STATISTICS IN MEDICINE
Statist. Med. 2007; **26**:4687–4702
Published online 10 April 2007 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/sim.2894

Bayesian and frequentist two-stage treatment strategies based on sequential failure times subject to interval censoring

Peter F. Thall[1,*,†], Leiko H. Wooten[1], Christopher J. Logothetis[2], Randall E. Millikan[2] and Nizar M. Tannir[2]

[1]*Department of Biostatistics, University of Texas, Houston, TX, U.S.A.*
[2]*Department of Genitourinary Medical Oncology, M.D. Anderson Cancer Center, University of Texas, 1515 Holcombe Boulevard, Houston, TX 77030, U.S.A.*

SUMMARY

For many diseases, therapy involves multiple stages, with the treatment in each stage chosen adaptively based on the patient's current disease status and history of previous treatments and clinical outcomes. Physicians routinely use such multi-stage treatment strategies, also called dynamic treatment regimes or treatment policies. We present a Bayesian framework for a clinical trial comparing two-stage strategies based on the time to overall failure, defined as either second disease worsening or discontinuation of therapy. Each patient is randomized among a set of treatments at enrollment, and if disease worsening occurs the patient is then re-randomized among a set of treatments excluding the treatment received initially. The goal is to select the two-stage strategy having the largest average overall failure time. A parametric model is formulated to account for non-constant failure time hazards, regression of the second failure time on the patient's first worsening time, and the complications that the failure time in either stage may be interval censored and there may be a delay between first worsening and the start of the second stage of therapy. Four different criteria, two Bayesian and two frequentist, for selecting a best strategy are considered. The methods are applied to a trial comparing two-stage strategies for treating metastatic renal cancer, and a simulation study in the context of this trial is presented. Advantages and disadvantages of this design compared to standard methods are discussed. Copyright © 2007 John Wiley & Sons, Ltd.

KEY WORDS: adaptive design; Bayesian design; clinical trials; dynamic treatment regime; interval censoring; simulation; treatment policy

1. INTRODUCTION

For many diseases, therapy involves multiple stages. Each stage begins with evaluation of the patient's disease status, and the physician chooses a treatment based on this evaluation and

*Correspondence to: Peter F. Thall, Department of Biostatistics, University of Texas, Houston, TX, U.S.A.
†E-mail: rex@mdanderson.org

Contract/grant sponsor: NCI; contract/grant number: RO1-CA-83932

Received 28 April 2006
Accepted 28 February 2007

the patient's history of previous treatments and therapeutic outcomes. Diseases for which multi-stage treatment strategies are used include cancer, psychological disorders and high blood pressure. In a given stage, a treatment may be a combination of drugs, radiotherapy and surgery in oncology; drugs and an intervention such as psychotherapy for a behavioural disorder; or a drug plus an exercise and diet regime for treating high blood pressure. There is a growing statistical literature on outcome-adaptive strategies for choosing a patient's treatments sequentially based on the patient's history of treatments and outcomes. Such algorithms have been referred to variously as 'dynamic treatment regimes' [1, 2], 'clinical strategies' [3], 'adaptive treatment strategies' [4], 'treatment policies' [5] or 'multi-course treatment strategies' [6, 7].

Physicians often must make treatment decisions in the face of great uncertainty. What works for one patient may not work for another, and often it is unknown why a treatment succeeded or failed for a given patient. Since this may be due to effects of latent variables specific to the patient, a common multi-stage strategy is to repeat a treatment that previously achieved a desired clinical outcome for the patient, and otherwise switch to a different treatment. This sort of 'switch away' strategy reflects the idea that the most useful information for predicting whether a given treatment will succeed for a given patient is whether that treatment has been successful previously for that patient. For many types of cancer, therapy begins with one or more courses of a 'frontline' treatment. The patient's disease is evaluated at the end of each course or according to a planned schedule. If an evaluation shows that, compared to a pre-treatment baseline evaluation, the patient's cancer is stable or a partial remission has been achieved, then the frontline treatment is continued. If the cancer has worsened compared to baseline ('progressed'), then it is common medical practice to switch to a different 'salvage' treatment. The terminology 'frontline,' 'salvage' and 'progression' from oncology corresponds to analogous treatments and events in many similar medical settings. For example, if an anti-psychotic drug A given initially fails due to the patient experiencing a psychotic episode and drug B is then given, one may identify A as the frontline treatment, B as the salvage treatment, and the psychotic episode is analogous to progression of cancer since both are a worsening of the disease state.

In this paper, we present a Bayesian framework for a clinical trial comparing two-stage outcome-adaptive strategies in which each patient receives an initial treatment and is switched to a different treatment when disease worsening is first observed. In particular, here 'outcome-adaptive' refers to treatment decisions that are made within patients, rather than between patients. While such within-patient treatment strategies may be part of a trial that also uses between-patient adaptive decision rules, such as group-sequential tests or rules to drop within-patient strategies found to perform poorly, the designs that we will discuss here do not include between-patient adaptive decision rules. We consider trials in which each patient is randomized among a set of treatments at enrollment and this treatment is continued until it fails due to disease worsening. The patient is then re-randomized among a set of treatments that excludes the treatment (s)he received initially. This ensures ignorability [8] since, given the stage 1 data, the stage 2 treatment assignment does not depend on possible future outcomes. Our underlying model accounts for several practical complications. Due to the use of scheduled examinations, unless disease worsening is discovered by frank signs or symptoms when it occurs, the worsening time is interval censored since it is only known that it occurred between the examination when it was discovered and the previous examination. The first or second worsening time may be interval censored in this way. The model also includes non-constant failure time hazards and regression of the second failure time on the patient's first worsening time.

DOI: 10.1002/sim

We establish a probability model in Section 2. Criteria for selecting a best two-stage strategy are presented in Section 3. Section 4 describes numerical methods for carrying out the necessary computations. The motivating application is presented in Section 5, including details of prior elicitation. A simulation study evaluating the selection methods is presented in Section 6. We close with a discussion in Section 7.

2. PROBABILITY MODELS

2.1. Two-stage strategies

For clinical trials evaluating two-stage strategies in which the second treatment must differ from the first, the sample size allocated to each strategy may be severely limited. For example, our application considers six two-stage strategies, so the overall sample size of 240 yields only 40 patients per strategy. Moreover, because on average therapy is discontinued without a second stage for about 20% of patients, the expected number of patients for whom the effects of each two-stage treatment strategy may be evaluated is only 32. Lavori and Dawson [2] note the combinatorial explosion in the number of possible strategies in the context of a trial examining multi-stage strategies for treating clinical depression, and they use this to motivate consideration of a limited but meaningful set of possibilities. We take this general approach here by considering only two stages and assuming a parametric model to borrow strength across strategies. Denoting by (A, B) the strategy in which treatment A is given initially and B is given if A fails, we will exploit the fact that patients given (A, B) and those given (A, C) both provide stage 1 data on treatment A.

2.2. Outcomes and likelihoods

Let T_1 denote the time to first treatment failure, defined as either first disease worsening or discontinuation of therapy. Discontinuation may be due to several reasons, including an adverse event so severe that therapy cannot be continued, the patient choosing not to receive further therapy ('dropout'), the physician deciding the patient's disease has worsened to an extent that further therapy is futile, or death. We combine these possibilities into one event to obtain a reasonably tractable probability model. An additional rationale is that, aside from death, any combination of the above reasons may cause discontinuation and often it is unclear which were the actual causes.

