

A Hybrid Phase I-II/III Clinical Trial Design Allowing Dose Re-Optimization in Phase III

Andrew G. Chapple

Department of Statistics, Rice University, Houston, Texas, U.S.A.

email: agc6@rice.edu

and

Peter F. Thall

Department of Biostatistics, The University of Texas M.D. Anderson Cancer Center, Houston, Texas, U.S.A

email: rex@mdanderson.org

SUMMARY: Conventionally, evaluation of a new drug, A , is done in three phases. Phase I is based on toxicity to determine a “maximum tolerable dose” (MTD) of A , phase II is conducted to decide whether A at the MTD is promising in terms of response probability, and if so a large randomized phase III trial is conducted to compare A to a control treatment, C , usually based on survival time or progression free survival time. It is widely recognized that this paradigm has many flaws. A recent approach combines the first two phases by conducting a phase I-II trial, which chooses an optimal dose based on both efficacy and toxicity, and evaluation of A at the selected optimal phase I-II dose then is done in a phase III trial. This paper proposes a new design paradigm, motivated by the possibility that the optimal phase I-II dose may not maximize mean survival time with A . We propose a hybridized design, which we call phase I-II/III, that combines phase I-II and phase III by allowing the chosen optimal phase I-II dose of A to be re-optimized based on survival time data from phase I-II patients and the first portion of phase III. The phase I-II/III design uses adaptive randomization in phase I-II, and relies on a mixture model for the survival time distribution as a function of efficacy, toxicity, and dose. A simulation study is presented to evaluate the phase I-II/III design and compare it to the usual approach that does not re-optimize the dose of A in phase III.

KEY WORDS: Bayesian Design, Clinical Trial, Dose Finding, Phase I-II Clinical Trial, Phase III Clinical Trial

This paper has been submitted for consideration for publication in *Biometrics*

1. Introduction

After a new treatment agent, A , is identified in pre-clinical studies, conventional clinical drug development and evaluation is carried out in three phases (Cancer.org, 2018). In phase I, the aim is to identify a dose, called the “maximum tolerable dose” (MTD), having acceptable toxicity probability. Phase I trials typically are small, with a wide variety of designs, including the 3+3 algorithm (Storer, 1989), continual reassessment method (O’Quigley et al., 1990), and escalation with overdose control (Babb et al., 1998). Efficacy of A at the MTD then is evaluated in phase II using the estimated probability π_E of a short-term event (“response”), such as 50% solid tumor shrinkage or complete remission of leukemia. Most phase II designs compare $\pi_E(A)$ with A at the MTD to an assumed $\pi_E(C)$ of a conventional therapy, C . Phase II trials often are small, and may include an early stopping rule if $\pi_E(A)$ is poor compared to $\pi_E(C)$. If A is found to be promising in phase II, this may motivate a randomized phase III trial of A versus C based on a long term outcome, such as survival time.

Many phase II designs have been published. Simon et al. (1985) proposed a randomized selection design for two or more experimental treatments. For single-arm phase II trials, two-stage designs were proposed by Simon (1989) based on response, and by Bryant and Day (1995) based on response and toxicity. Bayesian sequential designs were proposed by Thall and Simon (1994) for a binary response, and by Thall et al. (1995) for monitoring multiple outcomes. Lee and Liu (2005) used predictive probabilities for futility rules, and Yin et al., (2012) used adaptive randomization to favor empirically better treatment arms.

It now is recognized widely that the conventional phase I \rightarrow phase II \rightarrow phase III paradigm has many flaws, and has led to many negative phase III trials. Two studies (Bio, 2016; Arrowsmith, 2011) showed that only about 50% of phase III trials yield an improvement over standard therapy. Seruga, et al. (2015) discussed causes of failure in phase III, including insufficient evidence of anti-disease activity in early phase trials, disagreements about how

phase II trials should be designed, and reliance on phase II efficacy events or other surrogates not associated with longer survival. Yuan, Nguyen, and Thall (2016, Chapter 1) discuss problems with the conventional phase I \rightarrow phase II paradigm, mainly due to limited sample sizes and ignoring efficacy when determining an MTD in phase I.

Many alternatives have been proposed that create hybrid designs by combining conventional phases, most commonly phase I-II or phase II-III. Thall (2008) reviewed phase II-III designs and discussed problems with the conventional phase II \rightarrow phase III paradigm. “Select-and-test” phase II-III designs, where two or more experimental agents are chosen in phase II and randomized against C in phase III while maintaining desired overall type I and type II error rates, are given by Thall et al. (1988), Schaid et al. (1990), Stallard and Todd (2003), and many others. A phase II-III design proposed by Inoue, Thall and Berry (2004) uses both an early efficacy (response) indicator, Y_E , and survival time, Y_S . Denote $\pi_E = \Pr(Y_E = 1)$, the probability density function (pdf) of Y_S by $f_S(t)$, and the conditional pdf of $[Y_S | Y_E]$ by $f_S(t | Y_E = y)$ for $y = 0, 1$. Their approach relies on a mixture model of the general form

$$f_S(t) = f_S(t | Y_E = 1)\pi_E + f_S(t | Y_E = 0)(1 - \pi_E). \quad (1)$$

Denote the indicator of early toxicity by Y_T and $\pi_T = \Pr(Y_T = 1)$. Because phase I designs use Y_T but ignore Y_E when choosing a MTD, they are likely to choose a dose having reasonable π_T but ineffectively low π_E . For example, consider a dose-finding scenario with five doses, true toxicity probabilities $(.05, .10, .20, .30, .35)$, and true efficacy probabilities $(.05, .10, .20, .30, .60)$. If the CRM is used with target toxicity probability $.30$, this most likely will select dose 4 as optimal. By ignoring Y_E , however, dose 5 is chosen less frequently, despite the fact that it has only a $.05$ higher π_T than dose 4 but doubles π_E from $.30$ to $.60$. Phase I-II designs are motivated, in part, by the desire to overcome this sort of problem. Examples include the two-stage design of Hoering et al. (2011), studying combination therapies (Huang et al., 2006), using the odds ratio between π_E and π_T (Yin et al, 2006), and

basing decisions on elicited numerical utilities of the possible elementary events determined by efficacy and toxicity (Thall and Nguyen, 2012). Thall and Cook (2004) proposed, and Thall et al. (2014) refined, the so-called “Eff-Tox” phase I-II design based on maximizing an estimate of an efficacy-toxicity trade-off function, $\phi(\pi_E, \pi_T)$. The function $\phi(\pi_E, \pi_T)$ increases in π_E , decreases in π_T , and quantifies the desirability of each probability pair (π_E, π_T) .

This article presents a new Bayesian hybrid design that combines a phase I-II design followed by a modified phase III design, based on both early and late outcomes. We will call this a phase I-II/III trial design. For simplicity, we will use survival time, Y_S , as the long term outcome, although progression-free survival (PFS) time will work in precisely the same way. Our approach relies on a mixture model for the distribution of Y_S that generalizes the model (1) by including both efficacy and toxicity indicators, (Y_E, Y_T) , to characterize early outcome. After phase I-II and an initial stage of phase III have been completed, the phase I-II/III design may re-optimize the dose of the experimental agent A based on mean survival time, μ_S . This approach hybridizes the phase I-II \rightarrow phase III paradigm, in which dose-finding for A is done using (Y_E, Y_T) in phase I-II, rather than using only Y_T for dose-finding as in the more conventional phase I \rightarrow phase II \rightarrow phase III paradigm.

Our proposed phase I-II/III design has $K \geq 3$ stages. In stage 1, a phase I-II trial is conducted based on the short term binary indicators (Y_E, Y_T) , including adaptive randomization (AR) among doses of A based on the dose desirability criterion ϕ . The use of AR reduces the risk of getting stuck at a suboptimal dose in phase I-II. It addresses the “exploration versus exploitation” or “stickiness” problem, which is well known in sequential analysis (Sutton and Barto, 1998; Azriel, et al., 2010). AR improves the reliability of our proposed phase I-II/III design because it obtains more data on doses that may be sub-optimal in terms of the phase I-II criterion ϕ based on (Y_E, Y_T) but optimal in terms of μ_S .

