In *Biometrics* **59** (2003), pp. 229–236

Directed Walk Designs for Dose Response Problems with Competing Failure Modes

Janis Hardwick[1] **Mary C. Meyer**[2] **Quentin F. Stout**[1]

1. EECS Department, University of Michigan, Ann Arbor, MI 48109, USA
2. Statistics Department, University of Georgia, Athens, GA 30602, USA

SUMMARY: We examine adaptive allocation designs for the problem of determining the optimal therapeutic dose for subjects in early phase clinical trials. A subject can fail due to lack of efficacy or due to a toxic reaction. Successful subjects will have both a positive response and no toxic side effects. Thus, we seek to maximize the product of the non-toxicity and efficacy dose response curves. We are interested in sampling rules that perform well along several criteria, including the ethical criterion that, as often as possible, experimental subjects be treated at or close to the maximum in question. Statistically, we wish to identify the optimum dose with high probability at the close of the experiment. Here, we propose designs that combine new allocation policies, *directed walks*, with new smoothed shape constrained curve fitting techniques. These are compared with a variety of other curve fitting techniques and with up-and-down and equal allocation rules.

KEY WORDS: Bayesian; Convex-concave; CRM; Efficiency; Experimental design; Isotonic; Nonparametric; Phase I/II; Random walk; Sequential; Unimodal.

1 Introduction

Classical dose response problems for phase I clinical trials focus on locating specified quantiles of the relationship between dose and probability of toxic response to drug therapy. In the simplest case, a subject's response at a given dose, s, is modeled as a Bernoulli random variable with a probability of non-toxic outcome $p(s) = 1 - Q(s)$, where $Q(s)$ is assumed to be continuous and non-decreasing in $s \geq 0$. Considerable attention has been paid to such problems, and, in recent years, adaptive formulations of dose response problems have taken center stage. Two-stage designs, stochastic approximation, continual reassessment, dose escalation and up-and-down methods have all been proposed. One reason for this is an increased emphasis on addressing multiple experimental criteria. Nonadaptive or fixed allocation designs tend to focus on optimizing a single criterion such as the variability of an estimator. Potential simultaneous goals such as reducing risk to subjects, overall costs or time to decision have typically not been well incorporated in fixed designs.

When a range of "acceptable" doses has been established during a phase I study, then a phase II clinical trial will often follow. The aim of the phase II trial is to examine the efficacy of the therapy as it relates to adverse outcomes. In a variety of circumstances, it may be desirable to integrate phase I and II trials (known as phase I/II trials). One reason for developing such designs is to accelerate the process of getting a new drug to market. Ideally, the integrated trial will involve fewer subjects and less time. Another goal, addressing ethical concerns, is to allocate more subjects at or near doses that are both safe and efficacious.

An increasing amount of attention is being paid to phase I/II clinical trials with the competing failure modes of toxicity and lack of efficacy. Murtaugh and Fisher (1990) and Jennison and Turnbull (1993) have addressed the problem by assuming parametric bivariate response functions reflecting efficacy and toxicity. Bryant and Day (1995) propose two-stage designs that control error rates for the two failure modes while allowing for possibly correlated outcomes. Conaway and Petroni (1996) elicit tradeoff curves to characterize null hypotheses and develop stopping criteria to limit subject accrual. Kpamegan and Flournoy (2001) describe a modified random walk to locate the dose that optimizes the product of the efficacy and non-toxicity response curves. Tackling the same goal, Hardwick and Stout (2001) use modified multi-armed bandit rules.

Focusing more on curve estimation, Thall and Russell (1998) and O'Quigley, Hughes and Fenton (2001) use designs related to the continual reassessment method of O'Quigley, Pepe, and Fisher (1990). The procedure proposed in O'Quigley et al. (2001) also utilizes sequential probability ratio tests. Gooley et al. (1994), taking a nonparametric approach, and Thall and Russell (1998) specify error rates to define acceptable doses, incorporating elicited information about the expected response functions as a starting point for their designs. In all of these papers, simulation is used as a design aid.

1.1 *The Problem*

We examine a competing failure model in which the goal is to maximize the probability that a subject being treated exhibits both a *non-toxic* response and is *cured*. Throughout, we assume that observations may be taken at one of D fixed dose levels, $0 < s_1 < \cdots < s_D$. The dose levels need not be equidistant, although those used in our examples are.

Let $R(s)$ be a non-decreasing response curve that models the probability that dose s is effective (cures the patient). Then we take $F(s) = R(s)\{1 - Q(s)\}$ to be the probability of a "successful" outcome at dose s. We wish to locate the dose(s), s^*, that maximizes $\{F(s_k) : 1 \leq k \leq D\}$. For convenience, we often refer to dose level s_k as "dose k", $k = 1, \ldots, D$, and to the optimal dose s^* as dose k^*.

Some authors model this problem by defining the efficacy response curve only when there has been a nontoxic result. Call this $R'(s)$, and let $F'(s) = R'(s)\{1 - Q(s)\}$ be the success function for this dependent scenario. (See Thall and Russell, 1998; Kpamegan and Flournoy, 2001). In these cases, accompanying assumptions or model definitions are that R' is non-decreasing and that F' is unimodal. This conditional model is particularly appropriate when toxicities are severe and censor efficacy. It may also be preferable when there are clinical reasons to assume efficacy given non-toxicity is non-decreasing, whereas here the marginal efficacy, R, is assumed non-decreasing.

