Published in final edited form as:

J Stat Plan Inference. 2012 April 1; 142(4): 944–955. doi:10.1016/j.jspi.2011.10.016.

A Hybrid Geometric Phase II/III Clinical Trial Design based on Treatment Failure Time and Toxicity

Peter F. Thall^{1,*}, Hoang Q. Nguyen¹, Xuemei Wang¹, and Johannes E. Wolff²

¹Department of Biostatistics, M.D. Anderson Cancer Center, Houston, Texas, USA

²Department of Pediatrics, Tufts Medical Center, Boston, Massachusetts, USA

Abstract

The problem of comparing several experimental treatments to a standard arises frequently in medical research. Various multi-stage randomized phase II/III designs have been proposed that select one or more promising experimental treatments and compare them to the standard while controlling overall Type I and Type II error rates. This paper addresses phase II/III settings where the joint goals are to increase the average time to treatment failure and control the probability of toxicity while accounting for patient heterogeneity. We are motivated by the desire to construct a feasible design for a trial of four chemotherapy combinations for treating a family of rare pediatric brain tumors. We present a hybrid two-stage design based on two-dimensional treatment effect parameters. A targeted parameter set is constructed from elicited parameter pairs considered to be equally desirable. Bayesian regression models for failure time and the probability of toxicity as functions of treatment and prognostic covariates are used to define two-dimensional covariate-adjusted treatment effect parameter sets. Decisions at each stage of the trial are based on the ratio of posterior probabilities of the alternative and null covariate-adjusted parameter sets. Design parameters are chosen to minimize expected sample size subject to frequentist error constraints. The design is illustrated by application to the brain tumor trial design.

Keywords

Bayesian design; clinical trial; phase II/III design; treatment selection

1. Introduction

The problem of comparing several experimental treatments, $E_1, \dots, E_J$, to a standard treatment, S , is common in clinical trials. Reviews are given by Simon, Thall and Ellenberg (1994), Rubinstein, et al. (2004), Bretz, Schmidli, Konig, et al. (2006), Jennison and Turnbull (2006), and Thall (2008). For a one-dimensional outcome, the focus is on $J + 1$ parameters, $\theta = (\theta_0, \theta_1, \dots, \theta_J)$, where θ_0 corresponds to S and θ_j corresponds to E_j for $j = 1, \dots, J$. For example, each entry of θ may be a probability of tumor response, or median survival time.

© 2011 Elsevier B.V. All rights reserved.

*Corresponding author, rex@mdanderson.org.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

In this setting, Thall, Simon and Ellenberg (TSE) (1988) proposed a 2-stage phase II/III design for binary outcomes combining ideas from selection (cf. Bechhofer, Santner, and Goldsman, 1995) and group sequential testing (Pocock, 1977; O'Brien and Fleming, 1979; Lan and DeMets, 1983; Wang and Tsatis, 1987; Jennison and Turnbull, 2000). The TSE design randomizes patients among $E_1, \dots, E_J$ and S in stage 1, and proceeds to stage 2 only if the empirically best experimental treatment in stage 1, E_v , is promising compared to S . In stage 2, patients are randomized between E_v and S and a final test is conducted to decide whether E_v provides an improvement over S . Design parameters are chosen to minimize expected sample size subject to constraints on overall Type I and Type II error. Schaid, Wieand and Therneau (1990) proposed a similar 2-stage design for time-to-event outcomes that allows more than one E_j to move forward to stage 2, allows termination of the trial in stage 1 for either futility or superiority, and does pairwise comparisons with the possibility of concluding that more than one E_j provides an improvement over S . Many extensions of these designs have been proposed, including designs with more than two stages (Stallard and Todd, 2003; Stallard and Friede, 2008), a design that continues accrual between stages and uses the stage 1 data to determine the stage 2 sample size adaptively (Lin and Pledger, 2005), and an algorithm for computing decision boundaries (Cheung, 2008).

We address the problem of selecting among $E_1, \dots, E_J$ and comparing the selected treatment to S based on (T, Y) , where T is survival or event-free survival (EFS) time and Y indicates a nonfatal severe adverse event (SAE). Thus, death is accounted for by T and not by Y . Our motivation was the desire to provide a feasible design for a trial of four chemotherapy combinations for treating pediatric patients with choroid plexus tumors (CPTs). CPTs include three histological subtypes: choroid plexus carcinoma (CPC), atypical plexus papilloma (APP) and choroid plexus papilloma (CPP). A typical CPT patient is a very young child, with 61% less than three years of age. Because these tumors are too rare for standard protocol designs, to date no clinical trial specific for CPTs has been completed (Wolff et al., 2002). In the planned trial, the standard treatment is S = carboplatin + cyclophosphamide + etoposide + vincristine, and the three experimental treatments are E_1 = doxorubicin + cisplatin + actinomycin + etoposide, E_2 = high dose methotrexate and E_3 = temozolomide + irinotecan. The outcomes are T = EFS time, with an event defined as disease progression or death, and Y = I (severe toxicity). The baseline covariates are Age and indicators of metastatic disease, whether surgery achieved a complete resection (CR), and the unfavorable CPC histology. The goal is to increase T on average while controlling the probability of toxicity, compared to standard therapy, while also accounting for patient heterogeneity.

We propose a hybrid two-stage select-and-test design based on (T, Y) and baseline covariates. The logic of our design mimics that of the TSE design, with design parameters chosen to minimize expected sample size subject to overall Type I and Type II error constraints. The main differences between our procedure and the TSE design are that (i) we account for patient heterogeneity, (ii) we base the tests on 2-dimensional covariate-adjusted treatment effect parameters, and (iii) we take a Bayesian geometric approach in which test statistics are constructed from posterior probabilities of 2-dimensional parameter sets, rather than using conventional Z-scores.

2. Overview of the methodology

2.1 The Basic Idea

To establish the ideas underlying the geometric construction, for simplicity we temporarily consider the special case of two treatments, with $j = 0$ and 1 for S and E , we and ignore covariates. The parameters $\theta_{Y,j}$ and $\theta_{T,j}$ represent average behavior of two different outcomes. In the CPT trial, the physician decided to use $\theta_{Y,j}$ = probability of toxicity and $\theta_{T,j}$

$= \Pr(T > 24 \text{ months})$ with treatment j as the parameters, where T = time to treatment failure, defined as disease progression or death. Alternatively, the physician might use the mean or median of T as $\theta_{T,j}$. Thinking of $\theta_j = (\theta_{Y,j}, \theta_{T,j})$ as a two-dimensional parameter that describes the average behavior of the outcomes (Y, T) with treatment j , the problem of comparing E to S becomes that of comparing θ_1 to θ_0 .

In such settings, there are many methods for dimension reduction to obtain a one-dimensional decision criterion. Our approach exploits a Bayesian formulation and subjective goals of the physician, using the following construction. Let $\mu_0 = (\mu_{Y,0}, \mu_{T,0})$ be the prior mean of θ_0 based on experience with S . We elicit a set of fixed target pairs, $\{\theta^{e,1}, \dots, \theta^{e,m}\}$, for θ_E that the physician considers equally desirable improvements over μ_0 . We then form the polygonal line $\mathcal{L}$ connecting the elicited values, and define the target parameter set, $\Theta_{\mathcal{L}}$, of all θ pairs at least as desirable as a pair on $\mathcal{L}$. The test is based on the posterior probability of $\Theta_{\mathcal{L}}$ compared to the posterior probability of the null set Θ_0 where $\theta_Y \geq \mu_{Y,0}$ and $\theta_T \leq \mu_{T,0}$. Accounting for covariate effects requires the additional steps of fitting regression models, obtaining covariate-adjusted treatment effects, and transforming Θ_0 and $\Theta_{\mathcal{L}}$. This adds technical complications to the construction, but the essential idea is still to base tests on posterior probabilities of two-dimensional parameter sets. For the CPT trial, the target line and target parameter set in terms of $(\theta_{Y,j}, \theta_{T,j})$ are given in Figure 1. In other settings with two outcomes, this construction may be carried out in the same way, but using different parameters. For example, in a phase I/II dose-finding trial based on indicators of toxicity and response, $\theta_{T,j}$ would be replaced by the response probability.