Denote A given at x by $A(x)$. At the end of phase I-II (stage 1), an optimal dose $\hat{x}_{ET}^{\text{opt}}$ of

A based on ϕ is determined. In stage 2, phase III begins with patients randomized fairly to C and $A(\hat{x}_{ET}^{\text{opt}})$. Phase I-II patients are followed to observe their times of death or follow up. After a pre-specified number of deaths, n_2^* , have been observed from patients receiving C or $A(\hat{x}_{ET}^{\text{opt}})$ in stage 2, all (Y_E, Y_T) and survival data of patients treated with A in stages 1 and 2 are used to determine an optimal dose \hat{x}_S^{opt} such that $A(\hat{x}_S^{\text{opt}})$ maximizes μ_S , for use in the rest of phase III. The re-optimized dose \hat{x}_S^{opt} may or may not be the same as $\hat{x}_{ET}^{\text{opt}}$. Stages 3, . . . , K are a randomized group sequential trial with up to $K - 2$ tests comparing the mean survival times of $A(\hat{x}_S^{\text{opt}})$ versus C . To provide a concrete illustration, for stage 1 we use the Eff-Tox phase I-II design of Thall et al. (2014), extended to include AR.

The rest of the paper is organized as follows. In Section 2, the data structure, models, and decision criteria are presented. Section 3 presents details of trial conduct. Section 4 describes possible decisions, outcomes, and potential consequences of re-optimizing dose versus the conventional approach of using $\hat{x}_{ET}^{\text{opt}}$ in phase III. Section 5 presents results of simulation study to compare the phase I-II/III design to the phase I-II \rightarrow phase III paradigm. Section 6 concludes with a discussion. A computer program to implement the phase I-II/III design is available on CRAN in the package *Phase123*.

2. Data Structure, Models, and Decision Criteria

Given raw doses $\mathbf{d} = (d_1, \dots, d_J)$ of the experimental agent A , denote the standardized doses by $x_j = (d_j - \bar{\mathbf{d}})/sd(\mathbf{d})$ for $j = 1, \dots, J$. Let Y_S^o denote the observed time to death or administrative censoring and $\delta = I(Y_S = Y_S^o)$. Denote the parameters for the distribution of $[Y_E, Y_T|x]$ by $\boldsymbol{\theta}_{ET}$, and the parameters for the distribution of $[Y_S|Y_E, Y_T, x]$ by $\boldsymbol{\theta}_S$.

The Eff-Tox design is reviewed in Web Appendix Section A. Briefly, for each $m = E, T$, and x_j , it is assumed that $\pi_m(x_j, \boldsymbol{\theta}_{ET}) = P(Y_m = 1 | x_j, \boldsymbol{\theta}_{ET}) = \text{logit}^{-1}\{\eta_m(x_j, \boldsymbol{\theta}_{ET})\}$, with $\eta_T(x_j, \boldsymbol{\theta}_{ET}) = \tau_{T,1} + \tau_{T,2}x_j$ and $\eta_E(x_j, \boldsymbol{\theta}_{ET}) = \tau_{E,1} + \tau_{E,2}x_j + \tau_{E,3}x_j^2$, with $\tau_{T,2} > 0$, so that $\pi_T(x_j, \boldsymbol{\theta}_{ET})$ increases with x_j , but $\pi_E(x_j, \boldsymbol{\theta}_{ET})$ may be non-monotone. An association

parameter ψ determines the joint distribution of (Y_E, Y_T) from their marginals using a copula, so $\boldsymbol{\theta}_{ET} = (\tau_{T,1}, \tau_{T,2}, \tau_{E,1}, \tau_{E,2}, \tau_{E,3}, \psi)$. These parameters are assumed to be independent with priors $\psi \sim N(0, 1)$, $\tau_{E,2} \sim N(0, .20)$, and $\tau_{m,r} \sim N(\tilde{\mu}_{m,r}, \tilde{\sigma}_{m,r}^2)$ for $m = E, T$ and $r = 1, 2$. Numerical values of $(\tilde{\mu}_{m,r}, \tilde{\sigma}_{m,r})$ for $m = E, T, r = 1, 2$ are determined from elicited means of $\pi_m(x_j, \boldsymbol{\theta}_{ET})$, for $j = 1, \dots, J, m = E, T$, and a desired prior effective sample size. Adaptive dose-finding decisions are based on a trade-off function $\phi(\pi_E, \pi_T)$ for $\boldsymbol{\pi} \in [0, 1]^2$.

Denote $\phi_j = \phi[E\{\pi_E(x_j, \boldsymbol{\theta}_{ET}), \pi_T(x_j, \boldsymbol{\theta}_{ET})\} | data]$ for each dose x_j at any point during phase I-II based on the current data. The estimated optimal dose in an Eff-Tox trial is $\hat{x}_{ET}^{\text{opt}} = \max_j \{\phi_j\}$. To extend this design to include AR, rather than choosing $\hat{x}_{ET}^{\text{opt}}$ for each cohort during phase I-II, we adaptively randomize the next cohort to dose x_j with probability

$$\frac{\exp\{(\phi_j - \bar{Q})/sd(\mathbf{Q})\}}{\sum_{r:\phi_r \in \mathbf{Q}} \exp\{(\phi_r - \bar{Q})/sd(\mathbf{Q})\}},$$

where \mathbf{Q} is the current set of posterior mean desirabilities, ϕ_j , of doses that are acceptably safe and efficacious. This shrinks the selection probability of less desirable doses toward 0 while allowing selection of doses that are suboptimal in terms of ϕ . After the (Y_E, Y_T) outcomes of all N_{ET} patients in phase I-II have been evaluated, $A(\hat{x}_{ET}^{\text{opt}})$ based on the final phase I-II data is moved forward to stage 2, which is the first portion of phase III.

In the phase I-II/III design, we define two different types of truly optimal doses of A . Let $\boldsymbol{\theta}_m^{\text{true}}$ denote an assumed true value of $\boldsymbol{\theta}_m$ for $m = ET$ or S . The truly optimal dose that maximizes $\phi\{\pi_E(x, \boldsymbol{\theta}_{ET}^{\text{true}}), \pi_T(x, \boldsymbol{\theta}_{ET}^{\text{true}})\}$ is x_{ET}^{opt} . The truly optimal dose that maximizes the mean survival time $\mu_S(x, \boldsymbol{\theta}_S^{\text{true}})$ is x_S^{opt} . Let $k = 1, \dots, K$ index the stages of the phase I-II/III trial. Thus, $k = 1$ indexes the phase I-II trial, $k = 2$ indexes the first portion of phase III at the end of which the dose of A may be re-optimized based on μ_S , and $k = 3, \dots, K$ index the subsequent group sequential stages in phase III for comparing $A(\hat{x}_S^{\text{opt}})$ to C . Thus, there are up to $K - 2$ group sequential comparisons in phase III. Let $\mathcal{D}_{I-II,k}$ and $\mathcal{D}_{III,k}$ denote the data for patients at the end of stage k , from the phase I-II and phase III portions of the trial,

respectively. Therefore, $\mathcal{D}_{I-II,1}$ consists only of the (x, Y_E, Y_T) data from phase I-II patients, while $\mathcal{D}_{I-II,2}$ also includes these patients' survival time data (Y_S^0, δ) , up to the time at which the decision of whether to switch the dose based on mean survival time is made. $\mathcal{D}_{III,1}$ does not exist because phase III has not begun in stage 1. The re-optimized dose \hat{x}_S^{opt} is chosen based on $\mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2}$, which includes all (x, Y_E, Y_T) and (Y_S^0, δ) data at the end of stage 2.

