NIH Public Access
Author Manuscript
Stat Med. Author manuscript; available in PMC 2014 May 30.

Published in final edited form as:
Stat Med. 2013 May 30; 32(12): . doi:10.1002/sim.5735.

A Bayesian decision-theoretic sequential response-adaptive randomization design

Fei Jiang[a], **J. Jack Lee**[a,1], and **Peter Müller**[b]

[a]Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030

[b]Department of Mathematics, The University of Texas at Austin, 1 University Station C1200, Austin, TX 78712

Abstract

We propose a class of phase II clinical trial designs with sequential stopping and adaptive treatment allocation to evaluate treatment efficacy. Our work is based on two-arm (control and experimental treatment) designs with binary endpoints. Our overall goal is to construct more efficient and ethical randomized phase II trials by reducing the average sample sizes and increasing the percentage of patients assigned to the better treatment arms of the trials. The designs combine the Bayesian decision-theoretic sequential approach with adaptive randomization procedures in order to achieve simultaneous goals of improved efficiency and ethics. The design parameters represent the costs of different decisions, e.g., the decisions for stopping or continuing the trials. The parameters enable us to incorporate the actual costs of the decisions in practice. The proposed designs allow the clinical trials to stop early for either efficacy or futility. Furthermore, the designs assign more patients to better treatment arms by applying adaptive randomization procedures. We develop an algorithm based on the constrained backward induction and forward simulation to implement the designs. The algorithm overcomes the computational difficulty of the backward induction method, thereby making our approach practicable. The designs result in trials with desirable operating characteristics under the simulated settings. Moreover, the designs are robust with respect to the response rate of the control group.

Keywords

sequential method; response adaptive randomization; Bayesian decision–theoretic approach; backward induction; forward simulation

1. Introduction

Phase II clinical trials are designed to evaluate the efficacy of new treatments and find the correct dose of new drugs, as well as to address potential safety problems. In this paper, we focus on the goal of providing an initial test of the efficacy of new treatments. The purpose is to screen out the inefficacious treatments before launching a large-scale phase III study. Qualified trial designs should be able to achieve the prespecified goals stated in the protocols, such as satisfying the type I and II error rate requirements. Additionally, designs with smaller average sample size requirements and those that allow more patients to be

Copyright: © 2013 John Wiley & Sons, Ltd

[1]Correspondence to: Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030. jjlee@mdanderson.org.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

assigned to the more efficacious treatment arms are desirable. In this paper, we evaluate treatment efficacy by assessing a binary endpoint, e.g., response or nonresponse to a treatment. Hence an efficacious treatment has a higher response rate. The performance of the design is evaluated by assessing the operating characteristics of the trial, including the type I error rate (α), statistical power (1-type II error rate (β), average sample size, and average percentage of patients assigned to the better treatment arm.

Sequential stopping procedures that allow for interim decision making to achieve smaller average sample sizes have been well developed for clinical trial designs [1]. In addition, response-adaptive randomization (RAR) procedures have been proposed to allow more patients to be treated with the more efficacious treatments [2]. We combine these two procedures in order to achieve an overall goal of designing more efficient and ethical randomized phase II trials by reducing the average sample size and increasing the assignment of patients to the better treatment.

In the literature, sequential stopping and RAR procedures have been proposed under the frequentist setting. The classical frequentist group sequential methods include the Pocock method [3], the O'Brien and Fleming method [4], and the Lan and DeMets error spending function method [5]. The frequentist methods define stopping boundaries for controlling the overall type I error rate. By varying the boundary shape, the design can achieve different average sample sizes, α, and $1 - \beta$.

RAR procedures proposed under the frequentist setting include randomized play-the-winner, randomized Pòlya urn, birth and death urn, drop-the-loser , sequential maximum likelihood, and doubly-adaptive biased coin design (DBCD) methods. For comparing frequentist methods with the proposed method, we apply the DBCD procedure because it allows the choice of suitable tuning parameters to provide results that are comparable to those of the other designs [2, 6].

Under the frequentist setting, the sequential stopping and RAR procedures are usually applied separately because of complications from combining the data dependence structures, which make the asymptotic properties difficult to justify. The frequentist methods have additional deficiencies in such applications. The frequentist approach does not conform to the likelihood principle, which implies that statistical inference should be a function of the likelihood of the observed data only and should not depend on unobserved data under a certain design.

For example, the p-value in the frequentist analysis is defined as the probability of obtaining an event as or more extreme than the observed one under the null hypothesis. The more extreme events are unobserved but are used for inference. Moreover, as mentioned above, frequentist inferences are often based on asymptotic results and the specific assumed design. Hence, they may not be valid if the trial conduct deviates from the design [7].

Parallel to the frequentist methods, Bayesian methods have been proposed for both sequential stopping and adaptive randomization procedures. For recent reviews, see for example [8] and [9].

Under the Bayesian setting, Lee and Liu [10] developed a sequential design based on the predictive probability approach. They constructed stopping rules based on the predictive probability, which corresponds to the probability of rejecting the null hypothesis at the end of the study, providing the same trend as that observed continues. The trial is stopped if the predictive probability exceeds a chosen efficacy cutoff value or is lower than a futility cutoff value. By choosing the proper cutoff value, the design preserves the type I and II error rates

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

and improves the trial efficiency. Additionally, the design is flexible and easy to implement. It is also robust when the study conduct deviates from the original design.

Bayesian RAR procedures have also been discussed in the literature. Thall and Wathen developed an adaptive allocation rule based on the posterior distributions of p_i, where $i = 1, 2$, denote the response rates for treatment arm i. They defined the probability of allocation to arm 2 as a function of $Pr(p_2 > p_1/data)^c$, where c is a tuning parameter used to reduce the variability across the allocation probabilities [11]. The resulting trials treat patients more effectively by assigning more patients to the better treatment arm. We refer to this procedure as the generalized W. Thompson's (GWT) procedure in this paper because it generalizes the method proposed by W. Thompson [12] by adding the tuning parameter c, as discussed by Berry and Eick [13]. We adopt the GWT procedure under the Bayesian setting.

Compared with the frequentist methods, Bayesian methods are better alternatives for constructing combined designs involving both sequential stopping and adaptive randomization. This is because Bayesian inference relies on the posterior distributions of the parameters, which automatically consider the dependence among the observations.

In order to avoid the limitations of the frequentist approach, and take advantage of combining the sequential stopping and adaptive randomization procedures, we implement our combined design through a Bayesian decision-theoretic approach. This method is based on the Bayesian estimation procedure; therefore the estimation procedures are essentially adaptive, and the final results reflect the utility of the clinical trial, which is more meaningful. Moreover, the tuning parameters used in the method represent the costs of the decisions, which signify a direct connection between the decision-making procedures and the goals of the trial.

The Bayesian decision-theoretic sequential approach was described by Lewis and Berry [14]. They defined a $0 - K$ terminal loss function and showed that by adjusting the cost value K, the Bayesian decision–theoretic sequential designs result in smaller average sample sizes, and hence are more efficient compared with the frequentist designs. The loss function that is chosen is meaningful from the frequentist perspective. When K is a constant, it is essentially the same as the frequentist hypothesis test [15].

Berry and Eick discussed Bayesian decision-theoretic response-adaptive randomization designs [13]. They analyzed four adaptive allocation procedures under the Bayesian decision-theoretic framework, and compared them with an equal randomization procedure. The results show that the adaptive randomization designs allow patients to be treated more effectively than the balanced designs.

The second goal of this paper is to develop a method to address the computational issues of the Bayesian decision-theoretic approach. Traditionally, a Bayesian decision problem that involves sequential decisions is solved by the backward induction method. This method is computationally intensive and thus is not practically applicable. In order to reduce the computational demands, we make two adjustments while implementing the design under the Bayesian decision-theoretic framework. First, we limit our design goal to preserve the error rates and minimize the required sample size under the prespecified frequentist or Bayesian RAR procedures; i.e., we only apply the Bayesian decision-theoretic method to define the sequential stopping rules, and the decision rules are based on the losses of error decisions and enrollment costs. Second, we extend the constrained backward induction method introduced by Mueller et al. [16, 17] to obtain reasonable yet suboptimal solutions for Bayesian decision-theoretic problems. Further, we use the forward simulation method to approximate the suboptimal solutions. As the exact solutions of Bayesian decision problems

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

are not attainable, we evaluate our method by assessing the resulting operating characteristics in simulated trials.

In Section 2, we formally define the Bayesian decision-theoretic sequential method, and discuss how to incorporate an RAR procedure. In Section 3, we discuss the backward induction algorithm and the forward simulation and constrained backward induction methods. In Section 4, we show the simulation results and compare our designs with the frequentist power family boundary-DBCD design introduced by Morgan and Coad [18]. In Section 5, we provide our conclusions and discussions.

2. Bayesian decision-theoretic sequential and response-adaptive randomization method

Lewis and Berry [14] introduced a framework of the Bayesian decision-theoretic method, and illustrated its application to an animal study and clinical trials. Their simulation studies showed that the average sample sizes under the Bayesian decision-theoretic framework are smaller than those for trials using classical frequentist methods. The Bayesian decision-theoretic method has been used in more recent clinical trials. Gausche et al. [19] applied a Bayesian decision-theoretic method to evaluate the outcomes associated with the use of endotracheal intubation in pediatric patients in out-of-hospital emergency settings. They compared the effectiveness of two interventions: the use of only bag-valve-mask ventilation versus that form of ventilation plus endotracheal intubation. They assessed the short-term survival and neurological outcomes of the patients. The initial concern was that the two treatments would have a large difference in efficacy, so interim monitoring was desired. As the classical frequentist boundary method allows for relatively infrequent interim analyses, Gausche et al. chose to use a Bayesian decision-theoretic method in their study. In another example, Young et al. [20] applied the method to a clinical trial on the prophylactic use of phenytoin to prevent early posttraumatic seizures in children who experience blunt head trauma. Comparing phenytoin to a placebo, the Bayesian decision-theoretic method was used to assess the probability of pediatric patients remaining free from early posttraumatic seizures. Since Bayesian decision-theoretic designs are optimal with respect to the defined utility functions, they perform better than other designs in maximizing the utility functions. These two studies illustrate that the Bayesian decision-theoretic designs are easy to interpret. Such studies motivate us to extend the use of Bayesian decision-theoretic approaches to incorporate adaptive randomization and group sequential monitoring in clinical trials.

