

Screening Designs for Drug Development

DAVID ROSSELL^{1,2}, PETER MÜLLER¹, GARY L. ROSNER¹

*¹Department of Biostatistics & Applied Mathematics,
The University of Texas, M. D. Anderson Cancer Center
Houston, TX 77030, U.S.A*

²Department of Statistics, Rice University, Houston, TX 77005

rusi@stat.rice.edu

SUMMARY

We propose drug screening designs based on a Bayesian decision-theoretic approach. The discussion is motivated by screening designs for phase II studies. The proposed screening designs allow consideration of multiple treatments simultaneously. In each period new treatments can arise and currently considered treatments can be dropped. Once a treatment is removed from the phase II screening trial, a terminal decision is made about abandoning the treatment or recommending it for a future confirmatory phase III study. The decision about dropping treatments from the active set is a sequential stopping decision. We propose a solution based on decision boundaries in the space of marginal posterior moments for the unknown parameter of interest that relates to each treatment. We present a Monte Carlo simulation algorithm to implement the proposed approach.

We provide an implementation of the proposed method as an easy to use R library available for public domain download (<http://www.stat.rice.edu/~rusi/> or <http://odin.mdacc.tmc.edu/~pm/>).

1 INTRODUCTION

We develop a Bayesian decision-theoretic approach to screening designs for drug development. The proposed process is appropriate for a sequence of phase II studies targeting the same disease area, carried out at the same institution, competing for the same pool of potentially eligible patients, and subject to common resource constraints. For example, at large institutions dedicated to clinical research in cancer, such as the University of Texas M. D. Anderson Cancer Center, a large number of new agents or new combinations of anticancer agents undergo evaluation for activity. The process is typically carried out through separate phase II studies with only informal learning between studies – even if the studies draw patients with similar disease characteristics. We develop an approach that considers such a sequence of studies as one large encompassing screening design and borrows information between studies. An easy to use implementation as an R library allows interested readers to implement the proposed algorithm with minimal effort.

Most screening designs for culling active therapies from the many new agents that are in development consider each study in isolation, even though investigators recognize the need for reproducibility of results (Simon, 1987). After several similar phase II studies have appeared, one is left to combine the information informally and arrive at a decision whether to move ahead with the treatment or not. The question of how many repeat studies to complete is also left informal. In particular, one intuitively would think that the number of replicate phase II studies might depend on the strength of evidence already available concerning the activity of the new agent. Currently, however, decision

making does not incorporate such quantitative information in a formal way.

Yao et al. (1996) proposed a formal way to screen multiple agents for activity in a series of phase II vaccine trials. For each treatment being considered, a single-arm clinical study is carried out. The decision concerns choosing the sample size for each phase II study and a threshold to minimize the overall expected sample size (or time) needed until an active agent is identified. The decision problem is discussed in the frequentist paradigm, in which the type I and type II error probabilities are prespecified and preserved over the sequence of experiments. The formal setup in Yao et al. (1996) considers one treatment at a time and assumes independent binary outcomes. In later work, Yao & Venkatraman (1998), Wang & Leung (1998) and Leung & Wang (2001) consider a variety of extensions leading to 2-stage designs and fully sequential designs in the same setup. Strauss & Simon (1995) consider a generalization based on two-armed randomized trials for each new treatment. One arm is the new treatment, and the other arm is the best treatment found so far. At the end of this sequence of randomized studies, one chooses the “winner” that will be compared to a standard regimen in a randomized comparative trial. Stout & Hardwick (2005) discuss the above mentioned approaches as special cases of a more general setup.

In this paper, we build on these methods to develop a sequential decision-theoretic design for drug screening. We introduce two important directions of generalization. First, we allow for multiple treatments to be considered at any given time. New treatments can arise and existing ones can be dropped at any time if the current evidence suggests that it is optimal to abandon further development of them, or that it is optimal to move them to phase III. Second,

we cast drug screening as a decision problem. Using a simulation-based solution allows us to consider essentially arbitrarily complex utility functions and probability models. Also, the proposed approach includes the possibility to restrict the action space, e.g. by considering only designs with certain type-I and type-II error probabilities.

We propose a probability model that allows borrowing information between treatments, which is appropriate when treatments target the same disease and are likely to be based on similar mechanisms. We consider a utility function that includes terms related to sampling cost and a final payoff that is realized if the future phase III trial shows a statistically significant improvement over the standard of care. The use of a utility function that is based on the sampling cost and the final payoff means that we focus on the perspective of the drug developer or the investigator who is carrying out the trial. We propose to accommodate the interests of regulators and patients by restricting consideration to rules that satisfy constraints on type I and type II error probabilities. For comparison, we also consider a utility of the form proposed by Yao et al. (1996) who seek to minimize the number of patients before the first treatment is recommended for phase III. The decision criterion for the screening trial is the expected utility, appropriately marginalizing with respect to the unknown true success probability and the future outcomes in the phase III study. For an extensive discussion of utility functions for clinical trials see Gittins & Pezeshk (2002).

In Section 2, we formally state the drug screening process as a decision problem by defining a probability model, an action space, and a utility function that serves as the decision criterion. In Section 3, we discuss a simulation-based

approach for solving the decision problem. In Section 4 we show results for a simulated example. In Section 5 we assess the uncertainty and robustness of these results. In Section 6 we compare our approach with that of Yao & Venkatraman (1998) in a clinical immunology problem. Finally, we conclude in Section 7 with a final discussion of features and limitations of the proposed approach.