In such cases, the Directed Walk Algorithm (DWA), introduced in Section 2.2, could be modified so that efficacy information is updated only when a nontoxic result has occurred. This case is not pursued here. Instead, we assume that $R(s)$ and $Q(s)$ are observed for all patients. We also take the two response curves to be independent, although this is somewhat stronger than is needed for our results. A broader approach would be to fully model the dependency structure between the efficacy and toxicity. Two such models are proposed in Murtaugh and Fisher (1990), but they are highly parameterized and thus not appropriate for our purposes here. Further, in these models, the Gumbel and Cox bivariate binary models, it can arise that the success function, F', is not unimodal, which violates an assumption imposed here and by most other authors. Unimodality of the success curve helps assure consistency of estimators. In cases in which F' is unimodal, but F is used in its place, perhaps due to lack of knowledge of the dependency structure, then the DWA still correctly targets the location of the mode as long as $F'(s)$ is a monotonic function of $F(s)$. Also note that simple changes to the DWA exploration rule could be used to guarantee consistency even when success is not a unimodal function, albeit at the expense of reduced efficiency.

As an illustration of the present set-up, consider Figure 1 in which a toxicity function, $Q(s)$, and an efficacy function, $R(s)$ are plotted along with the resulting success curve, $F(s)$. The value $Q^{-1}(.3)$ is the "MD(30)", the dose at which 30% of the subjects are expected to become toxic. This value is often used to define a smaller range of doses for a subsequent efficacy study. Note that, in Figure 1, s^*, the dose that optimizes the success function, is lower than $Q^{-1}(.3)$. This suggests that the "phase I then phase II" trial sequence would involve placing more subjects at higher doses than would a phase I/II study.

In what follows, we work with designs in which, at each stage, the toxicity and efficacy response curves are estimated independently, and estimates of the optimum dose are obtained based on all data accrued to date. We restrict consideration to designs in which the next dose to be sampled is moved *one* dose in the direction of the updated optimum. While related to random walks, these sampling rules, which we call "directed walks", are not constrained by the Markov assumption, and thus they use all available information to determine each allocation. Such walking designs are to be contrasted with "jump to goal" sampling designs in which the next dose to be sampled is the current estimate of the optimum. Clearly there are times when allowing jumps would provide a more efficient design. Although jump to goal designs have been developed in Li, Durham, and Flournoy (1995) and Hardwick and Stout (2001), we restrict to walking rules here for space purposes.

To specify a sampling or allocation design, it is necessary to define a *sampling rule* and a *terminal decision rule*. The sampling rule determines which dose to sample at each stage. In general, sampling rules can include a curve fitting scheme as well as special rules to handle start-up and ending processes, boundaries and premature convergence. The terminal decision rule determines the dose declared best at termination. We assume a fixed, small number, n, of subjects — generally 25 to 50. Furthermore, we assume a subject population that is uniform in prognosis throughout the phase I/II experiment.

To evaluate the sampling designs, we estimate both a *sampling* and a *decision* error. These measures are based on the probability of sampling at or deciding on a sub-optimal dose weighted by the difference between the success probabilities at that dose and the optimal one. See Section 4.

In the next section, we review an up-and-down random walk design proposed in Kpamegan and Flournoy (2001), and outline our directed walk designs. Also discussed in this section are various special considerations, such as starting and ending processes. Following this, in Section 3, we specify a number of curve estimation approaches that are used in conjunction with the directed walks. Some of these are classical parametric techniques, both Bayesian and frequentist, while others are nonparametric. Included are a couple of new smoothed shape constrained curve fitting techniques. In Section 5, the various designs are compared via simulation, and in Section 6 we offer some final comments.

2 Sampling Rules

We first review "up-and-down" designs in which movement is a random walk and observations are obtained only from the success curve, F. Next, we consider a class of two component directed walk designs. In this case, outcomes are observed for both the toxicity and cure response curves, and both are used to update the allocation rule.

2.1 *An Up-and-down Design*

Previous up-and-down designs targeted quantiles of the toxicity probability curve. For the opposing failure mode problem, Kpamegan and Flournoy (2001) propose the following *optimizing up and down* design. At each Stage i, $i = 1, \ldots, \frac{n}{2}$, the rule requires that two consecutive doses be observed; say at s_k and s_{k+1}. Once the responses at Stage i have been observed, allocation for Stage $i + 1$ is as follows:

1. Move up to $(k+1, k+2)$ if there is a success at dose $k + 1$, a failure at k and $k + 1 < D$.
2. Move down to $(k-1, k)$ if there is a success at k, a failure at $k + 1$ and $k > 1$.
3. Otherwise, stay at dose levels $(k, k+1)$.

There are no parametric assumptions utilized in the up-and-down scheme and no form of curve estimation is being used. For these reasons, this design is not expected to perform as well as the others considered. For the evaluations herein we select the dose with the highest observed success rate. Doses not sampled were considered to have a success rate of 0.