For patients who experience a first worsening, let T_2 denote the time from first worsening to second treatment failure, defined as either second disease worsening or discontinuation. Thus, $T_1 + T_2$ is the second failure time. Let T_D denote the time of discontinuation and S_j the time from the start of the stage j treatment to the jth disease worsening, for $j = 1, 2$. Development of a probability model is complicated by several factors, including the sequential disease examination process, the fact that failures may be of two types (discontinuation or disease worsening), the possibility of interval censoring, the possibility of a delay between first worsening and start of the second stage of therapy, and non-constant failure time hazards. To account for all of these complications with a reasonably tractable model we define T_1 and T_2 in terms of S_1, S_2 and T_D and model T_1 and $[T_2 \mid T_1]$. Thus, the first failure time is $T_1 = \min(S_1, T_D)$. Denoting by $Y_{1,W} = I(T_1 < T_D) = I(T_1 = S_1)$, the indicator that the first failure is disease worsening rather than discontinuation, T_2 is defined only if $Y_{1,W} = 1$. In many settings there is a delay between first worsening and the start of the second stage of therapy. To accommodate this, denoting the length of the delay by δ, we define the second failure time as $T_2 = \min(\delta + S_2, T_D - S_1)$. We

DOI: 10.1002/sim

Table I. The four possible cases for first treatment failure time T_1, second treatment failure time T_2, and discontinuation time T_D, when there is a fixed delay δ between first disease worsening time S_1 and the start of stage 2 therapy.

Case	Discontinuation	T_1	$Y_{1,W}$	T_2	$Y_{2,P}$	f_2
1	$T_D < S_1$	T_D	0	—	—	—
2	$S_1 < T_D < S_1 + \delta$	S_1	1	$T_D - S_1$	0	$f_{2,1}(T_2)$
3	$S_1 + \delta < T_D < S_1 + \delta + S_2$	S_1	1	$T_D - S_1$	0	$\mathscr{F}_{2,1}(\delta) f_{2,2}(T_2 - \delta)$
4	$S_1 + \delta + S_2 < T_D$	S_1	1	$\delta + S_2$	1	$\mathscr{F}_{2,1}(\delta) f_{2,2}(T_2 - \delta)$

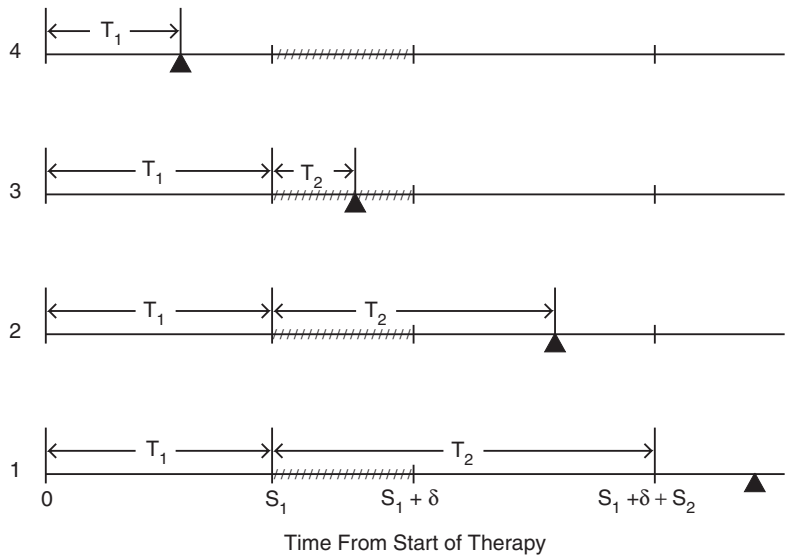

Figure 1. The four cases that determine the times T_1 and T_2 to first and second treatment failure. In each case, the time of discontinuation is indicated by a triangle, and the period of delay between first disease worsening at S_1 and the start of the second stage of therapy at $S_1 + \delta$ is denoted by hashmarks.

denote the indicator that the second failure is a disease worsening rather than a discontinuation by $Y_{2,W} = I(T_1 + T_2 < T_D)$. The cases for determining T_1 and T_2 are summarized in Table I and Figure 1, including the possibilities of discontinuation before first worsening, during the delay interval $(S_1, S_1 + \delta)$ between first worsening and the start of the second stage of therapy, during the interval $(S_1 + \delta, S_1 + \delta + S_2)$ after the second stage treatment has begun but before second worsening, or after second disease worsening.

Denote the patient's evaluation times by $0 = \tau_0 < \tau_1 < \cdots < \tau_k$. If the patient experiences a first worsening ($Y_{1,W} = 1$) and T_1 is observed rather than interval censored then we include T_1 in the vector of τ_j's. Thus, the τ_j's include the patient's examination times, which typically are scheduled but often deviate from the schedule due to patients being early, late or missing a scheduled visit, as well as the time of a first disease worsening that was discovered due to frank signs or symptoms rather than at a scheduled examination. Denoting the last follow-up time by T^0, we also write

DOI: 10.1002/sim

$\tau_{k^0} = T^0$ as a notational convenience, so that $\tau_k \leqslant \tau_{k^0}$. This allows the possibilities that the last follow-up time was the last examination time, $\tau_{k^0} = \tau_k$, or that the patient was known to have lasted to a time $T^0 = \tau_{k^0} > \tau_k$ after the last examination without a treatment failure, or that the patient died at a time $T^0 > \tau_k$. Let $\mathscr{A}_j = (\tau_{j-1}, \tau_j]$ denote the jth interval for $j = 1, \ldots, k^0$, with $\mathscr{A}_{k^0+1} = [\tau_{k^0}, \infty)$. Thus, $\tau_k < \tau_{k^0}$ only if administrative right censoring or failure occurs after the final examination, and by definition interval censoring of either T_1 or $T_1 + T_2$ may occur in any $\mathscr{A}_j$ for $j = 1, \ldots, k$ or $j = k^0 + 1$, but not between τ_k and τ_{k^0}.

We temporarily suppress dependence of event time distributions on their model parameters. Let f_1 be the probability density function (pdf) and $\mathscr{F}_1$ the survivor function (sf) of T_1, and denote the conditional pdf and sf of $[T_2 \mid T_1]$ by $f_2(\cdot \mid T_1)$ and $\mathscr{F}_2(\cdot \mid T_1)$. The first treatment failure is observed at T_1 if either worsening is first discovered by frank signs or symptoms or therapy is discontinued before disease worsening. Otherwise, worsening is discovered at one of the τ_j's and T_1 is interval censored. Let $Y_{1,0}$ be the indicator that the exact value of T_1 is observed and let $Y_{1,j}$ indicate that it is only known that $T_1 \in \mathscr{A}_j$ for $j \leqslant k$, so that $\mathbf{Y}_1 = (Y_{1,0}, Y_{1,1}, \ldots, Y_{1,k})$ has at most one entry equal to 1 and all other entries 0. Denoting interval censoring by $Y_{1,+} = \sum_{j=1}^{k} Y_{1,j}$, the indicator of right censoring at T^0 is $1 - Y_{1,0} - Y_{1,+}$. The likelihood for the first failure time data is

$$\mathscr{L}_1(T_1, \mathbf{Y}_1) = f_1(T_1)^{Y_{1,0}} \mathscr{F}_1(T^0)^{1-Y_{1,0}-Y_{1,+}} \prod_{j=1}^{k} \pi_{1,j}^{Y_{1,j}} \tag{1}$$

where $\pi_{1,j} = \Pr(T_1 \in \mathscr{A}_j) = \mathscr{F}_1(\tau_{j-1}) - \mathscr{F}_1(\tau_j)$.

To define $f_2(t \mid T_1)$, we first deal with the complication that there may be a delay of duration δ between first disease worsening at $T_1 = S_1$ and the start of stage 2 treatment at $S_1 + \delta$. Let $f_{2,1}(t \mid T_1)$ be the pdf of the time $T_{2,1}$ to second failure during the delay interval $[S_1, S_1 + \delta]$ and, if a second failure does not occur during $[S_1, S_1 + \delta]$, let $f_{2,2}(t \mid T_1)$ denote the pdf of the time $T_{2,2}$ from the start of the stage 2 treatment to second failure. We then define $T_2 = T_{2,1} I(T_{2,1} \leqslant \delta) + (\delta + T_{2,2}) I(T_{2,1} > \delta)$. Thus, for a patient randomized to strategy (A, B), $f_{2,1}$ depends on A but not B while $f_{2,2}$ depends on (A, B). The pdf of $[T_2 \mid T_1]$ thus takes the piecewise form

$$f_2(t \mid T_1) = f_{2,1}(t \mid T_1)^{I(T_{2,1} \leqslant \delta)} \{\mathscr{F}_{2,1}(\delta) f_{2,2}(t - \delta \mid T_1)\}^{I(T_{2,1} > \delta)} \tag{2}$$

Since the second failure may be due to either a discontinuation or a second disease worsening, $T_2 = T_D - S_1$ if $S_1 \leqslant T_D \leqslant S_1 + \delta$ and $T_2 = \min\{T_D - S_1, \delta + S_2\}$ if $T_D > S_1 + \delta$. In the simpler case where the second stage treatment is begun without delay at S_1, so that $\delta = 0$, the pdf of $[T_2 \mid T_1]$ reduces to $f_2(t \mid T_1) = f_{2,2}(t \mid T_1)$. The forms of f_2 under the various discontinuation cases are given in the last column of Table I. Although one may include the additional elaboration that δ is considered to be random, we did not do this in our application because any variability in δ was very small and we found that including it had no substantive effect on the trial's results.