2 DRUG SCREENING

Our approach is based on casting the screening process as a formal decision problem. The basic ingredients of a decision-theoretic setup are an action space \mathcal{A} of possible decisions $d \in \mathcal{A}$, a probability model $p(\theta, y)$ for all relevant random variables, including parameters θ and future data y , and a utility function $u(d, \theta, y)$. The probability model is conveniently factored into a prior probability model $p(\theta)$ and a sampling model $p(y | \theta)$. It can be argued (DeGroot, 2004) that a rational decision maker should choose an action in \mathcal{A} to maximize the expectation of u . The expectation is with respect to p , conditioning on all data observed at the time of decision making, and marginalizing over all parameters and all future data. Sometimes the action space is restricted to decisions that satisfy certain constraints, for example prespecified bounds on type-I and type-II errors (false positive and false negative rates). In such cases the maximization is carried out over the restricted set.

2.1 Action Space and Probability Model

Let y_{ti} be the outcome at time $t = 1 \dots T$ for treatment $i \in A_t$, where A_t is the set of treatments being considered at time t . We assume a finite time

horizon T for the entire screening process and we allow for a random number of treatments at any given time t .

After observing the outcomes y_{ti} , $i \in A_t$, we make a sequential stopping decision d_{ti} for each treatment. We denote with $d_{ti} = 0$ the action of removing treatment i from A_t and with $d_{ti} = 1$ the action of continuing recruitment for treatment i . If we decide $d_{ti} = 0$, then a terminal second step decision a_i indicates whether to abandon treatment i ($a_i = 0$), or whether to recommend to proceed with a confirmatory phase III study ($a_i = 1$).

Finally, before the next decision at time $t + 1$, new treatments might be proposed and added to the set A_{t+1} . Let Δn_t denote the number of new treatments arising in period t and denote with $\pi_j = Pr(\Delta n_t = j)$ for $j = 0, 1, \dots$ its probability distribution. In the last period, T , continuation is not possible. That is, $d_{Ti} = 0$ for all $i \in A_T$. Figure 1 illustrates the sequence of decisions and observations.

The formal definition of the decision problem requires a probability model for all involved random variables. We assume binomial sampling. We make this assumption mainly for ease of exposition. With minor modifications the proposed approach can be adapted to other sampling models. Thus, without major loss of generality we assume

$$y_{ti} \sim Bin(N_{ti}, \theta_i), \quad i \in A_t, \tag{1}$$

with known N_{ti} . In particular, accrual rates can vary across treatments. The unknown success probabilities arise from a common prior distribution, possibly involving a regression on treatment-specific covariates. We use a Beta prior, $\theta_i \sim Be(u, v)$, with random hyperparameters (u, v) that allow borrowing of in-

formation between treatments. As prior distribution on these hyperparameters we assume Gamma distributions, subject to a bound on $u + v$,

$$(u, v) \sim \text{Gamma}(A_u, B_u) \cdot \text{Gamma}(A_v, B_v) \cdot I(u + v \leq 10). \quad (2)$$

The restriction limits the extent of borrowing of strength across treatments. That is, no matter how many treatments and patients we have observed, the data will never provide more information about a new treatment than the equivalent of 10 patients. The choice of 10 is arbitrary. Any alternative bound, or no bound, could be used without any change in the following discussion. In the context of phase II trials with typically small sample sizes we consider 10 to be a reasonable choice.

Finally, we include a bound N^* for the number of eligible patients that can be recruited for enrollment at time t . Setting $N^* = \infty$ defines the problem without recruitment limits. We assume without loss of generality that N^* remains the same across t . When a new treatment arises and no patients are available, data collection for the new treatment has to wait until one of the existing treatments is dropped. We do not consider adaptive allocation to treatments.

2.2 Utility Function

Let n_T be the overall number of treatments considered in the screening process and let $d \equiv (d_{ti}, t = 1, 2, 3, \dots, T; i \in A_t)$ and $a \equiv (a_i, i = 1, \dots, n_T)$ denote the sequence of decisions. Recall that d_{it} denotes the stopping decision and a_i denotes the terminal decision upon stopping enrollment for treatment i . Let $y = (y_{ti}, t = 1, 2, \dots, T, i \in A_t)$, and let $\theta = (\theta_i; i = 1, \dots, n_T)$ denote the

parameters of the sampling model for y .

The utility function $u(d, a, \theta, y)$ formalizes preferences across possible outcomes corresponding to assumed responses y , parameters θ , and decisions (d, a) , i.e. it reports the value of a hypothetical realization (y, θ, a, d) of the entire trial. An important advantage of the proposed simulation-based solution is that we are free to specify a utility function that reflects the scientific problem, without constraints to convenient analytic properties.

In our implementation, we use a utility function that includes sampling cost plus a payoff for every treatment that is recommended for phase III and is approved at the end of a future confirmatory phase III study that compares the experimental therapy versus the standard of care. The payoff is weighted by the size of the advantage over the standard of care. Regulatory approval is formalized as a statistically significant treatment effect at the conclusion of the confirmatory trial. We build a utility function for the entire process in steps, leading eventually to the utility function stated in (3).

First, suppose that for treatment i we start recruitment at time t_{0i} and we stop recruitment at time t_{1i} , i.e., $d_{ti} = 0$ at time $t = t_{1i}$. If the treatment is abandoned ($a_i = 0$), then we only record a linear sampling cost $c_1 \cdot \sum_{t=t_{0i}}^{t_{1i}} N_{ti} = c_1 \cdot N_{.i}$. Here $N_{.i}$ is the total number of patients assigned to treatment i .