2.2 *Two Component Directed Walk Algorithms (DWA)*

As noted, the directed walk approach uses all outcomes observed to date to guide allocation of the next subject. The curves $Q(s)$ and $R(s)$ are estimated at each stage using observations of previous dose assignments, and the next dose selected is based on the estimated location of the optimal dose. Using this algorithm, one can work with parametric or nonparametric curve estimation methods and develop terminal estimators based on maximum likelihood, smoothed MLE, or Bayesian methods. We explore several such options.

Directed walk designs comprise a start-up procedure, a curve fitting and estimation routine, and a set of rules specifying behavior under special circumstances.

Start-up Procedure: Some methods require a certain amount of information to form a valid estimator. This is true for MLE's, for example, although not for Bayesian estimators. Thus, depending on the curve fitting method, the first few observations of the directed walk may follow an up-and-down scheme. Let an observation be a pair (X_e, X_t), where

$$X_e = \begin{cases} 0 & \text{if not efficacious} \\ 1 & \text{if efficacious} \end{cases} \qquad X_t = \begin{cases} 0 & \text{if toxic} \\ 1 & \text{if not toxic} \end{cases} \qquad \text{then when} \qquad (X_e, X_t) = \begin{cases} (1,1) & \text{stay} \\ (0,1) & \text{move up} \\ (1,0) & \text{move down} \\ (0,0) & \text{apply } \textit{exploration rule} \end{cases}$$

The exploration rules are discussed later in this section.

Curve Estimation: Given appropriate observations, we estimate Q and R using one of the curve estimation routines described in Section 3. We determine the dose, $\widehat{k^*}$, that has the highest estimated probability of success and move one step towards it. If we have just sampled that same dose then we utilize an exploration rule which, while usually reallocating to $\widehat{k^*}$, may also indicate a move away from it.

The DWA stops the experiment after n trials and estimates the optimum dose according to the curve estimation scheme employed. If the observations are not sufficient for the curve estimation method, then the dose with the highest observed success rate is chosen.

Starting Dose: The DWA can be started at any dose. Many investigators prefer taking the initial observation at the lowest available dose based on the physician's edict of *first, do no harm*. This perspective weighs toxicity failures more than those due to lack of efficacy. Here we weigh them equally. While one could explicitly incorporate weights in the objective function, to implement this viewpoint consistently one would also need to make adjustments elsewhere in the sampling decisions. Note that there are cases, such as with AIDS drug trials, in which issuing too small a dose is deemed more hazardous than a toxic response.

If one assumes that the mode of F is uniformly distributed among the doses, then starting at the lowest dose is on average $\approx D/2$ doses from the optimum, while starting at random is on average $\approx D/3$ doses away. In this case, the best choice would be to start in the middle, which is only $\approx D/4$ from the optimum. Another reason to prefer a middle start is that an investigator with any prior knowledge is likely to attempt to place the optimal dose near the center of the dose range.

Due to space limitations, we provide results only for the random start scenario. While this avoids the impact of a biased start in our simulations, we would not recommend it in practice, since some prior information is usually available.

Ending Processes: The DWA continues making observations until the sample size is reached. Since the optimizing up-and-down design in Kpamegan and Flournoy (2001) requires an even sample size, if n is odd then on the last stage we randomly choose between the pair of observations that the method indicates.

Note that in much of the literature related to this problem, authors have sought stopping rules that perform well with respect to trial goals, instead of using a fixed sample size. While a fixed sample size is used here to simplify comparisons, note that stopping rules could easily be incorporated into the DWA designs.

Boundary Considerations: We use the convention that whenever a rule indicates sampling at a dose outside the range $\{s_1, \ldots, s_D\}$, we instead sample at the closest endpoint. Note that for the up-and-down design, if we have sampled at (s_1, s_2) the rule stays at (s_1, s_2) even if a downward shift is indicated. For some purposes one might prefer to sample (s_1, s_1).

Exploration Rules: In some cases, sampling rules may get "stuck" allocating repeatedly to a suboptimal dose. To avoid this, we use an *exploration rule* to force occasional, but infinitely often (*i.o.*), sampling at neighbors of the dose in question. As long as the estimators employed are consistent, exploration rules guarantee that the optimal dose will be identified in the limit. Moreover, they are extremely important for problems with small sample sizes to ensure that sampling will eventually occur away from the starting dose region.

The exploration rule used for the results in Section 5 forces a move to a neighbor with probability p_{Ei}, $i = 1, \ldots n$, when 3 consecutive failures have occurred at the given dose, where $\Sigma_{i=1}^{\infty} p_{Ei} = \infty$ and $p_{En} \to 0$ as $n \to \infty$. There are an infinite variety of exploration options available and developing and analyzing them is an open area for research. For example, the rules can be created to ensure more exploratory sampling near the beginning of a trial. Furthermore, more sophisticated rules can be developed to improve or guarantee convergence rates of the sequence of the estimated optimal doses, $\widehat{k_n^*}$ to k^* as $n \to \infty$.

Delayed Responses: In situations where responses are delayed, the present techniques can still be used. Depending on the type of delay and the availability of subjects, a variety of start-up allocation procedures may be appropriate. Very little work has been done in this area, although Hardwick, Oehmke and Stout (2000) present an exact analysis for a two treatment problem with exponential delays. Their results suggest that moderate delays only mildly diminish efficiency. We would expect the same to hold here.

