally, there is some work being done now to formulate random walk designs under a delayed-response model (personal communication, Nancy Flournoy).

4.2 Continual Reassessment Method

If targeting a specific quantile is the primary goal of the study rather than estimation, then one may not need to spread the allocations unimodally around the target quantile. Instead, one might want to converge successively closer to the target quantile, and choose the last experimental design point as the optimal dose level. These goals closely relate to the stochastic approximation methods originated by Robbins and Monro (1951) and expanded upon by Wu (1985). While these methods are more relevant to large sample sequential experiments, O'Quigley, Pepe and Fisher (1990) have developed an adaptive Bayesian design called the *continual reassessment method* (CRM), which has performed quite well in small sample dose–response study simulations.

CRM RULE. In the CRM subjects are sequentially entered into the study, and, after each observation, Bayes's formula is used to estimate the dose level x^* associated with the target probability of response Γ. The original paper considers a one-parameter model $F_a(x_i)$, the one-parameter logistic distribution with parameter $a \in A$. (O'Quigley and Chevret, 1991, also explore a two-parameter logistic model where one parameter is known.) An initial prior $g(a)$ is assumed. (Chevret, 1993, explores various priors, although a gamma prior was suggested initially.) In the notation of O'Quigley, Pepe and Fisher (1990), let Ω_j denote the prior history of allocation and responses of the previous $j-1$ subjects, and let $f(a, \Omega_j)$ be a nonnegative function summarizing accumulated information about the parameter a. Associated prior probabilities of toxicity, p_i, are chosen at each prespecified dose level x_i. Experimentation begins at level x_s, $1 \leq s \leq K$, the level at which the chosen prior probability of toxicity is closest to Γ. Given the observed response on subject j, $f(a, \Omega_{j+1})$ is computed from $f(a, \Omega_j)$ using Bayes's formula. Estimates of response probabilities at dose level i, denoted Γ_{ij}, are reevaluated via $\int_A F_a(x_i) f(a, \Omega_j)\, da$. The jth patient is then assigned to the level x_i so that Γ_{ij} is closest to Γ. Absolute or squared error can be used as the loss function.

The recommended dose calculated after n patients is the estimate of x^*. Other estimators are explored by O'Quigley (1992).

A drawback of this design, as compared to the RWR rules already discussed, is that exact results are intractable, and the performance of the design must rely on simulation results. It also relies on a parametric model, unlike the previous nonparametric designs, although the design appears to be robust to other distributions (personal communication, John O'Quigley).

4.3 A Comparison of Three Designs for Dose–Response

It can be shown (thanks to Professor Stephen D. Durham for originally pointing this out) that any random walk rule has a corresponding tridiagonal GPU rule with the same stationary distribution (i.e., $\boldsymbol{\pi}$ for the random walk is the same as $\mathbf{v}$ for the GPU). In fact, if we choose ρ and r from GPU rule 1 so that $\Gamma = r/(r+\rho)$, then, asymptotically, dose levels are allocated unimodally around the unknown

quantile μ corresponding to the value of Γ (using the same argument as given by Durham and Flournoy, 1994, for the random walk rules). This is analogous to setting the bias of the coin to be $\Gamma/(1-\Gamma)$ for RWR rule 1. So we now have three rules which target any unknown quantile of the dose–response curve: GPU rule 1 and RWR rule 1, which asymptotically center the dose level assignments unimodally around μ, and the CRM rule, which sequentially gets closer to μ. It is important to note that, for years, practitioners have used ad hoc adaptive designs with no theoretical analysis of their consequences (see, e.g., Storer, 1989; Flournoy, 1993; and the "standard method" of Korn et al., 1994).

It would be interesting to compare the three methods discussed here for a given value of Γ. Depending on the goals of the study, many dose–response studies (i.e., phase I clinical trials) draw conclusions on the basis of very small numbers of patients. We performed a simple simulation study to compare GPU rule 1, RWR rule 1 and the CRM rule for very small samples. The simulation is similar to that described by Korn et al. (1994).