If we proceed with a phase III trial ($a_i = 1$), then we record the sampling cost $c_1 n_3$, where n_3 is the sample size of the future study, and we add a payoff for a significant phase III result, weighted by the estimated size of the advantage over the standard of care. Let θ_0 denote the success probability for the standard of care. Let $\hat{\theta}_i$ and $\hat{\theta}_0$ denote the maximum likelihood estimates for θ_i and θ_0 at the end of the phase III trial, and let B denote the event of

observing a significant result. Let c_2 denote the reward for recommending a treatment that shows a significant treatment effect in the confirmatory trial, i.e., the reward for a successful drug development. The reward is scaled by the estimated size of the advantage over placebo and the probability of B . We record $c_2 Pr(B | y_1, \dots, y_t) E(\hat{\theta}_i - \hat{\theta}_0 | B, y_1, \dots, y_t)$. Putting everything together, we have

$$u(d, a, \theta, y) = \sum_{i=1}^{n_T} -c_1 \cdot N_i + \sum_{i: a_i=1} \left[-c_1 \cdot n_3 + c_2 Pr(B | y_1 \dots y_{t_{1i}}) E(\hat{\theta}_i - \hat{\theta}_0 | B, y_1, \dots, y_{t_{1i}}) \right] \quad (3)$$

We now discuss the evaluation of n_3 , $Pr(B | y_1 \dots y_{t_{1i}})$ and $E(\hat{\theta}_i - \hat{\theta}_0 | B, y_1, \dots, y_{t_{1i}})$.

Let (m_{ti}, s_{ti}) denote the posterior mean and standard deviation for θ_i at time t , and let $(m_{.i}, s_{.i})$ denote their value at time t_{1i} . The phase III sample size n_3 is chosen for a test comparing $H_0, H_0 : \theta_i = \theta_0$, versus an alternative $H_1, H_1 : \theta_i = m_{.i}$, for a given significance level α_3 and power $1 - \beta_3$. Let $\bar{m} = (m_{ti} + \theta_0)/2$, and let z_p denote the $(1 - p)$ standard normal quantile. We assume that the final test is carried out as a z-test to compare two binomial proportions. Assuming known θ_0 , we approximate the phase III sample size as

$$n_3(m_{.i}, s_{.i}) = 2 \left(\frac{z_{\beta_3} \sqrt{m_{.i}(1 - m_{.i}) + \theta_0(1 - \theta_0)} + z_{\alpha_3/2} \sqrt{2\bar{m}(1 - \bar{m})}}{m_{.i} - \theta_0} \right)^2.$$

Next, we evaluate the posterior predictive probability $p(B | y_1 \dots y_{t_{1i}})$. The event B is defined by the z-statistic falling in the rejection region in favor of the experimental arm. Thus

$$P(B | y_1 \dots y_{t_{1i}}) = P \left(\hat{\theta}_i - \hat{\theta}_0 > z_{\alpha_3/2} \sqrt{2\bar{m}(1 - \bar{m})/n_3} | y_1 \dots y_{t_{1i}} \right).$$

Using a normal approximation to the posterior predictive distribution, $p(\hat{\theta}_i - \hat{\theta}_0 \mid y_1 \dots y_{t_{1i}})$ we can approximate $p(B \mid y_1 \dots y_{t_{1i}})$. Denote by μ_Δ and σ_Δ^2 the moments of this normal approximation,

$$\mu_\Delta = m_{.i} - \theta_0 \text{ and } \sigma_\Delta^2 = \frac{1}{n_3} m_{.i} (1 - m_{.i}) + s_{.i}^2 (1 - 1/n_3) + \frac{1}{n_3} \theta_0 (1 - \theta_0).$$

Finally, we evaluate $E(\hat{\theta}_i - \hat{\theta}_0 \mid B, y_1 \dots y_{t_{1i}})$, the posterior predictive expectation for the size of the advantage over standard of care, conditional on B . This conditional expectation is evaluated as the expected value of a normal random variable left-truncated at $k = z_{\alpha_3/2} \sqrt{2\bar{m}(1 - \bar{m})/n_3}$ (Jawitz, 2004).

2.3 Decision Boundaries

The described action set, probability model and utility function formally define the decision problem. We now proceed to find the optimal solution by maximizing the utility $u(d, a, \theta, y)$ as a function of the decisions, marginalizing with respect to θ and all future data that are unknown at the time of a decision and conditioning on all available data.

We first discuss the terminal decision a_i , the indicator for recommending a phase III trial. The terminal decision is carried out at time $t = t_{1i}$. From (3) we find that $a_i = 1$ is optimal if and only if

$$c_1 \cdot n_3 < c_2 Pr(B \mid y_1 \dots y_{t_{1i}}) E(\hat{\theta}_i - \hat{\theta}_0 \mid B, y_1 \dots y_{t_{1i}}) \quad (4)$$

If $m_{.i} < \theta_0$ it is not possible to achieve the desired power in the phase III trial, and we set $a_i = 0$. This solves the choice for the terminal decision a_i , once we have decided to stop enrollment in treatment i .

The continuation decision d_{ti} , is complicated by its sequential nature. To find the optimal solution at time t , we need to compare expected utilities

under $d_{ti} = 0$ and $d_{ti} = 1$. To find the expected utility under continuation, $d_{ti} = 1$, we need to know the decision for $t + 1$, etc. A full solution involves the use of backward induction. But the computational cost of backward induction makes a full solution infeasible even in fairly simple situations. DeGroot (2004), Brockwell & Kadane (2003) and Berry et al. (2001) discuss alternative, computationally intensive approaches that allow one to approximate full backward induction. Many Bayesian clinical trial designs avoid the difficult problem of optimal sequential decisions by stopping short of a formal decision-theoretic approach. Instead, many methods include a combination of posterior inference for the probability model with reasonable, but ad-hoc rules for the desired decisions. A typical example is the approach proposed in Thall et al. (1995). The method proceeds by evaluating posterior probabilities of clinically meaningful events. When these probabilities cross predefined boundaries certain decisions are indicated. The boundaries are fixed to achieve desired frequentist properties. Spiegelhalter et al. (2004) refer to such decision rules as proper Bayes. The main problem with such approaches is the large number of arbitrary choices. The major advantage is the ease of implementation.