Since the phase I-II \rightarrow phase III paradigm uses $\hat{x}_{ET}^{\text{opt}}$ throughout phase III, the primary motivation for our design is the possibility that $x_{ET}^{\text{opt}} \neq x_S^{\text{opt}}$, and that re-optimizing the dose of A may produce larger μ_S by comparing C to $A(\hat{x}_S^{\text{opt}})$ rather than $A(\hat{x}_{ET}^{\text{opt}})$ in the group sequential trial. To evaluate the effects of re-optimizing the dose of A based on μ_S during the first part of phase III, we require models for $[Y_E, Y_T|x]$ and $[Y_S|Y_E, Y_T, x]$, in order to formulate a mixture model for $[Y_S|x]$. This will include the effects of x on the indicators (Y_E, Y_T) , and the effects of (Y_E, Y_T) and x on the hazard function of Y_S . Let $\pi(y_E, y_T | x, \boldsymbol{\theta}_{ET})$ denote the probability distribution of (Y_E, Y_T) at dose x , where $(y_E, y_T) \in \{0, 1\}$. Let $f_{S|E,T}(y_S | y_E, y_T, x, \boldsymbol{\theta}_S)$ denote the conditional pdf of Y_S given the early binary outcomes and dose x of A . The mixture pdf of Y_S for patients treated with $A(x)$ is

$$f_S(y_S | x, \boldsymbol{\theta}_S, \boldsymbol{\theta}_{ET}) = \sum_{y_E=0}^1 \sum_{y_T=0}^1 f_{S|E,T}(y_S | y_E, y_T, x, \boldsymbol{\theta}_S) \pi(y_E, y_T | x, \boldsymbol{\theta}_{ET}). \quad (2)$$

The conditional mean survival time given (y_E, y_T) of a patient treated with $A(x)$ is

$$\mu_{S,A(x)}(y_E, y_T, \boldsymbol{\theta}_S) = \int_0^\infty y_S f_{S|E,T}(y_S | y_E, y_T, x, \boldsymbol{\theta}_S) dy_S. \quad (3)$$

At the end of stage 2 of phase I-II/III, we choose \hat{x}_S^{opt} based on all observed data, where \hat{x}_S^{opt} maximizes the posterior mean of the parametric mean survival time

$$\mu_{S,A(x)}(\boldsymbol{\theta}_S, \boldsymbol{\theta}_{ET}) = \sum_{y_E=0}^1 \sum_{y_T=0}^1 \mu_{S,A(x)}(y_E, y_T, \boldsymbol{\theta}_S) \pi(y_E, y_T | x, \boldsymbol{\theta}_{ET}) \quad (4)$$

at $A(x)$. Conventionally (Y_E, Y_T) are used as surrogates for Y_S in choosing a dose $\hat{x}_{ET}^{\text{opt}}$ in phase I-II, but (Y_E, Y_T) are ignored when modeling survival in phase III.

We assume that the distribution of $[Y_S | Y_E, Y_T, x]$ has the Cox type hazard function

$$h(t | Y_E, Y_T, x, \boldsymbol{\theta}_S) = h_0(t) \exp\{\beta_1 x + \beta_2 x^2 - e^{\beta_E} Y_E + e^{\beta_T} Y_T\}, \quad t > 0. \quad (5)$$

We assume that β_1 , β_2 , β_E , and β_T are independent with identical non-informative $N(0, 100)$ priors. For robustness, we assume that the baseline hazard is piecewise exponential with $h_0(t) = \exp(\lambda_l)$ for $t \in (t_l, t_{l+1}]$ under the partition $t_0 = 0 < t_1 < \dots < t_{L+1} = \max(\mathbf{Y}_S^o)$. We allow the dimension L of the baseline hazard to vary, with prior $L \sim Poi(\zeta_S)$ and assume that the locations of the split points \mathbf{t} vary according to the even order statistics with a uniform distribution of size $2L$, as in Lee et al. (2015) and Chapple et al. (2017). This prevents obtaining intervals in $h_0(t)$ having few events for estimating λ_l . We suggest values of $\zeta_S \in \{3, 4, 5, 6, 7\}$, since most hazard shapes can be approximated very accurately with 1 to 5 pieces. The resulting posterior distribution is not sensitive to the choice of ζ_S in this range for sample sizes greater than 50. We assume a normal prior with mean 0 and variance 25 for λ_1 , denoted $\lambda_1 \sim N(0, 25)$, and borrow strength when $L > 1$ for adjacent intervals via the prior $\lambda_l \sim N(\lambda_{l-1}, \sigma_\lambda^2)$, with the prior of σ_λ proportional to $1/\sigma_\lambda$. The variance of λ_1 ensures posterior hazard values seen in practice, while maintaining prior non-informativeness.

Denote $\hat{\mu}_{S,A(x)} = E(Y_S | x, \mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2})$, the posterior mean survival time for $A(x)$ given the data from phases I-II and III at the end of stage 2. We compute this quantity under the mixture model (2) by estimating the posterior mean survival time

$$\hat{\mu}_{S,A(x)}(y_E, y_T | \mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2}) = E\{\mu_{S,A(x)}(y_E, y_T, \boldsymbol{\theta}_S) | \mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2}\}$$

for each pair $(y_E, y_T) \in \{0, 1\}$, under the formula (3), and computing the posterior mean

$$\hat{\pi}(y_E, y_T | x, \mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2}) = E\{\pi(y_E, y_T | x, \boldsymbol{\theta}_{ET}) | \mathcal{D}_{I-II,2} \cup \mathcal{D}_{III,2}\}$$

of each bivariate probability under the Eff-Tox model given in Web Appendix A.

Since there will be limited survival time follow up information after n_2^* events, the design only evaluates the means until the maximum observed patient follow up time. The trial is continued after n_2^* patient events using the dose \hat{x}_S^{opt} of A having the highest posterior

mean $\widehat{\mu}_S^{\text{opt}} = \max_{x_j} \{\widehat{\mu}_{S,A(x_j)}\}$. After making this decision, the design does not use data from patients who were treated at doses $x_j \neq \widehat{x}_S^{\text{opt}}$. After obtaining values of n_3^*, \dots, n_K^* from East 6 statistical software (2016), n_2^* is chosen such that the design can switch doses with high accuracy, but can still yield high power for phase III trials, given the truly optimal dose x_S^{opt} has been selected. Suitable values of n_2^* can be determined using the function `SimPhase123` in the package *Phase123*. This approach will result in a larger sample size of patients in the C arm being compared to $A(\widehat{x}_S^{\text{opt}})$ if $\widehat{x}_{ET}^{\text{opt}} \neq \widehat{x}_S^{\text{opt}}$. We use Markov chain Monte Carlo to obtain posterior distributions for $\boldsymbol{\theta}_{ET}$ and $\boldsymbol{\theta}_S$, using 2000 iterations and 1000 discarded as burnin. This gives good convergence of the parameters, shown by the posterior of L settling on one or two values as well as traceplots for the parameters $\boldsymbol{\lambda}|L$, $\boldsymbol{s}|L$, and the coefficients in the linear terms of the Eff-Tox and survival hazard models. A detailed account of computational algorithms used to simulate posterior samples is given in Web Appendix B.

3. Trial Conduct

In this section, we give specific rules for conducting a phase I-II/III clinical trial. Each of the computer functions described below is contained in the R package *Phase123*, available on CRAN, including documentation of inputs and examples. Additional information on the trial parameters is given in Web Appendix C and a tutorial on several of the functions is given in Web Appendix D. When designing a phase I-II/III trial, the statistician should consult with the physician to establish design parameters, such as ϕ , maximum sample sizes N_{ET} and N_S , and the number of comparative tests $K - 2$ following dose re-optimization. The group sequential boundaries for stopping the trial due to futility \underline{u}_k or superiority \bar{u}_k may be obtained using East 6 statistical software (2016), specifying a null value of μ_C , desired improvement Δ , type I error, power under the alternative, maximum sample size N_S , and information proportions for determining n_k^* for $k = 3, \dots, K$. If no futility decision is desired

at look k then $\underline{u}_k = 0$. The information proportions used to determine n_k^* should be large enough ($> 30\%$) to avoid making unreliable decisions based on a small amount of patient data if a dose is re-optimized for A .