Details of the simulation. We simulated 2,000 replications of GPU rule 1 and RWR rule 1 and compared results on the CRM rule contained in Korn et al. (1994, Table I). Unfortunately, no measures of variability are given in Korn et al. (1994). In Table 3, we present results for $\Gamma = 0.25$ and six dose levels. For GPU rule 1, we therefore chose $r = 1$ and $\rho = 3$. For RWR rule 1, we set $b = 1/3$. We used the same start-up and stopping rules as did Korn et al. (1994): we start at level 1 and go up successively until there is a toxicity, and then begin implementing the rule; we stop when a patient is assigned to a dose level at which six previous patients have already been assigned. In Table 3, we report the percentages treated at each dose level, the average number of patients treated, the average number of toxicities and the observed proportion of toxicity.

GPU rule 1 appears to do slightly better with regard to requiring fewer patients to be treated and fewer toxicity per trial, on average, but the total toxicity is identical for the three rules. All three have unimodal assignment distributions around the true 25th percentile. We tried simulations to target other quantiles and found similar results. All three rules behave similarly, with the RWR rule generally requiring slightly more patients per trial.

Given that the three designs are comparable, which should be used? This depends on the goal of the experiment. The GPU has the disadvantage of potentially skipping dose levels. The CRM can be restricted (as it was in the simulations of Korn et al., 1994) to assigning only adjacent dose levels, and the RWR rule can only assign to adjacent levels. It is interesting to note that, even though very slow rates of convergence to $\mathbf{v}$ were reported for tridiagonal GPU rules in Section 2.6, the unimodality around the target quantile appears to become evident quite quickly from our simulations. Certainly, the random walk rules are easiest to implement. The CRM has the disadvantage of requiring numerical integration. It should be noted that simulation was not necessary for the random walk rules. For a fixed n, the exact distributional results of the total toxicity and the assignment distribution is worked out. For details, see Durham, Flournoy and Rosenberger (1996). Software is currently being developed which will give the exact distributional results for any random walk design. In terms of quantile estimation, we have discussed only maximum likelihood estimation in the context of the random walk rules and the GPU rules. Other estimators for the random walk rules are explored and compared in Durham, Flournoy and Rosenberger (1996). Finally, if selection bias is an issue, the GPU would be the best choice, as it is completely randomized. Selection bias could be a potential problem in particular with the CRM rule.

TABLE 3
Simulation results comparing GPU rule 1, RWR rule 1 and CRM rule; $\mathbf{p} = (0.05, 0.10, 0.25, 0.35, 0.50, 0.70)$, $\Gamma = 0.25$ (*true* $\mu = 3$), *2,000 replications*

	GPU	RWR	CRM
Percentage treated at each level			
Level 1	15	13	13
Level 2	22	24	22
Level 3	24	27	28
Level 4	21	21	21
Level 5	14	11	12
Level 6	4	4	3
Average number of patients treated per trial	12.5	15.1	13.4
(25th percentile, 75th percentile)	(10, 15)	(12, 18)	(12, 15)
Average number of toxicities per trial	3.3	3.9	3.5
Total toxicities (%)	26	26	26

5. CONCLUSION

We have discussed recent areas of research in adaptive designs for human experimentation, focusing in particular on randomized designs for dose–response studies and phase III clinical trials. A few admonitions are advisable at this point regarding adaptive designs for phase III trials. First, very successful treatments may also be toxic, and if preliminary toxicity studies have not been performed, it

may not be wise to adapt early until toxicity is ruled out. Perhaps an equal allocation strategy could be employed for the first m patients, until sufficient data are accrued; then adaptation can follow for the remaining $n - m$ patients. Second, there is the subtle potential for bias in that patients aware of the principles underlying adaptive designs may want to be randomized late in the trial, to have a better chance of being assigned to the better therapy. Such *accrual bias* may necessitate the blinding of the patient to his or her sequence in the randomization scheme, leading to some ethical concerns.

With respect to dose–response studies or phase I trials, we have focused on adaptive designs which target an unknown quantile of the underlying dose–response curve. There is a large recent literature on designs and methods for phase I clinical trials that would be relevant here if space allowed. The interested reader is referred to Flournoy (1993) (which is based on elegant results of Tsutakawa, 1980), Russek-Cohen and Simon (1994), Gooley, Martin, Fisher and Pettinger (1994), and Gastonis and Greenhouse (1992), to name a few.

Potential applications abound in other areas (see Flournoy and Rosenberger, 1995, a recent IMS monograph). Recent animal rights activism may make these designs timely for more efficient and ethical animal bioassay. Industrial applications have a big potential, because issues of cost-effectiveness and efficiency replace ethical considerations. For example, suppose in a quality control study different stress levels are applied to some items sequentially to find the stress–response relationship. For cost-efficiency, one would like to destroy as few of the items as possible. Finally, the author is currently using the GPU in neurophysiological threshold experiments (see Rosenberger and Grill, 1996), where a "hearing threshold" is determined in a given individual. Randomizing the sequence of hearing stimuli is desirable, while efficiently estimating the threshold-response curve.

