Harvard University
Harvard University Biostatistics Working Paper Series

Year 2014 *Paper* 182

Optimal Bayesian Adaptive Trials when Treatment Efficacy Depends on Biomarkers

Yifan Zhang* Lorenzo Trippa†
Giovanni Parmigiani‡

*Harvard University, yiz639@mail.harvard.edu
†Harvard University and Dana-Farber Cancer Institute, ltrippa@jimmy.harvard.edu
‡Harvard University and Dana-Farber Cancer Institute, gp@jimmy.harvard.edu

http://biostats.bepress.com/harvardbiostat/paper182

Optimal Bayesian Adaptive Trials when Treatment Efficacy Depends on Biomarkers

Yifan Zhang[1], Lorenzo Trippa[1,2], and Giovanni Parmigiani[1,2]

[1]Department of Biostatistics, Harvard University, Boston, Massachusetts, U.S.A.

[2]Department of Biostatistics and Computational Biology, Dana-Farber Cancer Institute, Boston, Massachusetts, U.S.A.

Abstract

Clinical biomarkers play an important role in personalized medicine in cancer clinical trials. A response adaptive trial design enables researchers to use treatment results observed from early patients to aid in treatment decisions of later patients. In this article, we describe the mathematical steps for computing the theoretical optimal biomarker-integrated adaptive trial designs. The optimal design maximizes the expected trial utility given any pre-specified utility function, though we focus here on maximizing responses within a given patient horizon. We describe the performance of the optimal design under different scenarios. Bayesian Adaptive Randomization (BAR) is emerging as a practical approach to develop adaptive trials. We compare the Bayesian Adaptive Randomization(BAR) and the theoretical optimum to quantify the loss of utility in different scenarios. We conclude that BAR is nearly optimal in a broad range of scenarios. We also compare the BAR to a frequentist play-the-winner rule where biomarkers are integrated. Our work provides absolute benchmark for the evaluation of trial designs in personalized medicine.

1 Introduction

Recent insight into the genetic drivers of cancer [Wood et al., 2007], the development of drugs whose action depends specifically on the activity of these targets, and the resulting heterogeneity of treatment responses have highlighted the need to systematically include information on a tumor's genetic characteristics into treatment decisions. This trend is generally referred to as "precision" or "personalized" cancer treatment [Savard, 2013, Schilsky, 2010]. A tremendous amount of resources is presently allocated to developing personalized cancer medicine [Chin et al., 2011, Compton, 2007, Hamburg and Collins, 2010, La Thangue

and Kerr, 2011, van't Veer and Bernards, 2008, Walther et al., 2009]. Since the Food and Drug Administration (FDA) approved trastuzumab [Baselga et al., 1998] in human epidermal growth factor receptor 2 (HER2)-positive breast cancer patients, there have been development of personalized medicine for the treatment of chronic myelogenous leukemia[Druker et al., 2001], colon cancer [Allegra et al., 2009], lung cancer [Paez et al., 2004], and others [Gonzalez-Angulo et al., 2010, McDermott and Settleman, 2009].

The development of targeted treatments is inspiring a change in the design of clinical trials. For instance, in the Biomarker-integrated Approaches of Targeted Therapy for Lung cancer Elimination (BATTLE) study [Herbst et al., 2010, Kim et al., 2011, Rubin et al., 2011], non-small cell lung cancer (NSCLC) patients had been classified into five subgroups defined by biomarkers reflecting tumor genetics. A more complex biomarker-adaptive design, also allowing treatments to enter and exit the trial is the I-SPY 2 [Barker et al., 2009]. The goal of these trial is to identify the most beneficial treatment separately for each patient subgroups.

Statistically, adaptive trial designs for cancer clinical trials have been studied extensively, including both frequentist [Eisele, 1994, Ivanova, 2003, Wei and Durham, 1978, Zelen, 1969] and Bayesian designs [Atkinson and Biswas, 2005, Berry and Eick, 1995, Rosner and Bekele, 2010, Thall and Wathen, 2007], although biomarkers are not often included. From a decision theoretic perspective, we lack results on optimal trial designs integrating biomarkers. In this paper, we fill this gap by deriving an optimal Bayesian biomarker-integrated adaptive trial design. Our optimality goal is to maximize the number of successfully treated patients over a given horizon, including the patients in the trial itself. Our results generalize earlier work by [Berry and Eick, 1995]. Our optimal design assigns treatment to each patient based on their own biomarker, and on accumulating results in the trial. At the end of the trial, patients are assigned to their optimal treatment, potentially depending on the biomarker, for the remainder of the horizon of interest.