We propose rules that are derived as optimal Bayes rules by maximizing expected utility. But we avoid the prohibitive computational cost of backward induction by appealing to an approximation. Instead of a full backward induction solution, we use decision boundaries in the space of marginal posterior moments $(\log(s_{ti}), m_{ti})$ to approximate the optimal sequential decision. See Figure 2 for an example. The decision boundary is defined by two line segments starting at (s_0, b_0) and going through (s_1, b_1) , with $b_1 > b_0$ and (s_1, b_2) , with $b_2 < b_0$, respectively. The two values s_0 and s_1 are fixed, leaving $b = (b_0, b_1, b_2)$

to identify the decision boundary. At the end of each period t , we compare the marginal moments $(\log(s_{ti}), m_{ti})$ with the decision boundaries. If m_{ti} lies between the two lines, then we continue to accrue patients for treatment i , that is $d_{ti} = 1$. If not, we drop treatment i ($d_{ti} = 0$). In summary,

$$d_{ti}(b) = I \left(b_0 + \frac{b_2 - b_0}{s_1 - s_0} (\log(s_{ti}) - s_o) < m_{ti} < b_0 + \frac{b_1 - b_0}{s_1 - s_0} (\log(s_{ti}) - s_o) \right) \quad (5)$$

We write $d = d(b)$ to highlight the nature of d as a rule determined by the decision boundary b .

Figure 2 shows the decision boundaries for a specific choice of (b_0, b_1, b_2) . The (fixed) offset s_o determines the $\log(s_{ti})$ value where the two half lines join. We always stop accruing patients when $\log(s_{ti}) < s_o$. This has the desirable implication of imposing an upper bound on the amount of information, as measured by posterior variance, before making a stopping decision. The stopping decision is followed by the terminal decision a_i , as described earlier.

Using decision boundaries as in (5) reduces the solution of the sequential decision problem to finding the optimal parameters $b = (b_0, b_1, b_2)$. The optimal choice is determined by maximizing expected utility

$$b^* = \arg \max_b E [u(d(b), a, \theta, y)] \quad (6)$$

The expectation is over $\theta \sim p(\theta)$ and $y_{ti} \sim p(y_{ti} | \theta)$, and plugging in the optimal terminal decisions a_i .

3 EXPECTED UTILITY MAXIMIZATION BY SIMULATION

We resort to forward simulation to evaluate expected utility

$$U(b) = E [u(d(b), a, \theta, y)] \quad (7)$$

using the optimal terminal rule a . Forward simulation was introduced in Carlin et al. (1998) to solve sequential decision problems that can be described by decision boundaries. We simulate once, up front, possible realizations, $j = 1, \dots, M$, of the screening process, keeping all treatments in the trial until a final horizon T . That is, we do not include stopping in the simulation. The arrival of new treatments is simulated using the multinomial probabilities π_j .

To evaluate expected utility $U(b)$ for a decision boundary described by b we look through the file of saved simulations. Let u_i denote the i -th term in (3). Whenever a treatment hits the decision boundary b , it is removed from the current set. When this happens we compute the optimal terminal decision a_i using (4) and record the realized utility u_i for this treatment. Summing u_i over all treatments we get a realization of the utility (3). Averaging over all simulated realizations, $j = 1, \dots, M$, we obtain an estimate $\hat{U}(b)$ of the expected utility $U(b)$. In other words, we use the Monte Carlo average $\hat{U}(b)$ to evaluate the expected utility integral (7). Similarly, we can evaluate the expected value of other summary statistics, such as the number of patients tested with each treatment or the probabilities of type I and II errors. Finally, evaluating $\hat{U}(b)$ over a grid on b , we find the optimal decision boundary b^* .

Evaluation of $U(b)$ as a sample average \hat{U} does not exploit assumed regularities of the expected utility surface as a function of b . That is, we ignore that we could learn about b also by looking at close-by designs b' . This is formalized by fitting a smooth surface $\tilde{U}(b)$ to the observed sample averages $\hat{U}(b)$ as a function of b . Such smoothing was proposed in Müller & Parmigiani (1996) as a generic method to improve expected utility evaluations. We propose to define a smooth surface $\tilde{U}(b)$ as a locally weighted linear regression of

$\hat{U}(b)$ on b , using only main effects for b_0 , b_1 and b_2 .

The described algorithm requires the evaluation of (m_{ti}, s_{ti}) for a large number of times, treatments and simulations. This can be very computationally intensive when no closed form is available, as is the case for the model defined by (1) and (2). We implemented instead an empirical Bayes approximation to (m_{ti}, s_{ti}) as proposed, for example, in Gelman et al. (1995).

4 SIMULATION EXAMPLE

We implemented the described method for the following problem. We assume a standard of care with success probability $\theta_0 = 0.5$, sampling in cohorts of $N_{ti} = 2$ patients, and multinomial probabilities $(\pi_0, \dots, \pi_3) = (0.7, 0.2, 0.05, 0.05)$ for the arrival of new treatments Δn_t . We specify no limit on the number of available patients, i.e., $N^* = \infty$.

We set the prior parameters to $A_u = 3$, $B_u = 1$, $A_v = 3$ and $B_v = 1$, corresponding to a prior mean $E(\theta_i) = 0.5$ and standard deviation $SD(\theta_i) = 0.27$. The experimenter expects new treatments to be as good as the standard of care on the average, but the actual performance of individual treatments can vary considerably. For the utility function, we use relative weights $c_1 = 1$ and $c_2 = 10,000$, i.e. the final payoff for a successful drug is 10,000 times the sampling cost for one patient. The value of c_2 was chosen to achieve a power of approximately 80%. See below for a definition of power and type-II error in the context of this simulation. The time horizon is assumed as $T = 100$. We investigated the impact of T on the solution by considering a doubling of the time horizon to $T = 200$. Comparing the reported optimal rules we found no significant change, leading us to interpret $T = 100$ as a reasonable

approximation for a process with infinite horizon.