2.2 Model and parameters

We define the distribution of T conditional on Y to allow the possibility that toxicity may affect the event time. Let ξ denote the vector of model parameters. For a patient with baseline covariates $\mathbf{Z} = (Z_1, \dots, Z_q)$ given treatment $j = 0, 1, \dots, J$, denote the probability of toxicity by $\theta_{Y,j}(\mathbf{Z}, \xi) = \Pr(Y = 1 | j, \mathbf{Z}, \xi)$ and the conditional probability density function and survivor function of T given $\mathbf{Z}$ and Y by $f_j(t | \mathbf{Z}, Y, \xi)$ and $\bar{F}_j(t | \mathbf{Z}, Y, \xi) = \Pr(T > t | j, \mathbf{Z}, Y, \xi)$. Denoting the time to the event or independent right censoring by T° and $\varepsilon = I(T^\circ = T)$, the likelihood function of the observed outcome $(T^\circ, \varepsilon | Y)$ given $\mathbf{Z}$ and j is

$$\mathcal{L}(T^\circ, \varepsilon, Y | j, \mathbf{Z}, \xi) = f_j(T^\circ | \mathbf{Z}, Y, \xi)^\varepsilon \bar{F}_j(T^\circ | \mathbf{Z}, Y, \xi)^{1-\varepsilon} \theta_{Y,j}(\mathbf{Z}, \xi)^Y \{1 - \theta_{Y,j}(\mathbf{Z}, \xi)\}^{1-Y}. \quad (1)$$

Let $\theta_{T,j}(\mathbf{Z}, Y, \xi)$ be a parametric function quantifying the effect of treatment j on T given Y and $\mathbf{Z}$. This may be a mean, a median, or $F_j(t^* | \mathbf{Z}, Y, \xi)$ for fixed time t^* . Thus, $\theta_j(\mathbf{Z}, Y, \xi) = (\theta_{T,j}(\mathbf{Z}, Y, \xi), \theta_{Y,j}(\mathbf{Z}, \xi))$ is a 2-dimensional parametric function. The treatment effect parameters that are basis of the test are defined as follows. Let $\mathbf{Z}_Y$ and $\mathbf{Z}_T$ be vectors of variables obtained from $\mathbf{Z}$, of dimensions q_Y and q_T . We assume parametric regression models under which $\theta_{T,j}(\mathbf{Z}, Y, \xi)$ is a function of the linear term

$$\eta_{T,j}(\mathbf{Z}_T, Y, \xi) = \gamma_{T,j} + \beta_{T,1} Z_{T,1} + \dots + \beta_{T,q_T} Z_{T,q_T} + \beta_{T,q_T+1} Y = \gamma_{T,j} + \beta_T \mathbf{Z}_T + \beta_{T,q_T+1} Y, \quad (2)$$

and $\theta_{Y,j}(\mathbf{Z}, \xi)$ is a function of the linear term

$$\eta_{Y,j}(\mathbf{Z}_Y, \xi) = \gamma_{Y,j} + \beta_{Y,1} Z_{Y,1} + \dots + \beta_{Y,q_Y} Z_{Y,q_Y} = \gamma_{Y,j} + \beta_Y \mathbf{Z}_Y, \quad (3)$$

where β_Y and β_T are covariate effects. Expression (2) specifies the conditional distribution of $[T|j, Y, \mathbf{Z}]$, with $\beta_{T,qT+1}$ the effect of toxicity on the distribution of T . Expression (3) specifies the marginal distribution of $[Y|j, \mathbf{Z}]$. The parameter $\gamma_{T,j}$ is the effect of treatment j on T after adjusting for Y and $\mathbf{Z}$, and $\gamma_{Y,j}$ is the effect of treatment j on Y after adjusting for $\mathbf{Z}$. Thus, the parameter vector $\gamma = (\gamma_0, \gamma_1, \dots, \gamma_J)$ may be considered a $2(J+1)$ -dimensional, covariate-adjusted generalization of the $(J+1)$ -dimensional parameter vector $\theta = (\theta_0, \theta_1, \dots, \theta_J)$ given earlier in the univariate setting. For each E_j , we define the 2-dimensional E_j -versus- S treatment effect $\delta_j = \gamma_j - \gamma_0 = (\gamma_{T,j} - \gamma_{T,0}, \gamma_{Y,j} - \gamma_{Y,0}) = (\delta_{T,j}, \delta_{Y,j})$, so that larger $\delta_{T,j}$ and smaller $\delta_{Y,j}$ are more desirable. The test statistics will be defined in terms of $2J$ -dimensional parameter vector $\delta = (\delta_1, \dots, \delta_J)$.

2.3. Steps of the construction

We first list the steps required to construct the test statistics. Additional details are given Sections 3 and 4. Let Θ denote the set of all possible values taken on by the $\theta_j(\mathbf{Z}, Y, \xi)$ pairs, and denote the set of all possible δ_j 's by Δ . Since $\theta_{T,j}(\mathbf{Z}, Y, \xi)$ is defined conditional on the toxicity outcome Y , in the following construction targets are elicited for each value $y^* = 0, 1$ of Y and averaged to obtain a set of targets that do not depend on Y .

1. Ask the physician to specify a reference covariate vector $\mathbf{z}^*$. For each value $y^* = 0, 1$ of Y , denote the prior mean of $\theta_0(\mathbf{z}^*, y^*, \xi)$ by $\mu_{\theta,0}(y^*) = (\mu_{\theta,T,0}(y^*), \mu_{\theta,Y,0}(y^*))$. For a fixed toxicity probability θ_y^* , denote $\mu_{\theta,T,0} = \mu_{\theta,T,0}(1)\theta_y^* + \mu_{\theta,T,0}(0)(1 - \theta_y^*)$ and $\mu_{\theta,0} = (\mu_{\theta,T,0}, \mu_{\theta,Y,0})$.
2. For each $y^* = 0$ or 1 , elicit a set of target θ pairs $\{\theta^{(e,1)}(y^*), \dots, \theta^{(e,m)}(y^*)\}$ in Θ from the physician that correspond to equally desirable improvements over $\mu_{\theta,0}(y^*)$. The two sets should be elicited so that $\theta^{(e,r)}(1)$ corresponds to $\theta^{(e,r)}(0)$ for each $r = 1, \dots, m$. Obtain a single set of mean target pairs by forming the weighted averages $\theta^{(e,r)} = \theta^{(e,r)}(1)\theta_y^* + \theta^{(e,r)}(0)(1 - \theta_y^*)$, for $r = 1, \dots, m$. Construct the polygonal target line, $\mathcal{L}$ that connects the mean elicited target pairs $\bar{\theta}^e = (\theta^{(e,1)}, \dots, \theta^{(e,m)})$.
3. The target set $\Theta_{\mathcal{L}} = \bigcup_{\theta' \in \mathcal{L}} \{\theta \in \Theta : \theta_T \geq \theta'_T \text{ and } \theta_Y \leq \theta'_Y\}$ consists of all θ pairs at least as desirable as a pair on $\mathcal{L}$. Figure 1 illustrates $\mu_{\theta,0}$, $\mathcal{L}$, and $\Theta_{\mathcal{L}}$ for the CPT trial.
4. In Δ , denote the line obtained from $\mathcal{L}$ by $\mathcal{D}_{\mathcal{L}}$, the 2-dimensional alternative treatment effect parameter set obtained from $\Theta_{\mathcal{L}}$ by $\Delta_{\mathcal{L}}$, and the null set of all $\delta = (\delta_T, \delta_Y)$ in Δ with $\delta_T \leq 0$ and $\delta_Y \geq 0$ by Δ_0 . Figure 2 illustrates $\mathcal{D}_{\mathcal{L}}$, Δ_0 , and $\Delta_{\mathcal{L}}$ for the CPT trial.
5. The test statistics used in each stage of the trial are the posterior probability ratios

$$R_j(data) = \frac{\Pr(\delta_j \in \Delta_{\mathcal{L}} | data)}{\Pr(\delta_j \in \Delta_0 | data)}, \quad j=1, \dots, J. \quad (4)$$

Under the construction, larger $R_j(data)$ corresponds to greater superiority of E_j over S .

2.4 Model and Parameters for the CPT Trial

The CPT trial design initially was motivated by the physician's desire to use (Y, T) as the outcome, and to study several experimental regimens. The physician chose to use the probability of toxicity and $\Pr(T > 24 \text{ months})$ as the reference parameters, and considered T to be independent of Y , so the formulation was simpler than that given above. Specifically, $\beta_{T,qT+1} \equiv 0$ in (2), and only one set of target pairs was elicited in step 2 of the construction. The covariates used in the CPT trial, given in Table 1, are well-known to be related to event-free survival time for this disease and were specified by the physician. Denoting $I_{\pm}(A) = +1$

if event A occurs and -1 if not, to obtain $\eta_{T,j}(\mathbf{Z}_T, Y, \xi)$ and thus $\gamma_{T,j}$, we defined the four binary covariates $Z_{T,1} = I(\text{Age} \geq 3)$, $Z_{T,2} = I(\text{metastatic disease})$, $Z_{T,3} = I(\text{surgery achieved a CR})$, and $Z_{T,4} = I(\text{histology} = \text{CPC})$. Using these, the physician specified the reference patient to be 2 years old, with non-metastatic disease, surgical CR, and CPC histology, denoted by $\mathbf{z}^* = (-1, -1, +1, +1)$. In contrast, the probability of toxicity only varied with Age, and was modeled using the logistic fractional polynomial function $\eta_{Y,j} = \text{logit}\{\theta_{Y,j}(\text{Age}, \xi)\} = \gamma_{Y,j} + \beta_{Y,1}\text{Age}^{1/2} + \beta_{Y,2}\log(\text{Age})$. Details of how this function was determined from elicited values are provided in Section 5.