The conditional pdf of T_2 given that T_1 is interval censored in $\mathscr{A}_j$ is

$$f_2(t_2 \mid T_1 \in \mathscr{A}_j) = \pi_{1,j}^{-1} \int_{\tau_{j-1}}^{\tau_j} f_2(t_2 \mid t_1)\, f_1(t_1)\, \mathrm{d}t_1 \tag{3}$$

For $T_1 \leqslant \tau_{r-1} < T_1 + T_2 \leqslant \tau_r$, we denote the conditional probability

$$\pi_2(\mathscr{A}_r \mid T_1) = \Pr(T_1 + T_2 \in \mathscr{A}_r \mid T_1) = \mathscr{F}_2(\tau_{r-1} - T_1 \mid T_1) - \mathscr{F}_2(\tau_r - T_1 \mid T_1)$$

DOI: 10.1002/sim

When the conditioning event is that $T_1 \in \mathscr{A}_j$ we denote

$$\pi_2^{(I)}(\mathscr{A}_r \mid \mathscr{A}_j) = \Pr(T_1 + T_2 \in \mathscr{A}_r \mid T_1 \in \mathscr{A}_j) = \pi_{1,j}^{-1} \int_{\tau_{j-1}}^{\tau_j} \pi_2(\mathscr{A}_r \mid t_1) f_1(t_1)\, \mathrm{d}t_1$$

for $\tau_j \leqslant \tau_{r-1}$. Let $Y_{2,0}$ indicate that the exact value of $T_1 + T_2$ is observed, and let $Y_{2,r}$ indicate that $T_1 + T_2$ is only known to have occurred in $\mathscr{A}_r$, with $Y_{2,+} = \sum_{r=1}^{k} Y_{2,r}$ and $\mathbf{Y}_2 = (Y_{2,0}, Y_{2,1}, \ldots, Y_{2,k})$. Since $Y_{2,1} = 0$ in any case, the conditional likelihood of T_2 may be expressed as follows:

$$\begin{aligned}\mathscr{L}_2(T_2, \mathbf{Y}_2 \mid T_1, \mathbf{Y}_1) = &\left[f_2(T_2 \mid T_1)^{Y_{2,0}} \left\{ \prod_{r=2}^{k} \pi_2(\mathscr{A}_r \mid T_1)^{Y_{2,r}} \right\} \mathscr{F}_2(T^0 - T_1 \mid T_1)^{1-Y_{2,0}-Y_{2,+}} \right]^{Y_{1,0}} \\ &\times \prod_{j=1}^{k} \left[f_2(T_2 \mid \mathscr{A}_j)^{Y_{2,0}} \left\{ \prod_{r=j+1}^{k} \pi_2^{(I)}(\mathscr{A}_r \mid \mathscr{A}_j)^{Y_{2,r}} \right\} \right. \\ &\left. \times \pi_2^{(I)}(\mathscr{A}_{k^0+1} \mid \mathscr{A}_j)^{1-Y_{2,0}-Y_{2,+}} \right]^{Y_{1,j}} \end{aligned} \tag{4}$$

Denoting $v_k = \Pr(Y_{k,W} = 1)$ for $k = 1$ and 2, the overall likelihood is

$$\begin{aligned}&\mathscr{L}(T_1, \mathbf{Y}_1, Y_{1,W}, T_2, \mathbf{Y}_2, Y_{2,W}) \\ &\quad = \mathscr{L}_1(T_1, \mathbf{Y}_1)\{v_1 \mathscr{L}_2(T_2, \mathbf{Y}_2 \mid T_1, \mathbf{Y}_1) v_2^{Y_{2,W}} (1 - v_2)^{1-Y_{2,W}}\}^{Y_{1,W}} (1 - v_1)^{1-Y_{1,W}} \end{aligned} \tag{5}$$

2.3. A parametric model

Based on historical experience, it is well known that the rate of disease progression or death for patients with metastatic renal cancer increases over time [9, 10]. Consequently, we assume that T_1 and $[T_2 \mid T_1]$ follow Weibull distributions. For real-valued α and ξ, we denote by Weib(α, ξ) the Weibull with $\log[-\log\{\mathscr{F}(t \mid \alpha, \xi)\}] = \mathrm{e}^{\xi}\{-\alpha + \log(t)\}$. This distribution has scale and shape parameters e^{α} and e^{ξ}, median $\zeta(\alpha, \xi) = \mathrm{e}^{\alpha}\{\log(2)\}^{\exp(-\xi)}$, mean $\mu(\alpha, \xi) = \mathrm{e}^{\alpha}\,\Gamma(1 + \mathrm{e}^{-\xi})$ and hazard function $h(t \mid \alpha, \xi) = \exp(\xi - \alpha \mathrm{e}^{\xi}) t^{\exp(\xi)-1}$. In other applications, it may be appropriate to use other event time distributions chosen based on experience with the particular disease and type of treatment regime being studied.

For a patient given initial treatment A, we assume first failure time distribution

$$[T_1 \mid A] \sim \mathrm{Weib}(\alpha_A, \xi_A) \tag{6}$$

Two patients who receive different strategies with the same frontline treatment, say (A, B) and (A, C), both contribute data for estimation of the stage 1 parameters (α_A, ξ_A), as in a conventional trial to compare only stage 1 treatments in terms of T_1. Since a randomized trial of multi-stage strategies has such a conventional trial of the stage 1 treatments embedded within it, the more complex design adds information to what would be obtained conventionally. This point will be discussed further in Section 7.

DOI: 10.1002/sim

Let $\text{Exp}(\alpha)$ denote an exponential distribution with mean e^{α}. For a patient given strategy (A, B) who had first progression at $T_1 = S_1$, we define the two variables

$$[T_{2,1} \mid (A, B), T_1] \sim \text{Exp}(\gamma_A + \beta_A \log(T_1)) \tag{7}$$

and

$$[T_{2,2} \mid (A, B), T_1] \sim \text{Weib}(\alpha_{A,B} + \beta_A \log(T_1), \xi_{A,B}) \tag{8}$$

In the notation of Section 2.1, $f_{2,1}$ is the exponential pdf (7) and $f_{2,2}$ is the Weibull pdf (8). This piecewise model accounts for the stage 1 treatment A effects if $T_2 \leqslant \delta$, or the two-stage strategy (A, B) effects if $T_2 > \delta$, as well as regression of T_2 on the first disease worsening time. Under this model, after adjusting for the effect of T_1, the hazard of a second failure is constant during the delay period $(S_1, S_1 + \delta)$, but once the stage 2 therapy with B has begun at $S_1 + \delta$ the hazard may be non-monotone. Many alternative models are possible. We have assumed the particular forms (6)–(8) to obtain a reasonably tractable model having the properties described above.