We add one more important feature to the decision problem. Let $1 - \beta$ denote the probability of an effective treatment, i.e., a treatment with simulation truth $\theta_i > \theta_0$, being recommended for phase III. The probability is over repeated simulations, and averaging with respect to the prior over all $\theta_i > \theta_0$. We refer to β as the false negative probability (type-II error), and $1 - \beta$ as power. Similarly, we define α as the false positive probability (type-I error). We constrain the set of allowable decision rules b to such rules that imply $\alpha \leq 0.05$ and $\beta \leq 0.20$, i.e., power $> 80\%$. The motivation for adding the constraint is that the utility function (3) could be criticized as being too narrowly focused on the perspective of the investigator and drug developer only. In the simulation the constraint is imposed by restricting the grid search for the optimal rule b^* to decision rules b that satisfy the conditions. To evaluate α , we find the relative frequency of simulated treatments with $\theta_i < \theta_0$ in the forward simulation that are recommended for phase III. To evaluate β , we count the treatments with $\theta_i > \theta_0$ that are not recommended for phase III. All results are based on $M = 1000$ forward simulations.

We evaluate the expected utility $U(b)$ in (7) over a 3-dimensional grid, as described in Section 3. Figures 3abc plot the surface $\hat{U}(b_0, b_1, b_2)$ for several values of b_0 . The flat nature of the surface with respect to b_1 indicates that a wide range of b_1 values yield similar expected utility.

Figure 2 shows the optimal decision boundaries subject to $\alpha \leq 0.05$ and $\beta \leq 0.2$, some simulated trajectories, and the terminal decisions. For each treatment in the simulated trial we plot the trajectory of (m_{ti}, s_{ti}) . We follow each treatment from right (high, prior variance) to left until the trajectory

crosses a decision boundary. This defines the stopping time $t = t_{i1}$ and the terminal decision a_i .

Rows 1 and 2 of Table 1 provide the solution to both, the unconstrained and the constrained optimization problems. The optimal unconstrained decision has higher expected utility than the optimal constrained decision, but it requires a larger average number of patients and it implies higher α and β .

Finally, we fit a smooth surface $\tilde{U}(b)$ to $\hat{U}(b)$ by locally weighted linear regression, as proposed in Section 3. The optimal bandwidth is selected by leaving 1/3 of the grid points out of the model fit and minimizing the mean square error of the predictions for those points. Figures 3def show the fit.

The optimal decisions, shown in rows 3 and 4 of Table 1, are very similar to those obtained without smoothing. The fact that the optimal design changed little confirms that the chosen Monte Carlo sample size, $M = 1000$, was sufficiently large for this optimal design problem. For smaller M , the advantages of smoothing should be more noticeable.

5 UNCERTAINTY AND SENSITIVITY OF THE OPTIMAL DECISION

5.1 *Uncertainty*

We consider two sources of uncertainty in the final solution b^* . First, numerical errors in evaluating the expected utilities imply uncertainty about the location of the maximum b^* . We refer to this uncertainty as *numerical uncertainty*. Second, even if we identify the correct mode of the expected utility surface, there may be other designs with almost equally high expected utility. We refer to a set of designs b with expected utility $U(b)$ within a small neighborhood of

$U(b^*)$ as *almost optimal designs*. While it is possible to reduce the first source of uncertainty by more extensive simulation, the latter uncertainty is inherent in the problem. We can only aim to honestly describe it.

To evaluate the numerical uncertainty in b^* we select designs b within a neighbourhood of b^* . We then approximate the expected utility $U(b)$ in that neighbourhood with a quadratic response surface, $U(b) = Q(b; \gamma) + \epsilon$. Here γ are the regression coefficients of the quadratic function and ϵ are independent normal residuals. The posterior distribution on γ implies a posterior distribution on the mode $b_Q^*(\gamma)$ of the response surface. We report 95% posterior intervals for b_Q^* to summarize the numerical uncertainty in the optimization. The results are shown in Table 2. The table is based on a neighborhood of b^* defined by $\|b - b^*\| \leq 0.01$. We judge the reported uncertainties to be negligible based on the comparison with the suboptimal set of designs discussed below. The small size of the reported numerical uncertainties confirms that the chosen Monte Carlo sample size $M = 1000$ was sufficiently large.

Next we find the set of almost optimal designs. In the forward simulation $20^3 = 8000$ triples $b = (b_0, b_1, b_2)$ were considered, i.e., we estimated the expected utility and type I and type II error rates for 8000 possible values of b . Of these, 55 designs satisfied the constraint $\hat{\alpha} \leq 0.05$ and $\hat{\beta} \leq 0.2$ and had an expected utility greater than 95% of the utility under the optimal design b^* , i.e., $\hat{U}(b) \geq 0.95 \hat{U}(b^*)$. Here b^* denotes the optimal rule for the constrained problem. We refer to these 55 designs as the almost optimal designs. They are suboptimal, but only by a negligible difference in expected utility. The range of almost optimal designs is reported in Table 2. Since the decision problem is invariant with respect to any additive shift of the utility function, it is not

possible to recommend a universal threshold like the 95% chosen here. The choice depends on the problem. The reported range of almost optimal designs is a useful diagnostic to help interpret, critique and modify the proposed solution. Typically the utility function is only a stylized description of the decision problem. The range of suboptimal designs allows the investigator to consider adjustments of the proposed solution to accommodate secondary goals and nuances of the decision problem that were not included in the formal utility function. For example, a large range on b_1 or b_2 might lead an investigator to propose designs with a narrower continuation region than b^* , i.e., shorter total time for each treatment under consideration.

5.2 Sensitivity Analysis

We assess the sensitivity of the solutions with respect to the choice of the main features of the decision problem: the utility function, the prior probability model, the parametrization of the decision boundary and the maximum number of patients enrolled across all trials at each time (N^*).