Having developed the theory behind the optimal design, it is now possible to evaluate

the efficiency of more heuristic approaches by comparing their utility to the optimum. While the optimal design maximizes the expected trial utility, the treatment assignment at each stage is deterministic and it imposes heavy computational burden [Powell, 2007]. Thus, a natural comparison is with the Bayesian Adaptive Randomization (BAR) design [Thall and Wathen, 2007], and with the frequentist play-the-winner rule. The comparison will show that using BAR does not generally lead to a significant loss in terms of the expected trial utility when compared to the theoretical optimum. BAR often shows higher expected trial utility than play-the-winner rule when the patient population is stratified by biomarkers. This comparison is the primary practical goal of this article.

In Section 2 we define the problem and provide a description of the optimal trial design. We assume a binary response, with no delay, and a general multivariate binary biomarker space. In section 3, we interpret how the biomarker affects the optimal expected trial utility, using synthetic data. In section 4 we address computational and design issues.

2 Biomarker Integrated Trials

2.1 Assumptions and Notation

We assume that n patients are recruited in a clinical trial to study the effects of I treatments, labeled by $i \in \{1, ..., I\}$. There are K biomarkers integrated in the trial, with binary biomarker indicators b_k's, $k = 1, ..., K$. For example if the biomarkers are somatic mutations in the tumor, $b_k = 1$ denotes the presence of a mutation in gene k, while if $b_k = 0$ the gene is unaltered. Patients' profiles can be defined by these biomarkers indicators and summarized by the vector $\mathbf{G} = (b_1, ...b_K)$. There are many options for dividing patients into subgroups according to their profiles. For example, 4 biomarkers are used in the BATTLE trial to define biomarker groups [Zhou et al., 2008]. Using our notation, $K = 4$. Also b_1: EGFR mutation or amplification, b_2: K-ras and / or B-raf mutation, b_3: VEGF and / or VEGFR expression, b_4: RXR and / or cyclin D1 expression. In the BATTLE trial, if a patient has any profile

with $b_1 = 1$, he/she belongs to the EGFR mutation/amplification group. Any patient with $b_1 = 0, b_2 = 1$ is binned in a second profile with EGFR negative and a K-ras/B-raf positive tumors, and so forth (see [Zhou et al., 2008] for the complete algorithm). A patient with $\mathbf{G} = (0, 0, 0, 0)$ is assigned to the "no mutation" group. However, this is not the only way of defining biomarker groups. For example, if one is interested in investigating the effects of drugs on patients with either an EGFR mutation or K-ras/B-raf mutation, then patients with profile $\mathbf{G} = (0, 1, 0, 0)$ could be assigned to the same group as patients with profile $\mathbf{G} = (1, 0, 0, 0)$. The optimal trial design proposed in this paper can be applied in both situations described above, as long as biomarker groups are mutually exclusive. We denote by j be the biomarker groups in the trial, with $j \in \{1, ..., J\}$.

Similarly to [Berry and Eick, 1995], we consider a patient horizon, including patients in the trial, as well as subsequent patients who will be treated based directly on the result of trial. The size of the patient horizon is an important element in adaptive designs [Upton and Lee, 1976], because it controls how the design balances the two potentially competing needs of learning the treatment effects and treating patients ethically within the trial.

We assume that the outcome of interest is binary and that it is known immediately after treatment for each patient. Let w_{ij} be the probability of observing a success, or response rate, for treatment i in subgroup j, $\forall i, j$. The goal of a clinical trial is to study the effects of different treatments in various biomarker subgroups, to choose the most effective treatment for each of these subgroup at the end of the trial, and to apply it for the remainder of the patient horizon. The values of w_{ij} are the parameters of interest.