Recall that v_k is the probability that the kth failure is disease worsening, for $k = 1, 2$. The parameter vector for strategy (A, B) is

$$\boldsymbol{\theta}_{A,B} = (v_{1,A}, v_{2,A,B}, \alpha_A, \xi_A, \gamma_A, \beta_A, \alpha_{A,B}, \xi_{A,B})$$

Accounting for all six strategies, the overall model parameter vector $\boldsymbol{\theta}$ contains two instances of each parameter $v_{1,a}$, α_a, ξ_a, γ_a, and β_a in $\boldsymbol{\theta}_{A,B}$ indexed by a single treatment $a = A$ or B and six instances each of $v_{2,A,B}$, $\alpha_{A,B}$ and $\xi_{A,B}$. Thus, $\boldsymbol{\theta}$ has 28 entries, made up of 10 single-treatment parameters and 18 interaction parameters.

Denote the mean overall failure time under strategy (A, B) by $\mu_{A,B}(\boldsymbol{\theta}) = E\{T \mid (A, B), \boldsymbol{\theta}\}$. Since $T = T_1 + Y_{1,W} T_2$ and $E\{Y_{1,W} \mid (A, B)\} = v_{1,A}$,

$$\begin{aligned}
\mu_{A,B}(\boldsymbol{\theta}) &= E(T_1 \mid A, \boldsymbol{\theta}) + v_{1,A} E\{T_2 \mid (A, B), \boldsymbol{\theta}\} \\
&= E(T_1 \mid A, \boldsymbol{\theta}) + v_{1,A} E[E\{T_2 \mid T_1, (A, B), \boldsymbol{\theta}\}] \\
&= \mu(\alpha_A, \xi_A) + v_{1,A} \int E\{T_2 \mid T_1 = t, (A, B), \boldsymbol{\theta}\} f_1(t \mid \alpha_A, \xi_A)\, \mathrm{d}t
\end{aligned} \tag{9}$$

where $f_1(t \mid \alpha_A, \xi_A)$ is the $\text{Weib}(\alpha_A, \xi_A)$ pdf of T_1 for a patient receiving A in stage 1. The form of $\mu_{A,B}(\boldsymbol{\theta})$ given in (9) shows that since effects of the interaction parameters appear in the conditional mean of $[T_2 \mid T_1]$, this must be averaged over $f_1(T_1 \mid \boldsymbol{\theta}_1)$. Substituting the exponential distribution (7) for $f_{2,1}$ and the Weibull (8) for $f_{2,2}$ into the general piecewise form for f_2 given by (2), and denoting $\lambda_A(t) = \exp\{\gamma_A + \beta_A \log(t)\}$, the conditional expectation inside the integral in (9) is

$$E\{T_2 \mid T_1 = t, (A, B), \boldsymbol{\theta}\} = (1 - e^{-\delta/\lambda_A(t)})\lambda_A(t) + e^{-\delta/\lambda_A(t)} \{\delta + t^{\beta_A} e^{\alpha_{A,B}} \Gamma(1 + e^{-\xi_{A,B}})\} \tag{10}$$

Thus, the expected time to second failure is an average of the expectations in the cases where the patient does or does not experience a second failure during the delay period before the second stage of therapy is begun. In the simpler case where there is no delay between first progression and start of the stage 2 treatment ($\delta = 0$), γ_A disappears from $\boldsymbol{\theta}_{A,B}$, the piecewise form (2) simplifies to $f_2(t \mid T_1) \equiv f_{2,2}(t \mid T_1)$, and (10) reduces to the Weibull mean $t^{\beta_A} e^{\alpha_{A,B}} \Gamma(1 + \xi_{A,B}^{-1})$.

DOI: 10.1002/sim

3. CHOOSING A BEST STRATEGY

We will evaluate four criteria, two Bayesian and two frequentist, for selecting a best strategy based on the mean overall failure times $\{\mu_{A,B}(\boldsymbol{\theta}) : (A, B) \in \mathscr{S}\}$. The two Bayesian methods select the strategy having, respectively, the largest posterior mean or the largest posterior median of $\mu_{A,B}(\boldsymbol{\theta})$. We denote these methods by B-Weib-Mean and B-Weib-Median. The first frequentist method selects the strategy having the largest maximum likelihood estimator (MLE) of $\mu_{A,B}(\boldsymbol{\theta})$ computed under the same likelihood (5) assumed by the Bayesian methods. We denote this method by F-Weib-MLE. The second frequentist method also uses maximum likelihood, but is much simpler than the other three methods in that it ignores the multi-course structure, interval censoring and delay between courses by assuming that the overall failure time T for strategy (A, B) follows a simple Exp($\mu_{A,B}$) distribution. For this method, which we denote by F-Exp-MLE, there are only six parameters, the $\mu_{A,B}$'s themselves. Denoting by T_i^0 the right-censored overall failure time of patient i, the F-Exp-MLE method uses the conventional exponential model MLE $\hat{\mu}_{A,B} = \sum T_i^0 / \sum I(T_i = T_i^0)$ and selects the strategy for which $\hat{\mu}_{A,B}$ is largest. The sums in $\hat{\mu}_{A,B}$ include all patients who received only A and discontinued, and patients who received B in stage two following disease progression with A in stage 1. This method borrows strength across strategies since the censoring times T_i^0 for patients who discontinue with A before stage 2 contribute to $\hat{\mu}_{A,B}$, $\hat{\mu}_{A,C}$ and $\hat{\mu}_{A,D}$. Note that, although the six parameters $\{\nu_{2,A,B} : (A, B) \in \mathscr{S}\}$ characterizing the proportion of second failures that are disease worsenings appear in the likelihood (5), they have no affect on $\mu_{A,B}$ and are of secondary interest.

4. NUMERICAL METHODS

To compute the posterior mean and median of $\mu_{A,B}(\boldsymbol{\theta})$, we used the following 3-step algorithm, which first approximates $\mathscr{L}(\text{data} \mid \boldsymbol{\theta})\text{prior}(\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$ with a multivariate normal $p_{\text{N}}(\boldsymbol{\theta} \mid \tilde{\boldsymbol{\mu}}, \tilde{\Sigma})$ having mean $\tilde{\boldsymbol{\mu}}$ and variance–covariance matrix $\tilde{\Sigma}$, and then uses defensive importance sampling to compute posterior quantities. The algorithm exploits the fact that $\tilde{\boldsymbol{\mu}}$ is the mode of p_{N}. In carrying out the following computations, the numerical integrations required to compute the integrals in $f_2(t_2 \mid T_1 \in \mathscr{A}_j)$, $\pi_2^{(I)}(\mathscr{A}_r \mid \mathscr{A}_j)$ and the overall mean $\mu_{A,B}(\boldsymbol{\theta})$ were evaluated numerically using the double exponential integration method of Takahasi and Mori [11].

Step 1: Approximate the mode of $\mathscr{L}(\text{data} \mid \boldsymbol{\theta})\text{prior}(\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$ using the simplex method of Nelder and Mead [12], and set $\tilde{\boldsymbol{\mu}}$ equal to this mode. This is carried out in the log domain, in terms of $l(\text{data}, \boldsymbol{\theta}) = \log\{\mathscr{L}(\text{data} \mid \boldsymbol{\theta})\text{prior}(\boldsymbol{\theta})\}$, to avoid numerical underflows as a result of multiplying very small values.

Step 2: Determine a quasi-random sample $\{\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(K)}\}$ of $\boldsymbol{\theta}$ values near $\tilde{\boldsymbol{\mu}}$ using the method of Halton [13], which efficiently distributes the sample over the $\boldsymbol{\theta}$ domain. Evaluate $l(\text{data}, \boldsymbol{\theta})$ and the quadratic form $Q(\boldsymbol{\theta}, \tilde{\mu}, A) = (\boldsymbol{\theta} - \tilde{\mu})^{\text{T}} A (\boldsymbol{\theta} - \tilde{\mu})$ at the quasi-sample points, treat $l(\text{data}, \boldsymbol{\theta}^{(j)})$ like the jth observed outcome variable, $Q(\boldsymbol{\theta}^{(j)}, \tilde{\mu}, A)$ like the jth observed predictor and the entries of A as parameters in a linear regression model, solve for the entries of A using conventional least squares, and set $\tilde{\Sigma} = -0.5\, A^{-1}$.

