



King's Research Portal

DOI: 10.1002/sim.7353

Document Version Peer reviewed version

Link to publication record in King's Research Portal

Citation for published version (APA):

Barrett, J. É. (2017). Information-adaptive clinical trials with selective recruitment and binary outcomes. Statistics in Medicine, 36(18), 2803-2813. https://doi.org/10.1002/sim.7353

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

•Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research. •You may not further distribute the material or use it for any profit-making activity or commercial gain •You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Information-adaptive clinical trials: a selective recruitment design

James E. Barrett

University College London, London, U.K.

Summary. We propose a novel adaptive design for clinical trials with time-to-event outcomes and covariates (which may consist of or include biomarkers). Our method is based on the expected entropy of the posterior distribution of a proportional hazards model. The expected entropy is evaluated as a function of a patient's covariates, and the information gained due to a patient is defined as the decrease in the corresponding entropy. Candidate patients are only recruited onto the trial if they are likely to provide sufficient information. Patients with covariates that are deemed uninformative are filtered out. A special case is where all patients are recruited, and we determine the optimal treatment arm allocation. This adaptive design has the advantage of potentially elucidating the relationship between covariates, treatments, and survival probabilities using fewer patients, albeit at the cost of rejecting some candidates. We assess the performance of our adaptive design using data from the German Breast Cancer Study group and numerical simulations of a biomarker validation trial.

Keywords: Adaptive trial design; Optimal allocation; Proportional hazards model

1. Introduction

Adaptive clinical trials offer a potentially more efficient and ethical way to conduct clinical trials. Covariate-adaptive designs try to ensure that the distributions of covariates across different arms are balanced, thus resulting in more comparable cohorts on each arm (Pocock and Simon, 1975; Taves, 1974). Response-adaptive randomisation attempts to allocate more patients to the effective treatment arms. As the trial progresses and more information is acquired on the efficacies of each treatment arm the allocation probabilities

shift towards the more effective treatments. Zhang and Rosenberger (2007) develop an optimal response-adaptive design under exponential and Weibull parametric models for time-to-event outcomes. See Yin (2012) for a good overview of adaptive designs.

We regard the primary goal of a clinical trial as establishing a statistical relationship between covariates, treatments, and survival outcomes. As we will show, not all patients on a trial provide the same amount of statistical information. Some covariate values are more informative than others. In addition, the informativeness of a covariate value will depend on what has been observed so far in the trial. As an example, consider two scenarios where a patient with particular covariate values is available for recruitment. In the first scenario another patient with precisely the same covariate values has already been recruited. In the second scenario suppose the candidate's covariates come from a region of covariate space that has not previously been sampled. Intuitively we expect the candidate to be more informative in the second scenario since they provide access to previously unobserved covariates values and outcomes.

Our aim in this paper is to address a practical question: given limited resources and the observation that not all patients are equally informative, what is the optimal way to conduct a clinical trial? We propose that it may be advantageous to selectively recruit and allocate patients on the basis of how much information they are likely to provide. Covariates are measured for candidate patients, and based on those values and what has been inferred from the trial up to that point a recruitment probability is computed. In other words, we *filter* out patients that are unlikely to significantly reduce the uncertainty surrounding model parameters.

Predictive biomarkers, which indicate whether a patient is likely to respond well to a particular treatment or not, are increasingly useful in the drive towards personalised medicine and targeted therapy. A potential application of our selective-recruitment design would be to validate a biomarker by looking at treatment-biomarker interaction terms in a proportional hazards model. We test this using numerical simulations. Sargent et al. (2005) discuss alternative adaptive designs for validating predictive biomarkers.

Our filtering approach is similar in spirit to some existing designs. Freidlin and Simon

(2005) propose a trial design which attempts to find a gene signature that will identify a subset of 'sensitive' patients who are more likely to respond to the treatment. In a randomised discontinuation design (Rosner et al., 2002) patients who fail to respond to a treatment in the first phase of the trial are dropped from the second part, thereby isolating a responsive subset of patients with a stronger statistical signal. Another type of trial known as 'enrichment designs' (Temple, 2010) enrich the recruited cohort with patients who are more likely to have the event of interest. For example, patients with a particular biomarker. Given that more events of interest are observed greater statistical power can be achieved within the enriched cohort.

