

*University of Texas, MD Anderson Cancer  
Center*

UT MD Anderson Cancer Center Department of Biostatistics  
Working Paper Series

---

*Year 2005*

*Paper 15*

---

Continuous Bayesian adaptive randomization  
based on event times with covariates

Ying Kuen Cheung\*

Lurdes Y. T. Inoue†

J. Kyle Wathen‡

Peter F. Thall\*\*

\*Columbia University, yc632@columbia.edu

†University of Washington, linoue@u.washington.edu

‡U.T.M.D. Anderson Cancer Center, jkwathen@mdanderson.org

\*\*U.T.M.D. Anderson Cancer Center, rex@mdanderson.org

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/mdandersonbiostat/paper15>

Copyright ©2005 by the authors.

## Continuous Bayesian adaptive randomization based on event times with covariates

Ying Kuen Cheung<sup>1,\*</sup>, Lurdes Y. T. Inoue<sup>2</sup>, J. Kyle Wathen<sup>3</sup> and Peter F. Thall<sup>3</sup>

<sup>1</sup>*Department of Biostatistics, Mailman School of Public Health, Columbia University, 722 West 168th Street, New York, NY 10032, U.S.A.*

<sup>2</sup>*Department of Biostatistics, University of Washington, Seattle, Washington 98195, U.S.A.*

<sup>3</sup>*Department of Biostatistics and Applied Mathematics, The University of Texas, M.D. Anderson Cancer Center, 1515 Holcombe Boulevard, Houston, TX 77030, U.S.A.*

### SUMMARY

In comparative clinical trials, the randomization probabilities may be unbalanced adaptively by utilizing the interim data available at each patient's entry time to favour the treatment or treatments having comparatively superior outcomes. This is ethically appealing because, on average, more patients are assigned to the more successful treatments. Consequently, physicians are more likely to enrol patients onto trials where the randomization is outcome-adaptive rather than balanced in the conventional manner. Outcome-adaptive methods based on a binary variable may be applied by reducing an event time to the indicator of the event's occurrence within a predetermined time interval. This results in a loss of information, however, since it ignores the censoring times of patients who have not experienced the event but whose evaluation interval is not complete. This paper proposes and compares exact and approximate Bayesian outcome-adaptive randomization procedures based on time-to-event outcomes. The procedures account for baseline prognostic covariates, and they may be applied continuously over the course of the trial. We illustrate these methods by application to a phase II selection trial in acute leukaemia. A simulation study in the context of this trial is presented. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: adaptive design; Bayesian statistics; censored data; clinical trials; ethics; historical data; survival

### 1. INTRODUCTION

Randomization is the established statistical method for obtaining unbiased estimates of comparative treatment effects, and thus is a key component of controlled clinical trials. In practice,

\*Correspondence to: Ying Kuen Cheung, Department of Biostatistics, Columbia University, 722 West 168th Street, Room 641, New York, NY 10032, U.S.A.

†E-mail: yc632@columbia.edu

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

most randomization methods are balanced, allocating on average an equal number of patients to each treatment arm. However, many physicians find balanced randomization objectionable because it may force them to use treatments that they believe are unlikely to provide their patients with the best chance of clinical benefit [1]. A practical alternative is outcome-adaptive randomization [2–5] which uses the data from patients treated previously in the trial to unbalance the randomization probabilities in favour of the treatment or treatments observed to have comparatively superior outcomes. Outcome-adaptive randomization provides a compromise between ethical concerns and the scientific goal of obtaining unbiased treatment comparisons. A different type of adaptive randomization, based on treatment assignments but not patient outcomes, aims to achieve balance within patient subgroups [6–8]. This is not our focus, and hereafter we use ‘AR’ to refer to outcome-adaptive randomization aimed at favouring empirically superior treatments.

Although there is a substantial literature on a wide variety of AR methods, until very recently such methods have seen almost no use in practice. Starting in 2001, we began to design and conduct clinical trials at M.D. Anderson Cancer Center (MDACC) and other medical institutions using a variety of Bayesian AR procedures [9]. We have found that physicians who are hesitant to enter patients onto conventionally randomized trials find AR more appealing because, on average, AR assigns more patients to the more successful treatments.

This paper is motivated by a practical problem that has arisen repeatedly when applying these methods in trials where response is defined as a particular event occurring within a fixed period of time from the start of treatment. This is illustrated in our application, a randomized trial comparing three treatments for acute myelogenous leukaemia (AML), where a ‘response’ is defined as the event that the patient survives at least 50 days. To implement an AR procedure, each time a new patient is enrolled the AR probabilities must be updated using the most recent data. An intrinsic problem when implementing AR based on the above binary response is that the responses of previous patients who are alive at a new patient’s accrual time, but who entered the trial less than 50 days earlier, are not known. If the randomization probabilities are computed using only the data from patients who have been observed and followed completely, then substantial information will be wasted. Moreover, the resulting data would be biased since, for example, this approach would ignore the data from a patient who has survived 45 days, but include one who died at day 45. An alternative is to randomize cohorts rather than single patients, while requiring that new patients must either be turned away or made to wait until the outcomes of all patients in the most recent cohort have been evaluated. This approach is impractical, however, since physicians usually are unwilling or unable to repeatedly suspend accrual, and the risks associated with delaying treatment or using an alternative therapy may outweigh the benefits of AR. An additional complication is that patient prognostic covariates may have a substantial effect on treatment outcome, and consequently patient heterogeneity may affect how an AR procedure behaves.

Nearly all AR methods in the statistical literature based on a binary outcome ignore the above problems by assuming that all outcomes are observed immediately. Exceptions include Eick [10], who considered a two-armed bandit treatment allocation procedure assuming geometric lifetimes, Louis [11], who considered exponential survival times, and Rosenberger and Seshaiyer [12], who studied a non-parametric AR method based on the logrank statistic. These methods assume that patients within a treatment arm are exchangeable, however. In this paper, we examine two Bayesian AR methods based on time-to-event outcomes with patient-specific baseline covariates in the context of the AML trial. The proposed methods

monitor the trial continuously and compute each new patient's randomization probabilities at the time of accrual without delay, utilizing all currently available data. The two methods differ in model formulation. The first approach assumes a completely specified Bayesian Weibull survival time regression model, with the prior distribution based on a preliminary fit to historical AML survival data. In the second approach, the survival time distribution is not fully specified, except for a stochastic ordering assumption and the constraint that the response probability is linked to a linear term accounting for treatment and baseline covariate effects. The second method utilizes an approximate posterior obtained by replacing nuisance parameters in a working likelihood with consistent estimates. This method generalizes the results of Cheung and Thall [13], who address the problem of early stopping in a phase II trial with exchangeable patients.

The remainder of the article is organized as follows. In Section 2, we describe the AML trial and present an analysis of the historical data. The exact and approximate Bayesian AR methods are presented in Sections 3 and 4, respectively. In Section 5, we present a simulation study of the two methods in the context of the AML trial, with each method implemented both with and without covariates. We close with some concluding remarks in Section 6.

## 2. HISTORICAL DATA FOR AN AML TRIAL

We consider a randomized phase II trial in acute myelogenous leukaemia (AML) aimed at selecting the most promising among three regimens: idarubicin+cytosine arabinoside (IA), IA+clofarabine (IAC), and IA+trioxacitabine (IAT). Since AML is a rapidly fatal disease and chemotherapy carries the risk of early regimen-related mortality, treatment success ('response') is defined as survival beyond day 50. We will use the rate of this event as the basis for comparing the three regimens.