We first consider changes to the utility function defined in (3). We leave the general form of the utility unchanged, but we now weight the payoff for a significant phase III result by the true advantage over placebo ($\theta_i - \theta_0$), rather than the estimated advantage. We define the utility function

$$u_2(d, a, \theta, y) = \sum_{i=1}^n -c_1 \cdot N_i + \sum_{i: a_i=1} [-c_1 \cdot n_3(m_{ti}, s_{ti}) + c_2 Pr(B | y_1, \dots, y_t) (\theta_i - \theta_0)]. \quad (8)$$

The optimal designs under the corresponding expected utility $U_2(b)$ are shown in rows 5-6 of Table 1 (smooth version only). Compared to the solution under

the original utility function, b_0 in the solution of the unconstrained problem decreases, the expected sample size \bar{N} increases and the type I error probability decreases slightly. The solution of the constrained problem is robust with respect to the change in the utility.

Next we consider changes in the prior probability model. In (2) we defined a hierarchical model with hyperparameters that allow the pooling of information between treatments. We investigate the change in the optimal design if at the time of analysis we ignore the hierarchy and use independent beta priors. We continue to use the hierarchical model as the simulation truth. For a meaningful comparison, we use U_2 since it does not depend on model-based estimates of θ_i . Table 1 presents the optimal decisions for both, the constrained and unconstrained problems. Without pooling information the expected number of patients per treatment is increased. The change is most extreme in the constrained problem. The expected utility of the optimal design changes only little. We conclude that by using the hierarchical prior we can gain the same payoff with fewer patients. Of course, this conclusion is only valid if the true sampling process does in fact include dependence across treatments.

Next, we consider changes in the parametrization of the decision boundaries. In (5) we imposed that the boundaries be linear in $\log(s_{ti})$. We investigate changing the boundaries to be linear in s_{ti} , i.e. in (5) we replace $\log(s_{ti})$ by s_{ti} . Results are shown in Table 1. The optimal design, its expected utility, and the expected number of patients per treatment are similar to the results for the log-scale parameterization. Lack of such robustness would be a concern. It would indicate that the optimal sequential rule is very poorly approximated by boundaries on the chosen grid.

Finally we consider changing N^* . We investigate the solution for only $N^* = 10$ eligible patients available to enroll at each time (across all treatments in A_t). The optimal rule, shown in Table 1, results in smaller sample sizes and reduced utility, especially under the unconstrained problem.

6 A CLINICAL IMMUNOLOGY EXAMPLE

We apply the proposed approach to a screening design for vaccines in a clinical immunology scenario. The same example was analyzed in Yao et al. (1996) and Yao & Venkatraman (1998) using an approach that minimizes the expected number of patients until the first effective treatment is identified. The design is restricted to a bound on type I and type II errors. They propose a two-stage design where an interim analysis is carried out after the first N_1 patients. The treatment is discarded if the number of successes is $\leq K_1$. Otherwise, N_2 more patients are accrued and the treatment is discarded if the overall number of successes is $\leq K_2$ and recommended otherwise. They then repeat the same process with the second treatment, and so forth. The decision parameters are K_1 , K_2 , N_1 and N_2 . The method also includes a truncation, i.e. stopping the accrual before observing N_1 patients if the number of failures is already $\geq N_1 - K_1$, and stopping before N_2 patients if we already have more than $N_2 - K_2$ failures. Truncation can significantly reduce the expected number of patients.

In our sequential approach, we define the utility function to be the average number of patients needed to recommend one treatment. This allows us to compare the results across methods. Specifically, we define the utility function to be the ratio of the total number of patients enrolled across all treatments

by the number of treatments recommended for phase III. Let $N_i = \sum_{t=t_{0i}}^{t_{1i}} N_{ti}$ denote the total number of patients on treatment i . We define

$$u_3(d, a, \theta, y) = - \sum_{i=1}^{n_T} N_i \bigg/ \sum_{i=1}^{n_T} a_i$$

Like Yao *et al.*, we set the success probability for the standard of care to be $\theta_0 = 0.5$ and we use a Beta prior with parameters $\hat{u} = 0.3188$, $\hat{v} = 0.5327$, chosen to match the moments based on historical data, $\hat{E}(\theta_i) = 0.3743$ and $\hat{V}(\theta_i) = 0.1265$. The prior gives a high probability to success probabilities close to 0 or 1.

We evaluate designs with $M = 1000$ simulations on a grid with 20 equally spaced values of b_0 in $[0.3, 0.7]$, b_1 in $[0.3, 0.8]$ and b_2 in $[0.2, 0.6]$. We use cohorts of $N = 2$ patients. After each batch the posterior moments are evaluated and the decision to stop is taken according to (5). For the terminal decision we use a fixed rule. Upon stopping the enrollment, a treatment is recommended when stopping was indicated by crossing the upper boundary, and a treatment is abandoned if stopping was indicated by crossing the lower boundary. We then select b to maximize \hat{U}_3 , the Monte Carlo sample average utility in the forward simulation. Again, the maximization is restricted to designs b that satisfy the constraints $\hat{\alpha} \leq \alpha_{max}$ and $\hat{\beta} \leq \beta_{max}$.

Table 3 shows the optimal decision boundaries for several values of α_{max} and β_{max} and compares them with the two-stage optimal design with truncation proposed in Yao & Venkatraman (1998). The fully sequential approach with the optimal decision boundaries yields a reduction between 27% and 57% in the expected number of patients necessary to recommend a treatment for phase III evaluation. In some cases the actual $\hat{\alpha}$ and $\hat{\beta}$ are lower than the

upper bound imposed by the constraints. The reduced sample size is a natural consequence of the fully sequential setup, and does not reflect on any deficiency in the other method.

7 DISCUSSION

We have proposed a Bayesian decision-theoretic approach to optimal screening designs for phase II studies. Its main strength is the generality of the simulation-based solution which allows for a wide range of probability models and essentially any utility function. Another advantage is the possibility of optimizing within a subset of rules that satisfy certain properties.

As a decision-theoretic approach, the proposed method inherits the usual limitations of expected utility maximization. In particular, it requires the specification of a utility function and a prior probability model.