We assume a proportional hazards model with a constant baseline hazard rate. The entropy of the posterior distribution is a useful way to quantify our uncertainty regarding the model parameters. As the trial progresses, and the space of plausible parameter values shrinks, the entropy decreases. The *informativeness* of a candidate is defined as the reduction in expected entropy in the hypothetical scenario where they are added to the cohort of existing recruits. The *ideal* candidate at time t is defined as the patient that would achieve the greatest possible reduction in expected entropy. By comparing the current candidate to the ideal candidate we can obtain a recruitment probability. The posterior is constructed using outcomes from all patients accrued up until time t. Patients who have not experienced any events are considered to be right-censored. Therefore, the recruitment probability changes dynamically as more events and patients are observed. An arm allocation probability can also be computed based on which arm has the lowest expected entropy. We also implement this in a more traditional setting where all candidates are recruited.

In Section 2 we provide the mathematical details and describe some approximations which are required. Results from experimental data generated by the German Breast Cancer Study group and numerical simulations are presented in Sections 3 and 4 respectively. Discussion on the practical applicability of our approach and concluding remarks are given in Section 5.

2. An information based adaptive protocol

2.1. Proportional hazards model

Suppose that N_t patients have been recruited onto the trial at time t. Observed data are denoted by $D_t = \{(\mathbf{x}_1, t_1, \Delta_1), \dots, (\mathbf{x}_{N_t}, t_{N_t}, \Delta_{N_t})\}$ where $\mathbf{x}_i \in \mathbb{R}^d$ is a vector of covariates for patient i (this vector may include biomarker values or treatment indicator variables). If patient i is censored then $\Delta_i = 0$ and t_i is the time of censoring, otherwise the primary event occurred at time t_i and $\Delta_i = 1$. Patients who have not experienced any event by t are considered right censored. We assume a proportional hazards model with a constant baseline hazard rate $\lambda \in (0, \infty)$:

$$h(t_i | \mathbf{x}_i, \lambda, \beta) = \lambda e^{\beta \cdot \mathbf{x}_i} \quad \text{for } i = 1, \dots, N_t$$
(1)

where $\boldsymbol{\beta} \in \mathbb{R}^d$ is a vector of regression coefficients. The covariates are assumed to be drawn from a known population distribution $p(\mathbf{x})$. The data likelihood is

$$p(D_t|\lambda,\beta) = \prod_{i=1}^{N_t} \left(\lambda e^{\beta \cdot \mathbf{x}_i}\right)^{\Delta_i} \exp(-\lambda t_i e^{\beta \cdot \mathbf{x}_i}) p(\mathbf{x}_i).$$
(2)

Using Bayes' rule we can write the posterior as

$$p(\lambda, \boldsymbol{\beta}|D_t, \boldsymbol{\theta}) = \frac{p(D_t|\lambda, \boldsymbol{\beta})p(\lambda|\boldsymbol{\theta})p(\boldsymbol{\beta}|\boldsymbol{\theta})}{p(D_t|\boldsymbol{\theta})}$$
(3)

where $p(D_t|\boldsymbol{\theta})$ is the marginal likelihood. The vector $\boldsymbol{\theta}$ contains hyperparameters that are required for the prior distributions. For the prior over λ we choose $\lambda \sim \text{Gamma}(\kappa_0, \chi_0)$, with shape and scale hyperparameters κ_0 and χ_0 respectively, and $\boldsymbol{\beta} \sim \mathcal{N}(0, \alpha_0^2 I)$. The value of $\boldsymbol{\theta} = (\kappa_0, \chi_0, \alpha_0^2)$ is fixed and we will henceforth drop the dependence on $\boldsymbol{\theta}$ for the sake of notational compactness.

2.2. Entropy as a measure of patient informativeness

At time t we have recruited N_t patients onto the trial. Suppose that a candidate patient with covariates \mathbf{x}^* has presented and we wish to estimate how much information we expect the candidate to provide if they are to be recruited. The information gain is defined as the

Information-adaptive clinical trials 5

reduction in the expected entropy of the posterior (3). The entropy is defined as

$$h(D_t) = -\langle \log p(\lambda, \beta | D_t) \rangle_{p(\lambda, \beta | D_t)}.$$
(4)

The notation $\langle \cdots \rangle_p$ denotes the expectation with respect to the density p. We then add the candidate to the existing cohort and take the expectation with respect to the unknown t^* :

$$H(\mathbf{x}^*|D_t) = \langle h(D_t \cup \{\mathbf{x}^*, t^*\}) \rangle_{p(t^*|\mathbf{x}^*, D_t)}$$

$$\tag{5}$$

where the argument of h is the union of D_t and the additional uncensored observation $\{\mathbf{x}^*, t^*\}$ and where

$$p(t^*|\mathbf{x}^*, D_t) = \langle p(t^*|\mathbf{x}^*, \lambda, \beta) \rangle_{p(\lambda, \beta|D_t)}.$$
(6)