In order to establish a Bayesian model that will be the basis for trial design and conduct, we first analysed historical data collected from 1146 AML patients treated at the MDACC between 1996 and 2001. A large variety of chemotherapy combinations were given to these historical patients. Most regimens contained cytosine arabinoside (A), combined with either idarubicin (I) or trioxacitabine (T). As there is no evidence from the historical data that the novel combinations IAC or IAT were either superior or inferior to the IA regimen, our goal is not to obtain information about the treatment effects from this data set. Rather, we hope to use the historical data to identify covariates that are predictive of survival so that we can incorporate them into the design of the proposed trial.

Let  $X$  denote survival time, and let  $T$  be a fixed time such that the patient's treatment is considered a success if  $X > T$ . In this example,  $T = 50$  days. Based on preliminary maximum likelihood fits of several models, we determined that a Weibull distribution gave a good fit to the historical data. This is illustrated by Figure 1, which gives a plot of  $\log[-\log\{\hat{S}_{KM}(t)\}]$  as a function of  $\log(t)$ , where  $\hat{S}_{KM}(t)$  is the Kaplan–Meier estimate of  $S(t) = \Pr(X > t)$ , and the straight line  $\log(\hat{\mu}) + \hat{\phi} \log(t)$  is determined by the maximum likelihood estimates of  $\mu$  and  $\phi$  under the Weibull model  $S(t) = \exp(-\mu t^\phi)$ .

For a Bayesian analysis, we included three baseline covariates that are well known to be prognostic of survival [14]:  $Z_1 = \text{age}$ ,  $Z_2 = \text{the indicator of the '−5/−7' cytogenetic abnormality wherein portions of the 5th or 7th chromosomes are missing}$ , and  $Z_3 = \text{the indicator of a poor}$

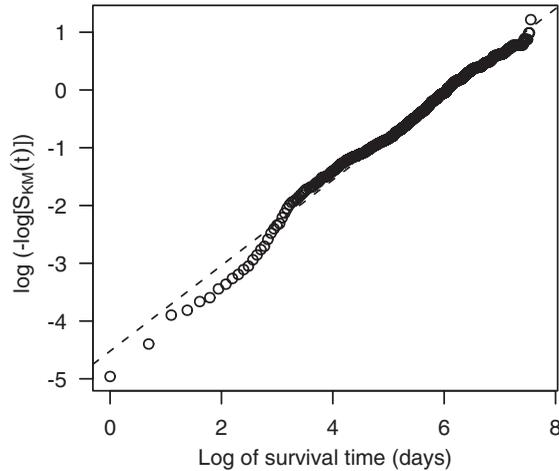


Figure 1. Plot of  $\log[-\log\{\hat{S}_{KM}(t)\}]$  as a function of  $\log(t)$  for the AML historical data. The dotted line is  $\log(\hat{\mu}) + \hat{\phi} \log(t)$  based on the maximum likelihood fit of the Weibull model ignoring covariates.

performance status (PS) wherein the patient is bedridden more than 50 per cent of the time. For a patient with prognostic covariates  $\mathbf{Z} = (Z_1, Z_2, Z_3)'$ , denote the survivor function of  $X$  by  $S(x|\mathbf{Z}, \boldsymbol{\theta}) = \Pr(X > x|\mathbf{Z}, \boldsymbol{\theta})$ , where  $\boldsymbol{\theta}$  denotes the vector of model parameters. Denoting

$$\boldsymbol{\beta}'\mathbf{Z} = \beta_1 \text{Age} + \beta_2 \mathbf{1}\{-5/-7 \text{ abnormality}\} + \beta_3 \mathbf{1}\{\text{PS} = \text{Poor}\}$$

the model including these covariates is given by

$$\log[-\log\{S(x|\mathbf{Z}, \boldsymbol{\theta})\}] = \log(\mu) + \phi \log(x) + \boldsymbol{\beta}'\mathbf{Z}$$

with  $\boldsymbol{\theta} = (\mu, \boldsymbol{\beta}, \phi)$ ,  $\mu > 0$  and  $\phi > 0$ . We assumed non-informative independent normal priors on  $(\log \mu, \beta_1, \beta_2, \beta_3, \log \phi)$ , each with mean 0 and variance 10.

The maximum likelihood estimates (MLEs) and Bayesian posterior values of the Weibull model parameters for the historical data are given in Table I. Posteriors were computed using the approximation given by Lindley and Smith [15]. The negative sign of the estimates of  $\log(\phi)$  shows that the hazard of death decreases over time. This corresponds to the well known fact that chemotherapy of AML carries a high risk of early death due to infection, toxicity, or the disease itself, and for patients who survive their induction therapy the risk of death continues to decrease. A corresponding Cox regression analysis (not shown) yielded virtually identical estimates of the three covariate parameters and variance estimates.

As expected, older age, the  $-5/-7$  cytogenetic abnormality, and poor PS were all negative risk factors, since  $\beta_j > 0$  is associated with a smaller value of  $X$  under the Weibull model. The column under 'Historical Data' in Table II gives the posterior median and 95 per cent credible interval (ci) of the 50-day survival probability,  $\Pr(X > T|\mathbf{Z}, \boldsymbol{\theta})$ , for each of the four prognostic subgroups determined by  $\mathbf{1}\{-5/-7 \text{ abnormality}\}$  and  $\mathbf{1}\{\text{PS} = \text{Poor}\}$ , for a 60-year-old patient. This illustrates the effects of these covariates on early survival in AML, as well as the variability inherent in the historical data.

Table I. Maximum likelihood estimates and Bayesian fit of the Weibull model with survivor function  $S(x|\mathbf{Z}, \boldsymbol{\theta}) = \exp\{-\mu x^\phi \exp(\boldsymbol{\beta}'\mathbf{Z})\}$  for the historical data.

Parameter	ML fit		Bayesian fit					
	Estimate*	SD	Mean	SD	Correlation with			log( $\phi$ )
					$\beta_1$	$\beta_2$	$\beta_3$	
log( $\mu$ )	-6.52	0.226	-6.49	0.225	-0.743	-0.203	-0.135	-0.703
$\beta_1$	0.021	0.0026	0.021	0.0026		-0.0547	-0.0099	0.0918
$\beta_2$	0.78	0.076	0.78	0.076			0.0035	0.203
$\beta_3$	1.13	0.112	1.13	0.112				0.139
log( $\phi$ )	-0.20	0.023	-0.20	0.023				

\*All  $P$ -values associated with the estimates under the ML fit were less than 0.0001.

Table II. Probabilities of 50-day survival by prognostic subgroup for a 60-year-old patient. Posterior medians are based on the fit of the Weibull model to the historical data, and prior medians are based on the prior used in the AR procedure in Section 3. Each median is followed by a 95 per cent credible interval, given in parentheses.

PS	-5/-7	Historical data	Prior used for trial conduct
		Median (95 per cent ci)	Median (95 per cent ci)
Good	No	0.88 (0.86, 0.89)	0.88 (0.30, 0.99)
Good	Yes <sup>†</sup>	0.75 (0.71, 0.78)	0.75 (0.06, 0.97)
Poor*	No	0.67 (0.60, 0.73)	0.67 (0.02, 0.96)
Poor	Yes	0.41 (0.33, 0.50)	0.41 (0.00, 0.92)

\*'Poor PS' means that the patient was bedridden more than 50 per cent of the time.