Step 3: Generate a sample of size 20 000 from the posterior using the iterative defensive importance sampling method of Owen and Zhou [14], with the normal distribution $p_{\text{N}}(\boldsymbol{\theta} \mid \tilde{\boldsymbol{\mu}}, \tilde{\Sigma})$ obtained in steps 1 and 2 used as the distribution approximating $\mathscr{L}(\text{data} \mid \boldsymbol{\theta})$

DOI: 10.1002/sim

prior($\boldsymbol{\theta}$). Use this posterior sample to compute $E\{\mu_{A,B}(\boldsymbol{\theta}) \mid \text{data}\}$ and median$\{\mu_{A,B}(\boldsymbol{\theta}) \mid \text{data}\}$ empirically.

For the frequentist methods, the MLE of $\boldsymbol{\theta}$ under each model was calculated using the same mode-finding method as Step 1 above, applied to the likelihood $\mathscr{L}(\text{data} \mid \boldsymbol{\theta})$ rather than to $\mathscr{L}(\text{data} \mid \boldsymbol{\theta})$ prior($\boldsymbol{\theta}$). All programming was done in C++. Computer programs are available from the second author on request.

5. APPLICATION

5.1. A Metastatic Renal Cancer Trial

We illustrate our method with the clinical trial that motivated this research. The goal of the trial is to compare two-stage treatment strategies for patients with metastatic renal cancer who have not been treated previously with systemic therapy. While many new agents with potential clinical anti-tumour activity currently are being produced at a rapidly increasing rate, the number of single agents that can be evaluated clinically is limited. Considering the number of possible multi-stage treatment strategies that may be of interest, the limitations are far greater. The trial described here was motivated by the desire to compare several new targeted agents in a randomized fashion, the belief that different agents given consecutively may have interactive effects, and the desire to provide a sound basis for selecting two-stage strategies for later evaluation in a large-scale phase III trial.

The strategies are based on four targeted therapies, which we denote by a, b, c, d. The four agents were chosen based on preliminary clinical and biological data. Three of the four agents target the vascular endothelial growth factor or its receptors, and the fourth inhibits the mTOR signalling pathway that controls messenger RNA and cell proliferation. When designing the trial, a and b recently had been approved by the U.S. Food and Drug Administration as frontline treatments for metastatic renal cancer but c and d had not. Thus, it was decided to include only strategies for which either a or b is given in stage 1. Since the second treatment of each strategy must be different from the first, the six strategies $\mathscr{S} = \{(a, b), (a, c), (a, d), (b, a), (b, c), (b, d)\}$ are evaluated. A total of 240 patients will be randomized fairly at enrollment between a and b. If a patient suffers a disease progression (s)he is then re-randomized among the three agents not given initially. At the end of the trial, the strategy having the largest posterior mean $\mu_{A,B}(\boldsymbol{\theta})$ will be selected.

While it is tempting to believe that the presence of a measurable target implies a high probability of response to an agent aimed at that target, unfortunately this usually is not so. For example, benefit from EGFR-directed therapy in colon cancer is completely independent of EGF or EGFR expression status. Likewise, expression data do not reliably predict benefit for EGFR therapy in lung cancer. In the latter case, a specific mutation in a subset of the population (females of oriental heritage), has been found to be associated with benefit from a particular agent but, even in this case, the effect is far below the level that would guide clinical practice, and the mutation has not been found in other contexts where EGFR inhibition is known to be biologically active, such as in colon or bladder cancer. Other examples include the lack of association between KIT expression and response to KIT inhibitors, and the lack of association between the androgen receptor and response to anti-androgens in the setting of prostate cancer. Thus, limiting a trial to patients who are 'target positive' would risk missing subsets of target negative patients in whom a drug, or

DOI: 10.1002/sim

in our trial a two-stage strategy, is effective. In metastatic renal cancer, there currently are no tests that predict clinical benefit from a particular drug better than classical histologic sub-type or clinical features. Identifying specific targets that can be measured at baseline and used as a basis for treatment assignment is a secondary goal of the metastatic renal cancer trial, however.

5.2. Establishing priors

Although $\boldsymbol{\theta}$ has 28 elements, assuming that the six subvectors $\{\boldsymbol{\theta}_{A,B} : (A, B) \in \mathscr{S}\}$ are exchangeable, it is only necessary to establish priors on the eight elements of one $\boldsymbol{\theta}_{A,B}$. To ensure suitably uninformative priors, large variances were used. Since historically 80 per cent of treatment failures in metastatic renal cancer are disease progression, we assumed all $\nu_{1,A}$'s and $\nu_{2,A,B}$'s were iid beta(0.80, 0.20). The remaining parameters $\alpha_A, \xi_A, \gamma_A, \beta_A, \alpha_{A,B}, \xi_{A,B}$ are real valued and were assumed to follow independent normal priors. To determine the prior means, we elicited expected values of medians and several other percentiles of T_1 and $T_2 \mid T_1$, and applied the least-squares method of Thall and Cook [15] to solve for the means of α_A, ξ_A, β_A, $\alpha_{A,B}$ and $\xi_{A,B}$. Setting $\text{var}\{e^{\xi_A}\} = \text{var}\{e^{\xi_{A,B}}\} = 100$ then yielded $\xi_A \sim \text{N}(-0.0516, 1.549^2)$ and $\xi_{A,B} \sim \text{N}(0.434, 1.395^2)$, and setting $\text{var}\{e^{\alpha_A}\} = 100$ gave $\alpha_A \sim \text{N}(-0.260, 1.613^2)$. Assuming the probability of discontinuation during the one month delay period is 0.02, substituting the prior mean $E(\beta_A) = 0.847$ and reference value $T_1 = 8$ months, and solving $1 - \exp(-1/e^{\gamma_A + 0.847 \log(8)}) = 0.02$ yielded $E(\gamma_A) = 2.141$. Equating $\text{var}(\beta_A) = \text{var}(\gamma_A) = \text{var}(\alpha_A)$ yielded $\beta_A \sim \text{N}(0.847, 1.613^2)$ and $\gamma_A \sim \text{N}(2.141, 1.613^2)$, and setting $\text{var}(\alpha_{A,B}) = \text{var}(\alpha_A)$ gave $\alpha_{A,B} \sim \text{N}(-2.100, 1.613^2)$. The resulting prior means of the scale and shape parameters were $E(e^{\alpha_A}) = 2.83$, $E(e^{\alpha_{A,B}}) = 0.45$, $E(e^{\xi_A}) = 3.15$ and $E(e^{\xi_{A,B}}) = 4.09$. Together, these priors yield $E(\mu_{A,B}) = 7.0$ and $\text{var}(\mu_{A,B}) = 167$ for the mean overall failure times, so the priors on the $\mu_{A,B}$'s were uninformative.

6. SIMULATION STUDY

6.1. Clinical scenarios

Each clinical scenario under which we will evaluate the design in the simulation study is characterized by a fixed value of $\boldsymbol{\theta}$. To facilitate interpretation, the scenarios are specified in terms of fixed values of $\zeta_1(A) = \text{median}\ (T_1 \mid A)$ and $\zeta_2(A, B) = \text{median}(T_{2,2} \mid T_1 = 8, A, B)$, the median time to second failure in stage 2 with treatment B following first progression with A at the reference value $T_1 = 8$ months. Given $\zeta_1(A)$ and $\zeta_2(A, B)$, to determine $\mu_{A,B}$ we must specify fixed values of the seven parameters α_A, $\xi_{A,B}$, ν_A, γ_A, β_A, $\alpha_{A,B}$, and $\xi_{A,B}$. To do this, in all the cases we fixed $\nu_{1,A} = 0.80$, $\gamma_A = 2.141$ and $\beta_A = 0.847$, their prior means. In each case, we equated the upper 95th percentiles of the distributions $\zeta_1(A)$ and $\zeta_2(A, B)$ to twice their specified fixed values, and these two additional constraints allowed us to solve for fixed values of α_A, ξ_A, $\alpha_{A,B}$, $\xi_{A,B}$.