The time-to-event density is $p(t^*|\mathbf{x}^*, \lambda, \beta) = \lambda e^{\beta \cdot \mathbf{x}^*} \exp(-\lambda t^* e^{\beta \cdot \mathbf{x}^*})$. This can be used to define an objective function E that will be used to determine the recruitment probability for the candidate

$$E(\mathbf{x}^*|D_t) = h(D_t) - H(\mathbf{x}^*|D_t).$$
(7)

2.3. Mathematical approximations

The expectation (4) is analytically intractable. Consequently, we develop a variational approximation of the the posterior $q(\lambda, \beta) \approx p(\lambda, \beta | D_t)$ with $q(\lambda, \beta) = q(\lambda)q(\beta)$. The purpose of a variational approximation is to approximate the posterior with a form that is more amenable to analytical integration (Bishop, 2006, Chapter 10). For the variational distributions q we choose a log-Normal distribution, $\log \lambda \sim \mathcal{N}(\mu_1, \sigma_1^2)$, and a multivariate Normal distribution for the regression coefficients, $\beta \sim \mathcal{N}(\mu_0, \Sigma_0)$ with $\Sigma_0 = \text{diag}(\sigma_{01}^2, \ldots, \sigma_{0d}^2)$. To achieve a 'good' approximation we minimise the Kullback-Leibler divergence between the distributions q and p with respect to the variational parameters $(\mu_1, \sigma_1^2, \mu_0, \sigma_{01}^2, \ldots, \sigma_{0d}^2)$:

$$KL(q||p) = \left\langle \log \left[\frac{q(\lambda)q(\boldsymbol{\beta})}{p(\lambda,\boldsymbol{\beta}|D_t)} \right] \right\rangle_{q(\lambda)q(\boldsymbol{\beta})}$$
$$= \left\langle \log q(\lambda) \right\rangle_{q(\lambda)} + \left\langle \log q(\boldsymbol{\beta}) \right\rangle_{q(\boldsymbol{\beta})} - \left\langle \log p(\lambda,\boldsymbol{\beta}|D_t) \right\rangle_{q(\lambda)q(\boldsymbol{\beta})}.$$
(8)

This is convenient since the first two terms give the entropy of the variational distribution which is required in (5). Equation (8) is explicitly calculated in Appendix A.

In addition, the expectations (5, 6) are analytically intractable. We make two further approximations:

(a)
$$p(t^*|\mathbf{x}^*, \lambda, \boldsymbol{\beta}) = \delta(t^* - \hat{t})$$
 where $\hat{t} = \langle t^* \rangle_{p(t^*|\mathbf{x}^*, \lambda, \boldsymbol{\beta})} = (\lambda e^{\boldsymbol{\beta} \cdot \mathbf{x}^*})^{-1}$.
(b) $p(\lambda, \boldsymbol{\beta}|D_t) = \delta(\hat{\lambda} - \lambda)\delta(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ where $(\hat{\lambda}, \hat{\boldsymbol{\beta}}) = \operatorname{argmax}_{(\lambda, \boldsymbol{\beta})} p(\lambda, \boldsymbol{\beta}|D_t)$.

The Dirac delta function $\delta(x)$ is loosely defined by $\delta(0) = \infty$ and is zero elsewhere. These approximations allow evaluation of the integrals (5, 6) and, additionally, it is computationally faster to obtain $(\hat{\lambda}, \hat{\beta})$ rather than numerically integrating (5, 6). Combining the above approximations we can write $\hat{t} = (\hat{\lambda}e^{\hat{\beta}\cdot\mathbf{x}^*})^{-1}$ and obtain

$$\hat{H}(\mathbf{x}^*|D_t) = \hat{h}(D_t \cup \{\mathbf{x}^*, \hat{t}\}) \tag{9}$$

$$\hat{h}(D_t) = -\langle \log q(\lambda) \rangle_{q(\lambda)} - \langle \log q(\beta) \rangle_{q(\beta)}.$$
(10)

These can be substituted into (7) to obtain an approximated objective function $\hat{E}(\mathbf{x}^*|D_t)$. Evaluation of these expressions require numerical optimisation of (3) and (8) in order to evaluate, but this is computationally feasible. Note that estimates of λ and β could be unstable at the early stages of the trial when few patients have been recruited. In this case, one could implement a 'burn in' phase where selective recruitment only begins after a certain number of patients have been recruited.