While eliciting the target pairs for the CPT trial from the physician, we used a graphical representation of preliminary versions of Figure 1, and modified the figure adaptively as he adjusted some values based on the figure. The process for obtaining the θ target pairs thus was straightforward. Initially, the shape of $\mathcal{L}$ was surprising. When we questioned this, the physician explained that this $\mathcal{L}$ very accurately reflects his strong belief that a very high risk of toxicity is an acceptable trade-off for a high probability of achieving a longer event-free survival time in CPT patients.

In general, T is not reduced to the binary indicator of the event ($T > t^*$), and all right-censored event time data (T^c, ε) are included in the likelihood and used to compute the posteriors that determine the test statistic. In the CPT trial, while $\theta_{T,j}(\mathbf{Z}, Y, \xi) = F_j(t^* | \mathbf{Z}, Y, \xi)$ for reference time $t^* = 24$ months is used as the basis for defining targets and hypotheses (described below in Section 3), T is not reduced to the binary indicator of the event ($T > 24$).

3. A geometry for comparing treatments

In this section, we provide details of the method for constructing 2-dimensional parameter sets that provide a geometric basis for treatment comparison. Given null pair $\mu_{\theta,0}$, a set of elicited target pairs $\tilde{\theta}^e$ and target set Θ_e are *admissible* if (i) $\mu_{\theta,0} \notin \Theta_e$ and (ii) θ_Y is a strictly increasing function of θ_T for $\theta \in \mathcal{L}$. We require these admissibility conditions to avoid cases that do not make sense. These conditions can be ensured easily during the target pair elicitation process by showing the physician a plot of $\mu_{\theta,0}$ and the elicited targets. The sample size, Type I error, and generalized power of the test, defined below in Section 4, depend on the shape of $\mathcal{L}$ and the distance between $\mu_{\theta,0}$ and $\mathcal{L}$. Similarly to a conventional test of a 1-dimensional parameter, in the present setting the sample size required to achieve given Type I error and generalized power would be larger for $\mathcal{L}$ closer to $\mu_{\theta,0}$. Condition (ii) rules out the possibility that some portions of $\mathcal{L}$ may be either horizontal or vertical line segments. For example, if $\mathcal{L}$ were allowed to contain a horizontal line segment, then there would be parameter pairs, θ and θ' , both on that horizontal segment such that $\theta_T < \theta'_T$ but $\theta_Y = \theta'_Y$. This would say that, despite the fact that θ' has greater efficacy than θ and the two pairs have identical toxicity, they are considered equally desirable, which is nonsense. In particular, “staircase” shaped $\mathcal{L}$ are inadmissible.

This construction is illustrated for the CPT trial by Figure 1. While we assume that T may depend on Y , it may be argued either that toxicity ($Y = 1$) decreases the mean EFS time since CPT patients who experience toxicity may have their chemotherapy dose reduced or, instead, that toxicity increases mean EFS time due to a positive association between toxicity and cancer cell killing by the chemotherapy. For the CPT trial, it was not necessary to to average two sets of targets, described earlier, since the physician T to be independent of Y . The historical mean was $\mu_{\theta,0} = (0.47, 0.11)$, and the five target pairs were $\theta^{e,1} = (0.65, 0.01)$, $\theta^{e,2} = (0.70, 0.05)$, $\theta^{e,3} = (0.80, 0.20)$, $\theta^{e,4} = (0.90, 0.40)$, $\theta^{e,5} = (0.99, 0.99)$, considered equally desirable improvements over $(0.47, 0.11)$. Figure 1 shows the targets, polygonal target line $\mathcal{L}$, target set Θ_e and null pair $\mu_{\theta,0}$. The trade-off between EFS and toxicity allows $\Pr(\text{toxicity} | \mathbf{z}^*)$ to increase substantially over the historical mean 0.11 provided that $\Pr(T >$

24 months $|z^*, y^*)$ is sufficiently large compared to the null value 0.47. The extreme case $\theta^{e,5} = (0.99, 0.99)$ was considered to be as desirable as $\theta^{e,2} = (0.70, 0.05)$ where $\Pr(T > 24 \text{ months})$ has an improvement from 0.47 to 0.70 and $\Pr(\text{toxicity} | z^*)$ drops from 0.11 to 0.05. This was due to the great importance placed on improving EFS time, and the fact that death is accounted for by T and not included in the definition of toxicity.

Figure 2 shows the structure obtained from that in Figure 1 by mapping each θ to a corresponding pair δ in the covariate-adjusted E_j -versus- S effect domain. The null pair $\mu_{\theta,0} = (0.47, 0.11)$, $\mathcal{L}$, and θ_e in the Θ domain in Figure 1 are mapped to $\mathbf{0} = (0, 0)$, $\mathcal{D}$, and Δ_c , respectively, in the Δ domain in Figure 2. While the vertical line from $\delta^{e,1}$ to $(\delta_T^{e,1}, -\infty)$ and the horizontal line from $\delta^{e,5}$ to $(+\infty, \delta_Y^{e,5})$ in Figure 2 both are portions of the boundary of the transformed target set Δ_c , these lines are *not* portions of the transformed target line $\mathcal{D} = \tau(\mathcal{L}) - \mu_{\gamma,0}$. In Δ , all pairs δ on the transformed target boundary $\mathcal{D}$ are equally desirable. However, pairs on either the vertical or horizontal lines on the boundary of Δ_c are *not* considered to be as desirable as pairs on $\mathcal{D}$, since this would not make sense.

4. A hybrid select-and-test design

4.1 Priors

Our formulation utilizes available information on covariate effect parameters, obtained by elicitation or from historical data. In contrast, we assume non-informative iid priors on the treatment effect pairs $\gamma_0, \gamma_1, \dots, \gamma_J$ because comparative E_j -versus- S decisions will be based on the posteriors of $\delta_1, \dots, \delta_J$. While it may be argued that the prior on γ_0 should be informative to reflect historical experience with S , we do not make this assumption in order to avoid confounding treatment effects based on the trial data with effects of latent variables that may differ in distribution between the trial and historical data. Under the assumption of no treatment-covariate interactions, any differences in the covariate effects (β_T, β_Y) between prior experience and the trial would be manifested approximately additively in the linear terms (3) and (2) for all $j = 0, 1, \dots, J$. Hence such differences would introduce additional variability but each δ_j will be approximately unbiased. Our formulation ensures that the treatment comparisons will be dominated by the data from the trial, with at most numerically trivial effects due to the non-informative priors on the γ_j 's.

4.2 Hypotheses and trial design

Figure 2 shows that the null set Δ_0 is the upper left quadrant of $\mathbb{R}^2$, and that Δ_0 and Δ_c are disjoint. The *global null hypothesis* is $H_0: \delta_1, \dots, \delta_J \in \Delta_0$. Since larger $\gamma_{T,j}$ and smaller $\gamma_{Y,j}$ are more desirable, under H_0 no E_j is better than S since $\gamma_{T,j} \leq \gamma_{T,0}$ and $\gamma_{Y,j} \geq \gamma_{Y,0}$. For each $j = 1, \dots, J$, we define the *jth alternative hypothesis* $H_{a,j}: \delta_j \in \Delta_c$. While $H_{a,j}$ says that E_j provides an improvement over S , it does not say anything about $\delta_{j'}$ for any $j' \neq j$, and in particular $H_{a,1}, \dots, H_{a,J}$ are not disjoint. The *global alternative hypothesis* that at least one E_j provides an improvement over S is $H_a = \cup_{j=1}^J H_{a,j}$.

During the trial, each patient's $\mathbf{Z}$ vector is recorded at enrollment, and a treatment j is randomly chosen using the Pocock-Simon method (1975) to balance on $\mathbf{Z}$. The treatment is administered and the patient's outcome (Y, T^o, ε) is observed. In stage 1, let $data_1$ denote the data available when the decision is made and denote the index of the empirically best experimental treatment at the end of stage 1 by $v = \arg\max\{j = 1, \dots, J : R_j(data_1)\}$. Let M_1 be the number of events observed and r_1 the decision cut-off for the test statistic $R_v(data_1)$. When the decision is made in stage 2, denote the data from all patients in arms E_v and S from both stages by $data_2$, let M_2 denote the total number of events observed from both

treatment arms E_v and S , and let r_2 denote the decision cut-off for the test statistic $R_v(data_2)$. Given (M_1, M_2, r_1, r_2) , the trial is conducted as follows:

Stage 1: Randomize patients among $E_1, \dots, E_J$ and S until M_1 events have been observed. The stage 1 decision is made when Y has been evaluated for all stage 1 patients enrolled up to time t_1 . If $R_v(data_1) > r_1$ then continue to stage 2; otherwise, terminate the trial and accept H_0 .