We apply a Bayesian decision-theoretic approach and an RAR procedure to define two-arm designs with binary endpoints. Without loss of generality, we assume that the two arms are the control (arm 1) and the experimental treatment (arm 2). We refer to the binary outcome as response (1) or nonresponse (0) to the treatments. We focus on the difference between the two response rates, $\delta = p_2 - p_1$, where p_1 and p_2 are the response rates for arm 1 and arm 2, respectively. The proposed design includes a preliminary first stage when $2n_{eq}$ patients are equally randomized to arm 1 and arm 2. After this initialization, the trial enrolls one additional patient sequentially at the ensuing stages until the sample size reaches the maximum value N. At each stage, we apply a Bayesian decision-theoretic approach to determine whether to continue the trial given the current state of the trial. If the decision is to continue the study, we use the RAR procedures to allocate the next patient to one of the treatments. The maximum number of analysis stages is $T = N - 2n_{eq} + 1$ under this setting. We index the analysis stages by $t = 1, \ldots, T$, and denote the sample size at time (stage) t for arm i, $i = 1, 2$, by n_{ti}, where $n_{11} = n_{12} = n_{eq}$ and $n_{T1} + n_{T2} = N$.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

2.1. Probability model

Let Δy_{ti}

denote the number of responses among m_{ti}

(cohort size at time t for treatment i) patients assigned to arm i at time t=1, … , T, and assume

$$\Delta y_{ti}|p_i \sim bin\,(m_{ti},\, p_i)\,,$$

where p_i is the true response rate of arm i, and the Δy_{ti}'s are assumed to be independent. In our implementation we use only m_{ti} = 1 or 0 depending on whether treatment i is selected or not. Generalization to different cohort sizes is straightforward.

Let $Y_{ti}=\sum_{j=1}^{t}\Delta y_{ji}$ be the total number of responses (achieving a favorable clinical response to the treatment) up to time t for arm i, and let $n_{ti}=\sum_{j=1}^{t}m_{ji}$ be the total number of patients. With the assumption that x_{ti} = 0 if m_{ti} = 0, i.e., no patients are assigned to arm i at stage t, we have

$$Y_{ti} \sim bin\,(n_{ti},\, p_i)\,.$$

The prior for parameter p_i is

$$f\,(p_i) = beta\,(\gamma_{0i})\,,$$

where $\gamma_{0i} = (\alpha_{0i}, \beta_{0i})$ are the beta distribution parameters. The posterior of p_i is $f(p_i|\gamma_{ti}) = beta(\gamma_{ti})$, with $\gamma_{ti} = (\alpha_{ti}, \beta_{ti}) = (\alpha_{0i} + y_{ti}, \beta_{0i} + n_{ti} - y_{ti})$.Here y_{ti} is the realization of Y_{ti}, the cumulative number of responses observed in arm i. The parameters are updated sequentially. The posterior at stage t is the prior for the next stage.

In our study, we implement two response-adaptive randomization procedures. Following the work of Thall and Wathen [11], we use the GWT procedure as the first allocation procedure with a Bayesian decision-theoretic sequential approach. We assign patients to arm 2 with probability proportional to the probability of arm 2 being superior raised to the power c. Letting $\delta = p_2 - p_1$, based on the posterior parameters, the allocation rate ψ_1 to arm 2 is defined as

$$\psi=\frac{Pr(\delta>0|\gamma_{ti})^c}{Pr(\delta>0|\gamma_{ti})^c+Pr(\delta<0|\gamma_{ti})^c}, \quad (2.1.1)$$

where c is a tuning parameter to adjust the extent of imbalance of the allocation and the variability of the allocation ratio across the trials [11].

Alternatively, we use the DBCD procedure under both the Bayesian decision-theoretic and frequentist settings in order to compare the two approaches under the same adaptive randomization scheme. The allocation procedure is defined by the allocation function

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

$$\mathbf{g}(\upsilon,\rho)=\begin{cases}1 & \text{if } \upsilon=0\\ \frac{\rho(\rho/\upsilon)^{\xi}}{\rho(\rho/\upsilon)^{\xi}+(1-\rho)((1-\rho)/(1-\upsilon))^{\xi}} & \text{if } 0<\upsilon<1\\ 0 & \text{if } \upsilon=1\end{cases} \quad (2.1.2)$$

where v is the currently observed allocation ratio, ρ is the target allocation ratio, and ξ is a tuning parameter that adjusts the convergence rates of the allocation ratios to the targets. We specify the target ratio on arm 2 to be

$$\rho=\frac{\sqrt{p_2}}{\sqrt{p_1}+\sqrt{p_2}}. \quad (2.1.3)$$

The choice of ρ has the aim of minimizing the expected number of treatment failures [21]. The value of ρ is generally unknown and can be estimated by plugging in the corresponding estimators $\hat{p}_1$, $\hat{p}_2$, such as Bayesian posterior means or maximum likelihood estimators under the Bayesian decision-theoretic and frequentist settings, respectively.

2.2. Decision-theoretic approach for clinical trials

Statistical decision-theoretic approaches can be applied to clinical trials to quantify various objectives of the clinical trials and to obtain optimal designs to achieve the specified objectives [12, 22]. In our setting, let S_t be the summary statistic based on the data Y_t, which represents the state of the trial at time t. The decision rule is a function $d(S_t, t)$ whose value is an action d to be taken, given the current state S_t. In general, we assume that S_t contains sufficient information about the data for making decisions. See later for more details and examples for the choice of S_t. In the sequential setting, the set of decision rules is referred to as policy π for the decision process, where $\pi = (d(S_1, 1), d(S_2, 2), \ldots, d(S_T, T))$. Letting Π denote the class of policies, the goal is to determine the action to take for each state by selecting a policy $\pi \in \Pi$ to minimize the expected total loss. The expected total losses are the objective functions whose values rely on the distributions of the sample paths and the costs of the decisions.

The choices of S_t vary with the underlying probability models. Under the beta-binomial assumption, we use a value of S_t that is equal to the sufficient statistic that contains the information about the number of patients who achieve a favorable clinical response to the treatment and the total number of observations on each arm, i.e., $S_t = (\gamma_{ti}, i = 1,2)$. The posterior parameter γ_{ti} is a natural choice because it contains all the relevant information. Note that under an equal randomization setting, a state is uniquely identified by a three-dimensional summary statistic because the allocation of patients is equal between the two arms. However, the adaptive randomization procedures enlarge the state space; therefore, in order to uniquely represent a state, S_t must be four-dimensional. As a result, RAR procedures increase computational complexity. We address the computational issues in the next section.

According to the observed state, we make a decision of stopping for efficacy, stopping for futility, or continuing the trial. It is convenient to describe the decisions as pairs $d = (d_1, d_2)$, including first a stopping decision $d_1 \in \{D_s, D_c\}$, and then a terminal decision $d_2 \in \{D_e, D_f\}$, in case we stop the trial. Here D_s and D_c represent the stopping and continuation decisions, respectively, and D_e and D_f represent the decisions of stopping for efficacy and futility, respectively. Let $L_\pi(S_t)$ be the expected total loss that is incurred by using policy π from stage t onwards, given the current state S_t. Its value depends on not only S_t but also the distribution of the sample path after stage t. Similarly, let $L_{d,\pi_{t+1}}(S_t)$ denote the expected loss under decision d at time t, assuming the use of policy π_{t+1} from $t + 1$ onward. When $d_1 = D_s$,

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

we drop $_{t+1}$ from the subindex because it is irrelevant. When $d_1 = D_c$, we drop d_2 from the notation because it is irrelevant. The optimal solution is the policy π^* that minimizes the expected total loss, i.e.,

$$L_{\pi^*}(S_0) = \min_{\pi} L_{\pi}(S_0)$$

and

$$\pi^* = \arg\min_{\pi} L_{\pi}(S_0),$$

where S_0 is the initial state.

In the rest of this section, we define the transition probability, the losses of the decisions, and the optimality equations. We will generically use $\ell(\cdot)$ to denote the realized loss, and $L(\cdot)$ to indicate the expected loss, after marginalizing with respect to some of the unknown variables. We start by describing the context and specific notations.

2.2.1. Transition Probabilities—We refer to $p(S_{t+1}|S_t, d(S_t))$ as the transition probability from the state S_t to S_{t+1} under the action $d(S_t)$. We divide the state space at time t into a continuation space $\mathscr{C}_t$ and a stopping space $\mathscr{S}_t$.

When $S_t = s \in \mathscr{S}_t$, the transition probability $p(S_{t+1} = s|S_t = s, d(s)) = 1$. The transition probability is nontrivial only if $S_t = s \in \mathscr{C}_t$. Therefore, we drop the action $d(S_t)$ from the conditioning subset in the notation, and write it as $p(S_{t+1}|S_t)$.

The probability $p(S_{t+1}|S_t)$ is characterized by the adaptive allocation rates and the posterior predictive distribution for the number of responses among the additional Δn_i patients who are allocated to each arm from the current to the next stages. Given S_t, S_{t+1} takes four values, depending on which arm is chosen for the next patient and whether or not the patient responds to the treatment. Therefore, S_{t+1} given S_t is determined by $\Delta y_{t+1,i} = (y_{t+1,i} - y_{ti})$, $i = 1,2$. For the beta (1,1) prior distribution for the response rate, conditional on Δn_i, the posterior predictive distribution for the number of responses $\Delta y_{t+1,\,i}$ is a product of two beta-binomial distributions, and can be simplified as a ratio of the gamma functions

$$\mathrm{y}_{\,t+1,i}|\Delta n_i,\ y_t = \prod_{i=1}^{2} \binom{\Delta n_i}{\Delta y_{t+1,i}} \left(\frac{\Gamma}{(n_{ti})\Gamma} (y_{ti}+\Delta y_{t+1,i})\,\Gamma(n_{ti} - y_{ti}+\Delta n_i - \Delta y_{t+1,i})\,\Gamma(n_{ti} - y_{ti})\,\Gamma(y_{ti})\,\Gamma(n_{ti}+\Delta n_i) \right). \quad (2.2.1)$$

For a fully sequential process, $\Delta n_i = 1$ or 0. $\Delta y_{t+1,\,i} = 1$ if the $(t + 1)^{th}$ patient is assigned to arm i and responds to the treatment; 0 otherwise. In fact, conditioning on the treatment selection, there is only one beta-binomial term left. To fully describe the probability of a patient achieving a favorable clinical response to the treatment on arm i at stage $t + 1$, we have to multiply the posterior predictive probability of the response rate on arm i by the probability of assigning a patient to arm i, which is defined in equation (2.1.1) or (2.1.2).