3 Curve Estimation Methods

Here we examine seven approaches to model the toxicity and efficacy response functions and to obtain estimators for the mode of the success function on $\{s_1, \ldots, s_D\}$. These include classical maximum likelihood and Bayes parametric methods as well as unsmoothed and smoothed shape constrained methods. Sampling is such that estimates are updated after every observation. Naturally, this represents an assumption that all subject responses are available prior to selecting a dose for the next subject in the study.

At any point in the experiment, if there have been m_k observations at dose k, resulting in x_k nontoxic responses and $m_k - x_k$ toxic ones, then the likelihood function for toxicity is

$$\mathcal{L} = \prod_{k=1}^{D} \binom{m_k}{x_k} q_k^{x_k} (1 - q_k)^{m_k - x_k}. \tag{1}$$

where $q_k = Q(s_k)$ for $k = 1, \ldots, D$. For the efficacy response curve, we substitute R for Q. Thus, in the following discussion we focus on estimation methods for Q.

Method 1: Two-parameter logistic maximum likelihood estimator

Perhaps the most common assumption for toxicity response functions is that they follow a logistic distribution with two unknown parameters

$$Q(s) = \frac{\exp(a + bs)}{1 + \exp(a + bs)}. \tag{2}$$

If b were allowed to be non-positive, we would occasionally get flat or even decreasing curve estimates, which would result in a choice of an extreme dose. If b were unbounded above, we could get an estimate with sharp corners and flat sections. Therefore we bound the parameter b by choosing a small $\epsilon > 0$, and require that $\epsilon \leq b \leq 1/\epsilon$. This produces smooth, strictly increasing estimates of Q. For the experiments herein, $\epsilon = 1/50$. Because there are no restrictions on a, a necessary and sufficient condition for existence of an MLE is that data be observed at two different doses, with at least one success and one failure.

Method 2: One-parameter Bayes

As a model for the continual reassessment method for locating quantiles, O'Quigley et al. (1990) use the one-parameter function $Q(s) = [\{\tanh(3s - 1.5) + 1\}/2]^a$ in examples. Since this function is increasing convex-concave, and has already been studied, we use it as an example here as well. As in O'Quigley et al. (1990), we assume that the unknown parameter, a, follows an exponential distribution $\pi(a) = e^{-a}$, $a > 0$, and use the mean of the posterior density $\hat{a}$ as an estimate for a to determine the next dose level. While not a true Bayesian design in that we do not determine the posterior expected values at each dose, this simplified semi-Bayesian approach is straightforward to implement. Note also that this method may be used to estimate the curve, regardless of the observations.

Method 3: Two-parameter Bayes

For this method we again use the logistic response function (2) in the likelihood (1), and a pseudo-conjugate (joint) prior for a and b.

$$\pi(a, b) \propto \prod_{k=1}^{D} \frac{\exp\{\gamma \mu_k m_k (a + b s_k)\}}{\{1 + \exp(a + b s_k)\}^{m_k \gamma}},$$

where μ_k is the prior guess for the probability of response at s_k, (see Meyer and Laud, 2000). The parameter γ is a weight given to the prior which represents the amount of "faith" in the prior relative to the data.

For the simulations in Section 5, we take $\mu_k = s_k$ since the dose ranges are conveniently between 0 and 1, and $\gamma = 0.1$, a small weight. The posterior is proportional to

$$\prod_{k=1}^{D} \frac{\exp\{(\gamma \mu_k m_k + x_k)(a + b s_k)\}}{\{1 + \exp(a + b s_k)\}^{m_k (1+\gamma)}}.$$

The mode of the posterior is used as the estimate of the curve. Note that the posterior at stage m is not the prior for the stage $m + 1$. Instead, the prior is redefined at each stage using the current design points.

Method 4: Nonparametric convex-concave shape constrained MLE

A shape constrained MLE maximizes the likelihood subject to shape assumptions on Q. Meyer (1999) gives an efficient algorithm for the non-decreasing convex-concave shape on $(q_1, \ldots, q_D)$.

For the nonparametric convex-concave MLE (both smoothed and unsmoothed), the minimum range for curve estimation is three. The MLE for the curve at the doses observed is determined, but for doses outside the observed range, in general their values are underdetermined. Here, they are given the value of the closest observed dose.

Method 5: Smoothed nonparametric convex-concave shape constrained MLE

One problem with nonparametric MLE's is that they tend to be flat at ends, since most of the information is near the estimated optimal dose. If there are several doses with equal estimated probabilities of survival, Method 4 chooses the smallest of these as "best". This leads to the choice of the lower extreme dose too often, especially with the smaller sample sizes. To get a smoothed nonparametric curve fit, we include a term in the likelihood to penalize flatness. The penalized likelihood is proportional to

$$\prod_{k=1}^{D} \binom{m_k}{x_k} q_k^{x_k} (1 - q_k)^{m_k - x_k} \prod_{k=2}^{D} \left(\frac{q_k - q_{k-1}}{s_k - s_{k-1}} \right)^{\lambda}. \tag{3}$$

Here λ is the smoothing parameter, with $\lambda = 0$ corresponding to the unsmoothed MLE. We choose λ to be small ($\lambda = 0.05$ in the simulations) because fidelity to the data is important, and because only a little smoothing is required for a better result.