Stage 2: Randomize patients between E_v and S until M_2 additional events from patients in these two arms have been observed. The final decision is made when Y has been evaluated for all patients. If $R_v(data_2) > r_2$ then accept $H_{a,v}$; otherwise accept H_0 .

Since at most one E_j may be selected in stage 1, the design allows the $J + 1$ possible conclusions $H_{a,1}, \dots, H_{a,J}$, that of the E_j 's is sufficiently promising for stage 2 evaluation, or H_0 , that no E_j is promising and the trial should be stopped. If the trial continues to stage 2, the data from patients enrolled in stage 1 are utilized when computing the stage 2 statistic $R_v(data_2)$ for the final decision. In particular, some of the M_2 additional events observed in stage 2 may come from patients enrolled in arms E_v or S in stage 1. Let t_v^* be the length of the time period from the start of therapy during which toxicity is monitored, which is 4 months for the CPT trial. Because Y is a binary variable based on a 4-month observation period, $Y = 1$ may be scored at any time toxicity occurs up to t_v^* but $Y = 0$ may be scored only after the patient has been observed for a period of length t_v^* without toxicity. The design makes decisions only after Y has been evaluated for all patients in each stage to avoid overestimating the $\theta_{Y,j}$'s. The waiting period needed to evaluate Y for all patients may be of length up to t_v^* , with the stage 1 decisions made at trial time $t_1 + t_v^*$. Let t_1 denote the time, from the start of the trial, when the M_1^{st} event is observed in stage 1. Ideally, patient enrollment should be suspended at trial time t_1 in stage 1 in order to evaluate Y for the last patient accrued, since it may be considered ethically undesirable to continue giving patients experimental treatments that soon may be considered inferior to S . If it is not feasible to suspend accrual between stages, accrual is continued and additional failures $T^\circ = T$ in that period may cause the planned M_1 to be overrun slightly by the time $data_1$ is used to make the stage 1 decision. This effect should be negligible in most settings, and it may be accounted for when computing the design's operating characteristics and deriving design parameters, which was done with the CPT trial. Computational details are provided below, in subsection 4.5. For the stage 2 decision, the values of (T°, ε) for stage 1 patients in arms E_v or S for whom $\varepsilon = 0$ at the stage 1 decision are updated at the final decision time to be either the patient's extended follow up time T° if the event still has not occurred ($\varepsilon = 0$), or the patient's observed event time if $T^\circ = T$ ($\varepsilon = 1$).

4.3 Operating characteristics

While the design uses Bayesian decision criteria, in practice it must have good frequentist properties. To define overall Type I and Type II error, we generalize the approach taken by TSE in the simpler case of one-dimensional parameters $\theta_0, \theta_1, \dots, \theta_J \in \mathcal{R}^1$, without covariates. As a frame of reference for what follows, we first briefly review the TSE formulation. The E_j -versus- S effects are $\delta_j = \theta_j - \theta_0$ for $j = 1, \dots, J$, the null set is $\Delta_0 = (-\infty, 0]$ and the alternative set is $\Delta_a = [\delta^*, +\infty)$ for fixed $\delta^* > 0$. The global null hypothesis is $H_0: \delta_1 = \dots = \delta_J = 0$, the j th alternative is $H_{a,j}: \delta_j \in \Delta_a$, and the global alternative is

$H_a = \bigcup_{j=1}^J H_{a,j}$. The Type I error is the probability of concluding H_a when H_0 is true. The *generalized power* (GP) is the probability that, for some j such that $H_{a,j}$ is true, one correctly concludes $H_{a,j}$. TSE show that, subject to the requirements (i) $\delta_1, \dots, \delta_J \in \Delta_0 \cup \Delta_a$ and (ii) at least one $\delta_j \in \Delta_a$, a parameter set $\{\delta_1, \dots, \delta_J\}$ minimizing the GP has exactly one $\delta_j = \delta^*$ and $\delta_{j'} = 0$ for all $j' \neq j$. This is called a *least favorable configuration* (LFC), and it may occur in

J ways, one for each E_j . The requirement (i) is imposed because no statistical procedure can reliably distinguish between fixed parameter values that are arbitrarily close. Since the indexing of the J experimental treatments is arbitrary, by symmetry the GP is the probability of concluding $H_{a,j}$ under the LFC where $\delta_1 = \dots = \delta_{J-1} = 0$ and $\delta_J = \delta^*$. The formulation given by TSE is slightly more general in that they consider two improvement values, a clinically insignificant value $\delta_1^* > 0$ and a clinically significant value $\delta_2^* > \delta_1^*$. The above account simplifies the formulation by setting $\delta_1^* = 0$ and writing $\delta_2^* = \delta^*$.

We define Type I error and GP similarly, but now the hypotheses refer to the parameter sets Δ_0 and Δ_c in $\mathbb{R}^2$. For each $j = 1, \dots, J$, the event that the design concludes $H_{a,j}$ is

$$A_j = \{v(\text{data}_1) = j, R_j(\text{data}_1) > r_1 \text{ and } R_j(\text{data}_2) > r_2\}, \quad (5)$$

and the event that the global alternative H_a is accepted is

$A = \bigcup_{j=1}^J A_j = \{R_v(\text{data}_1) > r_1 \text{ and } R_v(\text{data}_2) > r_2\}$. The event that the design accepts H_0 is

$$A_0 = \{R_v(\text{data}_1) \leq r_1\} \cup \{R_v(\text{data}_1) > r_1 \text{ and } R_v(\text{data}_2) \leq r_2\}. \quad (6)$$

Denote a $2J$ -vector of J fixed treatment effect pairs by $\vec{\delta}^* = (\delta_1^*, \dots, \delta_J^*)$. We make a strong distinction between such fixed parameter vectors, which will be used to evaluate the design's properties, and the random parameter vector $\vec{\delta}$ under the Bayesian model. In our setting, the GP is the probability $\varphi(\vec{\delta}^*)$ of correctly concluding $H_{a,j}$ for a true parameter vector $\vec{\delta}^*$ for which $\delta_j^* \in \Delta_c$. In order to define practical criteria for making decisions and deriving design parameters the following theorem is needed, which may be considered a two-dimensional version of the theorem in TSE (1988). Similarly to the one-dimensional case, we impose the requirement that all δ_j^* must be in either Δ_c or Δ_0 since no statistical procedure can reliably distinguish between parameters that are arbitrarily close.

Theorem. If (i) $\delta_1^*, \dots, \delta_J^* \in \Delta_c \cup \Delta_0$ and (ii) at least one $\delta_j^* \in \Delta_c$, then the GP is minimized if there is exactly one j such that $\delta_j^* \in \mathcal{D}$ and

A heuristic proof is given in the Appendix—The theorem does not specify the location of δ_j^* on $\mathcal{D}$ since this depends on the particular geometry of $\mathcal{D}$, which is determined by $\mu_{\theta,0}$ and $\vec{\theta}^e$. Since the indices of $E_1, \dots, E_J$ are arbitrary, let E_J be the treatment providing the improvement, so that $(0, \dots, 0, \delta^*)$ is the vector of J fixed parameter pairs with $\delta_J^* = \delta^*$ and all other pairs 0. The GP can be no smaller than $\min_{\delta_j^* \in \mathcal{D}} \varphi(0, \dots, 0, \delta^*)$, and we will call such a vector that minimizes the GP a LFC. Similarly, the Type I error is the maximum over $\delta^* \in \Delta_0$ of the probability of rejecting H_0 , given by

$\alpha^* = \sum_{j=1}^J \Pr(A_j | \vec{\delta}^* = \vec{0}) = J \times \Pr(A_J | \vec{\delta}^* = \vec{0})$, where $\vec{0}$ is the $2J$ -vector with all entries 0. Let $\bar{\alpha}$ be the design's largest allowed probability of a Type I error and $\bar{\varphi}$ the smallest allowed probability of making a correct decision under H_a . Let $\zeta(\vec{\delta}^*) = \Pr\{R_v(\text{data}_1) \leq r_1 | \vec{\delta}^*\}$ denote the probability of early termination after stage 1, and let $\bar{\zeta}$ be the maximum allowed value of ζ under an LFC. In order to feasibly obtain an optimal design, φ and ζ are evaluated at the m transformed target pairs $\{\delta^{e,1}, \dots, \delta^{e,m}\}$ that determine $\mathcal{D}$. The design requires the Type I error constraint $\alpha^* \leq \bar{\alpha}$, the incorrect early stopping probability constraint

$$\zeta^* =_{\text{def}} \max_{r=1, \dots, m} \{\zeta(\mathbf{0}, \dots, \mathbf{0}, \delta^{e,r})\} \leq \bar{\zeta} \quad (7)$$

and the GP constraint

$$\varphi^* =_{\text{def}} \min_{r=1, \dots, m} \{\varphi(\mathbf{0}, \dots, \mathbf{0}, \delta^{e,r}), r=1, \dots, m\} \geq \bar{\varphi}. \quad (8)$$

The design parameters are thus $\{\bar{\alpha}, \bar{\varphi}, \bar{\zeta}, M_1, M_2, r_1, r_2\}$.