The phase I-II/III design parameters must be calibrated to obtain good operating characteristics (OCs) under a reasonable array of possible scenarios. A smaller value of n_2^* ($\leq 20\%$ of the total information proportion) may be obtained by simulating the phase I-II/III trial under sets of different (a) Eff-Tox scenarios quantifying effects of x on (Y_E, Y_T) , (b) effects of (x, Y_E, Y_T) on survival, and (c) survival distributions. The stage 2 sample size n_2^* should be set by examining the design's OCs for several different values, to find n_2^* (1) large enough to give a high probability of selecting the optimal dose, but (2) small enough so, given that the design switches to a true optimal dose in stage 2, it has good generalized power figures. This can be done using the function *SimPhase123*. Specific rules for conducting a phase I-II/III trial are as follows:

- (1) Enroll the first cohort of patients in the phase I-II portion at the lowest dose. For each subsequent cohort until N^F patients have been treated, use the function *AssignEffTox* to obtain the next dose to give.
- (2) Once N^F patients have been enrolled in phase I-II, use the function *RandomEffTox* to adaptively randomize the next cohort of patients among acceptable doses, which allows doses that are empirically suboptimal in terms of $\phi(\pi_E, \pi_T)$ to be chosen.
- (3) After N_{ET} patients have been enrolled in phase I-II and their efficacy and toxicity outcomes have been evaluated, use the function *AssignEffTox* to obtain the dose $\hat{x}_{ET}^{\text{opt}}$ to continue to phase III.
- (4) Start phase III, randomizing patients equally between C and $A(\hat{x}_{ET}^{\text{opt}})$.
- (5) After n_2^* deaths have been observed, use the function *Reoptimize* to determine the dose \hat{x}_S^{opt} to continue with for the remainder of the trial.

- (6) Remove any patients treated with $\widehat{x}_{ET}^{\text{opt}}$ from consideration if the dose was switched and begin randomizing patients between C and $A(\widehat{x}_S^{\text{opt}})$.
- (7) For each stage $k = 3, \dots, K$, after n_k^* deaths occur, do two-sided tests for superiority or futility using the logrank test in R. Denoting the Z -score corresponding to the logrank statistic by Z , for futility bound \underline{u}_k and superiority bound \bar{u}_k , stop the trial if

$$|Z| > \bar{u}_k \text{ for superiority} \quad \text{or} \quad |Z| < \underline{u}_k \text{ for futility.}$$

- (8) Stop accrual after N_S patients have been enrolled in the phase III portion, including patients treated with a dose that is no longer considered optimal.

4. Possible Trial Outcomes

Before presenting our simulation results, we discuss possible design decisions and comment on each under different true states of nature. Because the phase I-II/III design may change the phase I-II selected dose of A in phase III before comparing A to C , the sequence of decisions that it makes may be correct and optimal, correct but suboptimal, wrong, or disastrously wrong, depends on x_S^{opt} and x_{ET}^{opt} , their estimates, and whether $\mu_{A(x_S^{\text{opt}})} = \mu_C$ or $\mu_{A(x_S^{\text{opt}})} \geq \mu_C + \Delta$. Since more than one dose of A may provide the desired improvement in μ_S of at least Δ over C , we denote the set of all such doses by $X^{\text{opt}} = \{x_j : \mu_{A(x_j)} \geq \mu_C + \Delta\}$. We define the generalized power (GP) to be the probability of (1) selecting a dose $x_j \in X^{\text{opt}}$ in stage 2 and (2) declaring $A(x_j)$ superior to C in one of stages $3, \dots, K$. The GP is the sum over $x_j \in X^{\text{opt}}$ of the probability of selecting x_j and declaring $A(x_j)$ superior to C . If X^{opt} contains more than one dose, then the GP is larger than the probability of the best possible decision, which is to select the optimal dose $x_S^{\text{opt}} \in X^{\text{opt}}$ that maximizes μ_S with A and declare $A(x_S^{\text{opt}})$ superior to C . We denote the probability of making this best decision by γ_1 and the GP by γ_2 . Thus, $\gamma_1 \leq \gamma_2$, with $\gamma_1 = \gamma_2$ if X^{opt} contains exactly one dose, which in this case must be x_S^{opt} .

To help sort this out, Table 1 provides explanatory comments on scenarios in stages $k = 1$ (phase I-II) and $k = 2$ (the first portion of phase III) regarding the true relationship between the optimal doses x_S^{opt} and x_{ET}^{opt} and their posterior estimates \hat{x}_S^{opt} and $\hat{x}_{ET}^{\text{opt}}$. If $\hat{x}_{ET}^{\text{opt}} = \hat{x}_S^{\text{opt}}$, then the phase I-II/III and phase I-II \rightarrow phase III designs make equivalent decisions. However, if $\hat{x}_{ET}^{\text{opt}} \neq \hat{x}_S^{\text{opt}}$, then switching provides a potential advantage. In this case survival data from patients who were treated with $\hat{x}_{ET}^{\text{opt}}$ during phase III are no longer relevant. Depending on the accrual rate, maximum sample size N_S , and number of patient events n_2^* needed to re-optimize dose, this may result in 30 to 100 patients being treated at doses no longer considered a part of the trial as phase III proceeds.

[Table 1 about here.]

[Table 2 about here.]

After choosing \hat{x}_S^{opt} , the phase I-II/III design makes group sequential decisions comparing $A(\hat{x}_S^{\text{opt}})$ to C , so the decisions in phase III depend on the selected \hat{x}_S^{opt} . But it may not be the case that $\hat{x}_S^{\text{opt}} = x_S^{\text{opt}}$. That is, the design may not choose the truly optimal dose in terms of mean survival time in stage 2. Table 2 lists possible decisions of a phase I-II/III design in stages $k = 2, \dots, K$ and how each decision may be viewed in terms of x_S^{opt} . Table 2 is ordered with the best outcomes listed first and the worst listed last, with outcomes 1, 2, and 3 being good and outcomes 4 and 5 being bad. In outcome 1, the design declares the dose that increases μ_S the most to be superior to C . In outcome 2, a dose of A is selected that provides a clinically meaningful improvement $\geq \Delta$ in μ_S compared to C , but the best dose of A is not chosen, so the decision is correct but the dose has not been truly optimized. Outcome 3 represents a correct decision, but it does not improve μ_S since it declares C superior to or equivalent to $A(\hat{x}_S^{\text{opt}})$. Outcome 4 gives a false positive result, including cases where the design wrongly chooses an inferior dose for which $\mu_{A(x_j)} < \mu_C$, which is worse than a conventional type I error. Outcome 5 represents the worst possible case, since not

only does the design wrongly conclude that the chosen dose gives $A(\widehat{x}_S^{\text{opt}})$ superior to C , but it might have obtained a successful trial result if it had correctly selected x_S^{opt} in stage 2.

These same decisions and interpretations are made in the conventional phase I-II \rightarrow phase III paradigm, with the difference that $\widehat{x}_S^{\text{opt}}$ is replaced with $\widehat{x}_{ET}^{\text{opt}}$. Compared to this conventional design, allowing the optimal dose to be switched in the phase I-II/III design makes selecting $\widehat{x}_S^{\text{opt}} = x_S^{\text{opt}}$ more likely, which increases the probabilities of outcomes 1 and 2 and decreases the probabilities of the disastrous outcome 5. Under outcome 3, the phase I-II/III design is likely to treat more patients because it is more likely to correctly pick the dose x_S^{opt} having the largest μ_S , thus making stopping the trial early for superiority of C or futility less likely. It will be more likely to switch to the dose having the longest mean survival time for outcome 4, however, which makes a false positive event more likely.

5. Simulation Study

To perform a simulation study comparing the phase I-II/III design to the phase I-II \rightarrow phase III paradigm, we first specify three different Eff-Tox scenarios, consisting of true efficacy and toxicity dose-probability vectors. We will use these to specify different relationships between (Y_E, Y_T) and Y_S . We evaluate the design with $J = 5$ doses using raw dose values $(d_1, \dots, d_5) = (1, 2, 3, 3.5, 5)$. For this study, each patient's (Y_E, Y_T) are evaluated in one month, and we assume for simplicity that no patients die before this month long window. For a dose x_j chosen in phase I-II, we test the null hypothesis $H_0 : \mu_C = \mu_{A(x_j)} = 24$ months versus $H_0 : \mu_C \neq \mu_{A(x_j)}$ with target $\mu_{A(x_S^{\text{opt}})} = 36$ months, a $\Delta = 12$ month improvement.