2.4. Obtaining a recruitment and allocation probability

Once a candidate patient presents with covariates \mathbf{x}^* we would like to define a recruitment probability $\rho(\mathbf{x}^*|D_t)$. In general, we can write $\mathbf{x}^* = [\mathbf{y}^*, \mathbf{z}]$ where \mathbf{y}^* are clinical covariates or biomarkers and \mathbf{z} indicates the allocated treatment arm. Suppose there are K arms in total and $\mathbf{z} \in {\mathbf{z}_1, \ldots, \mathbf{z}_K}$ where \mathbf{z}_k indicates allocation to arm k. The first step is to define the allocation probability to treatment arm k as

$$p(k|\mathbf{x}^*, D_t) = \frac{E(\mathbf{y}^*, \mathbf{z}_k | D_t)}{\sum_{j=1}^{K} \hat{E}(\mathbf{y}^*, \mathbf{z}_j | D_t)} \quad \text{for } k = 1, \dots, K.$$
(11)

Information-adaptive clinical trials 7

A treatment arm is chosen at random according to this distribution and is denoted by \mathbf{z}^* . Secondly, we define the *ideal* candidate as $\mathbf{y}_I = \operatorname{argmin}_{\mathbf{y}} \hat{E}(\mathbf{y}, \mathbf{z}^* | D_t)$. The ideal candidate would give us the greatest reduction in expected entropy. A recruitment probability is given by

$$\rho(\mathbf{x}^*|D_t) = f_0\left(\frac{\hat{E}(\mathbf{y}^*, \mathbf{z}^*|D_t)}{\hat{E}(\mathbf{y}_I, \mathbf{z}^*|D_t)}\right)$$
(12)

where f_0 is some function that remains to be specified. Since the argument of f_0 must lie in the interval [0, 1] we can choose f_0 to be the identity function in which case the closer the candidate is to the ideal patient the higher the probability of recruitment. Alternatively, we can choose $f_0(s) = \theta(s - p_0)$ for a specified threshold p_0 . The step function $\theta(s) = 0$ if $s \leq 0$ and $\theta(s) = 1$ otherwise. This results in deterministic recruitment. A more general option is $f_0(s) = (1 + \tanh(s/\beta_0 - p_0))/2$ which is equivalent to deterministic recruitment when $\beta_0 \to 0$. This allows the practitioner to implement a desired level of stringency in the recruitment process.

3. The German Breast Cancer Dataset

We applied our method to data obtained from the German Breast Cancer Study (GBCS) described in Hosmer et al. (2008, Section 1.3). Our goal is to infer the parameters for a single covariate in order to assess how our adaptive protocol performs. The data consist of time-to-event outcomes for 686 patients recruited between July 1984 and December 1989. There are eight covariates in total. We decided to use tumour size (mm) for a univariate analysis because a good spread (1st quartile = 20 mm, median = 25 mm, 3rd quartile = 35 mm) would make it suitable for filtering patients according to the covariate. Importantly, the dataset also contains the date at which each patient is diagnosed with primary node positive breast cancer so we can easily calculate the waiting-time between patients. This allows us to effectively 're-run' the trial. The primary event was recurrence.

To assess the information-adaptive design we decided to recruit a total of $N_T = 100$ patients. We used deterministic recruitment with a cutoff of $p_0 = 0.5$. The trial was terminated after 10 years. We compared this to a randomised clinical trial (RCT) in which



Fig. 1. Plot of the posterior entropy (10) for the RCT and ACT as a function of time. The vertical ticks indicate times at which a patient was recruited. The sharp drop at ≈ 0.75 years corresponds to the first primary event occurring.

Table 1. Inferred parameters and entropies of the full GBCS dataset (Full), the adaptive clinical trial (ACT), and the randomised clinical trial (RCT). In brackets are 95 percent confidence intervals and p is corresponding the p-value. N_{total} is the total number of recruits, N_{reject} is the number of rejected candidates, and t_R is the recruitment time in months.

| | N_{total} | N_{reject} | t_R | λ | β | entropy |
|------|-------------|--------------|-------|-----------|--|---------|
| Full | 686 | 0 | 67 | 0.13 | 0.36 (0.19,0.52), $p = 6.1 \times 10^{-6}$ | -4.54 |
| ACT | 100 | 278 | 31 | 0.11 | 0.44 (0.21,0.66), $p = 4.2 \times 10^{-5}$ | -3.49 |
| RCT | 100 | 0 | 11 | 0.14 | 0.11 (-0.27, 0.48), $p = 0.29$ | -2.83 |

the first 100 patients are recruited. The same proportional hazards model as Section 2.1 was used to analyse the RCT. The covariate values were median-centred and rescaled by 25 mm. The population density was assumed constant. We impose a uniform prior between ± 1 for the ideal covariate x_I . Hyperparameters were set to $(\kappa_0, \chi_0, \alpha_0^2) = (3, 1, 4)$.