4.4 Deriving optimal design parameters

While $\{\bar{\alpha}, \bar{\varphi}, \bar{\zeta}\}$ are predetermined, $\{M_1, M_2, r_1, r_2\}$ must be derived. As with any group sequential design based on event times where decisions are made when specified numbers of events are observed, the sample sizes N_1 and N_2 in stages 1 and 2 are random quantities determined by the design's required event counts, M_1 and M_2 , and the underlying accrual and event rates. The total sample size, N , is either N_1 or $N_1 + N_2$, with mean $E(N) = E(N_1) + (1 - \zeta)E(N_2)$, which may be computed under the null $\delta^* = \mathbf{0}$ or at an LFC. Subject to the constraints $\varphi^* \geq \bar{\varphi}$, $\alpha^* \leq \bar{\alpha}$ and $\zeta^* \leq \bar{\zeta}$, we derive the $\{M_1, M_2, r_1, r_2\}$ values that minimize the equally weighted mean overall sample size $\bar{E}(N) = \frac{1}{2}E_{H_0}(N) + \frac{1}{2}E_{LFC}(N)$, with the second expectation computed under the fixed parameter vector $(\mathbf{0}, \dots, \mathbf{0}, \delta^*)$ that minimizes the GP.

4.5 Numerical methods

For each fixed δ^* , we obtained the optimal $\{M_1, M_2, r_1, r_2\}$ for the CPT trial by searching in three nested loops, with M_1 in the outer loop, M_2 nested within M_1 , and (r_1, r_2) nested within M_2 . For each $\{M_1, M_2, r_1, r_2\}$ examined, given the assumed true value δ^* , we computed $\varphi^*, \zeta^*, \alpha^*$ and the objective function $\bar{E}(N)$ by generating 2,000 data sets from the likelihood with parameters corresponding to $(\mathbf{0}, \dots, \mathbf{0}, \delta^*)$. During the search, a design was considered possible if it satisfied the three error constraints. Initially, a small number of (M_1, M_2) pairs were considered. For each (M_1, M_2) pair, a local grid search was done in (r_1, r_2) starting with a coarse grid having increments .10 to identify possible solutions, then picking the possible (r_1, r_2) pair giving the smallest value of $\bar{E}(N)$, then localizing the grid around that pair and refining the grid increments to .001. This was repeated with the grid of (r_1, r_2) values refined to increments of .001. For each evaluation of $(\varphi^*, \zeta^*, \alpha^*, \bar{E}(N))$ at a pair of (M_1, M_2) values, if there were no (r_1, r_2) pairs giving a possible solution, M_2 was incremented to $M_2 + 1$, with this limited to at most 20 steps. Similarly, given M_1 , if no M_2 yielding a possible solution was found then M_1 was incremented to $M_1 + 1$, with this repeated until a possible solution was found. Once a set of possible designs were obtained, the (M_1, M_2) grid was localized at the pair $(M_1, M_2)^o$ minimizing $\bar{E}(N)$ and an exhaustive search was done in a contiguous set of pairs around $(M_1, M_2)^o$.

To simulate each data set when applying the above algorithm, each patient's $\mathbf{Z}$ was sampled from the historical covariate distribution ("Popn. Percent" in Table 1) and a treatment j was then randomly chosen based on $\mathbf{Z}$ using the Pocock-Simon method (1975). The toxicity indicator Y was simulated first using the probability $\theta_{Y,j}(\mathbf{Z}, \xi)$, substituting the prior means of $\gamma_{Y,0}$ and the covariate effects β_Y . The specified $\delta_{Y,j}^*$ determined the fixed intercept as $\gamma_{Y,j}^* = \delta_{Y,j}^* + E(\gamma_{Y,0})$. The event time $[T|\mathbf{Z}, Y]$ then was simulated under the lognormal model using the prior means of $\gamma_{T,0}$, β_T , $\beta_{T,qT+1}$ and σ_T^2 with $\gamma_{T,j}^* = \delta_{T,j}^* + E(\gamma_{T,0})$. For each simulated data set, the stage 1 posterior probability ratios $R_1(\text{data}_1), \dots, R_J(\text{data}_1)$ were computed, the

best stage 1 treatment index $v(data_1)$ was identified, the stage 2 test statistic $R_v(data_2)$ was computed, and the decision rules were then evaluated. Our construction involves a large number of nuisance parameters. While misspecifying these values may result in error inflation, we have used prior means, since these were the best available numerical values.

The values of $R_j(data_1)$ and $R_j(data_2)$ were computed numerically using Markov chain Monte Carlo with Gibbs sampling (Robert and Cassella, 1999). For prior $p(\xi)$, each generated chain $\{\xi^{(1)}, \dots, \xi^{(K)}\}$ had $K = 10,000$ samples distributed in proportion to the posterior integrand, $Lik(data, \xi) \times p(\xi)$. The first sample vector $\xi^{(1)}$ in each chain was set equal to the mode of the integrand, found using the Nelder-Mead algorithm (1965).

5. Establishing priors for the CPT trial

Since clinical experience with CPT (Wrede et al., 2005; Wrede, et al., 2007) indicates that the hazard function for EFS is certainly non-constant and very likely non-monotone, we assumed that $[T | Z, j]$ is lognormally distributed with $E\{\log(T) | Z, j\} = \gamma_{T,j} + \beta_T Z_T$ and $\text{var}\{\log(T) | Z, j\} = \sigma_T^2$. The parameter subvector for the distribution of T is thus $\xi_T = (\gamma_{T,0}, \gamma_{T,1}, \gamma_{T,2}, \gamma_{T,3}, \beta_T, \sigma_T)$.

To determine numerical values of the hyperparameters $\tilde{\xi}_T$ of the prior $p(\xi_T | \tilde{\xi}_T)$, we proceeded as follows. Denote a normal distribution with mean μ and variance σ^2 by $N(\mu, \sigma^2)$. Non-informative iid $N(0, 100)$ priors were assumed for the eight treatment effect parameters $(\gamma_{T,j}, \gamma_{Y,j})$ for $j=0,1,2,3$, and for the effect $\beta_{T,5}$ of Y on T . To obtain an informative prior on β_T , values of the EFS probabilities $\bar{F}(t | Z_T, \xi)$ were elicited for each of the 16 configurations of $(Z_{T,1}, \dots, Z_{T,4})$ and $t = .50, 1.0, 2.0$ and 3.0 years, with fixed $z_{T,5}^* = y^* = 0$. The elicited values are given in Table 1. We assumed that $\beta_{T,1}, \dots, \beta_{T,4}$ are normally distributed with common variance, and that $\log(\sigma_T)$ is normal. Indexing the elicitation times by $r = 1, 2, 3, 4$ and the covariate vectors by $l = 1, \dots, 16$, and regarding each prior mean EFS probability

$$\bar{F}_{r,l}(\tilde{\xi}_T) = E\{\bar{F}(t_r, Z_l | \xi_T) | \tilde{\xi}_T\} = \int \bar{F}(t_r, Z_l | \xi_T) p(\xi_T | \tilde{\xi}_T) d\xi_T \quad (9)$$

as a function of $\tilde{\xi}_T$, a generalization of the method of Thall and Cook (2004) was used to solve for $\tilde{\xi}_T$. This was done by treating the elicited probabilities $\bar{F}^{(e)}(t_r | Z_l)$ as observations with means $\bar{F}_{r,l}(\tilde{\xi}_T)$ in a nonlinear regression model parameterized by $\tilde{\xi}_T$ and solving for the values of $\tilde{\xi}_T$ that minimized $\sum_{r,l} \{\bar{F}^{(e)}(t_r | Z_l) - \bar{F}_{r,l}(\tilde{\xi}_T)\}^2$. The resulting normal priors of the covariate parameters have means $E(\beta_{T,1}, \beta_{T,2}, \beta_{T,3}, \beta_{T,4}) = (0.44, -0.41, 0.56, -0.53)$ with common variance 0.15^2 , and the prior on σ_T^2 is given by $\log(\sigma_T) \sim N(-0.08, 0.14^2)$. These were used as the priors for the model parameters when deriving the optimal designs.