We take a Bayesian approach to quantifying uncertainty in these parameters. We use a mixture prior, assigning a positive probability to the event that the effect of treatment i is the same across all biomarker groups, that is that $w_{i1} = ... = w_{iJ}$, as well as a positive probability to the event that there are some biomarker effects. Conditional on the latter, we define the prior distributions in the following way. The marginal priors of w_{ij}'s are Beta(1,1). The effects of two different treatments are independently distributed, whether

within or across subgroups. Namely, w_{ij} is independent of $w_{i'j'}$, for $i \neq i'; \forall j, j'$. When the effect of a treatment is allowed to vary across subgroups, it can take a different value in each subgroup. We do not consider the possibility that the effects may be equal in some subgroups and different in others because of the combinatorial complexity, though in practice this could be biologically justified. Consider now the effect of treatment i across biomarker groups, we allow for the possibility that treatment efficacy depends on biomarkers and use the following model as the joint prior of the treatment effects. Treatment i, with prior probability π_i, has the same effect across biomarker groups. If the effects of treatment i are not equal, they are independent and can be modeled with a multivariate uniform distribution. To summarize, the joint prior distribution of the unknown response rates for treatment i is

$$Pr(W_{i1} \leq w_{i1}, ..., W_{iJ} \leq w_{iJ}) = (1 - \pi_i) \prod_{j=1}^{J} w_{ij} + \pi_i \min_{j=1}^{J} \{w_{ij}\} \tag{1}$$

with the second term capturing the case when effects for treatment i are the same across all biomarker groups. The hyperparameters $\pi_i, i = 1, ..., I$ are fixed and represent the prior probability that the effect of treatment i depends on the biomarkers. The π_i's can be estimated from early phase trials or elicited from expert judgment. [Hammitt and Zhang, 2013].

In this paper, we study this specific prior in detail. However, our approach is more general, in that the algorithm we describe can easily accommodate other prior distributions. If reliable evidence on treatment effects is available from early phase trials, it can be summarized via the prior distributions of the $w'_{ij}s$, relaxing on assumption of uniformity.

Let X_n denote the treatment outcome of the n^{th} patient enrolled in the trial, and $\mathbf{G_n}$ the associated marker profile. The arrival of patient n also marks the n^{th} decision point in the optimal allocation: an adaptive trial design d assigns patient n to one of the I treatments based on results observed until decision point n. The notation $d(n|X_{n-1}, ..., X_1, \mathbf{G_n}, ..., \mathbf{G_1}) = i$ indicates that the design assigns patient n to treatment i after observing the treatment

outcomes of the first $n-1$ patients. The treatment assignment depends on the treatment outcomes of the previous $n-1$ patients summarized by $X_{n-1}, ...X_1, \mathbf{G_n}, ..., \mathbf{G_1}$, and on the marker profile $\mathbf{G_n}$ of the n^{th} patient. The best treatment for each biomarker subgroup is determined at the end of the trial and will be prescribed to future patients according to their marker profiles. After the trial, all patients in the same biomarker group will receive the same treatment.

We next define some notation that will be used in the calculation of the theoretical optimum. After the n^{th} patient has been assigned to a treatment and the outcome has been observed, we use matrices $\mathbf{M^{(n)}}$ and $\mathbf{S^{(n)}}$ to record the accumulated information up to that point. The entry $m_{ij}^{(n)}$, on the i^{th} row and j^{th} column of matrix $\mathbf{M^{(n)}}$, is the total number of patients in subgroup j assigned to treatment i up to that time point. The entry $s_{ij}^{(n)}$, on the i^{th} row and j^{th} column of the matrix $\mathbf{S^{(n)}}$, is the number among the $m_{ij}^{(n)}$ patients who responded to the treatment. Both matrices are of size $I \times J$. We assume that X_n's are conditionally independent given $\mathbf{G_n}$'s and w_{ij}'s. Thus, $\mathbf{M^{(n)}}$ and $\mathbf{S^{(n)}}$ together serve as the sufficient statistics for the treatment effects for the data up to decision point n.