It took approximately 1 year to recruit 100 patients onto the RCT. The adaptive clinical trial (ACT) took approximately 2.5 years, during which a total of 278 patients were rejected. In Figure 1 the posterior entropies for both the ACT and RCT are plotted. Initially the entropies are largely determined by the priors over λ and β but quickly drop as patients are recruited, although not monotonically. In the first 2.5 years of the trial the RCT has a lower entropy which is presumably due to the fact that more patients have been recruited compared to the ACT. Towards the end of the trial the ACT has a lower entropy due to a more informative cohort. Both entropies continue to decrease after recruitment has finished as more events are observed.

Table 1 shows the inferred model parameters (evaluated after 10 years) from the original dataset, the ACT, and the RCT. The ACT results in a significant non-zero value for β that is close to the value obtained using the full dataset (with N = 686). The RCT fails to infer any significant value.

In order to gain some intuition for how the recruitment probabilities are determined we have plotted the expected entropy as a function of the covariate x at various time points in Figure 2. We note that the function tends to have one maximum and two minima at $x = \pm 1$. This general shape is due to the nature of the proportional hazards model since extreme values of x will diminish the space of plausible parameter values more so than values close to zero, and consequently are more informative. The dashed line is the entropy below which a candidate will be recruited. In (a) the trial has started at t = 0 with two patients. There is a strong preference for individuals towards ± 1 . The next candidate (at t = 34 days) had $x^* = -0.52$ and so was recruited. In (b), some patients with covariate values > 1 have been recruited and this encourages recruitment of negative covariate values. At t = 267days no primary events have occurred. In (c), after t = 268 days the first primary event occurs for a patient with a positive covariate value. This additional piece of information further increases the benefit of recruiting negative covariate values over positive ones. Note that the vertical scale changes. This illustrates that the recruitment probability changes dynamically, and depends on the observed events and covariate values of the existing cohort. We conclude that in general we gain more information from covariate values that have been under-sampled or values where few primary events have occurred.

Individuals with covariates values far from zero will have the greatest reduction in expected entropy. This is because these terms will dominate the data likelihood in a proportional hazards model. Consequently, the covariate distribution in the ACT can differ considerably from the population distribution. Figure 3 shows the empirical covariate



Fig. 2. The expected entropy (9) as a function of x at various times during the ACT.

distributions for the original dataset and both trials. Due to the shape of the expected entropy function (see Figure 2) patients towards ± 1 were more likely to be recruited in the ACT. Consequently, almost no patients with $x \approx 0$ were recruited. The RCT density resembles the density of the full dataset.

4. Numerical simulation studies

Here we consider a scenario where the covariates consist of a two-dimensional biomarker $\mathbf{y}_i = (y_{i1}, y_{i2})$ and patients are given one of three treatments denoted by $\mathbf{z}_i = (z_{i1}, z_{i2}, z_{i3})$. A patient given treatment one would have $\mathbf{z}_i = (1, 0, 0)$, treatment two would have $\mathbf{z}_i = (0, 1, 0)$, and so forth. We are interested in whether there is any interaction between the biomarker and treatments, i.e. is the biomarker predictive. A proportional hazards model with interaction terms is assumed:

$$h(t|\mathbf{y}_{i}, \mathbf{z}_{i}, \lambda, \beta) = \lambda e^{\beta_{1}y_{i1}z_{i1} + \beta_{2}y_{i1}z_{i2} + \beta_{3}y_{i1}z_{i3} + \beta_{4}y_{i2}z_{i1} + \beta_{5}y_{i2}z_{i2} + \beta_{6}y_{i2}z_{i3}}.$$
(13)

This gives a total of six regression coefficients and the baseline hazard λ to be inferred. In all simulations we compared an adaptive trial to a randomised one.

To simulate survival data we generate a random vector $\mathbf{y} = (y_1, y_2)$ where $y_i \sim$ uniform(-1, +1) or $y_i \sim \mathcal{N}(0, 0.5)$ for i = 1, 2. A treatment arm \mathbf{z} is chosen (either randomly or according to (11)). A random number $w \sim$ uniform(0, 1) is generated, and an event time is given by the inverse of the cumulative distribution $t = -e^{-\beta \cdot \mathbf{x}} \log(1-w)/\lambda$ where $\mathbf{x} \in \mathbb{R}^6$ contains the same product terms between \mathbf{y} and \mathbf{z} as (13). Patients are



Fig. 3. Kernel smoothed empirical covariate densities (Gaussian kernel, bandwidth = 0.2) for (a) the full GBCS dataset, (b) the ACT, and (c) the RCT.