<sup>†</sup>-5/-7='Yes' indicates presence of the cytogenetic abnormality in which portions of chromosomes 5 or 7 are missing.

### 3. MODEL-BASED ADAPTIVE RANDOMIZATION

#### 3.1. Allocation criterion

Based on our fit of the historical data, we will assume the Weibull model in our development of the AR methodology. We extend the previous notation by introducing the treatment index  $k = 1, \dots, K$ , and denote the survivor function for treatment  $k$  by  $S_k(x|\mathbf{Z}, \boldsymbol{\theta})$ . Our AR criterion will be based on the fact that treatment  $k$  is superior to treatment  $k'$  if  $S_k(T|\mathbf{Z}, \boldsymbol{\theta}) \geq S_{k'}(T|\mathbf{Z}, \boldsymbol{\theta})$ . Let  $H_t$  denote the data accumulated from patients in the trial up to study time  $t$ . If patient  $i$  with covariates  $\mathbf{Z}_i$  enters the trial at study time  $t_i$ , we define the AR criterion

$$\gamma_i(k|\mathbf{Z}_i) = \Pr[S_k(T|\mathbf{Z}_i, \boldsymbol{\theta}) = \max_{1 \leq j \leq K} \{S_j(T|\mathbf{Z}_i, \boldsymbol{\theta})\} | H_{t_i}] \tag{1}$$

This is the posterior probability that treatment  $k$  is superior to all others, in terms of the probability of surviving beyond  $T$  given the covariate vector  $\mathbf{Z}_i$  of the patient accrued at  $t_i$ . Since  $\sum_{k=1}^K \gamma_i(k|\mathbf{Z}_i) = 1$  when the survival time distributions are continuous,  $\gamma_i(1|\mathbf{Z}_i), \dots, \gamma_i(K|\mathbf{Z}_i)$  may be used as the AR probabilities. This is similar to the AR criterion used by Thall

*et al.* [9], in a rather different setting, who adaptively randomize relapsed leukaemia patients among five different donor lymphocyte infusion times. For  $K=2$  and no covariate, (1) is similar to the randomization probability proposed by Thompson [16], who considered two independent binomial samples with probabilities following beta priors.

In this section, the AR procedure is based on the assumption that, in treatment group  $k$ ,  $X$  follows a Weibull distribution such that  $\log[-\log\{S_k(x|\mathbf{Z}, \boldsymbol{\theta})\}] = \log(\mu_k) + \phi \log(x) + \boldsymbol{\beta}'\mathbf{Z}$ , where  $\mu_k$  is the baseline rate parameter. Under this model, the AR criterion (1) is

$$\gamma_i(k) = \Pr[\mu_k = \min\{\mu_1, \dots, \mu_K\} | H_{t_i}] \quad (2)$$

which in particular does not depend on  $\mathbf{Z}$  or the time window  $T$ . Randomizing a patient to treatment  $k$  with probability  $\gamma_i(k)$  makes sense since a smaller  $\mu_k$  is associated with better survival under the Weibull model. Assuming a common shape  $\phi$  and covariate effects  $\boldsymbol{\beta}$  across treatments, the criteria (1) and (2) are equivalent to probabilistically favouring the arms with longer median survival times. In other words, pre-specification of a time interval  $T$  is not necessary for this method.

As an alternative to AR, one may view  $\gamma_i(k)$  as an allocation index and give a new patient the treatment having the largest  $\gamma_i(k)$ . This approach is similar to the idea of using a dynamic allocation index in a multi-armed bandit problem [17], although the latter is defined to maximize an expected reward specifically under a geometric discount sequence [18]. Operationally, such index-based strategies may be ‘frozen’ by an unbroken series of desirable outcomes in the most current patients. For example, with a binary treatment outcome, labelled ‘Success’ or ‘Failure’, a success in the most current patient will generate assignment of the next patient to the same treatment. In this sense, the index-based strategies imitate the deterministic play-the-winner (PTW) rule [19]. While Berry [20] shows that PTW is optimal in maximizing the number of successes under a model with independent beta priors with no covariates and delay, its deterministic feature has been criticized because it can prejudice the selection of patients admitted to the trial; see Reference [21] for example. Thus, outcome-adaptive randomization, as a control of selection bias in modern clinical research, provides a more practical solution than its deterministic counterpart.

The introduction of randomization into an outcome-adaptive design is not a novel concept. An example is the randomized PTW rule [2], whose characteristics can be established using an urn model. Our proposed method shares some of the philosophy of the randomized PTW rule, in that randomization is included in the adaptive treatment assignment to avoid the problem of selection bias. While implementation of our proposed method is considerably more complicated than randomized PTW, our method has the advantages that it accommodates censored data and adjusts for patient heterogeneity via explicit modelling. Moreover, since each treatment assignment decision is based on Bayesian posterior probability of an event (1), it is easily understood by non-statisticians. Explicit Bayesian modelling also facilitates prior elicitation.

### 3.2. Prior specification

To apply the AR procedure to the AML trial and compute (2), we must first specify a prior for  $\boldsymbol{\theta} = (\mu_1, \mu_2, \mu_3, \beta_1, \beta_2, \beta_3, \phi)$ . When using Bayesian methods to develop a clinical trial design, only one prior for  $\boldsymbol{\theta}$  may be specified. An important point in this regard is that the prior used for trial conduct need not be the only prior used to analyse the data at the end of the

trial. Indeed, it is conventional to carry out a Bayesian data analysis repeatedly using several priors, with each prior representing a different viewpoint by quantifying varying degrees of skepticism regarding treatment effects. This being said, a prior that is the basis for a Bayesian trial design must be fair in that it is not biased in favour of one treatment over the others, and also sufficiently vague so that artificial information is not introduced. However, the prior variances also must be calibrated so that the successive numerical computations of posteriors are sufficiently stable to allow the adaptive decision making criteria to be readily computed. To achieve these requirements for the AML trial, the prior parameter means were set equal to the historical posterior means, and the variances were specified by inflating the variances based on the historical data (Table I). Specifically, we set the variances of the  $\beta_j$ 's equal to 10 times the historical variances,  $\text{var}\{\log(\mu_j)\} = 1$ , which is roughly 20 times the historical variances of 0.051, and  $\text{var}\{\log(\phi)\} = 0.0008$ , which is roughly 1.5 times the historical variance of 0.00053. The inflation factor used to obtain  $\text{var}\{\log(\phi)\}$  was smaller since we found that the posterior computations were very sensitive to the distribution of  $\log(\phi)$  and larger values of  $\text{var}\{\log(\phi)\}$  destabilized the computations. This prior thus reflects average values in the historical data, but is sufficiently non-informative so that probability mass is assigned widely over the (0, 1) support. Table II gives the median and 95 per cent credible interval for each of the four prognostic groups defined by cytogenetics and PS, based on the prior used for trial conduct, and the corresponding prior distributions of  $\Pr(X > 50)$  are given in Figure 2. We assumed that all seven parameters were independent normals *a priori*. Although some of the historical correlations between  $\log(\mu)$  and other parameters are large (Table I), the historical treatments are different from those studied in the trial. For this reason, and to ensure that the MCMC algorithm for computing posteriors during the trial converges, we set all correlations equal to 0 in the prior used for trial design and conduct.