2.2.2. Loss functions—We first consider the loss of the decision $d_2 \in (D_e, D_f)$ when $d_1 = D_s$, i.e., the loss of the terminal decision upon stopping. We will use $\ell(d_2, \delta)$ for the loss of a terminal decision without considering the accrual costs.

We adopt the $0 - K$ loss function from Lewis and Berry [14] for one-sided testing in our study. Recall that the true response rates are denoted by p_i, $i = 1, 2$, and their difference is δ

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

$= p_2 - p_1$. Denoting the minimum difference for claiming efficacy by $\delta_0 (\delta_0 > 0)$, the definition of the losses is based on a hypothesis test of H_0: $\delta > 0$ versus H_1: $\delta > \delta > \delta_0$. The zone between 0 and δ_0 is considered to be an indifference region [14]. Using our notation and terminology, the terminal error loss function can be written as

$$\begin{aligned} \ell_{D_f} \quad &= \ell\,(d_2 = f, \delta) \quad &&= 0 \quad \text{if } \delta \le \delta_0 \\ & &&= K_1 \quad \text{if } \delta > \delta_0 \\ & &&\text{and} \qquad\qquad (2.2.2) \\ \ell_{D_e} \quad &= \ell\,(d_2 = e, \delta) \quad &&= K_2 \quad \text{if } \delta < 0 \\ & &&= 0 \quad \text{if } \delta \ge 0, \end{aligned}$$

where ℓ_{D_f}, ℓ_{D_e} represent the loss of wrongly stopping for futility and efficacy, respectively. According to the loss function, a loss would be incurred in the following two situations: we conclude the treatment is efficacious when the true difference $\delta \le 0$; and we conclude the treatment is inefficacious when the true difference $\delta > \delta_0$. There is no loss when the decision is consistent with the true parameters, or when the true parameter values are in the indifference region. According to the updated posterior distributions of the parameters, we are able to obtain $Pr(\delta \le \delta_0 | S_t)$ and $Pr(\delta > \delta_0 | S_t)$, the probabilities of the error decisions, and thus compute the expected losses of different actions, D_e, D_f, given the observed data.

There are two main reasons to select the $0 - K$ loss function, as in equation (2.2.2). First, the loss function is essentially the same as that defined for the standard hypothesis testing framework when $K_1 = K_2$, thus it has a clear interpretation from the frequentist perspective. Second, the cost assignments, K_1 and K_2, link directly with the error rates and can be interpreted as the costs of the type II and type I errors, respectively. Cheng and Shen [23] discussed the relationship between the α level and the value of K_2/K_1, and derived the upper bound of the ratio to achieve a certain α level. This property helps us to adjust the tuning parameters to control the error rates.

In the loss function (2.2.2) only terms for false decisions and for sampling cost appear. All are related primarily to efficacy. In contrast the adaptive allocation rule (2.1.1) is also related to ethical concerns by allocating more patients to better treatments. However, the two issues, ethics and efficacy are intertwined. A more efficient design is more ethical by exposing fewer patients to unnecessary risks.

The stopping losses and continuation losses are derived in appendix A.

3. Backward induction, constrained backward induction and forward simulation

The optimal policy π^* can be evaluated by an algorithm known as backward induction. In this section we introduce this algorithm and an approximate implementation that avoids the prohibitive computational cost of the full algorithm. As shown in equation (A.0.5), the optimal solution for stage t is obtained by evaluating an expectation over the expected total loss at the future time $t + 1$. Therefore, we should obtain the expected total loss at time $t + 1$ first. The standard way of implementing the procedure is by using the backward induction algorithm. However, as discussed before, the adaptive randomization procedures enlarge the state space, and therefore the usual backward induction algorithm becomes more complicated. As a result, we develop an alternative algorithm based on the constrained backward induction and forward simulation methods introduced by Mueller et al. [17]. The new algorithm aims to obtain a suboptimal solution for the decision problems. In this section, we review the backward induction algorithm, introduce the constrained backward

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

induction and forward simulation methods, and implement the methods under the response-adaptive randomization setting.

3.1. Backward induction

A decision problem is solvable by the backward induction algorithm only if there is a maximum stage at which the decision process must be stopped. After identifying the maximum stage, the procedure evaluates the expectations backward according to equations (2.2.2) through (A.0.3). Assuming the maximum stage is T, we obtain $L_{\pi_T^*}(S_T) = L_{(D_s, d_2)}(S_T)$, where $d_2(S_T) \in \{D_e, D_f\}$ as the final decision at S_T. We substitute this result into the optimal equation (A.0.4) for $t = T - 1$, and calculate $L_{\pi_{T-1}^*}(S_{T-1})$. The optimal decision at $t = T - 1$ selects $d_1(S_t)$ from D_c, D_s first, then $d_2(S_t)$ from D_e, D_f in the case when the decision D_s was selected. The selections are based on minimizing the expected losses of the actions. Then for $t = T - 2, T - 3, \ldots, 1$, we apply the procedure repeatedly to obtain the optimal decision for each S_t. The process is computationally intensive because we have to evaluate over all possible values of S_t, $t = 1, \ldots, T$. As the state S_t is four-dimensional, the computational effort grows as $O(4^N)$ in time, and $O(N^4)$ in storage, where N represents the maximum sample size of a trial.

Under the specified framework, we consider the following simplification to reduce computational effort. Recall that the unit enrollment cost, and hence the minimal expected continuation loss, is C. The trial will be stopped if the expected stopping loss at one state is less than C. Therefore we eliminate the states whose expected stopping losses are less than C before implementing the backward induction procedure.

3.2. Constrained backward induction

Brockwell and Kadane [24] and Mueller et al. [17] introduced the constrained backward induction method as an alternative to the full backward induction method to obtain results. The constrained backward induction method reduces the dimension of the state space and allows backward induction to be conducted based on a lower-dimensional summary statistic $\tilde{S}_t$.

Noting the reduced dimension of $\tilde{S}_t t$ it is clear that the results obtained under the constrained backward induction method might be suboptimal.

The goal is to choose a lower-dimensional statistic in order to reduce the computational burden, while also maintaining a certain level of accuracy. Mueller et al. argued that good choices of $\tilde{S}_t$ should be no more than three-dimensional [17]. We discuss the selection of summary statistics, forward simulation and constrained backward induction methods next.

3.3. Selecting the lower-dimensional summary statistics for constrained backward induction

We select a lower-dimensional summary statistic $\tilde{S}_t$ to best separate the continuation and stopping regions of the states based on the results from a small size, full backward induction, with the maximum sample size $N = 100$ and $\delta_0 = 0.4$. A sample size of $N = 100$ is chosen such that the exact computations are still feasible. Inspection of the exact solutions under different projections assists us in selecting the best summary statistic. We considered alternative summary statistics (not shown) and eventually identified the projection shown in Figure 1 as a reasonable separation. Figure 1 plots the posterior means versus the logarithms of the posterior variances of the difference $\delta = p_2 - p_1$ at certain stages. The middle regions are the continuing zones. The regions above and below correspond to the stopping zones for

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

efficacy and futility, respectively. Note that because we eliminate some stopping states by using the forward evaluation steps before the backward induction, the stopping regions are not fully displayed in the figure. Let ν_t, μ_t be the logarithm of the posterior variance, and the posterior mean of the δ at stage t. Based on the results, we define the lower-dimensional summary statistic $\tilde{S}_t = (\log(\nu_t), \mu_t)$ to partition the state space at time t. In other words, we assume the suboptimal policy is a sequence of decisions $d\left(\tilde{S}_t, t\right)$ that depend on the data only through the summary statistic $\tilde{S}_t = (\log(\nu_t), \mu_t)$. The dimension of the statistic $\tilde{S}_t$ is reduced compared with that of the original four-dimensional statistic S_t. It depends on the posterior mean, variance of the difference δ, and analysis stage t.

The number of states increases along with increases in the sample size; hence, overlapping of the states becomes serious at later stages. However, this is inevitable, since a value of $\tilde{S}_t$ does not correspond to a unique state represented by S_t, i.e., some closed states, S_t, may give the same values of $\tilde{S}_t$. This is the reason we claim our approach is suboptimal rather than optimal. Nevertheless, this compromise is necessary in order to make the backward induction algorithm practically feasible in a high-dimensional space.

3.4. Forward simulation

Upon selecting a summary statistic for constrained backward induction, we still face the daunting task of evaluating the expectation over all the sample paths whose states are represented by the reduced three-dimensional summary statistic.

Carlin et al. [25] introduced a Monte Carlo forward simulation method to evaluate the expected losses for different decision rules at each state and choose the optimal one to minimize the loss. This method replaces posterior integrals by their averages over selected, simulated sample paths. The selection of specific simulations accounts for conditioning. For example, the use of only simulations with $S_t = s$ amounts to conditioning on $S_t = s$.We implement this method in our study, and evaluate the desired expectation in equation (A.0.4) by averaging the expected losses over the simulated sample paths.

Detailed implementations of the above methods are presented in Appendix B, and plots based on the forward simulation and constrained backward induction method (which are similar to those in Figure 1) are presented in Figure 2.

4. Simulation example

We implement the methods under the following settings. We assume a two-arm design, with H_0: $p_1 = p_2$, H_A: $p_2 > p_1 + \delta_0$. The minimum difference in efficacy δ_0 is 0.2, the maximum sample size $N = 300$, and the equal randomization sample size is $2n_{eq} = 50$. We simulate 10000 repetitions of the trials under the setting where $p_1 = p_2 = 0.3$ and $p_1 = 0.3$, $p_2 = 0.5$. We also compare different methods under the settings where $p_1 \in B_1 := \{0.2, 0.3, 0.4, 0.5\}$ and $p_2 \in B_2 := \{\{0.1, 0.2, \ldots, 0.9\} \cap [p_1, 1)\}$. We evaluate the performances of the designs by comparing the resulting operating characteristics (OCs) of the trials. The OCs we consider are the type I error rate (α, the target is 0.05); power ($1 - \beta$, the target is 0.80); average sample size (ASN) under the null $\left(\bar{n}_\alpha\right)$, and under the alternative $\left(\bar{n}_\beta\right)$; the expected percentages of patients on arm 2 ($E(n_2/n)$), where n_2 is the sample size on arm 2 and n is the total sample size up to the stopping time; and the average loss (AVL) over the setting when $p_1 \in B_1$ and $p_2 \in B_2$. Note that the AVL is the sample realization of the expected total losses, which is the expectation with respect to the prior distribution of the parameters rather than with respect to any specific parameters. Therefore, we do not obtain

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

the AVLs simply by taking the simulated results under only specific settings, but by averaging the results across the simulation settings. We make comparisons of the OCs among the Bayesian decision-theoretic, sequential, response-adaptive randomization (BAR) designs and the frequentist, sequential, power family–DBCD (FAR) designs, as compared by Morgan and Coad [18].