Table 2. Mean square error between inferred and 'true' model parameters over 500 simulations. Comparison between both random and adaptive trials without selective recruitment and uniform and Gaussian distributed covariates.

| | β_1 | β_2 | β_3 | β_4 | β_5 | β_6 | λ |
|----------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| ACT (Uniform) | 0.348 | 0.374 | 0.361 | 0.384 | 0.418 | 0.352 | 0.00080 |
| RCT (Uniform) | 0.364 | 0.347 | 0.401 | 0.389 | 0.396 | 0.384 | 0.00084 |
| ACT (Gaussian) | 0.499 | 0.5120 | 0.487 | 0.438 | 0.445 | 0.430 | 0.00085 |
| RCT (Gaussian) | 0.470 | 0.494 | 0.504 | 0.471 | 0.518 | 0.435 | 0.00084 |

censored at random with probability $p_c \in [0, 1]$. If an individual is censored then the time-to-censoring is drawn from a uniform density between 0 and t. The first patient to be generated is recruited onto both the ACT and RCT. The waiting time until the next patient is drawn from an exponential density with rate parameter ξ . Hyperparameters were set to $(\kappa_0, \chi_0, \alpha_0^2) = (3, 1, 4)$.

4.1. Adaptive allocation without selective recruitment

In these simulations all patients were recruited. A total of N = 50 patients were recruited onto both trials. The trial was terminated after t = 100 arbitrary units of time. The rate parameter for waiting times was $\xi = 6$, and $p_c = 0.5$. Model parameters were set to $\beta = (0.8, -0.5, 1.1, -0.7, 0.6, 0.1)$ and $\lambda = 0.1$. In the ACT the expected entropy was used to determine which treatment arm each individual was allocated to as described in Section

2.4. In the RCT patients were allocated to one of the three arms at random.

A total of 500 simulations were run. We computed the mean square error between the inferred model parameters and the 'true' values used to generate the data. As shown in Table 2 we found essentially no difference between the randomised and adaptive trials for either uniformly or Gaussian distributed covariates. We found that the entropy at the end of the ACTs with uniform covariates was on average slightly lower than the RCTs (2.14 and 2.20 respectively), although the difference was statistically significant (p-value 0.017 with a one-sided paired t-test). For Gaussian distributed covariates the difference in entropies was insignificant. We also performed a chi-squared test to see if the allocation proportions of patients across arms differed from a uniform distribution. Each simulated trial was tested and we found no p-values less than 0.05 for either uniform or Gaussian distributed covariates. Since the chi-squared test was repeated for each trial the p-values were corrected for multiple hypothesis testing by controlling the false discovery rate (using the method of Benjamini and Hochberg (1995)) with the 'p.adjust' R function.

4.2. Adaptive allocation and recruitment

In these simulations the same parameters as above were used but patients were recruited onto the ACT selectively with a threshold of $p_0 = 0.66$. Over 500 simulations we found that the mean square error between the inferred and 'true' parameters was considerably lower in the ACTs than the RCTs as shown in Table 3. For uniformly distributed covariates 48.9% of the inferred parameter values were significant (at 0.05) in the ACT compared to 39.2% in the RCTs. Furthermore, the mean entropy at the end of the ACTs was 0.93, compared to 2.23 in the RCTs. On average 140.7 (standard deviation 42.9) individuals are rejected.

In the case of Gaussian distributed covariates the difference is more pronounced. 50.4% of parameters were significant in the ACT compared to 35.0% in the RCT. An average of 240.0 patients were rejected (standard deviation 61.9). Due to the Gaussian distribution there are more patients in the less informative region around zero. Therefore the number

| | β_1 | β_2 | β_3 | β_4 | β_5 | β_6 | λ |
|----------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| ACT (uniform) | 0.324 | 0.279 | 0.313 | 0.342 | 0.279 | 0.306 | 0.00079 |
| RCT (uniform) | 0.401 | 0.335 | 0.408 | 0.375 | 0.367 | 0.361 | 0.00081 |
| ACT (Gaussian) | 0.266 | 0.289 | 0.278 | 0.217 | 0.253 | 0.262 | 0.00085 |
| RCT (Gaussian) | 0.444 | 0.553 | 0.509 | 0.521 | 0.502 | 0.478 | 0.00082 |

Table 3. Mean square error between inferred and 'true' model parameters over 500 simulations.Comparison between random and adaptive trials with selective recruitment.

of rejections is higher and the benefit more substantial.