We are interested in choosing a design d from the set D of all possible biomarker-dependent patient allocations. A trial of size N is conducted following design d to study the unknown treatment effects $\mathbf{w} = (w_{ij})_{I \times J}$, and $\mathbf{w} \in \Omega$. When the allocation is d, $\mathbf{x} \in \mathcal{X}$ is observed with conditional density $f_d(\mathbf{x}|\mathbf{w})$, where $\mathbf{x} = (x_1, ...x_n)'$. At the end of the trial, the best treatment in each subgroup is selected and will be prescribed to future patients. We model this by considering a group of $N_h - N$ additional patients whose treatment decision will be collectively affected by the information accrued during the trial. This final treatment assignment is the last stage of the decision process and is called the "terminal decision". The quantity N_h is often termed "patient horizon" [Colton, 1963, Canner, 1970]. While d denotes the allocation of patients to treatments during the trial, d^H denotes the joint allocation of the remaining $N_h - N$ patients to treatment following the trial. In our work d^H is allowed to depend on the individual patient's marker profile.

The best treatment is defined with respect to a utility function $U(d, d^H, \mathbf{x})$ designed to capture the most important goals of conducting the trial. Comparing treatment efficacy is helpful insofar as it allows to improve treatment outcomes. Here we explicitly consider the outcomes of all N_h patients, including the N in the trial and the subsequent $N_h - N$ in clinical practice. A utility function is defined as the total number of favorable treatment outcomes:

$$U(d, d^H, \mathbf{x}) == \underbrace{\sum_{n=1}^{N} x_n}_{\text{inside the trial}} + \underbrace{\sum_{n=N+1}^{N_h} x_n}_{\text{outside the trial}} \tag{2}$$

This specific form of utility function has been used by several authors [Armitage, 1985, Berry and Eick, 1995] in the design of trials not incorporating biomarkers. We will use this utility function to illustrate the behavior of the optimal design under different biomarker effects, and to make comparisons among trial designs.

2.2 Optimal Solution

The optimal adaptive choice of d and d^H is obtained by dynamic programming [Bellman, 2003, Parmigiani and Inoue, 2009]. We begin with the optimal d^H, which is computed conditional on information $X_1, ..., X_N$, $\mathbf{G_1}, ..., \mathbf{G_N}$. The expected value of the utility in (2), given this information, is maximized by assigning each patient to the treatment with the highest expected success rate in his/her marker group. We denote this optimum terminal decision by $\tilde{d}^H(\mathbf{x})$. Now this optimal $\tilde{d}$ can be computed by maximizing expected utility

$$\mathcal{U}(d) = \int_{\mathbf{X},\mathbf{w}} U(d, \tilde{d}^H(\mathbf{x}), \mathbf{x}) dF_d(\mathbf{w}, \mathbf{x}),$$

with respect to d.

The optimal design achieves two goals: (1) identifying the best treatment for each biomarker group at the end of a trial, (2) maximizing the expected number of favorable

treatment outcomes. Let $R^{(n)}_{d,\tilde{d}^H}[\mathbf{M^{(n)}}, \mathbf{S^{(n)}}]$ denote the expected number of favorable treatment outcomes from the $(n+1)^{th}$ to the N_h^{th} patient conditional on information up to stage n. Also, let $\mu^{(n)}_{ij}$ be the posterior mean of w_{ij} calculated conditional on the same information, that is $\mu^{(n)}_{ij} = E[w_{ij}|X_n, ..., X_1; \mathbf{G_n}, ..., \mathbf{G_1}]$.

This notation can be used to write $\widetilde{d}^H$ explicitly. For example, if $\mu^{(N)}_{\tilde{i}j} = \max_i(\mu^{(N)}_{ij})$, future patients in subgroup j will be assigned to treatment $\tilde{i}$. At the end of a trial following the optimal design, the expected number of future favorable treatment outcomes is

$$R^{(N)}_{d,\tilde{d}^H}[\mathbf{M^{(N)}}, \mathbf{S^{(N)}}] = \sum_{n=N+1}^{N_h} \max_{i=1}^{I}(\mu^{(N)}_{i\mathbf{G_n}}).$$

where the marker profile of the n^{th} patient, $\mathbf{G_n}$, determines the biomarker subgroup.