4.1. Bayesian decision-theoretic, sequential, response-adaptive randomization (BAR) designs

Under the BAR design, the prior distributions are chosen to be beta(1, 1) for both arms. The cost is $C = 1$ for recruiting one patient. We adjust the K_1, K_2 values to control the error rates, and choose the exponent c in (2.1.1) between $\frac{1}{2}$ and $\frac{t}{2T}$ to control the allocation rates.

4.2. Frequentist, sequential, power family–DBCD randomization (FAR) designs

The frequentist, sequential, power family method defines stopping rules through specifying the lower and upper stopping boundaries as

$$low_t = \delta_0 \sqrt{I_t} - \lambda_2 (t/T)^{\Delta} - 1/2 \text{ and } up_t = \lambda_1 (t/T)^{\Delta} - 1/2, \quad (4.2.1)$$

where I_t is the information level at time t [18].

The design applies the DBCD adaptive procedure as the allocation procedure. The allocation rates are defined by the functions in equations (2.1.2) and (2.1.3), where the $\hat{p}_i$ in equation (2.1.3) is the maximum likelihood estimator.

We choose $\Delta = 0.5$ or 0, and adjust the values of λ_1 and λ_2 based on the simulations to achieve the target error rates. Note that when $\Delta = 0.5$, the boundary resembles the Pocock boundary, and when $\Delta = 0$ it resembles the O'Brien-Fleming boundary. We use ξ in equation (2.1.2) to adjust the convergent rate of the allocation ratio to the target rate. Usually it is selected to be 2 [2]. However, because of early stopping, we require a higher convergent rate and select $\xi = 10$ in this application.

In equation (4.2.1), the estimated information level is

$$\hat{I}_t = \frac{\bar{p}_t \left(1 - \bar{p}_t\right)}{1/n_{t1} + 1/n_{t2}},$$

where

$$\bar{p}_t = \frac{\hat{p}_{t1} n_{t1} + \hat{p}_{t2} n_{t2}}{n_{t1} + n_{t2}} \text{ and } \hat{p}_{ti} = y_{ti}/n_{ti}.$$

Then the estimate of the standard Z statistic is

$$\hat{Z} = (\hat{p}_{t2} - \hat{p}_{t1}) \sqrt{\hat{I}_t}.$$

4.3. Results

We evaluate the designs by assessing their operating characteristics (OCs) through simulation studies. First, we generate the settings to demonstrate the influence of K_1 and K_2

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

on α, $1-\beta$ and the average sample size (ASN), as shown in Table 1. Then we compare the results from the BAR and FAR designs (see Table 2). In both tables, we assume $p_1 = p_2 = 0.3$ under the null hypothesis, and $p_1 = 0.3$, $p_2 = 0.5$ under the alternative hypothesis. Additionally, we create plots with different p_1, p_2 values to examine the robustness of the designs. As a reference, the non-sequential, fixed sample size, equal randomization design requires a total sample size of 148 to achieve 80% power with a one-sided 5% type I error rate.

To generate the data shown in Table 1, we choose an adaptive randomization procedure, the generalized W. Thompsons (GWT) procedure with $c = t/(2T)$, and vary the costs K_1 and K_2 in order to evaluate their effects on $1-\beta$ and the ASN. Given $K_1 = 1000$, we assign the type I error rate cost K_2 as 1500, 5000, and 8000 in BAR_a, BAR_b and BAR_c, respectively. These adjustments reduce the type I error rate by directly increasing the type I error costs. Further, the power is reduced because the type II error cost K_1 is fixed, which means that the relative importance of the type II error rate is less. Moreover, the ASN increases because the overall error cost becomes larger compared with the fixed enrollment cost C.

In BAR_d, BAR_e and BAR_f, we fix the ratio of K_1 and K_2, but enlarge their values. The results show that $1-\beta$ and the ASN are both increased, but α is decreased. For the same reason as stated previously, the increase in the ASN is due to an increase in the overall error cost. The type I and II error rates are both reduced because of the increases in K_1 and K_2. Since K_1/K_2 is fixed, the type I error rate does not change as much as when the value of K_1/K_2 varies.

To sum up, there are two ways to control α, $1-\beta$ and the ASN. First, we can fix the values of K_1 and K_2 to reduce α substantially. But this also reduces $1-\beta$ and increases the ASN under the alternative hypothesis. Second, we can increase the values of $K_1 and K_2$ while keeping their ratio fixed. This reduces α and increases $1-\beta$; however, it increases the ASN at the same time. This result shows that we can calibrate the $K_1 and K_2$ values to achieve a target α level, a specific value for $1-\beta$, or to control the size of the ASN. At the same time, we have to carefully consider the trade-offs among the three.

Table 2 provides a comparison of the OCs among the BAR and FAR designs. We choose the $K_1 and K_2$ values that allow the trials to achieve the target $\alpha = 0.05$ and $1-\beta = 0.8$ values. BAR_2 and BAR_3 are the Bayesian decision–theoretic–GWT (BAR-GWT) designs with $c = t/(2T)$ or $c = 1/2$, respectively. BAR_1 is the Bayesian decision–theoretic-DBCD design. FAR_1 and FAR_2 are the frequentist power family-DBCD designs with $\Delta = 0 and \frac{1}{2}$, respectively. BEQ is the Bayesian decision–theoretic equal randomization design. FEQ_1 and FEQ_2 are the frequentist equal randomization designs with $\Delta = 0 and \frac{1}{2}$, respectively. The results show that the sample sizes required by all the designs listed in the table are less than the sample size required by the non-sequential equal randomization design.

In order to make fair comparisons between the Bayesian decision-theoretic and frequentist power family sequential approaches, we have to first eliminate the confounding effects from the RAR procedures. Hence, we only compare the designs under similar adaptive randomization settings. Additionally, because of the trade-offs between the ASN and type I and II error rates, we are not able to make an overall conclusion based on merely one or two components of the OCs. One suggestion is to maintain the same error rates for the designs and compare their ASNs. However, that strategy does not reflect the relative importance of the different components of the OCs. Instead, we compare the resulting average losses (AVLs) from the Bayesian decision-theoretic perspective. Because the AVL takes into account not only the ASN and error rates, but also their relative importance, it is a reasonable score to use in ranking the methods.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

The first three rows in Table 2 show the results of the BAR_1, FAR_1 and FAR_2 designs, which apply the DBCD as the RAR procedure. Using the AVLs as comparison criteria, BAR_1 is the best of the three designs, as expected, since its aim is to minimize the AVL. It has the smallest average loss (245.90) compared to those of the other two frequentist methods (256.02 and 249.12) when the error costs are $K_1 = 1300 and K_2 = 2700$. The results also demonstrate the trade-offs among the OCs, as BAR_1 has smaller ASN but lower power than the FAR_1 or FAR_2 designs. The allocation of patients is slightly more balanced under the BAR_1 design compared to the FAR_1 design. The lower PBA (number of patients assigned to the better treatment arm) under BAR_1 is due to its higher probability of stopping a trial early, which limits the effects of the RAR procedure. The observation confirms that the comparisons based on a single OC are not sufficient to draw a conclusion about the overall performance of the design.

The middle section of Table 2 lists the results of a comparison of the BEQ, FEQ_1 and FEQ_2 designs. All three designs apply equal randomization throughout the trial. We implement the BEQ design by using a full backward induction algorithm. This design, which aims to minimize the AVL, has the smallest AVL and requires the lowest ASN under the alternative hypothesis.

A comparison of the resulting AVLs shows that, under the same adaptive randomization settings, the Bayesian decision-theoretic designs minimize the expected total losses, and therefore are preferable over the frequentist designs. As a consequence, the BAR designs are better than the FAR designs. Additionally, we favor the BAR design over the BEQ design because the BAR designs have the advantage of assigning more patients to the better treatment.

The bottom two rows of Table 2 list the results of a comparison of two additional robust designs, BAR_2 and BAR_3, which we based on the GWT adaptive randomization procedure.

Both the BAR_2 and BAR_3 designs are better than the equal randomization designs at attaining the required error rate, stopping the trials early at the cost of slightly increasing the ASN, and decreasing the required power. The BAR_2 design has larger power (0.845) and smaller ASN (105.16), but smaller PBA (55% on the better arm) compared with those respective numbers (0.815, 109.84, 62%) for the BAR_3 design under the alternative hypothesis. We conclude that the BAR_2 design operated more efficiently, but the BAR_3 design assigned more patients to the better treatment arm. The trade-off between trial efficiency and treating individual patients better in a clinical trial always exists. The trial investigators should choose the designs that accomplish the specific purposes of the trial. For instance, if the trial requires a small ASN, the BAR_2 design is desirable; whereas the BAR_3 design is desirable if the trial requires a higher PBA. Note that we give the AVLs of the two designs just for reference. The values are not directly comparable because the cost assignments are different for the designs.

In addition to the settings of fixed values for p_1 and p_2, we evaluate the performances of the designs under settings of varying values where $p_1 \in B_1$, $p_2 \in B_2$. The tuning parameters are chosen to satisfy the error rate requirements under the settings where $p_1 = 0.3$ and $p_2 = 0.3$ or 0.5. We use these settings to compare the BAR_1, BAR_2, BAR_3, FAR_1 and FAR_2designs.

Figure 3 shows the power curves obtained from all the designs by fixing p_1 at 0.2, 0.3, 0.4, and 0.5, but varying $\delta = p_2 - p_1$. The power curves increase as δ increases. The rates of increase are greater in the indifference region compared with the rates in the region where $p_2 - p_1 > \delta_0$. The designs satisfy the type I and II error rate requirements under all settings. The

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

results agree with those shown in the table, in which the BAR_3 design has relatively less power; whereas the power levels of the other designs are similar.

Figure 4 shows the average sample sizes (ASNs) under the different design settings. When the p_1 values are small ($p_1 = 0.2$), the BAR_2*and*BAR_3 designs have relatively smaller ASNs compared to those of the FAR_1, FAR_2, and BAR_1 designs. Due to a more imbalanced allocation, the BAR_3 design generally has a larger ASN than the BAR_2 design under the alternative hypothesis. However, when $p_1 = 0.4,0.5$, the BAR_1 design outperforms all the other designs. Overall, the BAR designs perform better than the FAR designs in every setting.