Method 6: Nonparametric monotone shape constrained MLE

The unsmoothed monotone or isotonic shape constrained MLE method differs from Method 4 only in the shape assumption. The monotone constraint is weaker, and computationally much simpler. For each dose k such that $m_k > 0$, let $\widehat{q_k} = x_k / m_k$ be the observed success rate. Then the monotone shape constrained MLE is the weighted least squares monotone regression of $\hat{q}_k$, where $\widehat{q_k}$ is weighted by m_k, $k = 1, \ldots, D$.

Method 7: Smoothed nonparametric monotone shape constrained

A semi-Bayesian approach is used to smooth the monotone shape constrained MLE. The toxicity at each dose is given a beta prior. At each stage, a weighted least squares monotone regression is fit to the posterior distributions, using the posterior mean as the value and the sum of the posterior beta parameters as the weight. The prior used in Section 5 was the same for all doses, Be(0.45,0.05). This is a weak prior, so that the data dominates the behavior, and it has a high expected value to initially encourage exploration to doses not yet sampled. In some settings one might have more information and use a stronger prior or perhaps one that is not uniform on the doses.

4 Evaluation Criteria

Although we use various curve fitting techniques to guide the DWA, we do not evaluate designs via measures of closeness to the entire curve. Attempts to optimize in this manner would conflict with the ethical goal of optimizing subject well-being. As mentioned, we seek designs that behave well along two performance measures — a sampling error to assess experimental losses and a decision error to predict future losses based on the terminal decision.

Given any decision rule and sampling design, there exists a probability measure on the doses that reflects the chance, $\xi_n(k)$, that dose k is selected as best at the end of an experiment of size n. While one could take $\xi_n(k^*)$, the *probability of correct selection* or P(CS), as a measure of decision efficacy, this slightly misses the goal of diminishing harm to patients because it does not differentiate between selecting a dose with near-optimal success rate versus one with a meager success rate. Thus, we define *decision efficiency* as: $\mathcal{D}_n = \left\{\Sigma_{k=1}^{D}\xi_n(k)p_k\right\}/p^*$, where $p_k = F(s_k)$ for $k = 1,\ldots,D$ and $p^* = p_{k^*}$.

The sampling error is the normalized expected loss incurred when sampling from doses other than k^*. Letting $\mathbf{E}_n(n_k)$ denote the expected number of observations on dose k in an experiment of size n, we define *sampling efficiency* as: $\mathcal{S}_n = \left\{\Sigma_{k=1}^{D}\mathbf{E}_n(n_k)\cdot p_k\right\}/(n\cdot p^*)$. This is closely related to the well known measure *expected successes lost:* $np^* - \sum_{k=1}^{D} p_k\mathbf{E}(n_k)$.

Note that $p^*\mathcal{S}_n$ is the expected success rate of patients in the experiment, and $p^*\mathcal{D}_n$ is the same rate for subsequent patients if they are given the treatment selected as best.

Large Sample Performance

We say that a design is *efficient* if $\mathcal{D}_n$ and $\mathcal{S}_n$ converge to 1 as $n \to \infty$. For each method, our analysis of efficiency assumes that the model assumptions in the curve estimation for Q and R are met and that F is strictly unimodal.

To establish decision efficiency, two requirements must be met. First, the estimators of $(p_1,\ldots,p_D)$ must be consistent, which they all are. Secondly, the allocation algorithm must sample the best dose, k^*, and its two neighbors *i.o.* For the DWA, this is guaranteed by the exploration rule. For the up-and-down and equal allocation rules all doses are sampled *i.o.*, so no exploration rule is required.

To obtain sampling or experimental efficiency, it is necessary and sufficient that the rate at which dose k^* is sampled goes to 1 as $n \to \infty$. This is the case for the DWA although not for the optimizing up-and-down nor equal allocation rules.

5 Data Models and Performance

The seven different curve estimation methods were evaluated and compared along with the up-and-down rule and an equal allocation scheme, using estimated values of $\mathcal{D}_n$ and $\mathcal{S}_n$ obtained via simulation. Data were generated for every combination of the following: three models of probability curves for toxicity and efficacy, run lengths (n=25 and n=50), and number of dose levels (D=6 and D=12). For each combination, at least 1000 simulated experiments were performed.

In Figure 2 (a), Model 1 is shown. The true probability curves are

$$Q(s) = [\{\tanh(5s - 3.5) + 1\}/2]^4 \text{ and } R(s) = \exp(-2 + 5s)/\{1 + \exp(-2 + 5s)\}.$$

For Model 2, Figure 2 (b), the curves are

$$Q(s) = \exp(-0.5 + 2s)/\{1 + \exp(-0.5 + 2s)\} \text{ and } R(s) = \exp(-1 + 10s)/\{1 + \exp(-1 + 10s)\}.$$

Note that the success probability curve is rather flat compared to that of Model 1.

Model 3 is nonparametric. The success curve has approximately the same shape as that of Model 2, but both the toxicity and efficacy curves stay away from 0 and 1, making the individual curves harder to estimate. Figure 2 (c) illustrates Model 3 with D=12.