For toxicity, based on clinical experience in treating CPTs only $Z_{T,1} = \text{Age}$ was used as a covariate in the definition of $\theta_{Y,j}$. Toxicity data taking the specific form of Y were not available, however. Consequently, to establish a model for $\theta_{Y,j}(Z_{T,1}, \gamma_{Y,0}, \beta_Y)$ and an informative prior on β_Y , toxicity probabilities for each of nine patient ages were elicited independently from each of three oncologists with extensive experience treating patients with CPTs. The sets of elicited toxicity probabilities, shown in Figure 1, were remarkably consistent in that all three physicians gave values that described the same non-monotone function of Age in which the probability of toxicity is high for infants, very low for patients 3 to 50 years of age, and moderately higher for older patients. While most CPT patients are very young, the age domain used in the elicitation is appropriate for the trial because its eligibility is not restricted to children, and CPT patients can be as old as 80 years. Denoting

the elicited probability at the r^{th} age from the l^{th} physician by $\theta_{Y,l,r}^{(e)}$, these values were fit as a function of Age by considering mixed models with fractional polynomial linear terms (Royston and Altman, 1994) taking the general form

$\text{logit}\{\theta_{Y,l,r}^{(e)}\} = \gamma_{Y,0} + \beta_{Y,1} \text{Age}^{p_1} + \beta_{Y,2} \text{Age}^{p_2} + \rho^{(j)} + \varepsilon_{l,r}$ where $\gamma_{Y,0}$, $\beta_{Y,1}$, $\beta_{Y,2}$ are fixed parameters, $\rho^{(1)}$, $\rho^{(2)}$, $\rho^{(3)}$ are iid $N(0, \sigma_\rho^2)$ physician effects, and the $\varepsilon_{l,r}$'s are iid $N(0, \sigma_\varepsilon^2)$ residuals. One best fitting model to use as a basis for constructing a design was determined following the approach recommended by Saurbrei and Royston (1999). This was done by varying the exponents p_1 and p_2 exhaustively over the set $\{-2, -1, -0.5, 0, 0.5, 1, 2, 3\}$, where the exponent "0" corresponds to $\log(\text{Age})$. The 28 models of the above form with $p_1 \neq p_2$, and the 8 models not including the term $\beta_{Y,2} \text{Age}^{p_2}$, were considered. Each of these 36 models was fit both with and without physician effects, for a total of 72 models. The Bayesian Information Criterion (BIC, Schwarz, 1978) values ranged from 85.5 to 140.2 for the 72 models, with the best fitting model given by the linear term including $\text{Age}^{1/2}$ and $\log(\text{Age})$. Figure 1 gives plots of the three sets of elicited toxicity probabilities, along with the posterior mean fitted curve under the best fitting fractional polynomial model. To obtain the informative distributions on $\beta_{Y,1}$ and $\beta_{Y,2}$ needed for the trial design, a Bayesian version of this model was fit, dropping the physician effects, with $\gamma_{Y,0}$, $\beta_{Y,1}$, $\beta_{Y,2}$ following $N(0, 100)$ priors and σ_ρ^2 following an inverse gamma prior with mean 1 and variance 1000. The estimated linear term of the fitted model obtained by replacing each parameter with its posterior mean, each subscripted with its standard deviation, is

$$\text{logit}\{\hat{\theta}_r(\text{Age}, \gamma_{Y,0}, \beta_{Y,1}, \beta_{Y,2})\} = -2.26_{0.46} + 0.78_{0.19} \text{Age}^{1/2} - 1.33_{0.27} \log(\text{Age}).$$

The posteriors $\beta_{Y,1} \sim N(0.78, 0.19^2)$ and $\beta_{Y,2} \sim N(-1.33, 0.27^2)$ obtained from this fit were used as their priors in formulating the trial design. The fitted models from the elicited data for Y and T gave prior means $\mu_{\theta,0} = (0.47, 0.11)$. That is, the mean probability that $T > 24$ months for the reference patient z^* was 0.47, and based on the elicited age-dependent toxicity probabilities the mean probability that a 2-year-old patient would suffer toxicity was 0.11.

6. Design parameters for the CPT trial

6.1 Optimal designs

The optimal designs with $\bar{\alpha} = 0.05$, $\bar{\phi} = 0.80$ and $\bar{\zeta} = 0.05$ were derived for $J = 1, 2$, and 3. The case $J = 1$ is a two-arm, two-stage design that uses the same geometric construction but only compares one $E = E_1$ to S based on $\delta = \delta_1$. In this case, the null hypothesis is $H_0: \delta \in \Delta_0$ and there is one alternative, $H_{a,1} = H_a: \delta \in \Delta_c$. This case was included as a basis for comparison to assess the effects of having to perform a selection in stage 1.

The optimal designs are given in Table 2. It is important to bear in mind that the test cutoffs r_1 and r_2 pertain to the posterior probability ratios $R(\text{data}_1)$ and $R(\text{data}_2)$, and not to conventional Z-scores. The numbers of events required to make decisions in each stage, M_1 and M_2 , increase substantially with J due to the facts that, with larger J , more treatments are evaluated and the stage 1 decision includes selection among the E_j 's as well as comparison to S . This produces increases in overall expected sample size and trial duration with J , although these values also vary greatly with the case under which they were computed. The GP, ϕ^* , follows a U-shaped pattern as the true θ^* moves along $\mathcal{L}$, as illustrated by Figure 2, with ϕ^* smallest at $\theta^{e,2} = (0.70, 0.05)$ and the GP values of $\theta^{e,2}$, with $\theta^{e,3}$ very close to each

other in each case. This makes intuitive sense in terms of the centered targets $\delta^{e,j} = \theta^{e,j} - (0.47, 0.11)$, shown in Figure 2, since $\delta^{e,2}$ is closest to (0,0).

In addition to the GP values φ^* at the null and the five target pairs, Figure 2 also shows values of φ^* for δ pairs on lines connecting (0,0) to three of the targets, and for also values of δ in the interior of Δ_c . Considered together, the numerical values of φ^* in Figure 2 illustrate that $\varphi^*(\delta)$ increases monotonically as δ moves away from the null pair (0,0) toward any target pair. The response surface determined by φ^* as a function of δ may be considered a generalization of a conventional power curve based on a single parameter.

6.2 Robustness

To examine robustness of the method to the assumed lognormal distribution for T , we evaluated the design's behavior with T contaminated by multiplicative noise. For the design with $J = 3$, in each case of true θ^* in Table 2 we simulated samples of the event times with each T_i replaced by $G_i T_i$, where the G_i 's were independent gamma random variables with mean 1 and variance σ_G^2 . This had the effects of both changing the distribution of T_i and increasing its variability. Denoting $E(G_i) = \mu_G$ and $E(T_i) = \mu_T$, since

$$\text{var}(G_i T_i) = \sigma_T^2 \sigma_G^2 + \sigma_T^2 \mu_G^2 + \sigma_G^2 \mu_T^2 = \sigma_T^2 (\sigma_G^2 + \mu_G^2 + \sigma_G^2 \mu_T^2 / \sigma_T^2),$$

the effect of multiplying T_i by G_i is to increase the standard deviation σ_T by the

multiplicative factor $\sigma_{GT}/\sigma_T = (\sigma_G^2 + \mu_G^2 + \sigma_G^2 \mu_T^2 / \sigma_T^2)^{1/2}$. This factor depends not only on the variances σ_T^2 and σ_G^2 but also the values of μ_G^2 and μ_T^2 . Table 3 gives the values of φ^* , PET, and $E(N)$ for $\sigma_G = 0$, corresponding to the assumed lognormal distribution, and $\sigma_G = .2, .4$, and $.6$, corresponding to successively larger amounts of contamination. As shown in Table 3, for $\sigma_G = .4$ or $.6$ the value of σ_{GT}/σ_T may be quite large, and it varies substantially depending on the value of the true target pair θ^* where the design is evaluated. The Type I error and GP hold up quite well even when there is substantial contamination, with φ^* decreasing substantially only at Targets 2, 3, or 4 for large values $\sigma_{GT}/\sigma_T \geq 1.39$. Thus, the method appears to be robust to all but extreme departures from the lognormal assumption.