The general steps for calculating the expected trial utility of the optimal design starts with $R^{(N)}_{d,\tilde{d}^H}[\mathbf{M^{(N)}}, \mathbf{S^{(N)}}]$ and $\mu^{(N)}_{ij}$. In a backward manner, each $R^{(n)}_{d,\tilde{d}^H}[\mathbf{M^{(n)}}, \mathbf{S^{(n)}}]$ is calculated with $R^{(n+1)}_{d,\tilde{d}^H}[\mathbf{M^{(n+1)}}, \mathbf{S^{(n+1)}}]$ and $\mu^{(n)}_{ij}$ for $n = N-1, ...0$. Specifically, the theoretical optimum can be calculated with the following steps:

1. Initialize $R^{(N)}_{d,\tilde{d}^H}[\mathbf{M^{(N)}}, \mathbf{S^{(N)}}] = \sum_{n=N+1}^{N_h} \max_{i=1}^{I}(\mu^{(N)}_{i\mathbf{G_n}})$
2. Update $R^{(n)}_{d,\tilde{d}^H}[\mathbf{M^{(n)}}, \mathbf{S^{(n)}}]$ with $R^{(n+1)}_{d,\tilde{d}^H}[\mathbf{M^{(n+1)}}, \mathbf{S^{(n+1)}}]$ and $\mu^{(n)}_{ij}$
3. Repeat the previous step for $n = N-1, ..., 0$
4. $\mathcal{U}(\tilde{d}) = R^{(0)}_{d,\tilde{d}^H}[\mathbf{M^{(0)}}, \mathbf{S^{(0)}}]$

The updating step (Step 2) demonstrates the decision-making process of treatment selection for patients enrolled in the trial. When it is time to assign a treatment to the n^{th} patient who is in subgroup j, the probability of observing a favorable treatment outcome is $\mu^{(n-1)}_{ij}$ if this patient is assigned to treatment i. Let $\{R^{(n-1)}_{d(n)=i}[\mathbf{M^{(n-1)}}, \mathbf{S^{(n-1)}}]\}$ denote the expected number of future favorable treatment outcomes if the n^{th} patient is treated with treatment i, then $\{R^{(n-1)}_{d(n)=i}[\mathbf{M^{(n-1)}}, \mathbf{S^{(n-1)}}]\}$ can be calculated with $\mu^{(n-1)}_{ij}$, $R^{(n)}_{d}[\mathbf{M^{(n)}_{d(n)=i}}, \mathbf{S^{(n)}_{d(n)=i}}]$ and

$R_d^{(n)}[\mathbf{M}^{(\mathbf{n})}_{\mathbf{d(n)=i}}, \mathbf{S}^{(\mathbf{n-1})}]$.

$$\{R^{(n-1)}_{d(n)=i}\ [\mathbf{M}^{(\mathbf{n-1})}, \mathbf{S}^{(\mathbf{n-1})}]\}$$
$$= \mu_{ij}^{(n-1)} \times R_d^{(n)}[\mathbf{M}^{(\mathbf{n})}_{\mathbf{d(n)=i}}, \mathbf{S}^{(\mathbf{n})}_{\mathbf{d(n)=i}}] + (1-\mu_{ij}^{(n-1)}) \times R_d^{(n)}[\mathbf{M}^{(\mathbf{n})}_{\mathbf{d(n)=i}}, \mathbf{S}^{(\mathbf{n-1})}], \text{where,}$$

$$\mathbf{M}^{(\mathbf{n})}_{\mathbf{d(n)=i}} = \begin{bmatrix} m_{11}^{(n-1)} & \cdots & \cdots & \cdots & m_{1J}^{(n-1)} \\ \vdots & \ddots & & \iddots & \vdots \\ \vdots & & 1+m_{ij}^{(n-1)} & & \vdots \\ \vdots & \iddots & & \ddots & \vdots \\ m_{I1}^{(n-1)} & \cdots & \cdots & \cdots & m_{IJ}^{(n-1)} \end{bmatrix}, \mathbf{S}^{(\mathbf{n})}_{\mathbf{d(n)=i}} = \begin{bmatrix} s_{11}^{(n-1)} & \cdots & \cdots & \cdots & s_{1J}^{(n-1)} \\ \vdots & \ddots & & \iddots & \vdots \\ \vdots & & 1+s_{ij}^{(n-1)} & & \vdots \\ \vdots & \iddots & & \ddots & \vdots \\ s_{I1}^{(n-1)} & \cdots & \cdots & \cdots & s_{IJ}^{(n-1)} \end{bmatrix}$$

$$R^{(n-1)}_{d(n)=\tilde{i}}\ [\mathbf{M}^{(\mathbf{n-1})}, \mathbf{S}^{(\mathbf{n-1})}] = \max_i^I(\{R^{(n-1)}_{d(n)=i}\ [\mathbf{M}^{(\mathbf{n-1})}, \mathbf{S}^{(\mathbf{n-1})}]\})$$
$$\text{and } R^{(n-1)}_{\tilde{d}}\ [\mathbf{M}^{(\mathbf{n-1})}, \mathbf{S}^{(\mathbf{n-1})}] = R^{(n-1)}_{d(n)=\tilde{i}}\ [\mathbf{M}^{(\mathbf{n-1})}, \mathbf{S}^{(\mathbf{n-1})}]$$