5.1 *Results*

The tradeoffs between the two efficiency measures according to method are illustrated in Figure 3. For each model the estimates of $\mathcal{D}_n$ are plotted against those for $\mathcal{S}_n$ when $n = 50$ and D=6. Methods falling in the top-right corner of the plots are the best, and points that fall together in groups should be considered roughly equal.

Numbers represent the methods as follows: (**1**) 2–param MLE; (**2**) 1–param Bayes; (**3**) 2–param Bayes; (**4**) Convex-concave; (**5**) Smoothed convex-concave; (**6**) Monotone; (**7**) Smoothed monotone; (**8**) Up-and-down; and (**9**) Equal allocation (EA). EA is included to represent fixed allocation techniques.

In Model 1, the underlying curves are the tanh and logistic and, as expected, the parametric models (1–3) performed extremely well. Somewhat surprisingly, however, so did the smoothed shape constrained methods (5,7), as did even the unsmoothed monotone method (6). The unsmoothed convex-concave (4) performed poorly.

Model 2 results are similar to those of Model 1, with an exception being the poor performance of the 1–parameter Bayes model (2). Both underlying curves in Model 2 are logistic and thus match the models assumed in the 2–parameter MLE (1) and Bayes (3) methods. While these methods perform well, however, they do no better than the smoothed shape constrained methods (5,7). Again, method (6) is close to the leaders.

The excellent performance of the 1–parameter Bayes method (2) is a bit unanticipated in Model 3. Recall that Model 3 was included to test the robustness of the parametric methods. Note also the poor performance of the 2–parameter MLE (1).

Since the up-and-down method (8) involves no curve fitting, the results for all three models are quite poor. As expected, however, the decision efficiency of the up-and-down rule is much better than the sampling efficiency. This also holds for the fixed allocation, which appears in the plots only for Model 2. The rest of the points are out of range.

The fourth plot in Figure 3 contains efficiencies averaged over the three models with D=6 and n=50. In this plot the 2–parameter Bayes and smoothed monotone methods appear to be preferable. In the equivalent plot (not shown) when D=12, the smoothed shape constrained methods outperform all the others, with the monotone version being best.

Generally speaking the efficiency data for cases in which D=12 and/or n=25 are quite similar to the results presented here and have been omitted due to space constraints.

6 Discussion

In this paper, we have focused on two aspects of the design of phase I/II clinical trials. The first relates to rules for movement of an adaptive sampling design on a discrete dose set. Along with the basic walk, there are a number of factors such as starting sequence, endpoint conventions and exploration rules, that play an important role in design performance. Ignoring these factors can lead to highly flawed designs.

The second emphasis in this research has been to examine curve fitting methods and assumptions. We took seven curve fitting techniques and two methods with no fitting at all and compared them according to their sampling and decision efficiency. It is difficult to draw general conclusions based on the close examination of only three sets of curves, although these results strongly hint that parametric models have powerful competitors in smoothed shape constrained methods.

With the parametric designs, two of the three cases had underlying structure that matched the curve schemes. However, these methods appear very sensitive to structure, whereas the simple assumption of monotonicity seems effective for estimating a variety of response curves. Needless to say, a more variable set of curves, perhaps even a space of curves, should be examined.

Because this is a design paper, we have not concentrated much on the analysis phase of the trials. It should be pointed out, however, that the designs can be analyzed from a variety of perspectives. Many of the methods described here use Bayesian concepts to some extent. This is natural when sampling adaptively, since updating prior information is an inherent part of the methodology. Nevertheless, we support the idea, used, for example, in Thall and Russell (1998), of evaluating designs for frequentist characteristics even if they have resulted from a Bayesian approach. The main difficulty in the analysis phase lies in gaining an understanding of the distribution of the estimators since these are affected by the sequential design. Simulation studies and exact computational methods can assist in this analysis.

Next, there are a number of useful extensions of this work. For example, the *type* of failure observed (toxicity or lack of efficacy) affects the DWA only during the starting sequence. It is possible that continuing to use this information after the estimation step kicks in would improve the algorithm. Other enhancements involve step size. It is likely that the greatest improvement will be garnered when variation of step size is allowed. Large steps at the beginning are needed for exploratory purposes, but smaller steps towards the end will typically improve both sampling and decision efficiency. One version of this process arises when the algorithm moves directly to the best estimate of the mode (Hardwick and Stout, 2000; and Li et al., 1995). Another version would be to define a sequence that reduces step size at a given rate. Note also that one can easily extend the present results to include polychotomous responses.

In conclusion, while it was expected that the parameterized designs would outperform all others when the underlying models matched those of the sampling model, this is apparently not the case. In particular, the smoothed shape constrained methods performed roughly as well as the parametric techniques for each model, while requiring fewer assumptions.

Acknowledgments

Research supported in part by National Science Foundation grants DMS–9504980 and DMS–0072910. Research supported in part by University of Georgia Faculty Research Grant. We thank the referees and the associate editor for their help in improving this manuscript.

References

[1] Bryant J. and Day R. (1995). Incorporating toxicity considerations into the design of two-stage phase II clinical trials. *Biometrics* **51**, 1372–1383.