ACKNOWLEDGMENTS

Some of the ideas and philosophies presented herein derived from discussions with Professor Stephen D. Durham and Professor Nancy Flournoy. The author thanks Professors Durham and Flournoy for their mentoring and enjoyable collaboration. Mr. Kavosh Soltani programmed the simulations. This research was wholly supported by Grant R29 DK51017-01 from the National Institute of Diabetes and Digestive and Kidney Diseases, which also supported Mr. Soltani's programming.

APPENDIX: RATES OF CONVERGENCE FOR THE GPU

Some preliminary theoretical results on rates of convergence for the GPU are as follows. As before, let ρ be the maximal eigenvalue for $\mathbf{E}$, and let λ be the second largest eigenvalue, with corresponding right eigenvector ξ. By Athreya and Karlin (1968), we have that if $\rho > 2\,\mathrm{Re}\,\lambda$, then

$$n^{-1/2}\xi \cdot \mathbf{Y_n} \to_L \text{Normal}(0, c), \tag{A.1}$$

where c is a constant. If $\rho = 2\,\mathrm{Re}\,\lambda$, then the same result holds with a normalization of $n \ln n$. If $\rho < 2\,\mathrm{Re}\,\lambda$, then

$$n^{-\lambda/\rho}\xi \cdot \mathbf{Y_n} \to W \quad \text{a.s.}, \tag{A.2}$$

where W is an unknown (presumably nonnormal) random variable (determining the distribution of W is an open problem). Note that the rate of convergence in (A.2) is the ratio of the second and first eigenvalues, so their relative magnitude is driving the speed. (It is also suggestive that the limit, being nonnormal, may be less stable.) The relative magnitude of ρ and $2\,\mathrm{Re}\,\lambda$ depends on how much mass the matrix has on or near the diagonal (to see this, perturb any diagonal matrix and notice the effect on the eigenvalues). As discussed in Section 2.6, the more spread out the mass is throughout the matrix, the faster the rate of convergence.

The limit laws in (A.1) and (A.2) are rates on the urn composition vector. How does this relate to the vector of splits $\mathbf{N_n} = \{N_1(n), \ldots, N_K(n)\}$? Smythe (1996) has shown, under certain conditions on the eigenvectors of $\mathbf{E}$, that (A.1) holds when $\mathbf{Y_n}$ is replaced by $\mathbf{N_n}$ (with different asymptotic variance) and that $n^{-1/2}\mathbf{N_n}$ is jointly asymptotically normal when $\rho > 2\,\mathrm{Re}\,\lambda$.

REFERENCES

ANDERSEN, J., FARIES, D. and TAMURA, R. (1994). A randomized play-the-winner design for multi-arm clinical trials. *Comm. Statist. A—Theory Methods* **23** 309–323.

ANDERSON, T. W., MCCARTHY, P. J. and TUKEY, J. W. (1946). "Staircase" method of sensitivity testing. Naval Ordinance Report 65-46, Statistical Research Group, Princeton.

ANSCOMBE, F. (1963). Sequential medical trials. *J. Amer. Statist. Assoc.* **58** 365–384.

ATHREYA, K. B. and KARLIN, S. (1967). Limit theorems for the split times of branching processes. *Journal of Mathematics and Mechanics* **17** 257–277.

ATHREYA, K. B. and KARLIN, S. (1968). Embedding of urn schemes into continuous time Markov branching processes and related limit theorems. *Ann. Math. Statist.* **39** 1801–1817.

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.

BEGG, C. B. (1990). On inferences from Wei's biased coin design for clinical trials, with discussion. *Biometrika* **77** 467–484.

BYAR, D. P., SIMON, R. M., FRIEDEWALD, W. T., SCHLESSELMAN, J. J., DEMETS, D. L., ELLENBERG, J. H., GAIL, M. H. and WARE, J. H. (1976). Randomized clinical trials—perspectives on some recent ideas. *New England Journal of Medicine* **295** 74–80.

CHEVRET, S. (1993). The continual reassessment method in cancer phase I clinical trials: a simulation study. *Statistics in Medicine* **12** 1093–1108.