3 Numerical Results and Comparison

3.1 Properties of the Theoretical Optimum in Two-Treatments Comparisons

Assume that a new cancer treatment (treatment 1) is to be compared with a conventional treatment (treatment 2). Also assume that early phase studies show that patients with a specific biomarker mutation may respond differently to the treatments, though there is no certainty about a difference in treatment effects. The patients can be divided into two biomarker groups: patients with mutation and patients without mutation. Patients in the two biomarker groups may respond differently to a same treatment. The treatment response rate w_{ij} is the probability of observing a favorable treatment outcome for a patient in subgroup j, $j = 1, 2$, who is assigned to treatment i, $i = 1, 2$. The effects of different treatments are independent, i.e., the response rates of the new treatment are independent of the response rates of the conventional treatment in the two biomarker subgroups. However, the response

rates of a same treatment in the two biomarker groups may either be equal, or independent when they are not equal. In other words, the effect of a treatment in the mutation group may or may not be different from the effect of the same treatment in the no mutation group. These assumptions, reflecting knowledge from the early phase trials, can be summarized by the following definition of the prior distribution of the treatment response rates, which is a special case of (1).

$$\begin{aligned} &w_{ij} \sim Beta(1,1); \quad w_{is} \perp w_{i'j'}, \quad i \neq i', \quad \forall j, j'; \\ &Pr(W_{i1} \leq w_{i1}, W_{i2} \leq w_{i2}) = (1-\pi_i) w_{i1} w_{i2} + \pi_i \min\{w_{i1}, w_{i2}\}, \quad i = 1 \text{ or } 2; \end{aligned} \tag{3}$$

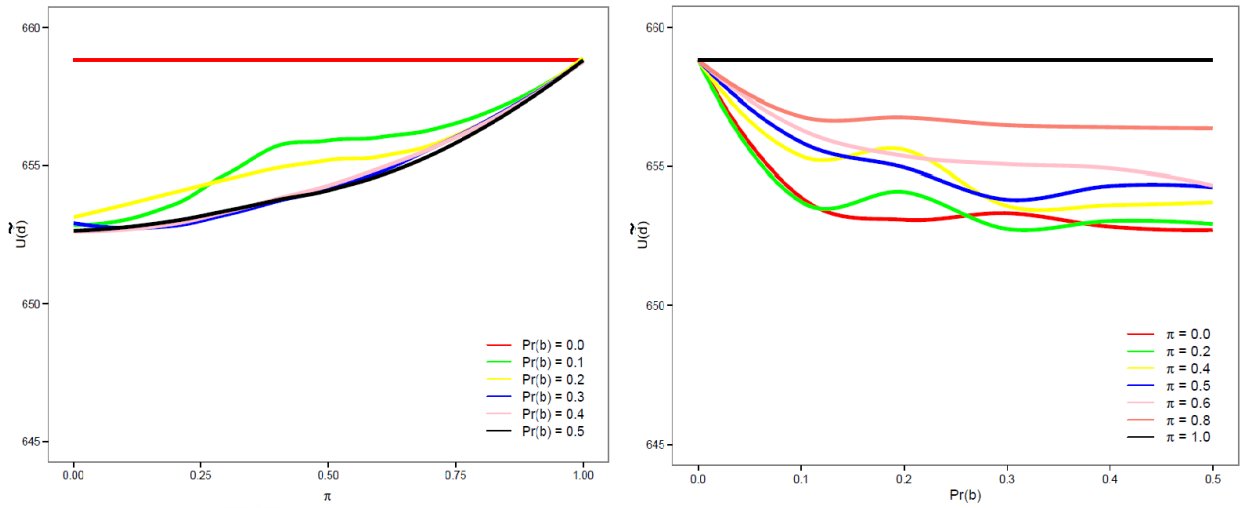

Figure 1: Model parameter influence on the theoretical optimum. Patient horizon $N_h = 1000$, trial size $N = 50$.