To construct a set of scenarios reflecting a reasonable array of possibilities, based on clinical experience we first specified the null values $\zeta_1 = 8$ and $\zeta_2 = 3$ months for the median failure times. In the other scenarios, $\zeta_1 = 12$ months is considered good frontline, and $\zeta_2 = 6$ months and 9 months are considered good and very good salvage. The scenarios are summarized in Table II. In all scenarios, we assume that 80 per cent of first failures are disease progressions, hence on average 32 patients are treated with each strategy. This underscores the importance of using the stage 1 data to learn about $\nu_{1,A}$, α_A and ξ_A, for $A = a, b, c, d$, and borrowing across strategies to learn about the regression parameters β_a and β_b.

DOI: 10.1002/sim

Table II. The scenarios studied in the simulations.

Scenario	Description	True parameter values: Medians	True parameter values: Overall means
1	Null case	All $\zeta_1 = 8$, all $\zeta_2 = 3$	All $\mu_{a,b} = 12.5$
2	a is good frontline	$\zeta_1(a) = 12$	$\mu_{a,b} = \mu_{a,c} = \mu_{a,d} = 17.7$
3	d is good salvage after either a or b	$\zeta_2(a,d) = \zeta_2(b,d) = 6$	$\mu_{a,d} = \mu_{b,d} = 15.0$
4	(a,d) is a good strategy	$\zeta_2(a,d) = 6$	$\mu_{a,d} = 15.0$
5	a is good frontline, (a,d) is a good strategy	$\zeta_1(a) = 12$ $\zeta_2(a,d) = 6$	$\mu_{a,b} = \mu_{a,c} = 17.7$ $\mu_{a,d} = 21.3$
6	(a,d) is a good strategy, (b,d) is a very good strategy	$\zeta_2(a,d) = 6$ $\zeta_2(b,d) = 9$	$\mu_{a,d} = 15.0$ $\mu_{b,d} = 17.5$

Note: In all the scenarios, unless otherwise specified, the medians are the null values $\zeta_1(A) = 8$ and $\zeta_2(A, B) = 3$ for all treatments A and B, corresponding to $\mu_{A,B} = 12.5$ months, and $\zeta_2(A, B)$ denotes median $(T_{2,2} \mid T_1 = 8)$ under (A, B). In all cases, 80 per cent of first failures are renal cancer progressions (worsenings).

In scenario 2, since a is a good frontline treatment the three strategies (a,b), (a,c), (a,d) are equally desirable with mean overall failure time 17.7, while (b,a), (b,c), (b,d) all have null mean 12.5. In scenario 3, since d is a good salvage therapy after either a or b, (a,d) and (b,d) are equally desirable with mean overall failure time 15.0, a smaller advantage over the null value 12.5 than in scenario 2. In scenario 4 only one strategy (a,d) is superior, due to an interaction. In scenario 5, since a is good frontline and a and d interact, (a,d) is by far the best strategy with $\mu_{a,d} = 21.3$, while (a,b) and (a,c) are superior but provide the smaller improvements $\mu_{a,b} = \mu_{a,c} = 17.7$. Scenario 6 has two superior strategies, one good with $\mu_{a,d} = 15.0$ and the other very good with $\mu_{b,d} = 17.5$.

6.2. Simulation methods

The trial was simulated 1000 times under each scenario, and the same data were used for all four selection methods. To simulate (T_1, T_2) for strategy (A, B), we first generated $T_1 \sim \text{Weib}(\alpha_A, \xi_A)$. Under the piecewise model given by (7) and (8), we then generated $T_{2,1} \mid T_1$ from the exponential distribution (7). If $T_{2,1} \leqslant \delta$ we set $T_2 = T_{2,1}$, and if $T_{2,1} > \delta$ we generated $T_{2,2} \mid T_1$ from the Weibull distribution (8) and set $T_2 = \delta + T_{2,2}$.

6.3. Simulation results

The simulation results are summarized in Table III. Each tabled mean is the average value over the 1000 simulated trials. The four methods perform very similarly under each of Scenarios 1–3, although under scenario 3 the B-Weib-Median and F-Weib-MLE methods both always select one of the two strategies (a,d) or (b,d) that include the superior salvage treatment d. The better performance of these two methods compared to the others is more pronounced under each of Scenarios 4–6, which show that B-Weib-Median and F-Weib-MLE perform equally well,

DOI: 10.1002/sim

Table III. Simulation results.

Scenario	Method	Strategy					
		(a, b)	(a, c)	(a, d)	(b, a)	(b, c)	(b, d)
1	B-Weib-Mean	16, 11.9	17, 12.0	17, 11.9	17, 12.0	17, 12.0	17, 11.9
	B-Weib-Median	16, 11.9	16, 11.9	16, 11.9	18, 11.9	18, 11.9	16, 11.9
	F-Weib-MLE	17, 11.8	16, 11.8	16, 11.8	17, 11.9	17, 11.9	17, 11.9
	F-Exp-MLE	18, 11.9	17, 11.8	16, 11.8	16, 11.9	17, 11.9	16, 11.9
2	B-Weib-Mean	**29**, **17.3**	**37**, **17.5**	**34**, **17.4**	0, 11.9	0, 11.9	0, 11.9
	B-Weib-Median	**31**, **17.3**	**37**, **17.4**	**33**, **17.3**	0, 11.9	0, 11.9	0, 11.8
	F-Weib-MLE	**31**, **17.3**	**35**, **17.3**	**34**, **17.3**	0, 11.8	0, 11.8	0, 11.8
	F-Exp-MLE	**33**, **16.8**	**34**, **16.8**	**33**, **16.8**	0, 11.8	0, 11.9	0, 11.8
3	B-Weib-Mean	1, 12.0	1, 12.1	**51**, **14.9**	1, 12.0	0, 12.1	**45**, **14.9**
	B-Weib-Median	0, 11.9	0, 12.0	**52**, **14.9**	0, 12.0	0, 12.0	**48**, **14.8**
	F-Weib-MLE	0, 11.9	0, 11.9	**52**, **14.9**	0, 11.9	0, 11.9	**48**, **14.8**
	F-Exp-MLE	0, 11.8	1, 11.9	**55**, **14.2**	0, 11.9	1, 11.8	**44**, **14.0**
4	B-Weib-Mean	1, 12.0	1, 12.0	**94**, **14.8**	2, 11.9	2, 11.9	1, 11.9
	B-Weib-Median	0, 11.9	0, 11.9	**97**, **14.7**	1, 11.8	1, 11.8	1, 11.8
	F-Weib-MLE	0, 11.9	0, 11.9	**98**, **14.7**	1, 11.8	1, 11.8	1, 11.8
	F-Exp-MLE	2, 11.9	2, 11.8	**84**, **14.0**	5, 11.9	4, 11.8	3, 11.7
5	B-Weib-Mean	**0**, **17.4**	**1**, **17.6**	**98**, **22.1**	0, 11.9	0, 12.0	0, 11.9
	B-Weib-Median	**0**, **17.4**	**0**, **17.5**	**100**, **21.9**	0, 11.8	0, 11.9	0, 11.8
	F-Weib-MLE	**0**, **17.4**	**0**, **17.4**	**100**, **21.9**	0, 11.8	0, 11.8	0, 11.8
	F-Exp-MLE	**5**, **16.9**	**5**, **16.9**	**91**, **20.3**	0, 11.8	0, 11.8	0, 11.8
6	B-Weib-Mean	1, 12.0	1, 12.1	**5**, **14.9**	0, 12.1	0, 12.1	**93**, **18.2**
	B-Weib-Median	0, 11.9	0, 12.0	**4**, **14.8**	0, 12.0	0, 12.0	**96**, **18.1**
	F-Weib-MLE	0, 11.9	0, 11.9	**5**, **14.8**	0, 11.9	0, 11.9	**95**, **18.1**
	F-Exp-MLE	0, 11.8	0, 11.8	**13**, **14.1**	0, 11.8	0, 11.9	**87**, **16.5**

Note: Each pair of entries is the selection percentage and estimated mean overall failure time, $\mu_{A,B}$. Results are given for the Bayesian Weibull model-based methods using the posterior mean (B-Weib-Mean) and median (B-Weib-Median) of $\mu_{A,B}$, and for the frequentist methods based on maximum likelihood estimators of $\mu_{A,B}$ under the Weibull model (F-Weib-MLE) and under a simple exponential model for overall failure time (F-Exp-MLE). Entries for superior strategies are given in boldface type.