COAD, D. S. (1991). Sequential tests for an unstable response variable. *Biometrika* **78** 113–121.

COAD, D. S. (1992). A comparative study of some data-dependent allocation rules for Bernoulli data. *J. Statist. Comput. Simulation* **40** 219–231.

COLTON, T. (1963). A model for selecting one of two medical treatments. *J. Amer. Statist. Assoc.* **58** 388–401.

CONNOR, E. M., SPERLING, R. S., GELBER, R., KISELEV, P., SCOTT, G., O'SULLIVAN, M. J., VANDYKE, R., BEY, M., SHEARER, W., JACOBSON, R. L., JIMINEZ, E., O'NEILL, E., BAZIN, B., DELFRAISSY, J., CULNANE, M., COOMBS, R., ELKINS, M., MOYE, J., STRATTON, P. and BALSLEY, J. (1994). Reduction of maternal-infant transmission of human immunodeficiency virus type 1 with zidovudine treatment. *New England Journal of Medicine* **331** 1173–1180. (Report written for the Pediatric AIDS Clinical Trials Group Protocol 076 Study Group.)

CORNFIELD, J., HALPERIN, M. and GREENHOUSE, S. W. (1969). An adaptive procedure for sequential clinical trials. *J. Amer. Statist. Assoc.* **64** 759–770.

DERMAN, C. (1957). Nonparametric up and down experimentation. *Ann. Math. Statist.* **28** 795–798.

DIXON, W. J. and MOOD, A. M. (1948). A method for obtaining and analyzing sensitivity data. *J. Amer. Statist. Assoc.* **43** 109–126.

DURHAM, S. D. and FLOURNOY, N. (1994). Random walks for quantile estimation. In *Statistical Decision Theory and Related Topics V* (S. S. Gupta and J. O. Berger, eds.) 467–476. Springer, New York.

DURHAM, S. D., FLOURNOY, N. and MONTAZER-HAGHIGHI, A. A. (1993). Up-and-down designs. In *Computer Science and Statistics: Interface* (M. E. Tarter and M. D. Lock, eds.) **25** 375–384. Interface Foundation of North America, Berkeley.

DURHAM, S. D., FLOURNOY, N. and ROSENBERGER, W. F. (1996). A random walk rule for phase I clinical trials. *Biometrics*. To appear.

FAREWELL, V. T., VIVEROS, R. and SPROTT, D. A. (1993). Statistical consequences of an adaptive treatment allocation in a clinical trial. *Canad. J. Statist.* **21** 21–27.

FLEHINGER, B. J. and LOUIS, T. A. (1971). Sequential treatment allocation in clinical trials. *Biometrika* **58** 419–426.

FLOURNOY, N. (1993). A clinical experiment in bone marrow transplantation: estimating a percentage point of a quantal response curve. In *Case Studies in Bayesian Statistics* (C. Gastonis, J. S. Hodges, R. E. Kass and N. D. Singpurwalla, eds.) 324–336. Springer, New York.

FLOURNOY, N., DURHAM, S. D. and ROSENBERGER, W. F. (1995). Toxicity in sequential dose–response studies. *Sequential Analysis* **14** 217–227.

FLOURNOY, N. and ROSENBERGER, W. F., eds. (1995). *Adaptive Designs*. IMS, Hayward, CA.

GANTMACHER, F. R. (1959). *Matrix Theory* **2**. Chelsea, New York.

GASTONIS, C. and GREENHOUSE, J. B. (1992). Bayesian methods for phase I clinical trials. *Statistics in Medicine* **11** 1377–1389.

GOOLEY, T. A., MARTIN, P. J., FISHER, L. D. and PETTINGER, M. (1994). Simulation as a design tool for phase I/II clinical trials: an example from bone marrow transplantation. *Controlled Clinical Trials* **15** 450–462.

HARDWICK, J. (1989). Comment: recent progress in clinical trial designs that adapt for ethical purposes. *Statist. Sci.* **4** 327–336.

KALBFLEISCH, J. D. and PRENTICE, R. L. (1980). *The Statistical Analysis of Failure Time Data.* Wiley, New York.

KARLIN, S. and TAYLOR, H. M. (1975). *A First Course in Stochastic Processes.* Academic Press, New York.