[2] Conaway, M.R. and Petroni, G.R. (1996). Designs for phase II trials allowing for a trade-off between response and toxicity. *Biometrics* **52**, 1375–1386.

[3] 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.

[4] Hardwick, J. and Stout, Q.F. (2001). Optimizing a unimodal response function for binary variables. In *Optimum Design 2000,* A. Atkinson, B. Bogacka and A. Zhigljavsky (eds), 195–210. Kluwer Academic Publishers.

[5] Hardwick, J., Ohemke, R. and Stout, Q.F. (2001). Optimal adaptive designs for delayed response models: exponential case. In *MODA 6: Advances in Model Oriented Design and Analysis,* A. Atkinson, P. Hackl, and W. Müeller (eds), 127–134. Physica Verlag.

[6] Jennison, C. and Turnbull, B. (1993). Group sequential tests for bivariate response — interim analyses of clinical-trials with both efficacy and safety end-points. *Biometrics* **49**, 741–752.

[7] Kpamegan, E.E. and Flournoy, N. (2001). An optimizing up-and-down design. In *Optimum Design 2000,* A. Atkinson, B. Bogacka and A. Zhigljavsky (eds), 211–224. Kluwer Academic Publishers.

[8] Li, Y., Durham, S., and Flournoy, N. (1995). An adaptive design for maximization of a contingent binary response. *Adaptive Designs: Institute of Mathematical Statistics Lecture Notes* **25**, B. Rosenberger and N. Flournoy (eds), 179–196.

[9] Meyer, M. (1999). A comparison of nonparametric shape constrained bioassay estimators. *Statistics and Probability Letters* **42**, 267–274.

[10] Meyer, M. and Laud, P. (2002). Predictive variable selection in generalized linear models. To appear in the *Journal of the American Statistical Association* **97**.

[11] Murtaugh, P.A. and Fisher, L.D. (1990). Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics — Theory and Methods* **19**, 2003–2020.

[12] O'Quigley, J., Hughes, M., and Fenton, T. (2001). Dose-finding designs for HIV studies. *Biometrics* **57**, 1018–1029.

[13] 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.

[14] Thall, P.F. and Russell, K.E. (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54**, 251–264.

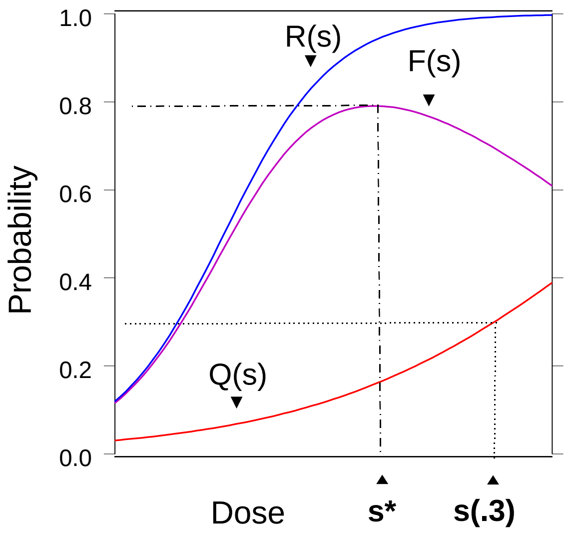

Figure 1: Example of non-decreasing toxicity (Q) and efficacy (R) response curves and the unimodal *success* curve (F). Here, s^* represents the mode of F, while s(.3) corresponds to the dose at which 30% of the subjects are expected to experience toxicity.

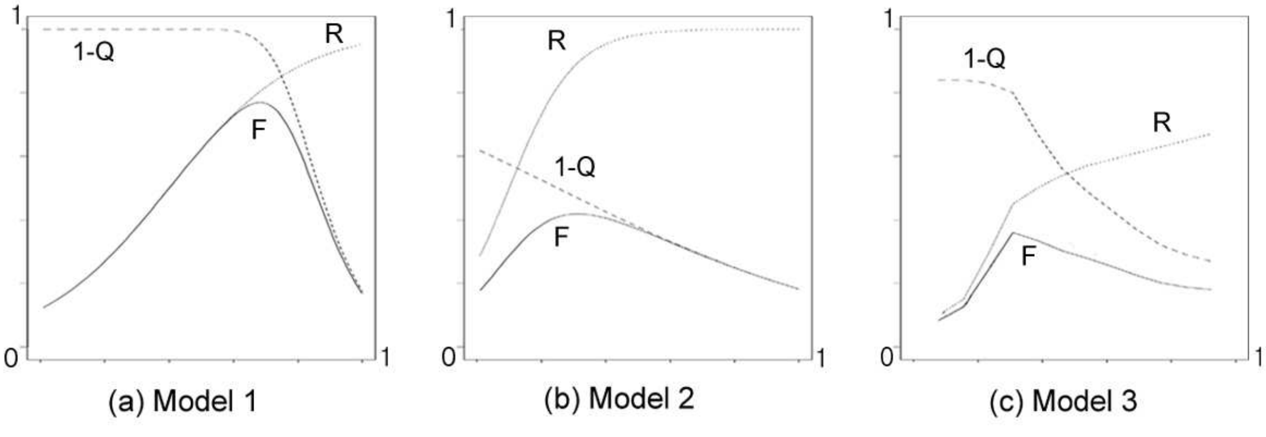

Figure 2: Model 1: $Q(s) = [\{\tanh(5s - 3.5) + 1\}/2]^4$ and $R(s) = \exp(-2 + 5s)/\{1 + \exp(-2 + 5s)\}$. Model 2: $Q(s) = \exp(-0.5{+}2s)/\{1{+}\exp(-0.5{+}2s)\}$ and $R(s) = \exp(-1{+}10s)/\{1{+}\exp(-1{+}10s)\}$. Model 3: Arbitrary curves