### 3.3. Likelihood and posterior computation

For study time  $t$ , let  $C_i(t) = \max(t - t_i, 0)$  denote the  $i$ th patient's time from entry to the trial. If  $X_i > C_i(t)$ , then the event time  $X_i$  has not been observed by  $t$ , hence  $C_i(t)$  is its current administrative censoring time. Let  $Y_i(t) = \mathbf{1}\{X_i > C_i(t)\}$  and  $\kappa_i$  be the  $i$ th patient's treatment. The likelihood at study time  $t$  is

$$\mathcal{L}(H_t | \boldsymbol{\theta}) = \prod_{k=1}^K \mathcal{L}_k(H_t | \boldsymbol{\theta}) = \prod_{k=1}^K \prod_{i=1}^{N(t)} \{f_k(x_i | \mathbf{Z}_i, \boldsymbol{\theta})\}^{1-Y_i(t)} \{S_k(C_i(t) | \mathbf{Z}_i, \boldsymbol{\theta})\}^{Y_i(t)} \mathbf{1}_{\{\kappa_i=k\}} \quad (3)$$

where  $f_k(x_i | \mathbf{Z}_i, \boldsymbol{\theta}) = -dS_k(x_i | \mathbf{Z}_i, \boldsymbol{\theta})/dx$  is the density of  $X_i$  under treatment  $k$  and  $N(t)$  is the number of patients who have been enrolled up to the study time  $t$ . Consequently, the joint posterior at study time  $t$  is

$$p(\mu_1, \dots, \mu_K, \phi, \boldsymbol{\beta} | H_t) \propto p(\phi, \boldsymbol{\beta}) \prod_k p(\mu_k) \mathcal{L}_k(H_t | \mu_k, \phi, \boldsymbol{\beta})$$

where  $\mathcal{L}_k(H_t | \mu_k, \phi, \boldsymbol{\beta})$  is the  $k$ th component of the likelihood given by (3). The priors  $p(\phi, \boldsymbol{\beta})$  and  $p(\mu_k)$  denote the normal densities described above.

We generated the posterior of  $(\mu_1, \mu_2, \mu_3, \phi, \boldsymbol{\beta})$  using Gibbs sampling by simulating from the full conditionals with adaptive rejection sampling [22]. The subvector  $\boldsymbol{\mu}$  was subsequently used to compute the randomization probabilities  $\gamma_i(k)$ 's. To calibrate the computations, we chose various combinations of number of chains, burn-in and chain length, and simulated

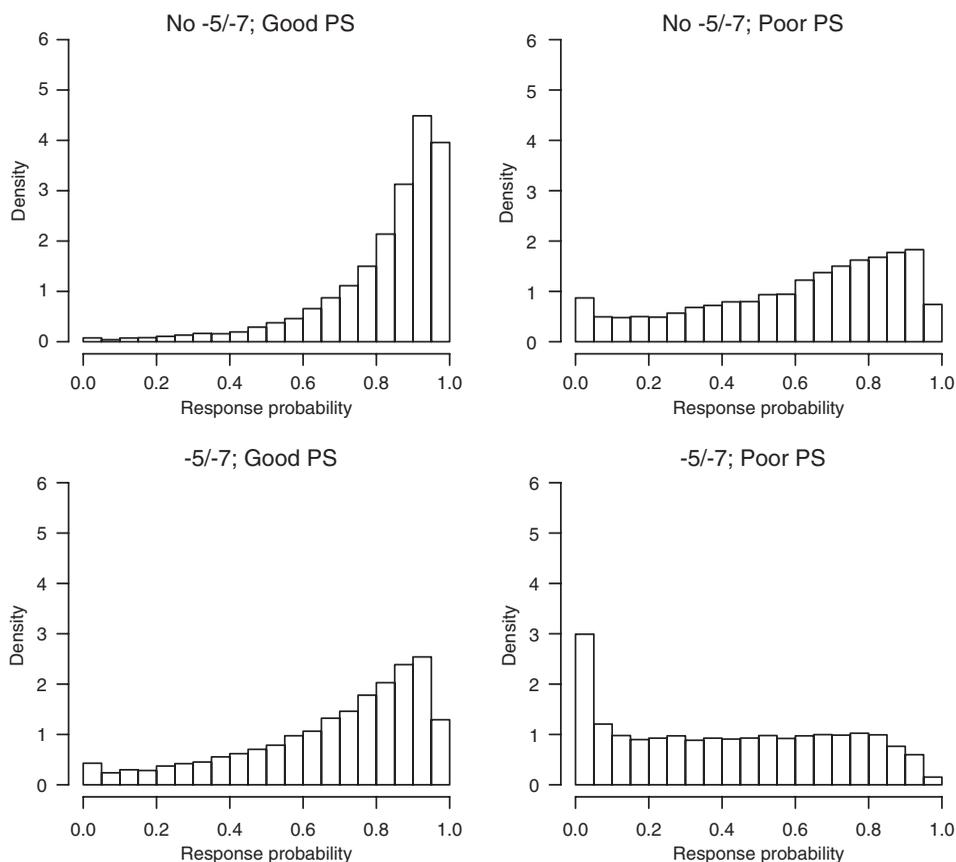


Figure 2. Distributions of the response probability,  $\Pr(X > 50)$ , for a 60-year-old patient, based on the prior used for trial conduct.

500 trials from each combination. The results indicated that three chains with a burn-in of 1000 followed by an additional 5000 samples gave adequate convergence. We assumed the chains had converged once the potential scale reduction was less than 1.1 as recommended by Gelman [22, pp. 131–143].

### 3.4. Ignoring prognostic covariates

To assess the value of including covariates in the AR criteria, we also considered an analogous model-based AR method that ignores covariates. Using this method, the survival times of all patients are assumed to follow the Weibull distribution given by  $S_k(x) = \exp(-\mu_k^* x^{\phi^*})$ . Under this model, the MLEs for  $\mu_k^*$  and  $\log(\phi^*)$  based on the historical data were  $-4.52$  and  $-0.30$ , respectively. In order to match the priors with those used in the covariate-adjusted method, we derived a prior for  $\mu_k^*$  and  $\phi^*$  based on the covariate-adjusted method for a ‘typical’ patient: aged 60, without the  $-5/-7$  abnormality, and with good PS. This implies that  $\log(\mu_k^*) + \phi^* \log(x) = \log(\mu_k) + \phi \log(x) + 60\beta_1$ , and we thus assumed  $\log(\mu_k^*)$  to follow

a normal prior with mean  $-5.23$  and variance  $0.52$  after adjusting for the negative historical correlation (Table I), and  $\phi^* = \phi \sim \text{LogNormal}(-0.20, 0.0008)$ . We assumed that  $\mu_k^*$  and  $\phi^*$  were independent *a priori* for computation stability. Posteriors were computed using Gibbs sampling, as for the method including covariates. The randomization probabilities were defined as in (2) with  $\mu_k$  replaced by  $\mu_k^*$ .