B-Weib-Mean performs well but has slightly smaller correct selection probabilities, and F-Exp-MLE clearly has the worst performance. It thus appears that, in terms of selecting a best strategy, it is worthwhile to account for the complexities of the data structure through an appropriately modelled likelihood. Moreover, it is interesting that B-Weib-Median has superior performance compared to B-Weib-Mean. The B-Weib-Median and F-Weib-MLE methods thus appear to do a very reliable job of selecting the best strategy in the cases studied. The greater simplicity of the frequentist method suggests that, once one has done the work of modelling the event time process in the likelihood, there may be little practical advantage to using a Bayesian model.

In terms of estimation, under each of the scenarios 1–4 all four methods slightly underestimate the $\mu_{A,B}$'s. In contrast, B-Weib-Mean, B-Weib-Median and F-Weib-MLE each slightly overestimates the means of the best strategies in scenarios 5 and 6, whereas F-Exp-MLE underestimates the $\mu_{A,B}$'s in all cases.

DOI: 10.1002/sim

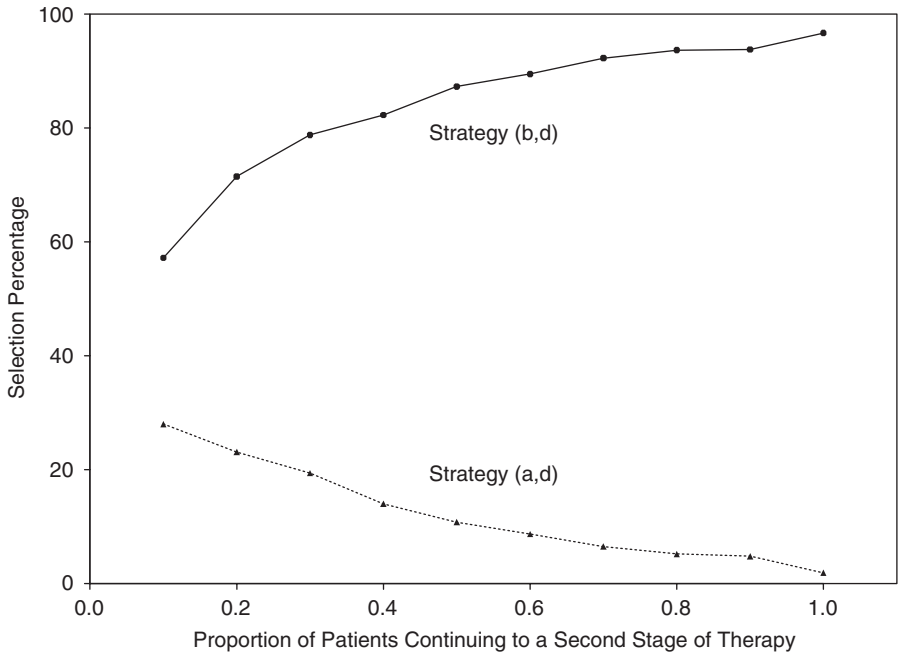

Figure 2. Sensitivity of the selection probabilities for the good strategy (a, d) with $\mu_{a,d} = 15.0$ months (dotted line with triangles) and very good strategy (b, d) with $\mu_{b,d} = 17.5$ months (solid line with circles) under Scenario 6 to the proportion of patients whose first treatment failure is a disease worsening and who thus receive a second stage of therapy.

The expected number of patients who receive each strategy is $40\,v_{1,A}$, since $v_{1,A}$ is the probability a patient first treated with A will suffer a disease progression and therefore receive a second stage of therapy. To examine the method's sensitivity to this probability, we repeated the simulations of B-Weib-Mean in scenario 6 while varying the fixed values of $v_{1,a} = v_{1,b}$ from 0.10 to 1.00, which has the effect of varying the expected sample size for evaluating each interaction from 4 to 40. The results are illustrated in Figure 2, which shows that the method's reliability is very sensitive to v_1, with the selection percentage for the best strategy (b, d) decreasing from 98 when $v_1 = 1.0$ to 57 when $v_1 = 0.10$.

7. DISCUSSION

While study of frontline treatments in randomized trials is routine, salvage treatments usually are evaluated in single-arm trials. This makes comparisons between salvage treatments problematic due to selection bias and effects of latent variables, since treatment effects are confounded with trial effects [16]. Comparisons between frontline treatments in terms of overall survival time also suffer from the biasing effects of non-randomly selected salvage therapies, since a patient's overall survival time depends on both the frontline and salvage treatments. Randomizing among strategies solves these problems by providing unbiased comparisons among strategies, among frontline treatments, and among salvage therapies given to patients who received the same initial

DOI: 10.1002/sim

treatment. Additionally, a trial with randomization among salvage treatments has the substantial logistical advantage that these treatments are all evaluated in one trial rather than in separate single-arm trials.

Although the primary aim of our design is to select the best two-stage strategy, a conventional randomized trial of *a versus b* evaluated in terms of T_1 alone is embedded in the design. If the primary goal were to perform a frequentist test of equal median first failure times, $\zeta_1(a) = \zeta_1(b)$, for the two frontline treatments then, assuming an accrual rate of nine patients per month and maximum trial duration of 38.5 months, the sample size of 240 would provide the basis for a two-sided level 0.05 test to detect a 50 per cent difference in ζ_1, from 8 to 12 months, with power 0.80. In this sense, one could think of the trial as a conventional phase III comparison of two frontline treatments that also includes comparisons of six two-stage strategies. Another interpretation is that the trial replaces six conventional single-arm phase II trials, namely two phase IIB trials of a and b as frontline therapies and four phase IIA trials of a, b, c and d as salvage regimens. However, randomization eliminates the bias that results from comparing treatments or two-stage strategies based on data from such single-arm trials. Finally, the use of time-to-event outcomes provides a much more reliable assessment of treatment effects than the binary response indicators conventionally used in phase II trials [17].

While it may appear unethical to continue randomizing patients to an initial treatment, a, if interim data indicate that it is inferior to initial treatment b in terms of T_1, this relies on the assumption that T_1 reliably predicts overall failure time, $T_1 + T_2$. In many clinical settings this simply is not the case. A regimen with a higher initial response rate or longer average time to first failure may not necessarily be the best choice when considered in the larger context of the overall therapeutic strategy. It may be the case that the treatment having the worst initial outcome, in our setting the shortest average T_1, is the first component of the best two-stage treatment strategy in terms of overall failure time. Although this may appear counter-intuitive, it is quite plausible due to the potential for the lack of cross-resistance between the first and second treatments, and the preservation of physiologic reserve with stage 1 therapy. For example, aggressive alkylator-based chemotherapy produces a high initial response rate in multiple myeloma but, in terms of overall long-term outcome, it is clear that a much better strategy in this disease is to use less aggressive therapies front-line, keeping chemotherapy in reserve as second line treatment. Thus, in general it is false that if b provides a longer average time to first failure then it necessarily follows that a strategy beginning with b rather than a will provide the greatest long-term clinical benefit. The main advantage of our proposed methodology, that it reliably identifies two-stage strategies that are superior in terms of overall failure time, would be lost if stage 1 treatments having comparatively worse first failure times were terminated early.

An important issue is accounting for possible effects of covariates $\mathbf{Z} = (Z_1, \ldots, Z_q)$, such as established prognostic variables or biomarkers thought to be related to outcome. This is a complex issue since effects of $\mathbf{Z}$ must be modelled within the multi-stage structure. An example in the context of two-stage strategies with discrete outcomes is given by Thall *et al.* [7]. This is particularly important when a measurable covariate indicating the presence of a putative target is available, since it generally is not the case that an agent aimed at a given target will improve clinical outcome in patients who are target positive. For example, benefit from EGFR-directed therapy in colon cancer or lung cancer is completely independent of EGF or EGFR expression status, KIT expression is not associated with response to KIT inhibitors, and presence of androgen receptors is not associated with response to anti-androgen agents for treating prostate cancer. Thus, including such biomarkers as covariates to assess possible effects, including biomarker-treatment

DOI: 10.1002/sim

interactions, is potentially useful. However, limiting a given treatment to patients who are positive for biomarkers that are thought, based on preclinical data, to be related to outcome using that treatment [18] risks doing away with the possibility of discovering that the treatment is beneficial in target negative patients, and of course it does away with the benefits of randomization.