7. Discussion

The proposed design is a hybrid in that Bayesian machinery is used to obtain decision rules, while the design's parameters are derived to ensure that it has specified frequentist properties. The methodology requires a substantial effort to elicit target parameter pairs and priors, and the process of deriving design parameters is computationally intensive. While these may be considered limitations, the two-dimensional covariate-adjusted target set Δ_c provides a much more refined goal than the conventional approach of increasing median EFS by a given amount while ignoring toxicity and patient covariates. From a practical perspective, the result of applying this design is a sample size small enough to allow the CPT trial to be completed in a realistic timeframe.

There are many other possible approaches to the problem we have addressed here. With regard to controlling error rates, a referee has pointed out that one may take an alternative approach since, if the hypothesis for treatments excluded at the interim analysis are retained, one can consider the disjunctive power to reject a false hypothesis and the familywise error rates to reject a true hypothesis. Aside from the preliminary selection, one might simultaneously use a noninferiority test for safety and a superiority test for efficacy, or use a positive quadrant test, as given by Jennison and Turnbull (1993). However, these methods

do not account explicitly for the relative importance of disease-free survival time and toxicity, which was our primary motivation for using the geometric construction, based on elicited target parameter pairs. Alternatively, one might use a utility score to combine the two outcomes, although this would require dealing with the fact that T has domain $[0, \infty)$. One might also use adaptive combination tests (cf. Bretz, et al., 2009), which allow a variety of interim decision rules, including Bayesian rules. Our proposed methodology controls the Type I error under the global null, H_0 , which is a strong requirement since it may be the case that some $\delta_j = 0$ but others not. In this regard, adaptive combination tests guarantee control of the Type I error for any configuration of nulls. As a future exercise, it would be worthwhile to investigate how such alternative approaches compare to that given here.

Our method relies on the strong assumption that there are no treatment-covariate interactions. We assume this in order to control the number of parameters so that we may obtain a feasible design that can actually be used. The model might be extended to include treatment-covariate interactions, for example in terms of subgroups $r = 1, \dots, K$ determined by $\mathbf{Z}$, by replacing the linear terms with the $2(J+1)K$ treatment-subgroup effects $\eta_{Y,j,r}(\xi)$ and $\eta_{T,j,r}(\xi)$. The number of parameter pairs would then be $K(J+1)$, far too large to implement the methodology practically, since the potential hypotheses and tests would be far more complex. Even with only $K = 2$ subgroups, with $J = 3$ there would be a total of 8 treatment-subgroup parameter pairs.

If the trial's entry criteria are sufficiently restrictive that it is reasonable to assume patients are homogeneous with respect to the outcome (Y, T) , then the parametric functions would take the simpler form $\theta_j(Y, \xi) = (\theta_{T,j}(Y, \xi), \theta_{Y,j}(\xi))$ for all patients in treatment arm j . In this case, $\eta_{Y,j}(\mathbf{Z}, \xi) = \gamma_{Y,j}$ and $\eta_{T,j}(\mathbf{Z}, \xi) = \gamma_{T,j} + \beta Y$, and the target pairs that determine the alternative region Θ_c would be elicited without reference to covariates. Finally, if it is realistic to assume that T is independent of Y , as in the CPT trial, then $\eta_{T,j}(\mathbf{Z}, \xi) = \gamma_{T,j}$.

References

- Bechhofer, RE.; Santner, TJ.; Goldsman, DM. Design and Analysis of Experiments for Statistical Selection, Screening and Multiple Comparisons. John Wiley & Sons; New York: 1995.
- Bretz F, Schmidli H, Konig F, Racine A, Maurer W. Confirmatory seamless phase II/III clinical trials with hypotheses selection at interim: general concepts. Biometrical Journal. 2006; 48:623–634. [PubMed: 16972714]
- Bretz F, Koenig F, Brannath W, Glimm E, Posch M. Adaptive designs for confirmatory clinical trials. Statistics in Medicine. 2009; 28:1181–1217. [PubMed: 19206095]
- Cheung YK. Simple sequential boundaries for treatment selection in multi-armed randomized clinical trials with a control. Biometrics. 2008; 64:940–949. [PubMed: 17970818]
- Jennison C, Turnbull BW. Group sequential tests for bivariate response: Interim analyses of clinical trials with both efficacy and safety endpoints. Biometrics. 1993; 49:741–752. [PubMed: 8241370]
- Jennison C, Turnbull B. Confirmatory seamless phase II/III clinical trials with hypotheses selection at the interim: opportunities and limitations. Biometrical Journal. 2006; 48:650–655. [PubMed: 16972717]
- Lan KKG, DeMets DL. Discrete sequential monitoring boundaries for clinical trials. Biometrika. 1983; 70:659–663.
- Liu Q, Pledger GW. Phase 2 and 3 combination designs to accelerate drug development. J American Statistical Association. 2005; 100:493–502.
- Nelder JA, Mead R. A simplex method for function minimization. Computer Journal. 1965; 7:308.
- O'Brien PC, Fleming TR. A multiple testing procedure for clinical trials. Biometrics. 1979; 35:549–556. [PubMed: 497341]
- Pocock SJ. Group sequential methods in the design and analysis of clinical trials. Biometrika. 1977; 64:191–199.

- Pocock SJ, Simon RM. Sequential treatment assignment with balancing for prognostic covariates in the controlled clinical trial. *Biometrics*. 1975; 31:103–115. [PubMed: 1100130]
- Robert, CP.; Cassella, G. Monte Carlo Statistical Methods. New York: Springer; 1999. p. 1999
- Royston P, Altman DG. Regression using fractional polynomials of continuous covariates: parsimonious parametric modelling. *Journal of the Royal Statistical Society, Series C*. 1994; 43:429–467.
- Rubinstein LV, Korn EL, Freidlin B, Hunsberger S, Ivy SP, Smith MA. Design issues of randomized phase II trials and a proposal for phase II screening trials. *Journal of Clinical Oncology*. 2004; 23:7199–7206. [PubMed: 16192604]
- Sauerbrei W, Royston P. Building multivariable prognostic and diagnostic models: transformation of the predictors by using fractional polynomials. *Journal of the Royal Statistical Society, Series A*. 1999; 162:71–94.
- Schaid DJ, Wieand HS, Therneau TM. Optimal two-stage screening designs for survival comparisons. *Biometrika*. 1990; 77:507–513.
- Schwarz GE. Estimating the dimension of a model. *Annals of Statistics*. 1978; 6:461–464.
- Stallard N, Todd S. Sequential designs for phase III clinical trials incorporating treatment selection. *Statistics in Medicine*. 2003; 22:689–703. [PubMed: 12587100]
- Stallard N, Friede T. Flexible group-sequential designs for clinical trials with treatment selection. *Statistics in Medicine*. 2008; 27:6209–6227. [PubMed: 18792085]
- Simon R, Thall PF, Ellenberg SS. New designs for the selection of treatments to be tested in randomized clinical trials. *Statistics in Medicine*. 1994; 13:417–429. (discussion p. 447–451). [PubMed: 8023026]
- Thall PF. A review of phase 2–3 clinical trial designs. *Lifetime Data Analysis*. 2008; 14:37–53. [PubMed: 17763973]
- Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004; 60:684–693. [PubMed: 15339291]
- Thall PF, Simon R, Ellenberg SS. Two-stage selection and testing designs for comparative clinical trials. *Biometrika*. 1988; 75:303–310.
- Wang SK, Tsiatis AA. Approximately optimal one-parameter boundaries for group sequential trials. *Biometrics*. 1978; 43:193–199. [PubMed: 3567304]
- Wolff JE, Sajedi M, Brant R, Coppes MJ, Egeler RM. Choroid plexus tumours. *British Journal of Cancer*. 2002; 87:1086–1091. [PubMed: 12402146]
- Wrede B, Liu P, Ater J, Wolff JE. Second surgery and the prognosis of choroid plexus carcinoma—results of a meta-analysis of individual cases. *Anticancer Research*. 2005; 25:4429–4433. [PubMed: 16334121]
- Wrede B, Liu P, Wolff JE. Chemotherapy improves the survival of patients with choroid plexus carcinoma: a meta-analysis of individual cases with choroid plexus tumors. *Journal of Neurooncology*. 2007; 85:345–51.

Appendix

For outcomes from the likelihood (1) with parameters corresponding to fixed $\vec{\delta}^*$, the posterior of $\vec{\delta}$ under the Bayesian model has mean $\vec{\delta}^*$, aside from negligible effects of the noninformative iid priors on the treatment effect pairs $\gamma_{j,0}$, $j = 0, 1, \dots, J$, which will be ignored for simplicity. If $\vec{\delta}^*$ contains more than one δ_j^* that is in Δ_c then $\varphi(\vec{\delta}^*)$ can be reduced by moving each δ_j^* for which this is the case from Δ_c to Δ_0 until exactly one $\delta_j^* \in \Delta_c$ remains. If this single δ_j^* is in the interior of Δ_c or falls on either the vertical or horizontal lines on $\partial \Delta_c$, then $\varphi(\vec{\delta}^*)$ can be reduced by either decreasing $\delta_{j,1}^*$ or increasing $\delta_{j,2}^*$. Consequently, $\delta_j^* \in \mathcal{L}$. Similar reasoning shows that $\varphi(\vec{\delta}^*)$ is minimized if the other $J - 1$ δ_j^* 's in Δ_0 all equal $\mathbf{0}$.