## 4. APPROXIMATE BAYES METHODS

### 4.1. A semi-parametric model

In the approximate Bayes approach, we do not fully specify a distribution of  $X$ , and require only two mild assumptions. The first assumption is that each  $S_k$  is stochastically ordered in  $\mathbf{Z}$ : for each pair  $\mathbf{Z}_1 \neq \mathbf{Z}_2$  and fixed  $\boldsymbol{\theta}$  we assume that either  $S_k(x|\mathbf{Z}_1, \boldsymbol{\theta}) \geq S_k(x|\mathbf{Z}_2, \boldsymbol{\theta})$  for all  $x$ , or  $S_k(x|\mathbf{Z}_1, \boldsymbol{\theta}) \leq S_k(x|\mathbf{Z}_2, \boldsymbol{\theta})$  for all  $x$ . In other words, the survival curves for different  $\mathbf{Z}$ 's do not cross. Since  $T$  is fixed, for convenience and brevity we denote  $S_k(T|\mathbf{Z}, \boldsymbol{\theta})$  by  $\zeta_k(\mathbf{Z})$  and suppress its dependence on  $\boldsymbol{\theta}$ . The second assumption is that

$$\zeta_k(\mathbf{Z}) = g^{-1}(\tau_k + \tilde{\boldsymbol{\beta}}' \mathbf{Z}) \quad (4)$$

where  $g$  is a suitable link function. We will assume the complementary log–log link,  $g(p) = \log\{-\log(p)\}$ , which is decreasing in  $p$ . This implies that treatment  $j$  is superior to treatment  $j'$  if  $\tau_j \leq \tau_{j'}$ , and consequently we define the randomization probability for arm  $k$  to be  $\tilde{\gamma}_i(k) = \Pr[\tau_k = \min\{\tau_1, \dots, \tau_K\} | H_{t_i}]$ . Under the Weibull model,  $\tilde{\boldsymbol{\beta}} = \boldsymbol{\beta}$  and  $\tau_k = \log(\mu_k) + \phi \log(T)$ , and the AR criterion  $\tilde{\gamma}_i(k)$  is identical to  $\gamma_i(k)$  defined in (2). Even if the Weibull assumption holds, however, the two AR methods are not identical due to differences in model formulation and methods for computing posteriors.

### 4.2. Prior specification

To facilitate comparison of the AR methods based on the full and semi-parametric models, we matched the prior specification for the parameters  $(\tau_1, \tau_2, \tau_3, \tilde{\boldsymbol{\beta}})$  to that used in the model-based AR method, as follows. First, the prior on  $\tilde{\boldsymbol{\beta}}$  is identical to that assumed for  $\boldsymbol{\beta}$ . In order to make the prior on the treatment effects,  $\tau_k$ , be similar to that in the model-based approach, we sampled  $\mu_k$  and  $\phi$  from their priors under the parametric model and calculated the sample mean and variance of  $\tau_k = \log(\mu_k) + \phi \log(50)$  which, under the assumption of normality, yielded  $\tau_k \sim N(-3.28, 0.88)$ . Note that although  $\phi$  had a lognormal prior in the model-based method, the simulation results verified that  $\tau_k$  thus defined was approximately normal because  $\text{var}\{\log(\phi)\}$  was small compared to  $\text{var}\{\log(\mu_k)\}$ . In order to ensure that the MCMC would converge under the wide range of data that could occur during the trial, we assumed  $\tau$  and  $\boldsymbol{\beta}$  were mutually independent.

As with the model-based method, we also examined the AR method based on model (4) without covariates, i.e. assuming  $\zeta_k = g^{-1}(\tau_k^*)$ . In this case, while it is convenient to assume that the response probability  $\zeta_k$  follows a beta prior, which will lead to a closed form posterior [13], we specified the prior on  $\tau_k^*$  so that it matched the priors of the other methods. We did this by noting that  $\tau_k^* = \tau_k + 60\beta_1$  for a ‘typical’ patient. Based on the priors of  $\tau_k$  and  $\beta_k$  under the Weibull model, we obtained  $\tau_k^* \sim N(-2.04, 0.42)$  for the prior under the semi-parametric model.

### 4.3. Working likelihood and posterior computation

The approximate Bayes method focuses on  $\xi_k(\mathbf{Z})$  and exploits the fact that, given  $X > T$ , the particular value of  $X$  contributes no additional information about  $\{\xi_1(\mathbf{Z}), \dots, \xi_K(\mathbf{Z})\}$ . Thus, with the approximate method we will consider information only up to time  $T$  for each patient. Let  $\tilde{C}_i(t) = \min\{C_i(t), T\}$  denote the  $i$ th patient's time on study up to a maximum of  $T$ . Define the patient's observable response status process  $A_i^o(t) = \{X_i > \tilde{C}_i(t)\}$ , and let  $\tilde{Y}_i(t) = \mathbf{1}\{A_i^o(t)\}$ . The current status likelihood for arm  $k$  at time  $t$  can be written as

$$\tilde{\mathcal{L}}_k(H_t|\boldsymbol{\theta}) = \prod_{i=1}^{N(t)} [\{P_k(A_i^o(t)|\mathbf{Z}_i, \boldsymbol{\theta})\}^{\tilde{Y}_i(t)} \{1 - P_k(A_i^o(t)|\mathbf{Z}_i, \boldsymbol{\theta})\}^{1-\tilde{Y}_i(t)}] \mathbf{1}_{\{\kappa_i=k\}} \quad (5)$$

where  $P_k(\cdot|\mathbf{Z}_i, \boldsymbol{\theta}) = \Pr(\cdot|\mathbf{Z}_i, \boldsymbol{\theta}, \kappa_i = k)$ . This is a partial likelihood for the probability measures  $\{P_k\}_{k=1}^K$ , which are of infinite dimension. If all currently enrolled patients have been completely followed, i.e.  $\tilde{C}_i(t) \equiv T$ , then (5) reduces to a binomial likelihood in the response probability  $\xi_k(\mathbf{Z}_i)$ . Because the current status likelihood contains nuisance parameters as well as the probabilities  $\{\xi_1(\mathbf{Z}_i), \dots, \xi_K(\mathbf{Z}_i)\}$  of interest, denoting  $\omega_{ki}(t|\mathbf{Z}_i) = P_k\{A_i^o(t)|X_i \leq T, \mathbf{Z}_i\}$ , we will exploit the decomposition  $P_k\{A_i^o(t)|\mathbf{Z}_i, \boldsymbol{\theta}\} = \xi_k(\mathbf{Z}_i) + \omega_{ki}(t|\mathbf{Z}_i)\{1 - \xi_k(\mathbf{Z}_i)\}$ , which follows from the fact that  $P_k\{A_i^o(t)|X_i > T, \mathbf{Z}_i\} = 1$ . Noting that  $1 - P_k\{A_i^o(t)|\mathbf{Z}_i, \boldsymbol{\theta}\} = \{1 - \omega_{ki}(t|\mathbf{Z}_i)\}\{1 - \xi_k(\mathbf{Z}_i)\}$ , we obtain a 'working' likelihood for arm  $k$  by replacing the  $\omega_{ki}$  with an estimate  $\hat{\omega}_{ki}$ :  $\tilde{\mathcal{L}}_k^{(w)}(H_t|\boldsymbol{\theta}) \propto$