To incorporate $\mathbf{Z}$ into the Weibull model given by (6), (7), and (8), for $A=a$ or b and $B=a, b, c$ or d, an extended model is given by $[T_1 \mid A, \mathbf{Z}] \sim \text{Weib}(\alpha_A+\sum_{j=1}^{q} Z_j\alpha_A^{(j)}, \xi_A+\sum_{j=1}^{q} Z_j\xi_A^{(j)})$, and $[T_{2,2} \mid (A, B), T_1, \mathbf{Z}] \sim \text{Weib}(\alpha_{A,B} + \sum_{j=1}^{q} Z_j\alpha_{A,B}^{(j)} + \beta_A \log(T_1), \xi_{A,B} + \sum_{j=1}^{q} Z_j\xi_{A,B}^{(j)})$. The $4q$ parameters $\{(\alpha_A^{(j)}, \xi_A^{(j)}) : A=a, b, j=1, \ldots, q\}$ characterize covariate effects on the distribution of T_1 in stage 1, and the $12q$ parameters $\{(\alpha_{A,B}^{(j)}, \xi_{A,B}^{(j)}) : (A, B) \in \mathscr{S}, j=1, \ldots, q\}$ characterize the covariate effects on the distribution of T_2 in stage 2. Thus, in addition to the 28 parameters of the original model, there are $16q$ additional covariate parameters, and such a model may easily become intractable. In practice a simpler formulation would be needed, such as assuming homogeneous covariate effects $\alpha_A^{(j)} = \alpha^{(j)}$ and $\xi_A^{(j)} = \xi^{(j)}$ for $A=a, b$, and $\alpha_{A,B}^{(j)} = \alpha^{(j)}$, $\xi_{A,B}^{(j)} = \xi^{(j)}$ for all $(A, B) \in \mathscr{S}$. An alternative approach would be to evaluate the predictive value of each Z_j separately, by comparing each 40-parameter model obtained by adding only the 12 parameters specific to Z_j to the 'null' 28-dimensional model with no covariates.

Our design has some similarities with the biased coin adaptive within subject (BCAWS) design proposed by Lavori and Dawson [3]. The BCAWS design starts all patients with the same initial treatment, and then randomly switches the patient to a salvage treatment, with the probability of switching based on whether a cumulative symptom score quantifying the success of the initial treatment over time exceeds a critical threshold. Lavori and Dawson apply the multiple imputation method of Rubin and Shenker [19] to obtain an approximately unbiased estimator of the mean final cumulative score of the optimal threshold. Our design is very different in that we randomize patients among treatments at both stages, randomize fairly in stage 2 rather than using the stage 1 data as a basis for choosing the randomization probabilities adaptively, and focus on selection among two-treatment strategies. An important similarity, however, is that both designs deal with strategies that switch the patient away from a treatment that has failed.

Our design is not outcome-adaptive between patients. It is outcome-adaptive within patients in that the second stage of therapy is begun after the stage 1 treatment has failed, and the stage 2 treatment must be different from that received in stage 1. However, as noted above the stage 2 randomization probabilities do not depend on T_1, but rather the patient is randomized with equal probability among the three treatments not given initially. Thus, patients could be randomized fairly among the strategies at baseline to avoid logistical problems during the trial. This avoids many of the difficulties, such as estimation bias, associated with designs that are outcome-adaptive between treatments, such as more elaborate versions of our design that unbalance the randomization probabilities interimly in favour of superior strategies, or that terminate inferior strategies early in a group-sequential fashion. Such designs introduce bias into estimators since the more successful treatments or strategies are over-represented, and are well known to be subject to potential biasing effects of systematic drift in patient prognosis over the course of the trial [20].

An important point regarding our definition of the composite discontinuation event is that the subevent in which the patient decides to discontinue therapy may be considered to be related to subsequent disease worsening or death. In this case, methods for utilizing such informative dropout information may be appropriate, such as that of Lunceford *et al.* [5]. Similarly, a deviation of the τ_j's from a planned schedule also may be considered to contain information about the progression

DOI: 10.1002/sim

time distributions. Although we will not deal with these issues here, they are important areas for future research in the context of multi-stage treatment regimes.

ACKNOWLEDGEMENTS

Peter Thall's work was partially supported by NCI grant RO1-CA-83932. The authors thank two referees and an Associate Editor for their constructive comments.

REFERENCES

1. Murphy SA. Optimal dynamic treatment regimes (with Discussion). *Journal of the Royal Statistical Society*, *Series B* 2003; **65**:331–366.
2. Lavori PW, Dawson R. Dynamic treatment regimes: practical design considerations. *Clinical Trials* 2004; **1**:9–20.
3. Lavori PW, Dawson R. A design for testing clinical strategies: biased adaptive within-subject randomization. *Journal of Royal Statistical Society*, *Series A* 2000; **163**:29–38.
4. Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* 2005; **24**:1455–1481.
5. Lunceford JK, Davidian M, Tsiatis AA. Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics* 2002; **58**:48–57.
6. Thall PF, Millikan R, Sung H-G. Evaluating multiple treatment courses in clinical trials. *Statistics in Medicine* 2000; **19**:1011–1028.
7. Thall PF, Sung H-G, Estey EH. Selecting therapeutic strategies based on efficacy and death in multi-course clinical trials. *Journal of the American Statistical Association* 2002; **97**:29–39.
8. Rubin D. Practical implications of modes of statistical inference for causal effects and the critical role of the assignment mechanism. *Biometrics* 1991; **47**:1213–1234.
9. Elaraj DM, White DE, Steinberg SM, Haworth L, Rosenberg SA, Yang YC. A pilot study of antiangiogenic therapy with bevacizumab and thalidomide in patients with metastatic renal cell carcinoma. *Journal of Immunotherapy* 2004; **27**:259–264.
10. Motzer RJ, Michaelson MD, Redman BG, Hudes GR, Wilding G, Figlin RA, Ginsberg MS, Kim ST, Baum CM, DePrimo SE, Li JZ, Bello CL, Theuer CP, George DJ, Rini BI. Activity of SU 11248, a multitargeted inhibitor of vascular endothelial growth factor receptor and platelet-derived growth factor receptor, in patients with metastatic renal cell carcinoma. *Journal of Clinical Oncology* 2006; **24**:16–24.
11. Takahasi H, Mori M. Double exponential formulas for numerical integration. *Publication RIMS Kyoto University* 1974; **9**:721–741.
12. Nelder JA, Mead R. A simplex method for function minimization. *Computer Journal* 1965; **7**:308.
13. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical Recipes in* C++. Cambridge University Press: Cambridge, 2002.
14. Owen A, Zhou Y. Safe and effective importance sampling. *Journal of the American Statistical Association* 1999; **95**:135–140.
15. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 2004; **60**:684–693.
16. Estey EH, Thall PF. New designs for phase 2 clinical trials. *Blood* 2003; **102**:442–448.
17. Thall PF, Wooten LH, Tannir N. Monitoring event times in early phase clinical trials: some practical issues. *Clinical Trials* 2005; **2**:467–478.
18. Maitournam A, Simon R. On the efficiency of targeted clinical trials. *Statistics in Medicine* 2005; **19**:329–339.
19. Rubin D, Shenker M. Multiple imputation in health-care data bases: an overview and some applications. *Statistics in Medicine* 1991; **10**:5585–5598.
20. Karrison TG, Huo D, Chappell R. A group sequential, response-adaptive design for randomized clinical trials. *Controlled Clinical Trials* 2003; **24**:506–522.

DOI: 10.1002/sim