The use of decision boundaries to solve the sequential design problem greatly reduces the computational burden to find the optimal sequential decision. At the same time, however, it restricts the possible actions to those described by such decision boundaries. Instead of decision boundaries one could consider the optimal decision for all possible values of a suitable summary statistic, in our case (s_{ti}, m_{ti}) , on a finite grid. This is explored in Ding (2006).

The basic framework developed in this paper allows many generalizations. The prior model could easily be generalized to include a regression of success probabilities on treatment-specific covariates. For example, we might learn that treatments that target a specific molecular mechanism are more successful than others. Another important direction of generalization is the sampling

model. Little changes in the proposed algorithms if we replace the binary outcome by a continuous response or time to some clinically meaningful event, as long as we can define a single parameter upon which to base the inference. For example, when analyzing time until tumor progression one could define the summaries (m_{ti}, s_{ti}) as posterior moments of a log hazard ratio for treatment relative to the standard of care. The nature of the event time as a delayed response would cause no difficulty in the optimal design scheme. Delayed responses are accounted for in the definition of the posterior moments. In particular, the definition of the likelihood function would include different factors for censored observations and for observed event times, as usual in posterior inference for event time data.

Acknowledgments

Research was supported by NIH/NCI grants R33 CA97534-01 and R01 CA075981. We thank Raquel Montes Díaz and Roberto Carta for work on earlier prototypes of proposed method.

REFERENCES

BERRY, D. A., MUELLER, P., GRIEVE, A. P., SMITH, M., PARKE, T., BLAZEK, R., MITCHARD, N. & KRAMS, M. (2001). Adaptive bayesian designs for dose-ranging drug trials. In C. Gatsonis, R. E. Kass, B. Carlin, A. Carriquiry, A. Gelman, I. Verdinelli & M. West, eds., *Case Studies in Bayesian Statistics, Volume V*, Lecture Notes in Statistics. New York: Springer-Verlag.

- BROCKWELL, A. E. & KADANE, J. B. (2003). A gridding method for bayesian sequential decision problems. *Journal of Computational and Graphical Statistics* 12 566–584.
- CARLIN, B., KADANE, J. & GELFAND, A. (1998). Approaches for optimal sequential decision analysis in clinical trials. *Biometrics* 54 964–975.
- DEGROOT, M. (2004). *Optimal Statistical Decisions*. New York: Wiley-Interscience.
- DING, M. (2006). *Bayesian Optimal Design for Phase II Screening Trials*. Ph.D. thesis, M.D. Anderson Cancer Center and Rice University.
- GELMAN, A., CARLIN, J., STERN, H., & RUBIN, D. (1995). *Bayesian Data Analysis*. Chapman & Hall.
- GITTINS, J. & PEZESHK, H. (2002). A decision theoretic approach to sample size determination in clinical trials. *J Biopharm Stat.* 12 535–551.
- JAWITZ, J. (2004). Moments of truncated continuous univariate distributions. *Advances in Water Resources* 27 269–281.
- LEUNG, D. H. Y. & WANG, Y. G. (2001). Optimal designs for evaluating a series of treatments. *Biometrics* 57 168–171.
- MÜLLER, P. & PARMIGIANI, G. (1996). Optimal design via curve fitting of monte carlo experiments. *Journal of the American Statistical Association* 90 1322–1330.
- SIMON, R. (1987). How large should a phase ii trial of a new drug be? *Cancer Trt. Rpts.* 71 1079–1085.

- SPIEGELHALTER, D. J., ABRAMS, K. R. & MYLES, J. P. (2004). *Bayesian approaches to clinical trials and health care evaluation*. Chichester, UK: John Wiley and Sons.
- STOUT, Q. & HARDWICK, J. (2005). Optimal screening designs with flexible cost and constraint structures. *Journal of Statistical Planning and Inference* 132 149–162.
- STRAUSS, N. & SIMON, R. (1995). Investigating a sequence of randomized phase-ii trials to discover promising treatments. *Statistics in Medicine* 14 1479–1489.
- THALL, P., SIMON, R. & ESTEY, E. (1995). Bayesian sequential monitoring designs for single-arm clinical trials with multiple outcomes. *Statistics in Medicine* 14 357–379.
- WANG, Y. G. & LEUNG, D. H. Y. (1998). An optimal design for screening trials. *Biometrics* 54 243–250.
- YAO, T. J., BEGG, C. B. & LIVINGSTON, P. O. (1996). Optimal sample size for a series of pilot trials of new agents. *Biometrics* 52 992–1001.
- YAO, T. J. & VENKATRAMAN, E. S. (1998). Optimal two-stage design for a series of pilot trials of new agents. *Biometrics* 54 1183–1189.

Table 1: Optimal Decisions for the Simulation Example

Original problem								
	b_0^*	b_1^*	b_2^*	$\hat{E}(U)$	$\hat{\alpha}$	$\hat{\beta}$	\bar{N}	$\sqrt{\text{MSE}}$
$\max \hat{U}$	0.88	1.69	0.49	42320	0.03	0.23	18.2	10.9%
$\max \hat{U}, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.56	0.75	0.42	39410	0.02	0.20	17.1	12.0%
$\max \tilde{U}$	0.90	1.49	0.49	42268	0.03	0.23	17.8	10.9%
$\max \tilde{U}, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.56	0.75	0.42	39214	0.02	0.20	17.1	12.0%
Sensitivity to change in utility								
$\max \tilde{U}_2$	0.58	1.59	0.42	37130	0.01	0.23	27.4	10.7%
$\max \tilde{U}_2, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.56	0.70	0.36	33649	0.02	0.20	17.8	11.8%
Sensitivity to ignoring hierarchical model								
$\max \tilde{U}_2$	0.58	1.59	0.34	36853	0.01	0.24	32.6	10.0%
$\max \tilde{U}_2, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.51	0.75	0.26	34766	0.02	0.20	26.3	10.0%
Sensitivity to re-parameterizing boundaries								
$\max \tilde{U}$	0.90	1.65	0.55	42225	0.04	0.23	15.3	11.0%
$\max \tilde{U}, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.53	0.91	0.39	39931	0.02	0.20	22.1	10.5%
Limiting number of available patients to $N^* = 10$								
$\max \tilde{U}$	0.69	0.70	0.57	40448	0.06	0.31	4.5	16.1%
$\max \tilde{U}, \hat{\alpha} \leq .05, \hat{\beta} \leq .2$	0.48	0.70	0.42	37116	0.03	0.20	13.4	13.0%

\bar{N} is the average number of patients tested with each treatment. MSE is the average mean squared error in estimating θ_i . In the last two rows, N^* constrains the number of patients available for enrollment in each period.