$$\prod_{i=1}^{N(t)} [\{\xi_k(\mathbf{Z}_i) + \hat{\omega}_{ki}(t|\mathbf{Z}_i)(1 - \xi_k(\mathbf{Z}_i))\}^{\tilde{Y}_i(t)} (\{1 - \hat{\omega}_{ki}(t|\mathbf{Z}_i)\}\{1 - \xi_k(\mathbf{Z}_i)\})^{1-\tilde{Y}_i(t)}] \mathbf{1}_{\{\kappa_i=k\}} \\ \propto \prod_{i=1}^{N(t)} [\{\xi_k(\mathbf{Z}_i) + \hat{\omega}_{ki}(t|\mathbf{Z}_i)(1 - \xi_k(\mathbf{Z}_i))\}^{\tilde{Y}_i(t)} \{1 - \xi_k(\mathbf{Z}_i)\}^{1-\tilde{Y}_i(t)}] \mathbf{1}_{\{\kappa_i=k\}}$$

where the multiplier  $\{1 - \hat{\omega}_{ki}(t|\mathbf{Z}_i)\}$  does not depend on  $\theta$ , and  $\xi_k$  is specified by the parameters  $\tau_k$  and  $\tilde{\boldsymbol{\beta}}$  via the regression model (4). Thus, an approximate posterior based on the working likelihood at time  $t$  is

$$\tilde{p}(\tau_1, \dots, \tau_K, \tilde{\boldsymbol{\beta}}|H_t) \propto \tilde{p}(\tilde{\boldsymbol{\beta}}) \prod_{k=1}^K \tilde{p}(\tau_k|\tilde{\boldsymbol{\beta}}) \tilde{\mathcal{L}}_k^{(w)}(H_t|\tau_k, \tilde{\boldsymbol{\beta}})$$

where the priors  $\tilde{p}$  are the normal densities described in Section 4.2. The triplets  $(\tau_1, \tau_2, \tau_3)$  can be generated from this approximate posterior using the Gibbs sampler, and these are used to compute the  $\tilde{\gamma}_i(k)$ 's for implementing the AR.

### 4.4. Estimating nuisance parameters

A key step in computing the working likelihood is to estimate the nuisance  $\omega_{ki}(t|\mathbf{Z})$  at study time  $t$ . To facilitate exposition, we temporarily suppose that there is only one covariate,  $Z$ ,

restrict attention to treatment arm  $k$ , and assume that, for all  $x > 0$ ,

$$S_k(x|Z_1, \boldsymbol{\theta}) \geq S_k(x|Z_2, \boldsymbol{\theta}) \quad \text{if } Z_1 \leq Z_2 \quad (6)$$

that is,  $S_k(x|Z, \boldsymbol{\theta})$  is non-increasing in  $Z$ . Note that  $S_k(x|Z, \boldsymbol{\theta}) = E_k\{U_i(x)|Z, \boldsymbol{\theta}\}$  where  $U_i(x) = \mathbf{1}(X_i > x)$ , and thus  $S_k(x|Z, \boldsymbol{\theta})$  can be estimated by isotonic regression from the reduced sample  $\{U_i(x), Z_i\}$  for  $i$  such that  $S_i = k$  and  $\tilde{C}_i(t) = T$ ; see Appendix A. As a result, denoting the isotonic regression estimator by  $\hat{S}_k(x|Z)$ , the  $\omega_{ki}(t|Z_i)$ 's are estimated by

$$\hat{\omega}_{ki}^{\text{emp}}(t|Z_i) = \frac{\hat{S}_k(\tilde{C}_i(t)|Z_i) - \hat{S}_k(T|Z_i)}{1 - \hat{S}_k(T|Z_i)}$$

To avoid highly variable estimators with small samples, we estimated  $\omega$  using a weighted average of the empirical component  $\hat{\omega}^{\text{emp}}$  and a fixed prior component  $\hat{\omega}^{\text{prior}}$

$$\hat{\omega}_{ki}(t|Z_i) = \frac{m_k(t|Z_i)}{m_k(t|Z_i) + 1} \hat{\omega}_{ki}^{\text{emp}}(t|Z_i) + \frac{1}{m_k(t|Z_i) + 1} \hat{\omega}_{ki}^{\text{prior}}(t)$$

where  $m_k(t|Z_i) = m_k(t)\{1 - \hat{S}_k(T|Z_i)\}$  and  $m_k(t)$  is the number of patients in arm  $k$  who have been completely followed at time  $t$ . Using this weighted average, the prior component will become less influential as the sample size grows. For simplicity, we take  $\hat{\omega}_{ki}^{\text{prior}}(t) = 1 - \tilde{C}_i(t)/T$  (see also Reference [13]).

For the general case with more than one covariate, we apply isotonic regression under the stochastic ordering conditions (6) with  $\tilde{\boldsymbol{\beta}}'Z$  in place of the single  $Z$ . In actual implementation, we replace  $\tilde{\boldsymbol{\beta}}$  with the most current updated posterior mode  $\tilde{\boldsymbol{\beta}}_0$  of  $\tilde{\boldsymbol{\beta}}$ , assume stochastic ordering (6) for  $\tilde{\boldsymbol{\beta}}_0'Z$ , and apply this to obtain estimators of the  $\omega_{ki}$ 's as before.

## 5. APPLICATION TO THE AML TRIAL

### 5.1. Trial designs

We now apply each of the four AR methods to the AML trial. A total of 96 patients are randomized to the three treatment arms, with an anticipated accrual rate of 8 patients per month. Because the AR criteria tend to be somewhat unstable for small amounts of data, the first 30 patients are randomized evenly, with exactly 10 patients assigned to each arm. The AR procedure is applied thereafter, for the remaining 66 patients, with the AR criterion computed based on the most recent data available when each new patient arrives. For the approximate Bayes methods, each patient is followed for at most  $T = 50$  days. For the Weibull model-based method all patients are followed until either  $X$  is observed or the study ends. At the end of the study, defined as 50 days after the last patient is randomized, the treatment that has the largest posterior randomization probability criterion is selected.

### 5.2. Simulation study design

For the simulation study, each patient's prognostic covariates were drawn at random from the historical data. Patients' survival times were generated under the Weibull model, or a

Table III. Parameter values, median survival, and 50-day survival probabilities of the models used in the simulations. Only parameters that differ from the historical values are given.

Scenario	Arm	med*	Weibull		GOR	
			log $\mu_k$	Prob <sup>†</sup>	$\lambda_k$	Prob <sup>†</sup>
0	IA	1.00	-6.49	0.88	1070	0.80
	IAC	1.00	-6.49	0.88	1070	0.80
	IAT	1.00	-6.49	0.88	1070	0.80
1	IA	1.00	-6.49	0.88	1070	0.80
	IAC	1.00	-6.49	0.88	1070	0.80
	IAT	1.50	-6.82	0.91	1611	0.84
2	IA	1.00	-6.49	0.88	1074	0.80
	IAC	1.25	-6.67	0.90	1338	0.82
	IAT	1.50	-6.82	0.91	1611	0.84

\*Median survival relative to IA.

†50-day survival probability for a 60-year old patient with good PS and no -5/-7 abnormality.

generalized odds rate (GOR) model [23] characterized by the survivor function

$$S_k(x|\lambda_k, \phi, \zeta, \boldsymbol{\beta}, \mathbf{Z}, \boldsymbol{\theta}) = \{1 + \zeta(x/\lambda_k)^\phi e^{\boldsymbol{\beta}'\mathbf{Z}}\}^{-1/\zeta} \quad \text{for } x > 0$$

where  $\phi > 0, \zeta > 0$  and each  $\lambda_k > 0$ . The GOR model is a family that contains the Weibull distribution (as  $\zeta \rightarrow 0$ ), but in general does not satisfy the proportional hazards assumption. Thus, evaluation of the AR methods under the GOR model provides an assessment of robustness.