We also explored the effect of the threshold p_0 on the trial results. When $p_0 = 0.33$ we found that the MSE (averaged over the six beta values) was 0.287 in the ACT compared to 0.372 in the RCT with 44.0% of inferred parameters reaching statistical significance in the ACT compared to 39.6% in the RCT. An average of 22.0 patients were rejected (standard deviation 6.45). When the threshold was increased to $p_0 = 0.90$ the MSE was 0.358 versus 0.363, and the proportion of significant parameters was 41.7% versus 39.3%, in the RCT and ACT respectively. On average 237.3 (standard deviation 86.5) patients were rejected. This suggests that setting the threshold too high can be counterproductive.

5. Discussion

The practicality of our proposed design will depend on various economic and ethical considerations as well as the characteristics of each particular trial and the study population. For instance, if a covariate is relatively inexpensive to measure when compared to the costs of recruitment (treatment provision, follow-up, administration) then it may be sensible to selectively recruit informative patients. A large pool of patients can be inexpensively screened and then resources concentrated on those which are likely to provide the most information. In this case a selective recruitment design could result in significant cost reductions since fewer recruits are required overall.

Clinical trials are not primarily intended to be therapeutic, but rather as a means to generate medical evidence. Recruited patients may be exposed to treatments that are

ineffective (e.g. a placebo) or that are possibly even harmful. Our proposed design offers the possibility to conduct a trial using fewer patients than a traditional randomised design. This may be ethically attractive in some cases since ultimately fewer patients are offered treatment options with uncertain efficaciousness.

In a selective recruitment design the decision to recruit and allocate a patient can also take into account the probability of a successful response to treatment (although this was outside the scope of this paper). Patients can be recruited and allocated in a manner that balances the statistical informativeness of a decision against the potential benefit or harm to that individual. The decision making process must balance individual and collective benefits. Maximising statistical information offers a collective benefit to all patients outside the trial (both current and future) who could benefit from the trial findings. Naturally this must be offset by what is best for the trial participants. What our proposed design offers the practitioner is a framework to balance individual versus collective ethical considerations.

Selective recruitment designs suffer from a number of drawbacks, one of which is longer recruitment times. If the patient accrual rate is low it may render the overall recruitment period unfeasible. Selective recruitment designs are therefore only appropriate in situations where patients accrue relatively quickly or where longer recruitment periods are an acceptable compromise.

One of the consequences of a proportional hazards model is that the most informative patients tend to have extreme values of covariates. As a result the distribution of recruited patients may differ from the population distribution which might make it difficult to generalise results from the trial to the general population. Thus, some generalisability is sacrificed in return for greater statistical power. If this was deemed undesirable one could introduce a sufficient level of random sampling in addition to preferential accrual of informative patients. Each candidate patient has a minimum probability of recruitment with informative patients having a higher probability. Thus, selective recruitment need not be an all or nothing process; it can be used to enrich the trial with informative patients to a desired degree.

Finally, in the case of model misspecification undesirable biases may be introduced into

the dataset because the model choice influences the covariate distribution considerably. An additional limitation is that it is not yet clear how to estimate the sample size required for a certain level of statistical power — a calculation that is typically used when planning new trials.

In summary, our novel information-adaptive selective recruitment clinical trial design will reject non-informative patients. Individuals who are more likely to clarify the values of our model parameters are more likely to be recruited. We have demonstrated with both experimental and simulated data the feasibility of our approach. Statistically significant inferences can be achieved using fewer patients with a selective recruitment design than a randomised trial, although we found that treatment arm allocation using an entropy based measure (without selective recruitment) did not offer any improvement over a randomised design. Such a design may offer a more economical or ethically attractive route to discover the relationship between biomarkers, treatments, and survival outcomes.

It will be interesting to extend this work beyond the proportional hazards assumption to more complex survival models. Incorporation of response-adaptive protocols offer another promising extension. Throughout this work we have assumed a uniform population density. In the case of a non-uniform density it may be desirable to incorporate this into the definition of an ideal candidate such that an ideal candidate is both informative and likely to be observed. This will require further investigation. Further extensions of the model could include alternative outcomes such as binary or continuous measurements.