Table 2: Uncertainty in the Determination of the Optimal Decision

	original grid	optimal design b^*	numerical uncertainty	almost optimal designs
b_0	[0.4,0.9]	0.56	[0.56,0.56]	[0.43,0.58]
b_1	[0.7,1.7]	0.75	[0.75,0.80]	[0.70,0.81]
b_2	[0.1,0.6]	0.42	[0.36,0.42]	[0.23,0.42]

Table 3: Clinical Immunology Example

α_{max}	β_{max}	b_0^*	b_1^*	b_2^*	$-\hat{U}_3$	$\hat{\alpha}$	$\hat{\beta}$	N_{Yao}	% reduction
0.05	0.05	0.45	0.64	0.24	31.49	0.04	0.04	55.60	43.36
0.05	0.10	0.53	0.67	0.37	15.67	0.05	0.10	32.40	51.64
0.05	0.15	0.57	0.64	0.52	9.93	0.05	0.14	20.20	50.85
0.10	0.05	0.49	0.51	0.26	17.04	0.09	0.03	27.20	37.35
0.10	0.10	0.43	0.46	0.37	7.14	0.10	0.10	16.50	56.70
0.10	0.15	0.41	0.45	0.39	7.04	0.09	0.10	14.80	52.46
0.15	0.05	0.43	0.46	0.22	13.04	0.15	0.02	17.80	26.74
0.15	0.10	0.41	0.44	0.37	6.18	0.12	0.09	12.80	51.71
0.15	0.15	0.41	0.44	0.37	6.18	0.12	0.09	9.50	34.94

$(-\hat{U}_3)$: expected number of patients necessary to recommend one treatment.

N_{Yao} : expected number of patients necessary to recommend the first treatment when following the approach in Yao & Venkatraman (1998).

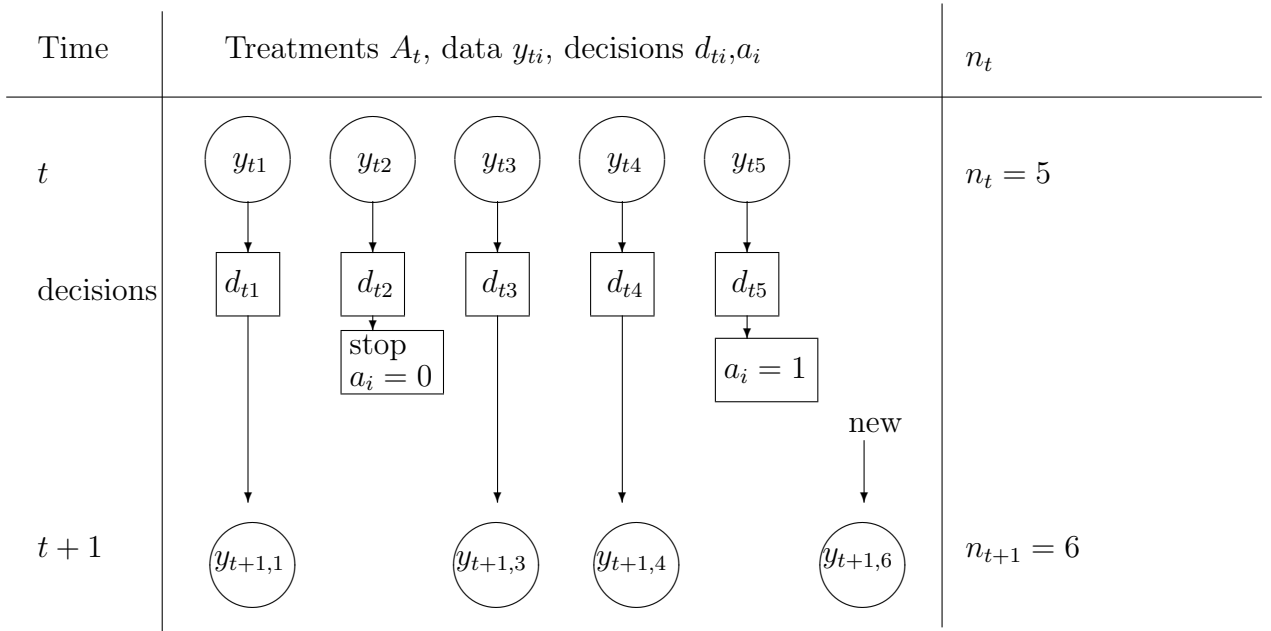


Figure 1: Multiple binomial experiments are available at time t . Some of them are dropped and some are introduced at the end of period t .







Figure 2: Forward simulation. For 20 treatments we plot $(m_{t_i}, \log s_{t_i})$ from the time t_{0i} treatment i arose until t_{1i} when it stopped. The two thick black half lines show a decision boundary b .




(a) $b_0 = 0.56$




(b) $b_0 = 0.66$




(c) $b_0 = 0.87$



(d) $b_0 = 0.56$



(e) $b_0 = 0.66$



(f) $b_0 = 0.87$

Figure 3: Heat map of unsmoothed expected utilities $\hat{U}(b)$ (first row) and smoothed $\tilde{U}(b)$ (second row) on a grid. The black star in each figure marks the optimal (b_1, b_2) combination for that b_0 .