Two key quantities describe a biomarker: the prevalence of mutations, $Pr(b = 1)$, and π as defined in (3). The remainder of this section considers the theoretical optimum in different scenarios for these quantities, throughout we assume that $Pr(b = 1)$ is known from previous studies. Figure 1 shows the relationship between π, $Pr(b = 1)$ and the utility associated with the theoretical optimum. We simulate 1000 times of a patient horizon $N_h = 1000$ and a trial size $N = 50$. For the purpose of visualization, we are showing the situations when $\pi_1 = \pi_2 = \pi$. The first panel of Figure 1 illustrates how the expected utility of a trial,

$\mathcal{U}(\tilde{d})$, changes with π, by fixing the mutation prevalence. Each curve corresponds to the expected number of favorable treatment outcomes under different biomarker effects. The two biomarker groups are exchangeable in the model assumption. Therefore, the two curves with $Pr(b = 1) = \delta$ and $Pr(b = 1) = 1 - \delta$, $\delta \in [0, 1]$ are identical. The curve with $Pr(b = 1) = 0$ is a horizontal line because when no patient have the mutation, and $\mathcal{U}(\tilde{d})$ is independent of π, π is a redundant parameter in the model.

When $Pr(b = 1) \in (0, 1)$, the patient population is a mixture of two subgroups. The expected utility of a trial increases with π. The monotonically increasing relationship can be explained by considering the two extreme cases with $\pi = 0$ and $\pi = 1$. When $\pi = 0$, the two biomarker groups are independent, and the number of positive patient responses in the total patient horizon is the sum of the numbers in the two biomarker groups. On the other hand, when $\pi = 1$ the treatment effects are always the same in the two subgroups, thus it is no longer necessary to differentiate the subgroups. The patient population can be modeled as a single group. Therefore, the comparison between these two situations is essentially comparing a larger size trial with the juxtaposition of two smaller trials. A bigger sample size provides more information to detect the best treatment. The two subgroups become more closely related as π increases. At each level of π, $\mathcal{U}(\tilde{d})$ is minimized when $Pr(b = 1) = 0.5$. When the subgroups are getting more balanced, the joint ability to find the best treatment in the two subgroups is the smallest.

The second panel of Figure 1 shows the expected trial utility as we vary the prevalence of mutation with fixed π. Because the two subgroups are exchangeable, we plot the results over the range $Pr(b = 1) \in [0, 0.5]$. Had the results been plotted over $Pr(b = 1) \in [0, 1]$, the curves would be symmetric about $Pr(b = 1) = 0.5$. When $\pi = 1$, the biomarker is known not to modify the treatment effects. The two subgroups can thus be combined and considered as a single group. Then the expected utility of the trial does not change with $Pr(b = 1)$, because the total sample size is fixed. Therefore, the curve corresponding to this situation is a horizontal line and is above other curves corresponding to situations when

the biomarker mutation modifies treatment outcomes. When $\pi < 1$, the expected number of favorable treatment outcomes is negatively correlated with the biomarker prevalence, but this dependence is very weak for prevalence larger than 10%.

3.2 Comparison of Bayesian Adaptive Randomization and the Theoretical Optimum

Bayesian Adaptive Randomization (BAR) is emerging as a practical approach to develop adaptive designs, and is computationally far more efficient than dynamic programming. Here we explore the extent to which BAR can achieve expected utility close to that of the optimal solution. To this end, we extend the approach by [Thall and Wathen, 2007] so that it can be applied to trials incorporating biomarkers. When there are two treatments and one biomarker, BAR assigns treatment 2 to the n^{th} patient who belongs to biomarker group j with probability $r_j^{(t)}$, and

$$r_j^{(n)} = \frac{[Pr^{(n)}(w_{1j} < w_{2j})]^c}{[Pr^{(n)}(w_{1j} < w_{2j})]^c + [Pr^{(n)}(w_{1j} > w_{2j})]^c}$$

where $Pr^{(n+1)}(w_{1j} < w_{2j}) = Pr(w_{1j} < w_{2j}|X_n, ..., X_1, \mathbf{G_n}, ..., \mathbf{G_1})$ is the posterior probability that treatment 2 is the better treatment for patients in subgroup j. With the same priors defined in section 2. We follow [Thall and Wathen, 2007] for the choice of tuning parameter c.