To implement phase I-II using the Eff-Tox design, the three equivalent (π_E, π_T) pairs used to establish the desirability function ϕ were $(.35, 0)$, $(.70, .40)$ and $(1, .75)$. The contour created by these three pairs is seen in Web Figure 1. The upper limit on π_T was $\bar{\pi}_T = .40$ and the lower limit on π_E was $\underline{\pi}_E = .30$. The threshold on the posterior probability that

$\pi_E > .30$ and $\pi_T < .40$ was set to be $p_E = p_T = .10$ for both acceptability rules. Patients were treated in cohorts of size 3, with up to $N_{ET} = 60$ patients enrolled in phase I-II (stage 1). We calibrated the phase I-II hyperparameters to have prior effective sample size .90 as suggested by Yuan, Nguyen and Thall (2016). We used prior mean toxicity probabilities of (.05, .10, .15, .20, .30) and mean efficacy probabilities (.20, .40, .60, .65, .70) for the five doses, to produce the hyperparameter means $(-4.23, 3.1, .02, 3.45, 0, 0)$ and standard deviations $(3.13, 3.12, 2.68, 2.69, .2, 1)$ for the prior of θ_{ET} . The *EffTox* program is freely available on the MDAnderson biostatistics software page.

For the phase I-II portion of the simulated trials, patients were treated in cohorts of size three and assigned doses after the previous cohort was fully evaluated, assuming an accrual rate of five patients per month with adaptive randomization begun after $N^F = 15$ patients. The three simulation scenarios' assumed true $\pi_E(x_j)$ and $\pi_T(x_j)$ are given in Table 3, with their selection percentages, true ϕ values, and numbers of patients treated, based on 5,000 simulated trials using the *EffTox* program.

[Table 3 about here.]

In the three Eff-Tox scenarios in Table 3, the respective optimal doses in terms of the tradeoff contour are doses = 3, 5, and 2. We only consider simulated phase I-II trials that advance to phase III, ignoring simulation replications where the trial stopped early. In scenario 1, doses 3 and 4 have nearly equivalent desirability, so we expect most patients in the phase I-II portion of the phase I-II/III trial to be treated at these two doses. In scenario 2, the highest dose 5 is considered optimal, most patients are treated at this dose, and it is selected in 49% of the simulations. In scenario 3, the dose 2 is optimal and doses 4 and 5 have unacceptably high toxicity probabilities, so we expect to treat fewer patients at these doses. The design treats the most patients at dose 2, which is selected with probability .51, but substantial numbers of patients are treated at doses 1 and 3. The use of AR assigns more

patients to doses 1 and 3, which allows the phase I-II /III design to better assess the functional relationship between dose and mean survival time. For the control group, we assume that the effects of toxicity and efficacy on overall survival are the same as those for the experimental group, and set the probabilities of toxicity and efficacy to be (.15, .40), (.10, .30), and (.35, .20) for the three Eff-Tox scenarios, respectively. For each of these simulation scenarios, we assume two different forms for the linear terms of the log hazard of Y_S . For $A(x)$, we assume $\eta_S(x, Y_E, Y_T) = \beta_0 + \beta_1 x + \beta_2 x^2 - \exp(\beta_E)Y_E + \exp(\beta_T)Y_T$. For the simulated data from the control group, we assume $\eta_S(C, Y_E, Y_T) = \beta_C - \exp(\beta_E)Y_E + \exp(\beta_T)Y_T$ and calibrate the additional parameter β_C so that we obtain the desired null value of 24 months for mean survival time. We first consider an exponential distribution with pdf $f(t|\rho) = (1/\rho) \exp(-t/\rho)$ where $\rho = \exp\{\eta_S(x, Y_E, Y_T)\}$, since the O'Brien Fleming group sequential bounds (O'Brien and Fleming, 1979) for the logrank test are based on this assumption. Later, we will consider several other distributions to evaluate the robustness of the methodology. Table 4 displays the six scenarios considered, which correspond to the different Eff-Tox scenarios listed in Table 3, as well as differing effects of dose, efficacy, and toxicity on survival time.

[Table 4 about here.]

These scenarios encompass several qualitatively and quantitatively different possible cases in connecting phase I-II to phase III. In scenario 1, the optimal dose in terms of μ_S is dose 3, which is selected with probability .29. In scenario 2, there is a large efficacy effect, leading to dose 5 being optimal in terms of μ_S , but this dose is only selected with probability .49 in phase I-II. Thus, we expect to see a large improvement in this scenario by using a phase I-II/III design. Similarly, in scenario 3, dose 3 is optimal in terms of μ_S , but is only selected with probability .18 in phase I-II. Scenario 4 represents a case with a large toxicity effect and small efficacy effect, making dose 1 optimal in terms of μ_S , but dose 1 is only chosen in 26% of the usual phase I-II trials. In scenario 5, dose 3 is the third best dose in terms of ϕ , but is

best in terms of μ_S , and it is only selected with probability .20. In this scenario, dose 4 also gives a significant improvement in μ_S compared to C , with $\mu_{A(x_4)} = 38.6$ months. In scenario 6, there is a large efficacy effect on overall survival, making dose 4 optimal in terms of both overall survival and ϕ , but dose 5 also has significantly improved survival compared to C . These two scenarios provide a basis for evaluating improvements in both γ_1 and the GP, γ_2 . To control the possibility of incorrectly switching due to chance outcomes, we do not allow the design to continue with a dose that had less than 6 patients treated. These scenarios also have varying effects of toxicity and efficacy on h_S , quantified by the coefficients β_E and β_T . This will evaluate the sensitivity of the method to these effects. Since the parameters (β_1, β_2) must be changed substantially to obtain similar μ_S values for different values of (β_E, β_T) , we do not perform a sensitivity analysis to these parameters within each scenario.

We assume that 10 patients, on average, are accrued each month during phase III, and that the phase III trial will begin 1 month after the phase I-II trial concludes. This waiting time could be increased to obtain longer survival follow up and thus improve the design's ability to re-optimize doses during stage $k = 2$. We enroll a maximum of $N_S = 500$ patients in phase III, which has up to three interim looks after $n_3^* = 200$, $n_4^* = 300$ and $n_5^* = 400$ deaths, with superiority decisions possible at each. We calibrated the stopping boundaries with East 6 statistical software (2016) using O'Brien-Fleming bounds (O'Brien and Fleming, 1979) with power .80 and type I error probability .05. We included a rule to determine if the trial should be stopped for futility, i.e. neither C nor $A(\hat{x}_S^{\text{opt}})$ is superior, after $n_4^* = 300$ deaths. The boundaries for declaring superiority of $A(\hat{x}_S^{\text{opt}})$ or C based on the standardized logrank statistics are $(\bar{u}_3, \bar{u}_4, \bar{u}_5) = (2.96, 2.53, 1.99)$, and the futility bound at the second look is $\underline{u}_4 = 1.001$. At the start of the phase III portion of the trial, we begin randomizing patients equally to $A(\hat{x}_{ET}^{\text{opt}})$ and C . After $n_2^* = 50$ deaths in the trial have occurred, we determine the dose \hat{x}_S^{opt} that patients receiving A should receive for the remainder of the

trial. This is the re-optimization step. Survival times for patients in phase I-II and phase III are generated after their toxicity and efficacy are scored, which does not allow the possibility that a patient may die before their short term indicators are seen.

[Table 5 about here.]

For each scenario and design, 5,000 simulation replications were performed. The simulation results are summarized in Table 5. In each of scenarios 1-4, $\gamma_1 = \gamma_2$, since there is one dose for which A is superior to C . Mean improvement in patient survival time with each design is denoted by \bar{W} , computed by averaging the differences between the true mean survival time with the selected dose of A and μ_C , if A is declared superior to C . In the simulations, \bar{W} is computed as the mean over $\{W^b, b = 1, \dots, 5000\}$, where

$$W^b = (\mu_{A(\hat{x}_S^{\text{opt}})} - \mu_C)I [A(\hat{x}_S^{\text{opt}}) \text{ is declared superior to } C \text{ in simulated trial } b].$$

Table 5 shows that, in general, the phase I-II/III design maintained type I error probability $\leq .05$ under H_0 and had a uniformly higher γ_1 and GP, γ_2 , compared to the conventional phase I-II \rightarrow phase III approach without dose re-optimization. The values of γ_1 , γ_2 , and \bar{W} are uniformly larger for the phase I-II/III design than for the conventional phase I-II \rightarrow phase III paradigm. The differences are extremely large in scenario 3, with an improvement of .73 in γ_1 , and a 9.68 month improvement in \bar{W} . The smallest advantage of the phase I-II/III design is seen in scenario 2, with an improvement of .09 for γ_1 and .90 for \bar{W} . In scenarios 5 and 6, where two doses of A give mean survival time larger than $\mu_C + \Delta = 36$ months, the phase I-II/III design provides respective improvements in γ_2 of .61 and .23, and improvements in γ_1 of .57 and .25. These scenarios illustrate the potential advantage of the phase I-II/III design compared to the conventional phase I-II \rightarrow phase III approach.