Figure 5 shows the percentages of patients assigned to the better arm (PBA) for the different designs. The BAR_3 design assigns the largest proportion of patients to the better treatment arm in most settings. The FAR_1 design, which uses the O'Brien-Fleming group sequential stopping boundary, also has a large PBA, especially under the alternative hypothesis as a result of the large ASNs. By contrast, the FAR_2 and BAR_1 designs generally have the smallest PBAs.

Figure 6 shows the standard deviations of the PBAs across the simulations. Based on these values, the BAR_1 design is the most stable, with the smallest standard deviation for the allocation ratio in every setting. Contrarily, the BAR_3 design generally has the largest standard deviation, followed by the BAR_2 design.

The comparisons above demonstrate that the BAR designs compare very favorably with the FAR designs. Considering each component of the OCs, we find that the best design is usually one of the BAR designs.

In most cases, p_1, the response rate of the control treatment arm, is an unknown parameter. To evaluate whether the designs are robust with respect to the values of p_1, we plot each component of the OCs across the p_1 values from 0.2 to 0.5 against p_2 in the supplemental figures S1, S2 and S3.

For the ASNs, the performances of the BAR_2, BAR_3, FAR_1 and FAR_2 designs are relatively invariant with respect to the changes in p_1. However, the performance of the BAR_1 design is affected by the value of p_1. The inconsistency in the results for the BAR_1 design is due to a combination of the Bayesian decision–theoretic sequential method and the DBCD adaptive randomization procedures. Recall that when computing the expected continuation loss, we consider the possible allocations of the next enrolled patient and their associated probabilities. Therefore, the stopping rules rely on the choice of the allocation procedure. Further, the allocation rates of the DBCD allocation procedure vary with the value of p_1 because the target ratio is a function of p_1. As a result, the stopping rules, and in turn the average sample sizes of the trials, are indirectly influenced by the value of p_1.

For the ASNs, the performances of the BAR_2, BAR_3, FAR_1 and FAR_2 designs are relatively invariant with respect to the changes in p_1. However, the performance of the BAR_1 design is affected by the value of p_1. The inconsistency in the results for the BAR_1 design is due to a combination of the Bayesian decision–theoretic sequential method and the DBCD adaptive randomization procedures. Recall that when computing the expected continuation loss, we consider the possible allocations of the next enrolled patient and their associated probabilities. Therefore, the stopping rules rely on the choice of the allocation procedure. Further, the allocation rates of the DBCD allocation procedure vary with the value of p_1 because the target ratio is a function of p_1. As a result, the stopping rules, and in turn the average sample sizes of the trials, are indirectly influenced by the value of p_1.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

As for the PBAs and their standard deviations, when δ is large, the PBAs of all the designs reduce to 0.5. This is because the trials in these settings would be terminated sooner following an equal randomization procedure. The BAR_1, FAR_1 and FAR_2 designs are not robust with respect to p_1 because the DBCD allocation rule is sensitive to changes in the value of p_1.

When comparing the operating characteristics as well as the robustness of the designs, Figures 3 to 6 show that the BAR designs outperform the FAR designs based on the OCs. The supplemental figures S1, S2 and S3 show that the GWT procedure surpasses the DBCD procedure from the perspective of robustness.

The above results also confirm that none of the BAR designs is uniformly better than the others based on every component of the OCs. Therefore, selecting among the BAR designs should be based on the specific trial considerations.

With a goal of controlling the ASN of a trial, the BAR_1 design is the best choice if the control arm (arm 1) has a known high response rate, e.g., $p_1 = 0.4 or 0.5$, as shown in Figure 4. This is because the resulting ASNs are the smallest among all the designs under these settings. However, the BAR_1 design is not robust with respect to p_1. Further, when p_1 is small, e.g., $p_1 = 0.2 or 0.3$, the ASNs of the BAR_1 design are larger than those of the BAR_2 and BAR_3 designs. Therefore, if p_1 is unknown or known to be small, we are in favor of selecting the BAR_2 design because it is robust and results in the smallest ASN when compared to the BAR_3, FAR_1 and FAR_2 designs in these settings.

When the goal is to improve the allocation of patients, it is desirable to have a design that assigns more patients to the better arm on average, but has low variability of allocating patients across the trial. Therefore, we consider both the PBA and the variability of PBA across the trial. According to the results from Figures 5 and 6, we exclude the BAR_1 design in the first place. Although the BAR_1 design has the smallest variability, it is the most balanced design, and has limited effect on improving patient allocation in the trials. Further, the BAR_1 design is not robust with respect to p_1. The $BAR_2 and BAR_3$ designs perform better than the BAR_1 design because they are robust with respect to p_1, and they are flexible, allowing for different PBAs and levels of variability by adjusting the c value.

Similarly, the choice between the BAR_2 and BAR_3 designs depends on the specific trial considerations. For example, if assigning more patients to the better treatment arm is the goal, the BAR_3 design is better, as shown in Figure 5. Conversely, if the ASN or the variability of PBA are more important, the BAR_2 design would be selected because it requires a smaller sample size and has lower risk of assigning patients to the inferior treatment arm, as shown in Figures 4 and 6.

5. Discussion

We propose a framework for clinical trial designs that combines a Bayesian decision–theoretic sequential method and response-adaptive randomization procedures. The goal of the design is to construct more efficient randomized phase II trials and assign more patients to better treatments. We describe the Bayesian decision problem and introduce the constrained backward induction and forward simulation methods to obtain reasonable and computationally feasible but suboptimal solutions for the problem under the adaptive randomization setting. The constrained backward induction method results in suboptimal decision rules by using the lower-dimensional summary statistic $\tilde{S}_t$ whose value might not uniquely represent a state. As a result, we can not obtain optimal solutions, even under extensive forward simulation. We can not compare solutions with an exact optimal solution

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

or even an improved suboptimal solution for lack of computationally feasible implementations. One could potentially compare the proposed designs with *k*-step look-ahead methods, a method used to evaluate expected losses under the Bayesian decision-theoretic setting [26].

In this study, we consider two-arm clinical trial designs with binary endpoints. We assume that the two arms are independent, an assumption that may be violated in some situations. Also, we use the average loss to compare the designs, acknowledging that this criterion favors the Bayesian decision-theoretic solution which is designed to minimize average loss, while the other methods do not.

Through simulation studies, we evaluate the performances of the designs by assessing their operating characteristic (OCs) in various scenarios. We first fix the allocation procedure and compare the average sample size and error rates between the Bayesian decision–theoretic and frequentist sequential approaches. We use the average loss as a weighted score of the average sample size and error rates to give an overall assessment of the performance of the designs. The results show that the Bayesian decision–theoretic approach results in a smaller average loss for the trial, as expected.

In terms of the average sample size, assignment of patients to the better treatment arm, and variability of patient assignment across the trial, our evaluations under different settings show that the BAR designs outperform the FAR designs in every component of the OCs. Further, the designs that use the GWT allocation procedures, the $BAR_2 and BAR_3$ designs, are more robust.

The results are favorable for the proposed BAR designs compared to the FAR designs. However, none of the BAR designs is uniformly better than the others. Therefore, the choice among the BAR designs should depend on the specific trial considerations. Moreover, the GWT procedure is more robust than the DBCD allocation procedures. Further, the BAR designs using the GWT procedure, the BAR_2 and BAR_3 designs, are amendable in order to achieve different PBAs and associated variables. Investigators can obtain a desirable BAR-GWT design by choosing a proper value of the parameter *c*.

The proposed method can be extended to design a trial with survival time as the endpoint. In this case, we have to specify the survival time distribution and carefully choose the summary statistic for the decision-making procedure. Further, our method could be extended to design a multi-arm clinical trial. Since the state space will be enlarged by considering multiple arms, we have to find another low-dimensional summary statistic. A higher-dimensional grid is necessary to evaluate the expected continuation losses. This may be a challenging task, but is feasible in practice using the forward simulation method.

The designs and evaluation procedures are implemented in R programming. The R code is available upon request.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported in part by grants CA16672 and CA97007 from the National Cancer Institute. The third author was partially supported by grant CA075981 from the National Cancer Institute. The authors thank Ms. LeeAnn Chastain for editorial assistance.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

A. The derivation of stopping and continuation losses

Terminal decision

According to the loss function and the probabilities of making erroneous decisions, we write the expected loss of decisions D_f,D_e as

$$\begin{aligned} E\left(\ell_{D_f}|s_t\right) &= K_1 Pr\left(\delta>\delta_0|S_t\right) \\ E\left(\ell_{D_e}|s_t\right) &= K_2 Pr\left(\delta<0|S_t\right). \end{aligned} \quad \text{(A.0.1)}$$

The probabilities of making erroneous decisions are based on the posterior probabilities of the difference in response rates. We can evaluate the value of $Pr(\delta > \delta_0|S_t)$ as

$$\begin{aligned} Pr\left(\delta>\delta_0|S_t\right) &= \int Pr\left(p_2>p_1+\delta_0|p_1,S_t\right) Pr\left(p_1|S_t\right) \mathrm{d}p_1 \\ &= \int_0^1 \int_{p_1+\delta_0}^1 \frac{p_2^{\alpha_{t2}-1}(1-p_2)^{\beta_{t2}-1}}{B(\alpha_{t2},\beta_{t2})} \mathrm{d}p_2 \frac{p_1^{\alpha_{t1}-1}(1-p_1)^{\beta_{t1}-1}}{B(\alpha_{t1},\beta_{t1})} \mathrm{d}p_1, \end{aligned}$$

where $B(\cdot)$ is the incomplete beta function. These values are computed based on the posterior parameters α_{ti}, β_{ti}, i = 1,2. In order to obtain the optimal solution for the decision problem, we choose the decision rule that selects the action with the smallest expected loss between D_e and D_f given the current state. Therefore, the terminal loss under the optimal decision $d_2=d_2^*$ is

$$L_{\left(D_s,d_2^*\right)}\left(S_t\right) = \min\left(E\left(\ell_{D_f}|S_t\right), E\left(\ell_{D_e}|S_t\right)\right). \quad \text{(A.0.2)}$$

The expectation is taken with respect to the unknown parameter δ in equation (2.2.2).