KORN, E. L., MIDTHUNE, D., CHEN, T. T., RUBINSTEIN, L. V., CHRISTIAN, M. C. and SIMON, R. M. (1994). A comparison of two phase I trial designs. *Statistics in Medicine* **13** 1799–1806.

LI, W. (1995). Sequential designs for opposing failure functions. Ph.D. dissertation, College of Arts and Sciences, American Univ., Washington.

MANTEL, N. (1966). Evaluation of survival data and two new rank order statistics arising in its consideration. *Cancer Chemotherapy Reports* **50** 163–170.

O'QUIGLEY, J. (1992). Estimating the probability of a toxicity at the recommended dose following a phase I clinical trial in cancer. *Biometrics* **48** 853–862.

O'QUIGLEY, J. and CHEVRET, S. (1991). Methods for dose finding studies in cancer clinical trials: a review and results of a Monte Carlo study. *Statistics in Medicine* **10** 1647–1664.

O'QUIGLEY, J., PEPE, M. and FISHER, L. (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46** 33–48.

ROBBINS, H. (1952). Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **58** 527–535.

ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **29** 400–407.

ROSENBERGER, W. F. (1992). Asymptotic inference problems arising from clinical trials using response-adaptive treatment allocation. Ph.D. dissertation, Graduate School of Arts and Sciences, George Washington Univ.

ROSENBERGER, W. F. (1993). Asymptotic inference with response-adaptive treatment allocation designs. *Ann. Statist.* **21** 2098–2107.

ROSENBERGER, W. F., FLOURNOY, N. and DURHAM, S. D. (1996). Asymptotic normality of maximum likelihood estimators from multiparameter response-driven designs. *J. Statist. Plann. Inference*. **55**.

ROSENBERGER, W. F. and GRILL, S. (1996). A randomized sequential design for threshold experiments. Revised for *Statistics in Medicine*.

ROSENBERGER, W. F. and LACHIN, J. M. (1993). The use of response-adaptive designs in clinical trials. *Controlled Clinical Trials* **14** 471–484.

ROSENBERGER, W. F. and SRIRAM, T. N. (1996). Estimation for an adaptive allocation design. *J. Statist. Plann. Inference*. **55**.

ROYALL, R. M. (1991). Ethics and statistics in randomized clinical trials (with discussion). *Statist. Sci.* **6** 52–62.

RUSSEK-COHEN, E. and SIMON, R. M. (1994). Selecting the best dose when a monotonic dose–response relation exists. *Statistics in Medicine* **13** 87–95.

SCHMOOR, C. and SCHUMACHER, M. (1992). Adaptive statistical procedures for the analysis of nonmonotone dose–response relationships. *Biometrie und Informatik in Medizin und Biologie* **23** 113–126.

SMYTHE, R. T. (1996). Central limit theorems for urn models. *Stochastic Process. Appl.* To appear.

STORER, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45** 925–937.

TAMURA, R. N., FARIES, D. E., ANDERSEN, J. S. and HEILIGENSTEIN, J. H. (1994). A case study of an adaptive clinical trial in the treatment of out-patients with depressive disorder. *J. Amer. Statist. Assoc.* **89** 768–776.

TEMPLE, R. (1981). Government viewpoint of clinical trials. *Drug Information Journal* **16** 10–17.

TSUTAKAWA, R. K. (1980). Selection of dose levels for estimating a percent point of a quantal response curve. *J. Roy. Statist. Soc. Ser. C* **29** 25–33.

WEI, L. J. (1979). The generalized Pólya urn design for sequential medical trials. *Ann. Statist.* **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. *J. Amer. Statist. Assoc.* **73** 840–843.

WEI, L. J., SMYTHE, R. T., LIN, D. Y. and PARK, T. S. (1990). Statistical inference with data-dependent treatment allocation rules. *J. Amer. Statist. Assoc.* **85** 156–162.

WU, C. F. J. (1985). Efficient sequential designs with binary data. *J. Amer. Statist. Assoc.* **80** 974–984.

YAO, Q. and WEI, L. J. (1996). Play the winner for phase II/III clinical trials. *Statistics in Medicine*. To appear.

ZELEN, M. (1969). Play the winner rule and the controlled clinical trial. *J. Amer. Statist. Assoc.* **64** 131–146.

ZELEN, M. and WEI, L. J. (1995). Foreword. In *Adaptive Designs* (N. Flournoy and W. F. Rosenberger, eds.). IMS, Hayward, CA.