The phase I-II/III design does have the drawback that it requires treating more patients and longer trial durations, on average, than the conventional paradigm. Part of this required increase is due to the design correctly switching to the best dose of A in terms of overall

survival, which decreases the likelihood that a trial will stop early by declaring C to be superior or due to futility. This increase in required sample size and trial duration are the price paid for the much larger probability of a successful phase III trial in cases where dose switching increases mean survival time with A .

[Table 6 about here.]

Since phase I-II trials may have sample sizes ranging from 24 to 90 in practice, we chose the Eff-Tox sample size $N_{ET} = 60$ in the simulations as a practical compromise that obtains a reasonable amount of information in stage 1. Web Tables 1 and 2, seen in Web Appendix E, summarize additional simulations with $N_{ET} = 90$. Values of γ_1 , γ_2 and \bar{W} for the phase I-II/III design all increased substantially with N_{ET} for all six scenarios. This is because more information at different doses in phase I-II makes switching to the best dose in stage 2 more likely.

To assess robustness of the phase I-II/III design to different event time distributions, we evaluated its performance for two lognormal distributions, with variances .25 and 1, a Weibull distribution with increasing or decreasing hazard, with shape parameters 4 or .5, and a gamma distribution with scale parameter 2. The true coefficients of Y_E and Y_T in the hazard function's linear term were kept constant for each distribution, and the remaining constant parameters β_0 , β_C , β_1 , β_2 were adjusted to obtain similar true means as in the exponential distribution simulation study. We exponentiated the linear term for the gamma and Weibull distribution rate parameters, but did not do this for the lognormal distribution. The means under the null and alternative hypotheses for each distribution are given in Web Table 3. Table 6 summarizes the robustness study, showing that under the alternative hypothesis, for each distribution, the phase I-II/III design has uniformly higher values of γ_1 , γ_2 , and \bar{W} , with substantially higher values for scenarios 1,3 and 5. For the Weibull distribution with decreasing hazards in each scenario, the decrease in γ_1 and γ_2 for both designs is due to the

assumptions of the logrank test being grossly violated by a high early failure rate. Because so many patients have early failures, patients are not followed as long before the final group sequential test. For this distribution, however, the phase I-II/III design still improved the probability of selecting the optimal dose of A compared to the conventional paradigm by .49, .21, .15, .59, .46 and .14 in the six scenarios, respectively. Similar improvements are seen under the other distributions. This shows that the logrank test is not robust to the Weibull distribution with decreasing hazard. An extension of the phase I-II/III design might incorporate a robust group sequential test in place of the logrank test, to reduce the loss in power under a Weibull with decreasing hazard. The type I error constraints are nearly met for each distribution. Some slight inflation in α above .05 may be attributed to the proportional hazards assumption being violated. For each distribution, the phase I-II/III design treats more patients, on average, under both H_0 and H_1 , and has slightly longer trial duration, but makes the correct decision much more often.

6. Discussion

We have proposed a new drug development strategy, which we call a phase I-II/III design, that re-optimizes the dose of an experimental agent A chosen in phase I-II during phase III based on mean survival time. We use information from all patients treated with A , including their short term efficacy and toxicity indicators, dose assigned, and survival time information, in order to more accurately select the dose of A that provides the highest posterior mean survival time. The design is based on an assumed a mixture model for the survival time distribution that averages over the possible short term phase I-II outcomes. While we have used the Eff-Tox trade-off based phase I-II design for stage 1 of the phase I-II/III design, one could replace the Eff-Tox design with any phase I-II design based on (Y_E, Y_T) that uses some dose optimality criterion ϕ and includes AR. However, the necessary modifications of the design parameters and computer software to accommodate such a change would be

non-trivial. Similarly, a complicated but straightforward extension of the methodology may address the problem of possible deaths before evaluation of (Y_E, Y_T) .

The simulations shows that, under a range of alternative cases, the generalized power γ_2 , and probability γ_1 of the best possible decision, both are greatly increased by the phase I-II/III design compared to the phase I-II \rightarrow phase III paradigm. The phase I-II/III design also has a much lower probability of making the least desirable decision, where a suboptimal dose is chosen and a true treatment advance is missed. A drawback of the phase I-II/III design is that it requires more patients and a slightly longer trial duration, on average, compared to the phase I-II \rightarrow phase III paradigm. This seems like a very reasonable price to pay for the much larger values of γ_1 , γ_2 , and \bar{W} , in cases where re-optimizing the dose of the experimental agent increases its associated mean survival time.

7. Supplementary Materials

Web Appendices and tables, referenced in Sections 2, 3, and 5, are available at the *Biometrics* website on Wiley Online Library.

Acknowledgements

Peter Thall's research was supported by NCI grants R01 CA 83932 and P30 CA 016672.

Andrew Chapple's research was partially supported by the NIH grant 5T32-CA096520-07.

References

- Arrowsmith, J. (2011). Trial watch: Phase III and submission failures: 2007-2010. *Nature Review Drug Discovery*. **10**, 87–87.
- Azriel, D., Mandel, M., and Rinott, Y. (2011). The Treatment Versus Experimentation Dilemma in Dose-Finding Studies. *Journal of Statistical Planning and Inference* **141**, 2759-2768.

- Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine*. **17**, 1103-1120.
- BIO, Biomedtracker, Amplion. Clinical Development Success Rates 2006-2015. <https://www.bio.org/bio-industry-analysis-published-reports>. (Accessed December 20, 2017).
- Bryant, J. and Day, R. (1995) Incorporating toxicity considerations into the design of two-stage phase II clinical trials. *Biometrics*. **51**, 1372–1382.
- Cancer.org. (2018). What are the phases of clinical trials? <https://www.cancer.org/treatment/treatments-and-side-effects/clinical-trials/what-you-need-to-know/phases-of-clinical-trials.html> (Accessed 17 Jan. 2018).
- Chapple, A.G., Vannucci, M., Thall, P.F., and Lin, S.H. (2017). Bayesian variable selection for a semi-competing risks model with multiple components. *Journal of Computational Statistics and Data Analysis*. **112**, 170–185.
- Chu, Y., Pan, H., and Yuan, Y. (2016) Adaptive Dose Modification for Phase I Clinical Trials. *Statistics in Medicine* **35**, 3497–3508.
- East 6 (2016). Statistical software for the design, simulation and monitoring of clinical trials. Cytel Inc., Cambridge, MA.
- Green, P. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, **4**, 711-732.
- Hoering, A., LeBlanc, M., and Crowley, J. (2011). Seamless phase I-II trial design for assessing toxicity and efficacy for targeted agents. *Clinical Cancer Research*. **17**,**4**, 640–646.
- Huang, X., Biswas, X., Oki, Y., Issa, J.P., and Berry, D. (2006). A parallel phase I/II clinical trial design for combination therapies. *Biometrics*. **63**,**2**, 429–436.
- Inoue L.Y.T., Thall P.F., and Berry D.A. (2002). Seamlessly expanding a randomized phase

- II trial to phase III. *Biometrics* **58**, 823-831.
- Jin, I-H, Liu, S., Thall, P.F., and Yuan, Y. (2014) Using data augmentation to facilitate conduct of phase I/II clinical trials with delayed outcomes. *Journal of the American Statistical Association* **109**, 525-536.
- Lee, J. and Liu, D. (2008). A predictive probability design for phase II cancer clinical trials. *Clinical Trials*. **5**, 93–106.
- Lee, K., Haneuse, S., Schrag, D., and Dominici, F. (2015). Bayesian semiparametric analysis of semicompeting risks data: investigating hospital readmission after a pancreatic cancer diagnosis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **64**, 253-273.
- O'Brien, P.C. and Fleming, T.R. (1979). A Multiple Testing Procedure for Clinical Trials. *Biometrics*. **35**, 549–556.
- O'Quigley J., Pepe M., and Fisher L. (1990) Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics*, **46**, 33-48.
- Schaid, D. J., Wieand, S., and Therneau, T. (1990). Optimal two-stage screening designs for survival comparisons. *Biometrika* **3, 1**, 507–513.
- Seruga, B., Ocana, A., Amir, E., and Tannock, I.F. (2015). Failures in Phase III: Causes and Consequences. *Clinical Cancer Research* **21**, 4551-4560.
- Simon, R. (1989). Optimal two-stage designs for phase II clinical trials. *Controlled Clinical Trials*. **10**, 1–10.
- Simon, R., Wittes R.E., and Ellenberg, S.S. (1985) Randomized phase II clinical trials. *Cancer Treatment Reports*. **69**, 137581.
- Stallard, N. and Todd, S. (2003) Sequential designs for phase III clinical trials incorporating treatment selection. *Statistics in Medicine*. **22**, 286–703.
- Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics*. **45**, 925-937.