Loss of trial continuation

To obtain the loss associated with continuing the trial after stage t, we have to consider the cost of continuing to accrue patients, and the possible expected losses characterized by S_{t+1} after enrolling the patients. We assume only one patient is enrolled at each stage; thus the cost of continued accrual for one stage is the unit cost C of enrolling one patient. Without loss of generality, herein we set C = 1 because only the relative magnitude of the error costs and the cost of recruiting one additional patient matter.

Under the policy π, the loss of continuation L_{D_c} is

$$L_{\left(D_c,\pi_{t+1}\right)}\left(S_t\right) = C + E\left(L_{\pi_{t+1}}\left(S_{t+1}\right)|S_t\right), \quad \text{(A.0.3)}$$

where

$$E\left(L_\pi\left(S_{t+1}\right)|S_t\right) = \sum_s Pr\left(S_{t+1}=s|S_t\right) L_\pi\left(S_{t+1}=s\right), \quad \text{(A.0.4)}$$

and where the summation in equation (A.0.4) is over the state space for S_{t+1}. The equation shows that we can obtain the expected continuation loss at stage t once we have knowledge about the expected total loss at stage t + 1. The recursive equation gives the rationality of using the backward induction technique to evaluate the expectations.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

Optimality equation

The optimal policy is $\pi^* = (d^*(S_1), d^*(S_2), \ldots, d^*(S_T))$, the collection of optimal decisions at each stage. To simplify the notation, we write $\pi_t = d(S_t, t)$ and $\pi_t^* = d^*(S_t, t)$, where $\pi_t^* = \arg\min_{\pi_t} L_{\pi_t}(S_t)$. We write the optimal programming equation as

$$\begin{aligned} L_{\pi_t}(S_t) &= \min\left(L_{D_s}(S_t), L_{D_c}(S_t)\right) \\ &= \min_{\pi_t}\left(E\left(\ell_{D_f}|S_t\right), E\left(\ell_{D_e}|S_t\right), C + \sum_s Pr\left(S_{t+1}=s|S_t\right) L_{\pi_{t+1}^*}(s)\right). \end{aligned} \quad \text{(A.0.5)}$$

The optimal decision solution π^* is a sequence whose elements satisfy equation (A.0.5) at any time t [27].

B. The implementation of the forward simulation and constrained backward induction methods

Implementation of forward simulation and constrained backward induction

The forward simulation and constrained backward induction methods are based on the following assumptions. First, with intensive simulations, we are able to generate the sample paths with moderate and high occurrence probabilities. This ensures that the sample average can well approximate the true expectation, as the samples with low occurrence probabilities contribute little when computing the expectations. Second, the closeness of the observations at each analysis stage can be defined by the distance between $\tilde{S}_t = (\log(\nu_t), \mu_t)$. Then, we can create the grid at each stage by the quantiles of $\tilde{S}_t$. The observations that are close to each other are considered to be in the same group; i.e., we assume the observations that fall into the same grid share the same future outcomes. As a result, the constrained backward induction is conducted on the unit of a grid rather than the single observation. This reduces the computational complexity substantially.

Implementation of forward simulation

We generate $M = 20000$ simulated sample paths. For the k^{th} sample path, we first simulate hypothetical response rates $\theta_k = (p_{1k}, p_{2k})$ from the prior, and simulate the observations $y_{tk} = (y_{1k}, \ldots, y_{Tk})$ based on the assumed sampling model $p(y_{tk}|\theta_k)$. According to the current simulated observations, we assign the new patients to the treatment arms by the chosen allocation rules. By doing this, we generate an extensive list of hypothetical trial histories. Given each realization in a sample path, we update the posterior parameters accordingly. We create a grid on $\tilde{S}_{tk} = (\log(\nu_{tk}), \mu_{tk})$, $t = 1, \ldots, T$, $k = 1, \ldots M$ at each analysis stage according to the simulated quantiles of $(\log(\nu_{tk}), \mu_{tk})$, $k = 1, \ldots, M$, where $(\log(\nu_{tk}), \mu_{tk})$ are the logarithm of the posterior variances and means of the difference of the response rates at time t in the k^{th} simulated sample path. Note that the simulated sample paths depend on the response adaptive randomization procedures. Therefore, we fix the allocation rules when performing the forward simulation.

The resulting plots, shown in Figure 2, have clear patterns. Hence, we believe $M = 20000$ simulated sample paths are more than enough. The patterns of the plots are invariant with the analysis stages. Therefore, we conclude that the choice of M is independent of the maximum sample size N. This graphical heuristic method can be used in practice to check whether the simulated sample size is large enough. We do not typically need to continue to generate samples when the patterns of the stopping and continuation regions are clear.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

Implementation of constrained backward induction

We implement backward induction on the state space represented by $\tilde{S}_t$, which is constrained to be three-dimensional. Let s^*_{tk} be the realization of $\tilde{S}_{tk}$, and A_{th} be the subset of 1, … , M (M is the simulation time) created according to the quantiles of (log(ν_{tk}), μ_{tk}). At each stage, we can identify into which grid s^*_{tk}, k = 1, … , M falls. We obtain $M_{th} = |A_{th}|$ as the number of *sstk* that falls in grid A_{th}.

As shown before, we start the constrained backward induction procedure from the final stage. Assuming the maximum analysis stage of a trial is T, we obtain the expected stopping loss for every simulated scenario $\tilde{S}_{Tk}$ at this stage. We take the average of the expected stopping loss over the scenarios with index $k \in A_{th}$ as the expected stopping loss of the grid h, i.e., for the h^{th} grid at analysis stage T, we compute the expected stopping loss given $\tilde{S}_{Tk}=s^*_{Tk}$ where $k \in A_{Th}$ as

$$L_{\pi_T}\left(s^*_{Tk}\right)=1/M_{Th}\underset{l\in_{Th}}{A}\sum L_{\left(D_s,d_2\left(s^*_{Tl}\right)\right)}\left(s^*_{Tl}\right).$$

The procedure then goes backward. When $\tilde{S}_{tk}=s^*_{tk}$ where $k \in A_{th}$ at stage t, $t < T$, we compute the expected stopping loss as

$$L_{\pi_t}\left(s^*_{tk}\right)=1/M_{th}\underset{l\in_{th}}{A}\sum L_{\left(D_s,d_2\left(s^*_{sl}\right)\right)}\left(s^*_{tl}\right). \quad \text{(B.0.6)}$$

Note that we apply equations (2.2.2) and (A.0.1) without changes for computing the expected stopping losses of the decision $d_2 \in \{D_e, D_f\}$. In other words, the four-dimensional state S_t is used for calculating the loss of d_2 for each scenario; whereas $\tilde{S}_t$ is used to determine the grid into which the current state falls.

The expected continuation loss of a sample path k, k = 1, … , M at time t is the unit accrual cost C plus the expectation of the expected total loss of the realization in the same path at time t + 1. We average the expected continuation losses over the index set A_{th} to obtain the expected continuation loss of grid h. Hence, the loss of continuation given $\tilde{S}_{tk}=s^*_{tk}$ where $k \in A_{th}$ is computed as

$$L_{\left(D_c,\cdot\right)}\left(s^*_{tk}\right)=1/M_{th}\underset{l\in_{th}}{A}\sum L_{\pi_{t+1}}\left(s^*_{(t+1)l}\right)+C. \quad \text{(B.0.7)}$$

Equations (B.0.6) and (B.0.7) show that the scenarios within a grid assume the same expected losses. We compare the resulting expected continuation loss and stopping loss for each grid, and choose the decision π^{**}_t between stopping and continuing to minimize the loss of the grid.

We consider this procedure to be valid because equation (B.0.7) is well defined as the expected losses of the scenarios at stage t + 1 are obtained by time t. In addition, the method is efficient because we use replications of the simulated sample paths to approximate the expected continuation loss rather than exactly evaluating the expected losses over all the possible outcomes that follow a state. Using this strategy, we only have to evaluate a fixed M number of scenarios for each stage. This strategy avoids rapid increases in the number of scenarios, thereby making our approach computationally feasible. Note that for each state, the intensive simulation provides the numbers of different future realizations for a state that have high probability of occurrence and are thus adequate for a good approximation. This

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

procedure may miss some distant future realizations whose occurrence probabilities are small. However, after weighting by the occurrence probabilities, the contributions of these scenarios to the expected losses are negligible, and therefore have little influence on the final results.

Since our summary statistics are three-dimensional, we create separated grids instead of using a single grid for every stage, as shown in the application of a two-dimensional statistic by Mueller et al. [17]. Assuming the simulated dataset contains almost all of the possible outcomes of $\tilde{S}_t$, we use quantiles of μ_t and $\log(\nu_t)$ to create an $R \times Q$ grid for the stage t, $t = 1, \ldots, T$. We plot the forward simulation and constrained backward induction results for several stages in Figure 2 when $K_1/K_2 = 1300/2700$, and $M = 20000$. After observing the range of μ_t and $\log(\nu_t)$ and the densities of the points, we decide to choose $R = Q = 30$. This produces grids with a horizontal interval length of around 0.06-0.1 and a vertical interval length of 0.06. We also try different R and Q values, ranging from 20–30. They lead to similar results. In practice, the grids can be chosen by observing the results from the forward simulation procedures. In principle, the chosen grid should be fine enough to distinguish the realizations with large differences, and each grid should be wide enough to contain sufficient samples.

During the trial, at stage t, we compute the posterior mean μ_t^* and log variance $\log(\nu_t^*)$ based on the realized observations, and find the grid into which it falls. We then make a decision based on the suboptimal decision for the grid. Problems occurs when the observed posterior log variance $\log(\nu_t^*)$ exceeds the maximum or minimum values in the simulated dataset at stage t. We use linear predictions to address this issue. According to the grid points on the horizontal line (on $\log(\nu_t)$), denoted by ν^{qt}, $q = 1, \ldots, Q$, we identify, in each interval $(\nu^{qt}, \nu^{(q+1)t})$, the minimum (maximum) value of μ_t in the upper (lower) stopping regions and the corresponding values of $\log(\nu_t)$ as the upper (lower) critical points denoted by $\left(\log(\nu)_u^{qt}, \mu_u^{qt}\right)$ $\left(\left(\log(\nu)_l^{qt}, \mu_l^{qt}\right)\right)$. Then, by using the method of least squares, we form upper and lower linear boundaries at each stage, which are the linear functions of $log(\nu_t)$ denoted by f_u, f_l. Once an observed value of $\log(\nu_t^*)$ falls into a grid range, we compare the value of this μ_t^* with the boundaries $f_u(\log(\nu_t^*))$, $f_l(\log(\nu_t^*))$ to make the decision to stop or continue the trial.