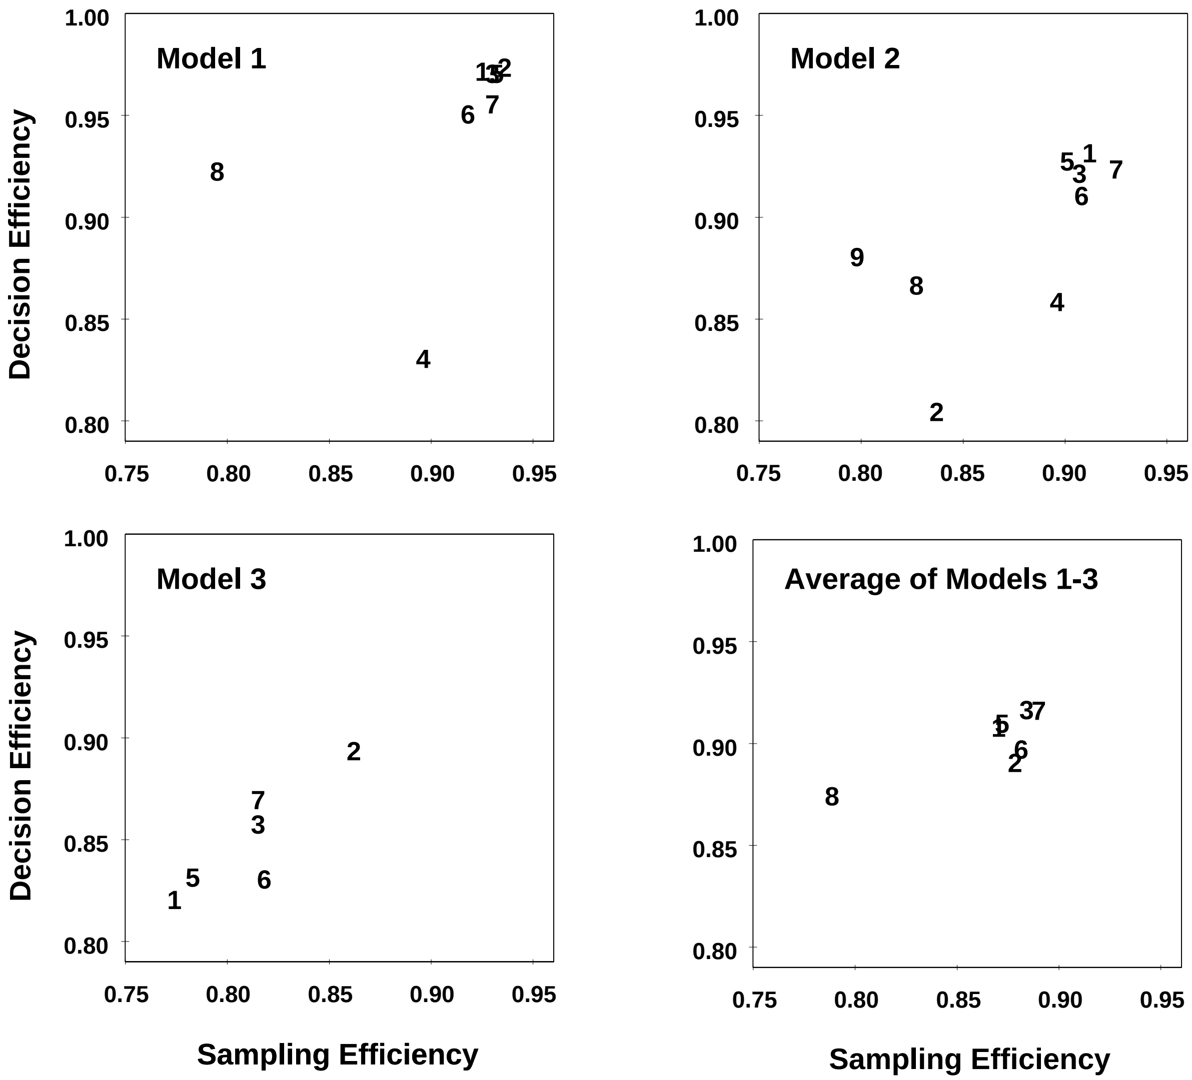

Figure 3: Efficiency Tradeoff Curves: Sampling versus Decision efficiency by Method and Model. $n = 50$ and $D = 6$. Methods that are missing from plots fall below the plot ranges. Methods are: **1.** 2–par MLE; **2.** 1–par Bayes; **3.** 2–par Bayes; **4.** Convex-concave; **5.** Smooth conv-conc; **6.** Monotone; **7.** Smooth monotone; **8.** Up-down; **9.** EA