- Sutton, R. S. and Barto, A. G. (1998), Reinforcement Learning: An Introduction. *Cambridge, MA: MIT Press*.
- Thall, P.F. (2008). A review of phase 2-3 clinical trial designs. *Lifetime Data Analysis* **14**, 37–53.
- Thall, P.F., Herrick, R.C., Nguyen, H.Q., Venier, J.J., and Norris, J.C. (2014). Effective sample size for computing prior hyperparameters in Bayesian phase I-II dose-finding. *Clinical Trials*, **11**, 657-666.
- Thall, P.F. and Cook, J.D. (2004). Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* **60**, 684–693.
- Thall, P.F. and Nguyen H.Q. (2012). Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics* **22**, 785–801.
- Thall, P.F., Simon, R., and Ellenberg, S.S. (1988). Two-stage selection and testing designs for comparative clinical trials. *Biometrika* **75**, 303-310.
- Thall, P.F. and Simon, R. (1994). Practical Bayesian guidelines for phase IIB clinical trials. *Biometrics*. **50**, 337–349.
- Thall, P.F., Simon, R.M., and Estey, E.H. (1995). Bayesian sequential monitoring designs for single arm clinical trials with multiple outcomes. *Statistics in Medicine*. **14**, 357–379.
- Yin, G., Chen, N., and Lee, J. (2012). Phase II trial design with Bayesian adaptive randomization and predictive probability. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. **61**, 219-235.
- Yin, G., Li, Y., and Ji, Y. (2006). Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics* **62**, 777–787.
- Yuan, Y., Nguyen, H., and Thall, P.F. (2016). Bayesian Designs for Phase I-II Clinical Trials. *Chapman and Hall/CRC Biostatistics Series*.

Received February 2018. Revised April 2018. Accepted April 2018.

Table 1

Possible relationships between \hat{x}_{ET}^{opt} , x_{ET}^{opt} , \hat{x}_S^{opt} , and x_S^{opt} , including comments related to the phase I-II \rightarrow phase III and phase I-II/III designs. $\hat{x}_{ET}^{opt} \stackrel{?}{=} x_{ET}^{opt}$ refers to whether or not the optimal dose is selected at the end of phase I-II based on ϕ , and $x_{ET}^{opt} \stackrel{?}{=} x_S^{opt}$ refers to whether the optimal dose based on μ_S is the same as that based on ϕ

$\hat{x}_{ET}^{opt} \stackrel{?}{=} x_{ET}^{opt}$	$x_{ET}^{opt} \stackrel{?}{=} x_S^{opt}$	Comments
$\hat{x}_{ET}^{opt} = x_{ET}^{opt}$	$x_{ET}^{opt} = x_S^{opt}$	The optimal dose in terms of μ_S was selected in phase I-II, so it is not desirable to switch doses at stage 2. In this scenario, the phase I-II/III design cannot provide an improvement over phase I-II \rightarrow III.
$\hat{x}_{ET}^{opt} = x_{ET}^{opt}$	$x_{ET}^{opt} \neq x_S^{opt}$	The dose selected in phase I-II is optimal in terms of ϕ but is not optimal in terms of μ_S . This illustrates the advantage of the phase I-II/III design over phase I-II \rightarrow III design.
$\hat{x}_{ET}^{opt} \neq x_{ET}^{opt}$	$x_{ET}^{opt} = x_S^{opt}$	The dose selected in phase I-II is suboptimal based on ϕ , but the optimal doses in terms of ϕ and μ_S are identical. This scenario illustrates the advantage of the phase I-II/III design over phase I-II \rightarrow III design.
$\hat{x}_{ET}^{opt} \neq x_{ET}^{opt}$	$x_{ET}^{opt} \neq x_S^{opt}$	The dose selected in phase I-II is suboptimal based on ϕ , but the optimal doses in terms of ϕ and μ_S are not identical. This scenario illustrates the advantage of the phase I-II/III design over phase I-II \rightarrow III design.

Table 2

Possible Phase I-II/III Trial Outcomes, O . x_S^{opt} is the truly optimal dose in terms of $\mu_{A(x)}$ and Δ is the desired improvement over μ_C . Column 2 gives the two trial decisions, the first row for selecting \hat{x}_S^{opt} and the second row for determining superiority, inferiority, or futility, with $A(\hat{x}_S^{\text{opt}}) \succ C$ indicating that $A(\hat{x}_S^{\text{opt}})$ is declared superior to C , and $C \succeq A(\hat{x}_S^{\text{opt}})$ indicating that the trial is stopped due to either superiority of C or futility.

O	Decision	Truth	Comments
1	$\hat{x}_S^{\text{opt}} = x_S^{\text{opt}}$ $A(\hat{x}_S^{\text{opt}}) \succ C$	$\mu_{A(x_S^{\text{opt}})} > \mu_C + \Delta$	This is the generalized power event at the optimal dose x_S^{opt} . The design correctly selects x_S^{opt} as optimal and declares $A(x_S^{\text{opt}})$ superior to C .
2	$\hat{x}_S^{\text{opt}} \neq x_S^{\text{opt}}$ $A(\hat{x}_S^{\text{opt}}) \succ C$	$\mu_{A(\hat{x}_S^{\text{opt}})} > \mu_C + \Delta$ $\mu_{A(x_S^{\text{opt}})} > \mu_{A(\hat{x}_S^{\text{opt}})}$	This is a generalized power event in a case where the design correctly concludes $A(\hat{x}_S^{\text{opt}})$ is superior to C but \hat{x}_S^{opt} is suboptimal, so it could have improved survival more had it chosen the truly optimal dose x_S^{opt} .
3	$\hat{x}_S^{\text{opt}} = \text{any } x_j$ $C \succeq A(\hat{x}_S^{\text{opt}})$	$\mu_{A(x_S^{\text{opt}})} \leq \mu_C$	This is a correct conclusion, but the phase I-II/III design will require an increased sample size compared to the phase I-II \rightarrow III design due to correctly switching to x_S^{opt} .
4	$\hat{x}_S^{\text{opt}} = \text{any } x_j$ $A(\hat{x}_S^{\text{opt}}) \succ C$	$\mu_{A(x_S^{\text{opt}})} \leq \mu_C$	This is a false positive conclusion. While the design may pick the best dose of A , it incorrectly concludes that A at that dose is superior to C .
5	$\hat{x}_S^{\text{opt}} \neq x_S^{\text{opt}}$ $A(\hat{x}_S^{\text{opt}}) \succ C$	$\mu_{A(\hat{x}_S^{\text{opt}})} \leq \mu_C$ $\mu_{A(x_S^{\text{opt}})} \geq \mu_C + \Delta$	This is a disastrous false negative conclusion. The design chooses a suboptimal dose based on μ_S and incorrectly concludes $A(\hat{x}_S^{\text{opt}})$ is inferior to C , instead of correctly selecting x_S^{opt} and declaring $A(x_S^{\text{opt}})$ superior to C .

Table 3

Eff-Tox Scenarios. True outcome probabilities $(\pi_E, \pi_T)^{TR}$, desirabilities $\phi\{(\pi_E, \pi_T)^{TR}\}$, and operating characteristics for the usual (non-adaptively randomized) EffTox phase I-II trial design.