References

1. Jennison, C.; Turnbull, BW. Group Sequential Methods with Applications to Clinical Trials. CRC Press Inc.; Boca Raton, FL: 2000.
2. Hu, F.; Rosenberger, WF. The Theory of Response-Adaptive Randomization in Clinical Trials. Wiley-Interscience; 2006.
3. Pocock SJ. Group sequential methods in the design and analysis of clinical trials. Biometrika. 1977; 64(2):191–199. URL http://www.jstor.org/stable/2335684.
4. O'Brien PC, Fleming TR. A multiple testing procedure for clinical trials. Biometrics. 1979; 35(3): 549–556. URL http://www.jstor.org/stable/2530245. [PubMed: 497341]
5. Lan KKG, DeMets DL. Discrete sequential boundaries for clinical trials. Biometrika. 1983; 70(3): 659–663. URL http://www.jstor.org/stable/2336502.
6. Rosenberger WF. Randomized play-the-winner clinical trials: review and recommendations. Control Clin Trials. Aug; 1999 20(4):328–342. [PubMed: 10440560]
7. Spiegelhalter, DJ.; Abrams, KR.; Myles, JP. Bayesian approaches to clinical trials and health-care evaluation. 1. Wiley; 2004.
8. Berry DA. Bayesian statistics and the efficiency and ethics of clinical trials. Statistical Science. 2004; 19(1):175–187.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

9. Holland JFT. Statistical innovations in cancer research. Cancer Medicine. 2003:411–425. chap. BC Decker. ch. 33.
10. Lee JJ, Liu DD. A predictive probability design for phase ii cancer clinical trials. Clin Trials. 2008; 5(2):93–106. doi:10.1177/1740774508089279. URL http://dx.doi.org/10.1177/1740774508089279. [PubMed: 18375647]
11. Thall PF, Wathen JK. Practical bayesian adaptive randomisation in clinical trials. European Journal of Cancer. 2007; 43(5):859–866. doi:DOI:10.1016/j.ejca.2007.01.006. URL http://www.sciencedirect.com/science/article/B6T68-4N2M6C7-1/2/a1d9a187f79022b0361e179bf70853c9. [PubMed: 17306975]
12. Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika. 1933; 25(3/4):285–294. URL http://www.jstor.org/stable/2332286.
13. Berry DA, Eick SG. Adaptive assignment versus balanced randomization in clinical trials: A decision analysis. Statistics in Medicine. 1995; 14:231–246. [PubMed: 7724909]
14. Lewis RJ, Berry DA. Group sequential clinical trials: A classical evaluation of bayesian decision-theoretic designs. Journal of the American Statistical Association. 1994; 89(428):1528–1534. URL http://www.jstor.org/stable/2291016.
15. Lehmann, EL.; Romano, JP. 3rd. Springer; 2005. Testing Statistical Hypotheses.
16. Müller P, Berry DA, Grieve AP, Krams M. A bayesian decision-theoretic dose-finding trial. Decision Analysis. Dec.2006 3:197–207. doi:10.1287/deca.1060.0079. URL http://dl.acm.org/citation.cfm?id=1235066.1235068.
17. Mueller P, Berry DA, Grieve AP, Smith M, Krams M. Simulation-based sequential bayesian design. Journal of Statistical Planning and Inference. 2007; 137(10):3140–3150. doi:DOI:10.1016/j.jspi.2006.05.021. Special Issue: Bayesian Inference for Stochastic Processes.
18. Morgan CC, Coad DS. A comparison of adaptive allocation rules for group-sequential binary response clinical trials. Stat Med. Apr; 2007 26(9):1937–1954. doi:10.1002/sim.2693. URL http://dx.doi.org/10.1002/sim.2693. [PubMed: 16981177]
19. Gausche M, Lewis RJ, Stratton SJ, Haynes BE, Gunter CS, Goodrich SM, Poore PD, McCollough MD, Henderson DP, Pratt FD, et al. Effect of out-of-hospital pediatric endotracheal intubation on survival and neurological outcome: a controlled clinical trial. JAMA: The Journal Of The American Medical Association. 2000; 283(6):783–790. URL http://ezproxy.rice.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=cmedm&AN=10683058&site=eds-live&scope=site. [PubMed: 10683058]
20. Young K, Okada P, Sokolove P, Palchak M, Panacek E, Baren J, Huff K, McBride D, Inkelis S, Lewis R. A randomized, double-blinded, placebo-controlled trial of phenytoin for the prevention of early posttraumatic seizures in children with moderate to severe blunt head injury. ANNALS OF EMERGENCY MEDICINE. 2004; 43(4):435–446. URL http://ezproxy.rice.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=edswsc&AN=000220501300001&site=eds-live&scope=site. [PubMed: 15039684]
21. Hu F, Rosenberger WF. Optimality, variability, power: Evaluating response-adaptive randomization procedures for treatment comparisons. Journal of the American Statistical Association. 2003; 98(463):671–678. URL http://www.jstor.org/stable/30045294.
22. Berry DA, Eick SG. Adaptive assignment versus balanced randomization in clinical trials: A decision analysis. Statistics in Medicine. 1995; 14:231–246. [PubMed: 7724909]
23. Cheng Y, Shen Y. Bayesian adaptive designs for clinical trials. Biometrika. 2005; 92(3):633–646. doi:10.1093/biomet/92.3.633. URL http://biomet.oxfordjournals.org/cgi/content/abstract/92/3/633.
24. Brockwell AE, Kadane JB. A gridding method for bayesian sequential decision problems. Journal of Computational and Graphical Statistics. 2003; 12(3):566–584. URL http://www.jstor.org/stable/1391039.
25. Carlin BP, Kadane JB. Approaches for optimal sequential decision analysis in clinical trials. Biometrics. Sep; 1998 54(3):964–975. [PubMed: 9750245]
26. Berger, JO. Statistical Decision Theory and Bayesian Analysis. 2nd. Springer; 1993.

27. Ross, SM. Introduction to stochastic dynamic programming / Sheldon Ross. Academic Press; New York: 1983.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

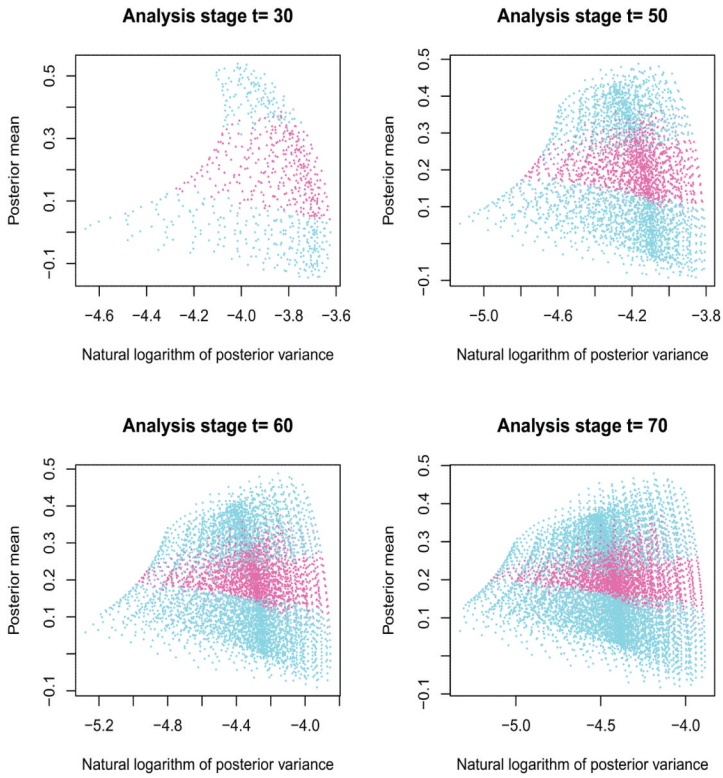

Figure 1.

Backward induction results: middle region–continuing region; lower region –stopping for futility; upper region –stopping for efficacy The results of exact backward induction. The results are based on the four-dimensional state space; the expected losses are calculated by taking into account all future possible outcomes. The axes represent the posterior means and logarithm of the posterior variances of the difference of the response rates

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

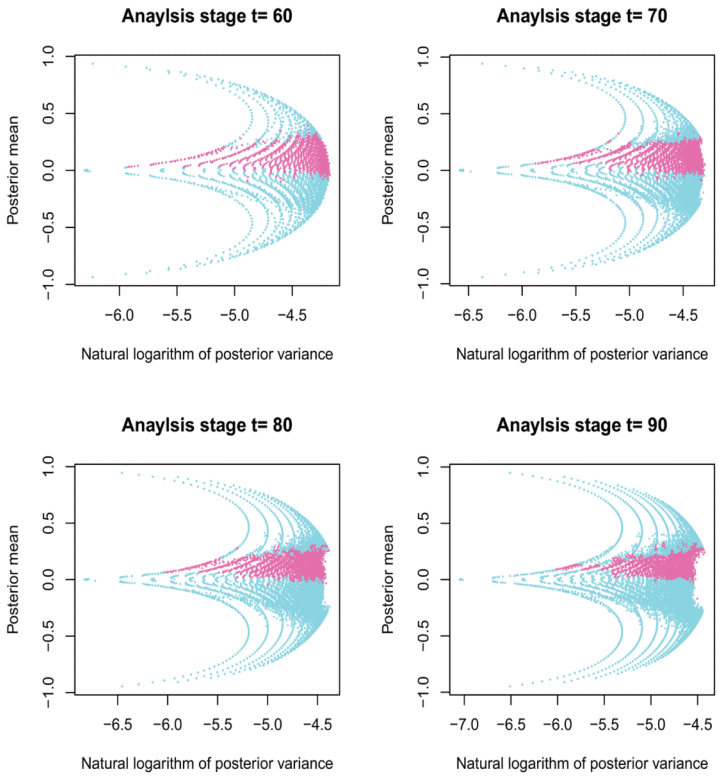

Figure 2.
Backward induction results: middle region–continuing region; lower region–stopping for futility; upper region–stopping for efficacy The results of the forward simulation and constrained backward induction method. The results are based on the simulated sample paths and the lowerdimensional state space. The axes represent the posterior means and logarithm of posterior variances of the difference of the response rates