A. Derivation of the Kullback-Leibler divergence

The first two terms of the Kullback-Leibler divergence (8) in Section 2.3 are simply minus the entropies of the variational distributions. These are $\langle \log q(\lambda) \rangle_{q(\lambda)} = -(1/2 + \log(2\pi\sigma_1^2)/2 + \mu_1)$ and $\langle \log q(\beta) \rangle_{q(\beta)} = -\sum_{\nu=1}^d \log(2\pi e \sigma_{0\nu}^2)/2$. The third term from (8) is

$$-N_{t}^{1} \langle \log \lambda \rangle_{q(\lambda)} - \mathbf{\Phi}_{t} \cdot \langle \boldsymbol{\beta} \rangle_{q(\boldsymbol{\beta})} + \langle \lambda \rangle_{q(\lambda)} \sum_{i=1}^{N_{t}} t_{i} \left\langle e^{\boldsymbol{\beta} \cdot \mathbf{x}_{i}} \right\rangle_{q(\boldsymbol{\beta})} - \left\langle \log p(\lambda | \kappa_{0}, \chi_{0}) \right\rangle_{q(\lambda)} - \left\langle \log p(\boldsymbol{\beta} | \alpha_{0}^{2}) \right\rangle_{q(\boldsymbol{\beta})}$$
(14)

where N_t^1 is the number of non-censored events up until time t and $\Phi_t = \sum_{i:\Delta_i=1} \mathbf{x}_i$. It is straightforward to show $\langle \log \lambda \rangle_{q(\lambda)} = \mu_1$, $\langle \lambda \rangle_{q(\lambda)} = e^{\mu_1 + \sigma_1^2/2}$ and $\langle \beta \rangle_{q(\beta)} = \mu_0$. The following result is needed (Coolen et al., 2005, Appendix D):

$$\int \mathrm{d}\mathbf{z} \, \frac{e^{-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})\cdot A^{-1}(\mathbf{z}-\boldsymbol{\mu})+\mathbf{b}\cdot\mathbf{z}}}{(2\pi)^{d/2}|A|^{1/2}} = e^{\boldsymbol{\mu}\cdot\mathbf{b}+\frac{1}{2}\mathbf{b}\cdot A\mathbf{b}} \tag{15}$$

from which it follows $\langle e^{\boldsymbol{\beta}\cdot\mathbf{x}_i}\rangle_{q(\boldsymbol{\beta})} = e^{\boldsymbol{\mu}_0\cdot\mathbf{x}_i + \frac{1}{2}\mathbf{x}_i\cdot\boldsymbol{\Sigma}_0\mathbf{x}_i}$. Note that (15) also defines the moment generating function for a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix A. The terms relating to the priors are $\langle \log p(\boldsymbol{\beta}|\alpha_0^2)\rangle_{q(\boldsymbol{\beta})} = -\sum_{\nu} (\sigma_{0\nu}^2 + [\boldsymbol{\mu}_0]_{\nu}^2)/2\alpha_0^2$ and $\langle \log p(\lambda|\kappa_0, \chi_0)\rangle_{q(\lambda)} = (\kappa_0 - 1) \langle \log \lambda \rangle_{q(\lambda)} - \chi_0^{-1} \langle \lambda \rangle_{q(\lambda)}$ where $[\boldsymbol{\mu}_0]_{\nu}$ denotes the ν th component of $\boldsymbol{\mu}_0$.

References

- Benjamini, Y. and Hochberg, Y. (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B* (Methodological), 289–300.
- Bishop, C. M. (2006) Pattern recognition and machine learning. Information science and statistics. Springer, 15.
- Coolen, A., Kühn, R. and Sollich, P. (2005) Theory of Neural Information Processing Systems. Oxford University Press.
- Freidlin, B. and Simon, R. (2005) Adaptive signature design: an adaptive clinical trial design for generating and prospectively testing a gene expression signature for sensitive patients. *Clinical Cancer Research*, **11**, 7872–7878.
- Hosmer, D. W., Lemeshow, S. and May, S. (2008) Applied Survival Analysis: Regression Modeling of Time-to-Event Data. John Wiley and Sons, second edn.
- Pocock, S. J. and Simon, R. (1975) Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics*, **31**, 103–115.

- Rosner, G. L., Stadler, W. and Ratain, M. J. (2002) Randomized discontinuation design: application to cytostatic antineoplastic agents. *Journal of Clinical Oncology*, 20, 4478– 4484.
- Sargent, D. J., Conley, B. A., Allegra, C. and Collette, L. (2005) Clinical trial designs for predictive marker validation in cancer treatment trials. *Journal of Clinical Oncology*, 23, 2020–2027.
- Taves, D. R. (1974) Minimization: a new method of assigning patients to treatment and control groups. *Clinical pharmacology and therapeutics*, 15, 443–453.
- Temple, R. (2010) Enrichment of clinical study populations. Clinical Pharmacology & Therapeutics, 88, 774–778.
- Yin, G. (2012) Clinical trial design: Bayesian and frequentist adaptive methods. John Wiley & Sons.
- Zhang, L. and Rosenberger, W. F. (2007) Response-adaptive randomization for survival trials: the parametric approach. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 56, 153–165.