The trial was simulated under three different scenarios, with the data generated under the Weibull with fixed parameter values. Under the null case, scenario 0, the fixed parameter values were set equal to the historical means of  $\boldsymbol{\theta}$  (Table I). For alternative scenario 1, we calibrated  $\log \mu_3$  so that IAT improved upon the historical median survival by 50 per cent. For alternative scenario 2,  $\log \mu_2$  and  $\mu_3$  were calibrated so that IAC and IAT provided 25 and 50 per cent improvements in median survival, respectively.

The three scenarios were also simulated under the GOR model, with  $\zeta = 2$  throughout and  $\lambda_k = (-2 \log 0.5/3\mu_k)^{1/0.819}$  chosen to match the median survival times in the Weibull scenarios. The covariate coefficients  $\boldsymbol{\beta}$  and the shape  $\phi$  were the same under the two models, set equal to the historical means in all scenarios. The simulation scenarios are summarized in Table III.

### 5.3. Simulation results

Under the null scenarios, all four methods have the same selection percentages and sample sizes by symmetry. However, when compared to a non-adaptive, balanced randomization, where on average each arm will receive 32 patients with a standard deviation (SD) of 4, the AR methods induce extra variability, with the corresponding SD varying from 11 to 15.

CONTINUOUS BAYESIAN ADAPTIVE RANDOMIZATION

Table IV. Selection probabilities and average sample sizes for the AML trial. Standard errors are given as subscripts.

Scenario	Model	Method	Covariates	Percent selected			Sample sizes		
				IA	IAC	IAT	IA	IAC	IAT
1	Weibull	Full model	Yes	15. <sub>50</sub>	15. <sub>50</sub>	70. <sub>65</sub>	27. <sub>16</sub>	27. <sub>16</sub>	42. <sub>20</sub>
			No	15. <sub>50</sub>	16. <sub>52</sub>	70. <sub>65</sub>	26. <sub>16</sub>	29. <sub>18</sub>	41. <sub>19</sub>
		Approximate	Yes	22. <sub>56</sub>	22. <sub>56</sub>	57. <sub>70</sub>	28. <sub>17</sub>	29. <sub>17</sub>	39. <sub>19</sub>
			No	22. <sub>56</sub>	23. <sub>66</sub>	56. <sub>70</sub>	29. <sub>16</sub>	29. <sub>16</sub>	38. <sub>19</sub>
2	Weibull	Full model	Yes	12. <sub>46</sub>	30. <sub>65</sub>	58. <sub>70</sub>	25. <sub>15</sub>	32. <sub>18</sub>	39. <sub>20</sub>
			No	19. <sub>55</sub>	29. <sub>64</sub>	52. <sub>71</sub>	27. <sub>18</sub>	32. <sub>18</sub>	37. <sub>20</sub>
		Approximate	Yes	18. <sub>54</sub>	33. <sub>66</sub>	49. <sub>71</sub>	27. <sub>16</sub>	32. <sub>18</sub>	37. <sub>19</sub>
			No	19. <sub>55</sub>	32. <sub>66</sub>	49. <sub>71</sub>	28. <sub>16</sub>	32. <sub>17</sub>	36. <sub>19</sub>
1	GOR	Full model	Yes	19. <sub>55</sub>	22. <sub>56</sub>	59. <sub>70</sub>	26. <sub>18</sub>	31. <sub>18</sub>	40. <sub>19</sub>
			No	21. <sub>58</sub>	21. <sub>58</sub>	58. <sub>70</sub>	30. <sub>18</sub>	29. <sub>16</sub>	37. <sub>20</sub>
		Approximate	Yes	20. <sub>57</sub>	20. <sub>57</sub>	60. <sub>69</sub>	28. <sub>16</sub>	29. <sub>16</sub>	39. <sub>18</sub>
			No	21. <sub>58</sub>	22. <sub>56</sub>	57. <sub>70</sub>	29. <sub>16</sub>	29. <sub>15</sub>	37. <sub>18</sub>
2	GOR	Full model	Yes	18. <sub>54</sub>	36. <sub>68</sub>	46. <sub>70</sub>	26. <sub>17</sub>	32. <sub>18</sub>	39. <sub>21</sub>
			No	20. <sub>58</sub>	35. <sub>67</sub>	45. <sub>70</sub>	27. <sub>15</sub>	31. <sub>18</sub>	38. <sub>19</sub>
		Approximate	Yes	20. <sub>57</sub>	31. <sub>65</sub>	49. <sub>71</sub>	28. <sub>16</sub>	31. <sub>17</sub>	37. <sub>19</sub>
			No	26. <sub>62</sub>	30. <sub>65</sub>	44. <sub>70</sub>	27. <sub>15</sub>	32. <sub>16</sub>	37. <sub>17</sub>

The simulation results for the alternative scenarios are summarized in Table IV. Under scenario 1 with Weibull data, where IAT provides a 50 per cent improvement in median survival compared to the other two arms, the model-based approaches have about a 70 per cent correct selection probability, with 41–42 patients treated on the IAT arm and 26–29 on each of the inferior arms. This is better than the approximate Bayes methods, which have a 56–57 per cent correct selection rate and allocate about 3 fewer patients to the superior arm. With balanced, non-adaptive randomization, selection of the best treatment arm at the end of the trial may still be based on the posterior criterion (2). Under scenario 1, if (2) is computed based on the full Weibull model, then equal randomization gives selection percentages (15,14,71); if (2) is computed based on the approximate posterior, the corresponding selection percentages are (22,21,57). Thus, despite the larger variability of the sample sizes with AR, the selection reliability of the AR methods is almost unaffected when compared to the balanced design. The behaviours of the methods are very similar under scenario 2 for Weibull data.

When the survival times follow a GOR distribution with shape parameter 2, the AR methods assuming a Weibull model lose about 7–12 per cent in correct selection rate, and have 2–4 fewer patients allocated to the best arm, depending on whether covariates are included. The approximate Bayes method with covariate adjustment suffers no loss in efficiency, despite the fact that the complementary log–log link is incorrect under the GOR model. This suggests that the approximate method is robust against a misspecified link, if heterogeneity is properly accounted for. Under the GOR model, the performances of all four AR methods are similar.

While it may seem somewhat surprising that ignoring covariates has no apparent effect on the selection percentages and sample sizes, this may be due in part to the increased variation in estimating additional parameters. However, a closer look reveals potential pathological behaviours for the unadjusted methods. The Weibull method without covariate adjustment selects the IA arm with probability 15 per cent under scenario 1, but with a larger probability of 19 per cent under scenario 2. This is counter-intuitive, since scenario 2 is 'easier' in that there are two better treatment arms than the IA arm, and thus any coherent method should select the IA arm with less probability in scenario 2 than in scenario 1. This pathology contributes to the drop in the selection percentage of the best arm, from 58 to 52 per cent. Similarly, the approximate Bayes method that ignores covariates selects the IA arm with a higher probability in scenario 2 (26 per cent) than in scenario 1 (21 per cent) under the GOR model. Therefore, accounting for covariates appears to lessen the likelihood of pathological behaviour.