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

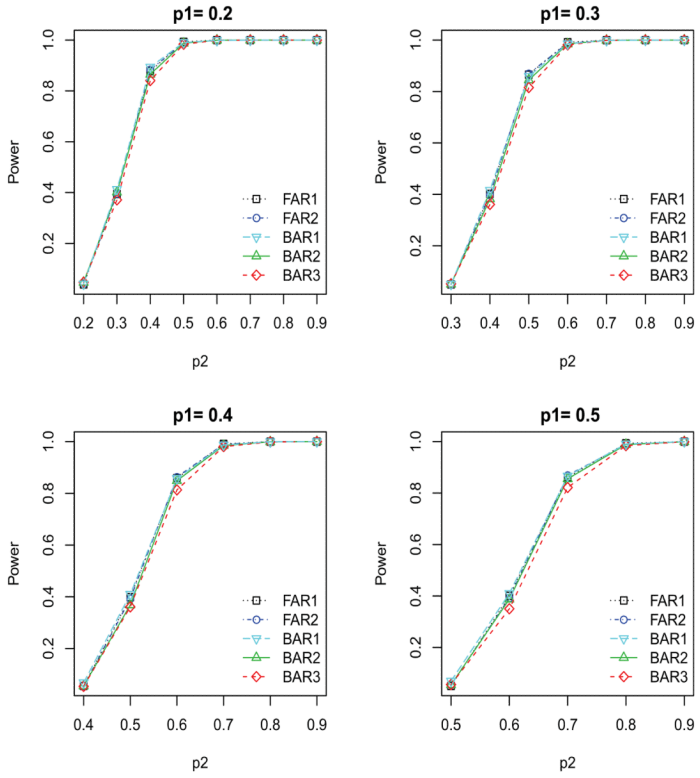

Figure 3.
Comparison of statistical powers.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

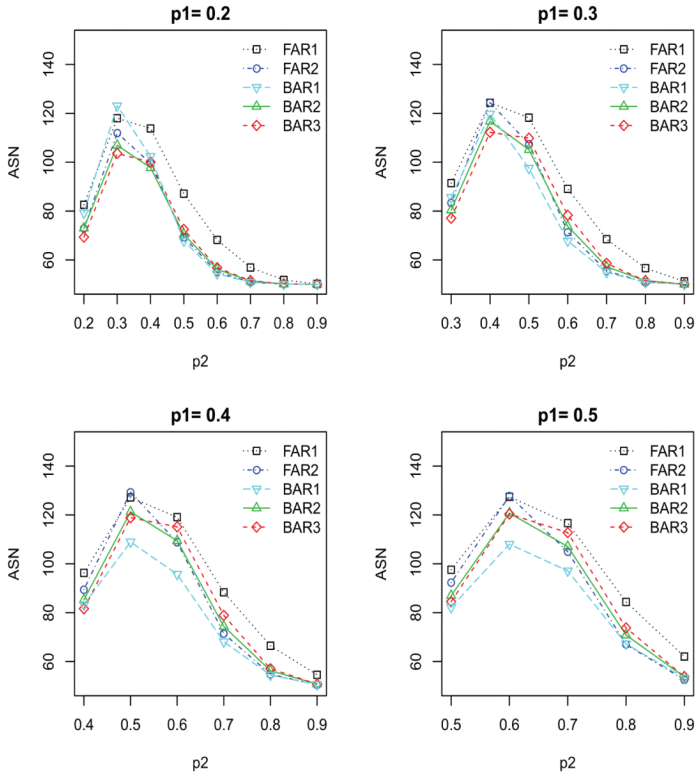

Figure 4.
Comparisons of average sample sizes (ASNs).

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

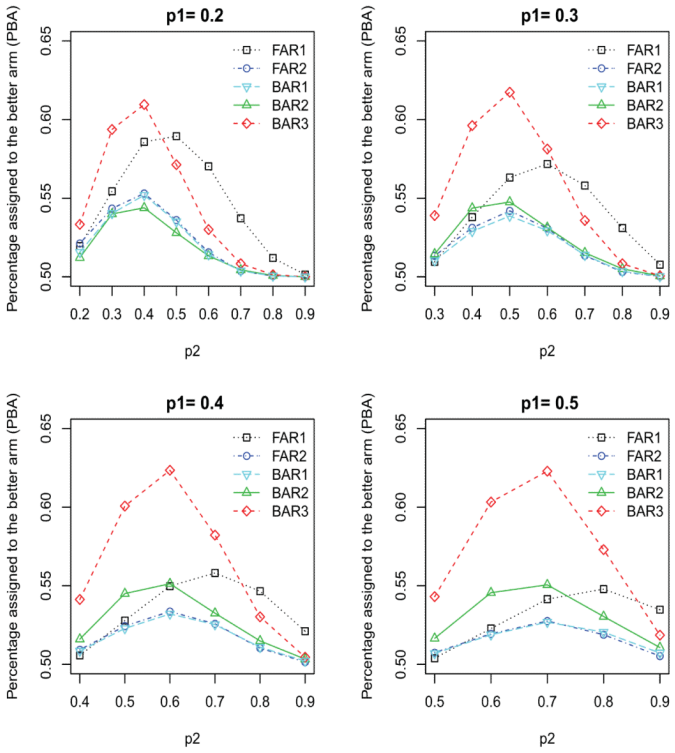

Figure 5.
Percentages of patients assigned to the better treatment arm.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

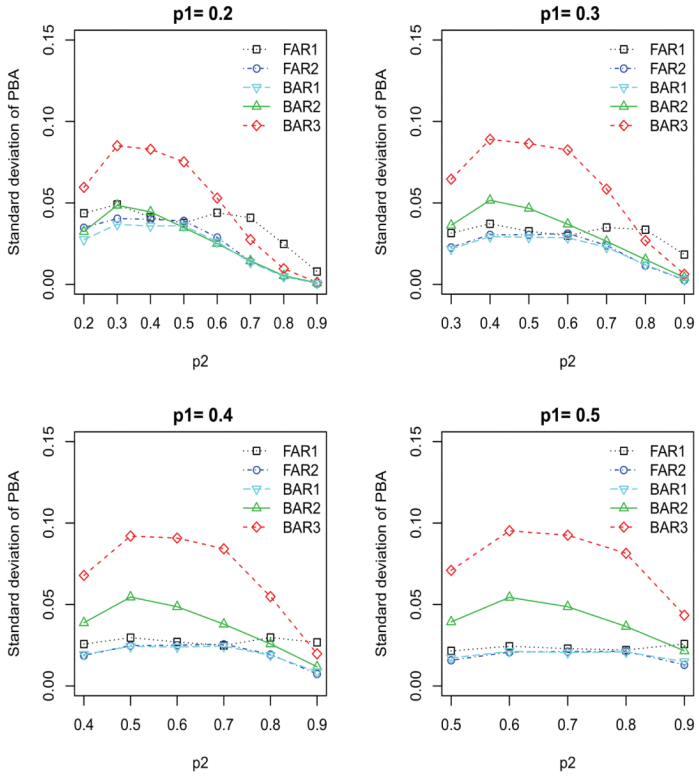

Figure 6.
Standard deviation of allocation ratio.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

Table 1

Comparison of the operating characteristics of BAR designs with different costs K_1, K_2 and $c = t/2T$. $p_1 = p_2 = 0.3$ under the null and $p_1 = 0.3$, $p_2 = 0.5$ under the alternative. Type I error rate (α), average sample size under the null hypothesis $\left(\bar{n}_\alpha\right)$, power ($1 - \beta$), average sample size under the alternative hypothesis $\left(\bar{n}_\beta\right)$, expected number of patients on arm_2 ($E(n_2/n)$) are recorded. The results are from 10,000 simulation runs.

		Operating Characteristics					
		Under the null			Under the alternative		
Design	K_1/K_2	α	$\bar{n}_\alpha$	$E(n_2/n)$	$1-\beta$	$\bar{n}_\beta$	$E(n_2/n)$
BAR_a	1000/1500	0.077	76.08	0.51	0.853	86.19	0.52
BAR_b	1000/5000	0.041	77.03	0.51	0.813	109.68	0.56
BAR_c	1000/8000	0.031	77.24	0.52	0.800	119.98	0.58
BAR_d	500/1500	0.064	65.84	0.51	0.779	79.21	0.52
BAR_e	1000/3000	0.053	77.11	0.51	0.830	100.42	0.54
BAR_f	2000/6000	0.043	90.33	0.52	0.875	122.76	0.57

NIH-PA Author Manuscript NIH-PA Author Manuscript NIH-PA Author Manuscript

Table 2

Comparing BAR designs with FAR designs. FAR designs with different values of γ_1 and γ_2. $p_1 = p_2 = 0.3$ under the null and $p_1 = 0.3$, $p_2 = 0.5$ under the alternative. Type I error rate (α), average sample size under the null hypothesis $\left(\bar{n}_\alpha\right)$, power ($1 - \beta$), average sample size under the alternative hypothesis $\left(\bar{n}_\beta\right)$, expected number of patients on arm_2 ($E(n_2/n)$), and the average loss (AVL) for the setting where $p_1 \in A_1$ and $p_2 \in A_2$ are recorded. ξ in equation 2.1.2 is chosen to be 10. The results are from 10,000 simulation runs. The BAR designs have better performance in reducing AVL on average.

					Operating Characteristics						
					Under the null			**Under the alternative**			
Design	***c***	**Δ**	K_1/K_2	λ_1/λ_2	α	$\bar{n}_\alpha$	$E(n_2/n)$	$1-\beta$	$\bar{n}_\beta$	$E(n_2/n)$	**AVL**
BAR_1	/		1300/2700		0.051	85.59	0.51	0.859	97.52	0.54	245.90
FAR_1		0		1.53/1.15	0.050	91.43	0.51	0.866	118.22	0.56	256.02
FAR_2		0.5		2.38/1.95	0.050	83.42	0.51	0.870	107.43	0.54	249.12
BEQ	/		4500/2000		0.050	83.12	0.50	0.864	104.14	0.50	343.79
FEQ_1	/	0		1.15/1.13	0.050	90.85	0.50	0.857	115.18	0.50	349.30
FEQ_2	/	0.5		2.42/1.90	0.049	82.63	0.50	0.856	110.82	0.50	351.38
BAR_2	*t*/2*T*		1200/3500		0.051	80.44	0.51	0.845	105.16	0.55	249.07
BAR_3	1/2		1000/4000		0.052	77.10	0.54	0.815	109.84	0.62	239.04

Stat Med. Author manuscript; available in PMC 2014 May 30.