Scenario	Value	1	2	3	4	5
1	$(\pi_E, \pi_T)^{TR}$	(.20, .10)	(.40, .15)	(.60, .25)	(.65, .35)	(.70, .50)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.37	-.13	.05	-.01	-.13
	% Selected	3	26	29	27	13
	# Treated	6.4	16.0	16.1	12.1	9.0
2	$(\pi_E, \pi_T)^{TR}$	(.2, .05)	(.25, .08)	(.35, .10)	(.40, .15)	(.55, .20)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.30	-.26	-.14	-.13	.04
	% Selected	5	11	18	16	49
	# Treated	8.4	9.1	10.1	9.8	22.4
3	$(\pi_E, \pi_T)^{TR}$	(.40, .10)	(.50, .15)	(.60, .35)	(.65, .60)	(.70, .70)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.06	.03	-.09	-.35	-.40
	% Selected	26	51	20	2	0
	# Treated	16.7	27.3	11.6	3.3	0.8

Table 4

Simulation Parameters. True Eff-Tox parameters β^{TR} corresponding to the phase I-II scenarios in Table 3, and true mean survival time of each $A(x_1), \dots, A(x_5)$ for each phase I-II/III scenario's null and alternative hypotheses.

Scenario	Eff-Tox Scen	Hyp	$(\beta_1, \beta_2, e^{\beta_E}, e^{\beta_T}, \beta_0)^{TR}$	$(\mu_{A(x_1)}, \mu_{A(x_2)}, \mu_{A(x_3)}, \mu_{A(x_4)}, \mu_{A(x_5)})^{TR}$
1	1	Null	(.1, -.5, .5, .5, 2.9)	(8.3, 17.9, 24, 22.5, 9.8)
		Alt	(.25, -2, .5, .5, 3.4)	(1, 14.5, 36.2, 28.3, 1)
2	2	Null	(.1, -.1, 1, .5, 2.6)	(14.0, 17.8, 21.9, 23, 24)
		Alt	(.5, 0, 1, .5, 2.3)	(7.1, 10.3, 16.0, 19.5, 36)
3	2	Null	(.1, -.5, .3, 1, 3.1)	(9.5, 18.5, 24, 22.5, 10.4)
		Alt	(.1, -1, .3, 1, 3.6)	(6.9, 24.7, 38, 33.1, 6.3)
4	3	Null	(-.3, .3, .3, 1, 2.3)	(24, 13.6, 8.9, 6.8, 7.8)
		Alt	(-.1, .3, .3, 1, 3.0)	(38, 24.6, 18.4, 15.0, 21.1)
5	3	Null	(.1, -.5, .3, .1, 3.0)	(9.3, 18.7, 24, 22.7, 10.4)
		Alt	(.1, -1, .3, .1, 3.6)	(7.8, 28.8, 44, 38.6, 7.4)
6	1	Null	(.75, -.5, .3, .25, 2.8)	(3.2, 10.1, 20.4, 24, 20.4)
		Alt	(1, -.6, .3, .25, 3.3)	(3.0, 12.9, 31.8, 40, 36.6)

Table 5

Simulation results. α is the probability of a type I error or concluding an inferior version of A is better than C under the null. γ_1 is the generalized power at $A(x_S^{opt})$ (probability of selecting the best dose x_S^{opt} and declaring it to be superior to C) under the alternative hypothesis. γ_2 is the generalized power (probability of selecting any truly superior dose of A and declaring it superior to C). \bar{W} is the mean improvement in patient survival under the alternative hypothesis, \bar{Dur} is the mean trial duration, and \bar{N} is the mean sample size.

Scenario	Design	Alternative Hypothesis					Null Hypothesis			
		\bar{W}	γ_1	γ_2	\bar{Dur}	\bar{N}	α	\bar{Dur}	\bar{N}	
1	Phase I-II \rightarrow phase III	4.04	.29	.29	4.05	431.1	–	.02	4.05	461.5
	Phase I-II/III	10.15	.83	.83	4.73	479.2	–	.03	4.32	492.0
2	Phase I-II \rightarrow phase III	7.87	.66	.66	4.29	459.2	–	.06	4.18	489.4
	Phase I-II/III	8.97	.75	.75	4.45	470.7	–	.02	4.18	489.9
3	Phase I-II \rightarrow phase III	1.83	.06	.06	3.10	355.0	–	< .01	3.28	385.7
	Phase I-II/III	11.51	.79	.79	4.56	476.9	–	.04	4.22	485.6
4	Phase I-II \rightarrow phase III	3.52	.25	.25	4.28	475.8	–	.05	3.48	407.8
	Phase I-II/III	5.86	.42	.42	4.30	472.0	–	.05	3.81	442.0
5	Phase I-II \rightarrow phase III	5.61	.21	.25	3.98	428.5	–	.01	3.83	440.4
	Phase I-II/III	16.71	.68	.88	4.24	464.4	–	.03	4.37	493.9
6	Phase I-II \rightarrow phase III	9.46	.34	.52	4.16	447.9	–	.02	4.16	466.5
	Phase I-II/III	12.67	.59	.75	4.53	472.7	–	.04	4.39	494.0

Table 6

Robustness simulation results. α = probability of a type I error or concluding that an inferior version of A is better than C under the null. γ_1 = generalized power at x_S^{opt} under the alternative hypothesis (probability of selecting the best dose x_S^{opt} and declaring $A(x_S^{opt})$ superior to C). γ_2 = generalized power (probability of selecting any superior dose and declaring $A(\hat{x}_S^{opt})$ superior to C). \bar{W} is the mean improvement in survival time under the alternative hypothesis. \bar{Dur} and \bar{N} are the mean trial duration and sample size, respectively.

Scenario	Distribution	Phase I-II \rightarrow phase III				Phase I-II/III Design			
		\bar{W}	γ_1	γ_2	α	\bar{W}	γ_1	γ_2	α
1	Lognormal, $\sigma = .5$	7.72	.30	.30	.02	14.56	.80	.80	.03
	Lognormal, $\sigma = 1$	6.18	.30	.30	.02	14.88	.85	.85	.03
	Weibull Increasing	6.37	.30	.30	.02	12.00	.77	.77	.03
	Weibull Decreasing	1.96	.13	.13	.02	4.14	.32	.32	.02
	Gamma	5.02	.30	.30	.02	11.72	.89	.89	.04
2	Lognormal, $\sigma = .5$	12.18	.73	.73	.05	15.14	.91	.91	.05
	Lognormal, $\sigma = 1$	11.70	.70	.70	.06	13.78	.82	.82	.03
	Weibull Increasing	9.08	.73	.73	.06	11.20	.90	.90	.06
	Weibull Decreasing	3.98	.32	.32	.05	4.75	.38	.38	.03
	Gamma	9.11	.73	.73	.06	11.30	.90	.90	.03
3	Lognormal, $\sigma = .5$	2.17	.06	.06	< .01	13.91	.85	.85	.04
	Lognormal, $\sigma = 1$	2.02	.06	.06	< .01	12.91	.81	.81	.03
	Weibull Increasing	2.14	.06	.06	< .01	13.63	.84	.84	.03
	Weibull Decreasing	.90	.03	.03	< .01	4.94	.32	.32	.03
	Gamma	2.07	.06	.06	< .01	13.23	.86	.86	.03
4	Lognormal, $\sigma = .5$	4.50	.32	.32	.05	7.23	.52	.52	.05
	Lognormal, $\sigma = 1$	3.86	.28	.28	.05	5.61	.40	.40	.06
	Weibull Increasing	4.67	.33	.33	.04	6.61	.47	.47	.04
	Weibull Decreasing	4.14	.28	.28	.06	6.21	.42	.42	.07
	Gamma	4.62	.31	.31	.04	6.93	.47	.47	.04
5	Lognormal, $\sigma = .5$	6.63	.21	.25	.01	18.19	.76	.96	.03
	Lognormal, $\sigma = 1$	5.74	.21	.25	.01	16.78	.69	.87	.03
	Weibull Increasing	7.41	.21	.25	.01	15.88	.64	.77	.03
	Weibull Decreasing	4.00	.16	.18	.01	12.27	.48	.63	.03
	Gamma	6.73	.21	.25	.01	18.73	.80	.92	.03
6	Lognormal, $\sigma = .5$	11.92	.34	.52	.04	15.56	.58	.83	.06
	Lognormal, $\sigma = 1$	11.50	.34	.52	.02	15.66	.60	.80	.04
	Weibull Increasing	10.29	.34	.52	.04	12.76	.54	.77	.05
	Weibull Decreasing	5.46	.21	.31	.02	6.63	.32	.38	.03
	Gamma	10.90	.34	.52	.02	14.80	.59	.83	.04