## 6. CONCLUDING REMARKS

We have presented exact and approximate Bayesian outcome-adaptive randomization procedures based on time-to-event outcomes, accounting for baseline prognostic covariates, that are applied continuously during a clinical trial. Our simulations indicate that, on average, all of the AR procedures allocate substantially more patients to superior treatment arms. When compared to the non-adaptive design with balanced randomization, the AR procedures have about the same correct selection rates. Thus, our AR procedures provide an ethically much more attractive family of clinical trial designs, compared to designs with conventional randomization, without sacrificing correct selection probability. A small price is that patient allocation is somewhat more variable with AR.

A feature of the AR procedures is that they combine Bayesian data analysis and trial design. Thus, the final data from the trial at the end can be analysed utilizing the same Bayesian model that is the basis for trial design and conduct. In this final data analysis, several priors may be used in a conventional Bayesian sensitivity analysis. Under the likelihood principle, the final inferences will depend on the AR procedure only via the imbalance in the treatment arms produced by the AR. The unbalanced allocation may result in an increased variability in the analysis. This may be regarded as the price paid for the fact that assigning more patients to the superior treatment arms is ethically more desirable.

Our simulations show that accounting for covariate effects reduces the probability that an AR method will behave pathologically. In particular, the approximate Bayes method with covariate adjustment appears to be robust to link mis-specification. In this article, we assume that the covariates affect survival under each treatment in an identical way. However, covariate-treatment interaction is an important issue in AR. In theory, an AR criterion similar to (1) may be extended to include covariate-treatment interaction. However, the issues of implementation and interpretation may become quite complex. This is an area for future investigation.

Finally, in comparing the fully parametric Bayesian approach to the semi-parametric approximate Bayes method, the latter had substantially lower correct selection probabilities and produced a slightly smaller imbalance in favour of superior treatment arms.

APPENDIX A: ESTIMATING  $S_k(x|Z, \theta)$  VIA ISOTONIC REGRESSION  $\hat{S}_k(x|Z)$ 

Define  $\mathcal{C}_k \equiv \{j : S_j = k \text{ and } \tilde{C}_j(t) = T\}$  as the index set for completely followed patients in arm  $k$ . Re-index the patients in  $\mathcal{C}_k$  so that  $Z_1 < Z_2 < \dots$  and let  $M_l$  denote the number of patients in  $\mathcal{C}_k$  with covariate  $Z_l$ . The isotonic estimator  $\hat{S}_k(x|Z)$  can be obtained by the pool-adjacent-violators algorithm [24]:

1. If  $U_1(x) \geq U_2(x) \geq \dots$ , then  $\hat{S}_k(x|z) = U_l(x)/M_l$  for  $z \in [Z_l, Z_{l+1})$ .
2. If  $U_l(x)/M_l \leq U_{l+1}(x)/M_{l+1}$  for some  $l$ , then combine  $Z_l$  and  $Z_{l+1}$  into one covariate level by letting  $\hat{S}_k(x|Z_l) = \hat{S}_k(x|Z_{l+1})$ . Replace the pair  $\{U_l(x), U_{l+1}(x)\}$  in the sequence  $\{U_r(x)\}$  by  $U_l(x) + U_{l+1}(x)$ , and replace the pair  $\{M_l, M_{l+1}\}$  in the sequence  $\{M_r\}$  by  $M_l + M_{l+1}$ . Thus, the two ratios  $U_l(x)/M_l(x)$  and  $U_{l+1}(x)/M_{l+1}$  are replaced in the sequence  $\{U_r(x)/M_r\}$  by the single ratio  $\{U_l(x) + U_{l+1}(x)\}/\{M_l + M_{l+1}\}$ .
3. Repeat (1) and (2) until an ordered, monotone non-increasing set of ratios is obtained.

## ACKNOWLEDGEMENTS

Peter Thall's work was partially supported by NCI Grant RO1-CA-83932. The authors are grateful to two referees for their constructive comments.

## REFERENCES

1. Emanuel EJ, Patterson WB. Ethics of randomized clinical trials. *Journal of Clinical Oncology* 1998; **48**:6–29.
2. Wei LJ, Durham S. The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association* 1978; **73**:840–843.
3. Ware JH. Investigating therapies of potentially great benefit: ECMO (with discussion). *Statistical Science* 1989; **4**:298–340.
4. Berry DA, Eick GE. Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. *Statistics in Medicine* 1995; **14**:231–246.
5. Kadane J. *Bayesian Methods and Ethics in a Clinical Trial Design*. Wiley: New York, 1996; 19–63.
6. Pocock S, Simon R. Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics* 1975; **31**:103–115.
7. Wei LJ. The adaptive biased coin design for sequential experiments. *The Annals of Statistics* 1978; **6**:92–100.
8. Therneau TM. How many stratification factors are 'too many' to use in a randomization plan? *Controlled Clinical Trials* 1993; **14**:98–108.
9. Thall PF, Inoue LYT, Martin T. Adaptive decision making in a lymphocyte infusion trial. *Biometrics* 2002; **58**:560–568.
10. Eick GE. The two-armed bandit with delayed responses. *The Annals of Statistics* 1988; **16**:254–264.
11. Louis TA. Sequential allocation in clinical trials comparing two exponential survival curves. *Biometrika* 1977; **33**:627–634.
12. Rosenberger WF, Seshaiyer P. Adaptive survival trials. *Journal of Biopharmaceutical Statistics* 1997; **7**: 617–624.
13. Cheung YK, Thall PF. Monitoring the rates of composite events with censored data in phase II clinical trials. *Biometrics* 2002; **58**:89–97.
14. Estey EH, Thall PF, Beran M, Kantarjian H, Pierce S, Keating M. Effect of diagnosis (RAEB, RAEB-t, or AML) on outcome of AML-type chemotherapy. *Blood* 1997; **90**:2969–2977.
15. Lindley DV, Smith AFM. Bayesian estimates for the linear model (with discussion). *Journal of the Royal Statistical Society, Series B* 1972; **34**:1–41.
16. Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 1933; **25**:285–294.
17. Gittins JC. Bandit processes and dynamic allocation indices (with discussion). *Journal of the Royal Statistical Society, Series B* 1979; **48**:148–177.
18. Berry DA, Fristedt B. *Bandit Problems: Sequential Allocation of Experiments*. Chapman & Hall: London, 1985.

19. Zelen M. Play the winner rule and the controlled clinical trial. *Journal of the American Statistical Association* 1969; **64**:131–146.
20. Berry DA. A Bernoulli two-armed bandit. *The Annals of Mathematical Statistics* 1972; **43**:871–897.
21. Simon R. Adaptive treatment assignment methods and clinical trials. *Biometrics* 1977; **33**:743–749.
22. Gelman. *Markov Chain Monte Carlo in Practice*, Gilks WR, Richardson S, Spiegelhalter DJ (eds). Chapman & Hall: London, 1996; 131–143.
23. Dabrowska DM, Doksum KA. Estimation and testing in a two-sample generalized odds-rate model. *Journal of the American Statistical Association* 1988; **83**:744–749.
24. Ayer M, Brunk HD, Ewing GM, Reid WT, Silverman E. An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics* 1955; **26**:641–647.