Figure 1.

Historical mean pair $\mu_{\theta,0} = (0.47, 0.11)$, elicited targets pairs $\{\theta^{e,1}, \dots, \theta^{e,5}\}$, target line $\mathcal{L}$, and target set $\Theta_{\mathcal{L}}$ in the natural parameter domain, with $\theta = \theta(Z)$ evaluated at the reference covariate vector, corresponding to a 2-year-old patient with non-metastatic disease, complete resection and CPC histology.

Figure 2.

Transformed elicited targets, target line and target set in the Δ domain of comparative, covariate-adjusted treatment effects. Values in this figure were obtained from those in Figure 2 by applying the transformation $\vec{\theta} \delta = \tau(\theta) - \mu_{\gamma,0}$. The null set is $\Delta_0 = \{\delta : \delta_T \leq 0, \delta_Y \geq 0\}$.

Figure 3. Elicited values of $\text{Pr}(\text{toxicity})$ as a function of age, with the posterior mean of the fitted curve under the Bayesian fractional polynomial model.

Table 1

Historical population percent and elicited prior event-free survival (EFS) time probabilities for each of 16 patient prognostic subgroups defined in terms of age, whether the tumor was metastatic at first diagnosis (MET), whether surgery achieved a complete resection (CR), and the unfavorable histology choroid plexus carcinoma (CPC).

Age ≥ 3	MET	CR	CPC	Popn. Percent	Elicited Prior Pr(EFS >t years)			
					t=0.5	t=1.0	t=2.0	t=3.0
No	No	Yes	Yes	18.9	0.68	0.58	0.50	0.45
No	No	Yes	No	0	0.89	0.85	0.83	0.81
No	No	No	Yes	26.2	0.47	0.23	0.18	0.15
No	No	No	No	1.7	0.68	0.66	0.60	0.50
No	Yes	Yes	Yes	3.9	0.65	0.42	0.32	0.25
No	Yes	Yes	No	2.4	0.86	0.65	0.54	0.50
No	Yes	No	Yes	6.3	0.44	0.12	0.08	0.04
No	Yes	No	No	1.2	0.65	0.45	0.41	0.38
Yes	No	Yes	Yes	20.1	0.87	0.82	0.79	0.70
Yes	No	Yes	No	0	0.94	0.91	0.89	0.82
Yes	No	No	Yes	9.9	0.73	0.65	0.50	0.43
Yes	No	No	No	2.2	0.77	0.69	0.56	0.50
Yes	Yes	Yes	Yes	0.5	0.84	0.53	0.50	0.46
Yes	Yes	Yes	No	2.2	0.90	0.70	0.60	0.55
Yes	Yes	No	Yes	3.9	0.70	0.48	0.35	0.28
Yes	Yes	No	No	0.7	0.74	0.62	0.36	0.29

Table 2

The optimal design parameters (M_1 , M_2 , r_1 , r_2) minimize $\bar{E}(N)$ subject to $\alpha^* \leq 0.05$, $\zeta^* \leq 0.05$ and $\varphi^*(\vec{\delta}^*) \geq 0.80$ for all $\vec{\delta}^* \in \mathcal{D}$. Each target has true $\vec{\theta}^* = (\theta_0^*, \theta_1^*, \dots, \theta_j^*)$ with $\theta_j^* = \theta^*$ and $\theta_j^* = \theta_0^* = (0.47, .11)$ for all $j < J$. PET = Pr(Early termination after stage 1) and E(Dur) = expected trial duration, in years, assuming 30 patients per year accrual.

$J = 1 : M_1 = 14, M_2 = 17, r_1 = 0.510, r_2 = 14.790$									
Case	True $\theta^* = \theta_1$	φ^*	PET	$E(N_1)$	$E(N_2)$	$E(N)$	$\bar{E}(N)$	E(Dur)	
Null	(0.47, 0.11)	0.05	0.61	57.4	17.2	64.1	—	2.5	
Target 1	(0.65, 0.01)	0.94	0.01	62.0	30.2	91.9	78.0	2.9	
Target 2	(0.70, 0.05)	0.80	0.05	63.0	30.3	91.7	77.9	2.9	
Target 3	(0.80, 0.20)	0.83	0.05	66.5	32.2	97.2	80.6	3.0	
Target 4	(0.90, 0.40)	0.99	<0.01	69.5	35.1	104.5	84.3	3.1	
Target 5	(0.99, 0.99)	>0.99	<0.01	77.1	41.2	118.0	91.2	3.4	
$J = 2 : M_1 = 19, M_2 = 45, r_1 = 0.431, r_2 = 16.199$									
Case	True $\theta^* = \theta_2$	φ^*	PET	$E(N_1)$	$E(N_2)$	$E(N)$	$\bar{E}(N)$	E(Dur)	
Null	(0.47, 0.11)	0.05	0.36	66.1	53.2	99.9	—	3.7	
Target 1	(0.65, 0.01)	0.96	0.01	69.4	67.2	135.7	117.8	4.3	
Target 2	(0.70, 0.05)	0.82	0.03	70.2	68.0	135.9	117.9	4.3	
Target 3	(0.80, 0.20)	0.82	0.05	72.2	70.1	139.0	119.5	4.3	
Target 4	(0.90, 0.40)	0.98	0.01	75.2	75.7	149.9	124.9	4.5	
Target 5	(0.99, 0.99)	>0.99	<0.01	79.3	87.2	166.5	133.2	4.8	
$J = 3 : M_1 = 28, M_2 = 67, r_1 = 0.605, r_2 = 17.376$									
Case	True $\theta^* = \theta_3$	φ^*	PET	$E(N_1)$	$E(N_2)$	$E(N)$	$\bar{E}(N)$	E(Dur)	
Null	(0.47, 0.11)	0.04	0.31	80.5	79.0	134.7	—	4.8	
Target 1	(0.65, 0.01)	0.94	0.02	83.5	90.8	172.1	153.4	5.4	
Target 2	(0.70, 0.05)	0.80	0.05	84.3	91.8	171.2	153.1	5.4	
Target 3	(0.80, 0.20)	0.80	0.05	86.0	95.5	176.8	155.7	5.5	
Target 4	(0.90, 0.40)	0.95	0.01	88.4	102.1	189.1	161.9	5.7	
Target 5	(0.99, 0.99)	>0.99	<0.01	92.2	118.0	210.0	172.4	6.1	

Table 3

Robustness of the optimal design for $J = 3$ experimental treatments to multiplicative contamination of T by a gamma noise variable G with mean 1 and standard deviation (sd) σ_G . The observed event time is GT , and the ratio $\sigma_{GT}/\sigma_T = \sigma_{GT}/\sigma_T$ is the multiplicative factor by which the sd of T is increased by G .

Case	σ_G	σ_{GT}/σ_T	ϕ^*	PET	E(N)
Null $\mu_T = .62$	0	1.00	.04	0.31	134.7
	.2	1.03	.05	0.31	134.0
	.4	1.11	.05	0.28	136.3
	.6	1.23	.06	0.29	132.5
Target 1 $\mu_T = 1.052$	0	1.00	.94	0.02	172.1
	.2	1.04	.93	0.02	171.3
	.4	1.17	.92	0.03	168.7
	.6	1.35	.90	0.03	165.4
Target 2 $\mu_T = 1.181$	0	1.00	.80	0.05	134.7
	.2	1.05	.79	0.05	134.0
	.4	1.19	.75	0.06	136.3
	.6	1.39	.73	0.07	132.5
Target 3 $\mu_T = 1.476$	0	1.00	.80	0.05	176.8
	.2	1.07	.79	0.04	176.5
	.4	1.25	.73	0.07	170.4
	.6	1.51	.65	0.08	166.1
Target 4 $\mu_T = 1.885$	0	1.00	.95	0.01	189.1
	.2	1.10	.95	0.05	188.0
	.4	1.35	.94	0.02	185.1
	.6	1.68	.89	0.03	179.2
Target 5 $\mu_T = 2.857$	0	1.00	>.99	<0.01	210.0
	.2	1.19	>.99	<0.01	208.9
	.4	1.63	>.99	<0.01	205.9
	.6	2.18	0.99	<0.01	201.0