While it is relatively easy to obtain an unbiased estimator of the prevalence of mutations in the patient population, it is often hard to choose π. In both designs, estimates of π are required to specify a design, and π needs to be elicited using the best knowledge. In this subsection, we will show the comparisons between the theoretical optimum and BAR in two types of simulations: first we generate treatment effects from the prior used for designing the trial; next we generate treatment effects from a prior with a different mixture weight, call it Π, to explore robustness to the choice of prior. In both cases we average across a collection

of values of **w**'s.

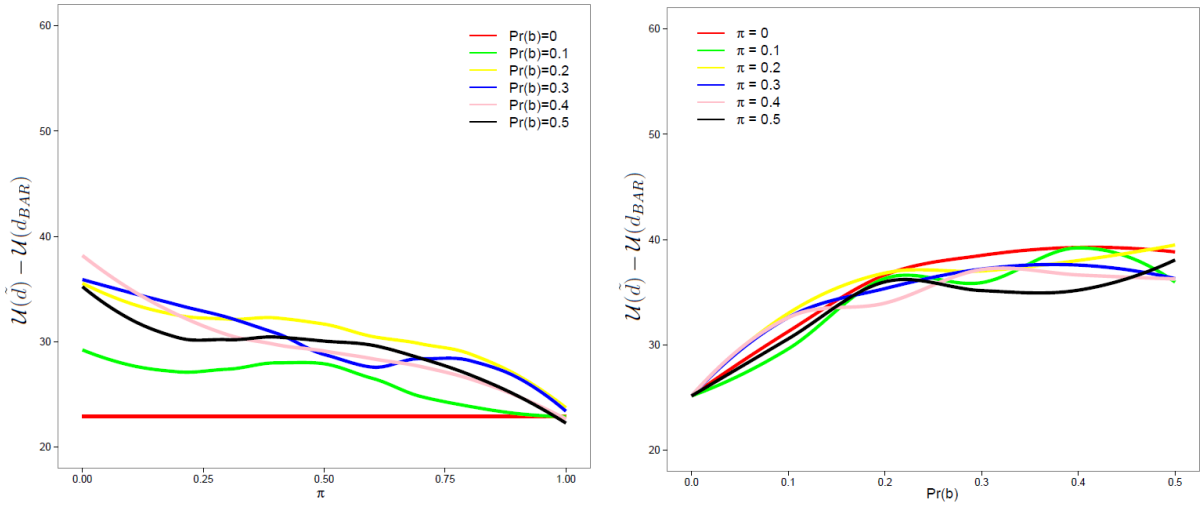

Figure 2: $\pi = \Pi$: comparison between the optimal design and the Bayesian Adaptive Randomization (BAR). Each curve plots the difference in the expected trial utility, $\mathcal{U}(\tilde{d}) - \mathcal{U}(d_{BAR})$. Patient horizon $N_h = 1000$, trial size $N = 50$.

In order to facilitate the display of the results graphically, we assumed that $\Pi_1 = \Pi_2 = \Pi$ and $\pi_1 = \pi_2 = \pi$. Figure 2 shows the comparison between the optimal design and BAR when generating from the design prior. The first panel shows how the difference in expected trial utility changes with π by holding $Pr(b = 1)$ constant. The second panel illustrates the relationship between $\mathcal{U}(\tilde{d}) - \mathcal{U}(d_{BAR})$ and $Pr(b = 1)$ with fixed π. For each combination of π and $Pr(b = 1)$, the two designs are used to implement trials with the same simulated patients, we then observe the numbers of favorable treatment outcomes for the two designs and compute the difference. Each curve in the two panels of Figure 2 is the expected utility lost by using BAR compared to the best achievable result. The difference between BAR and the theoretical optimum is generally small, between 2.5% and 4% of the size of the patient horizon.

In the first panel, the difference between the optimal design and BAR is minimized when the patient horizon consists of identical patients, namely, either all patients have the biomarker mutation or none has. The difference in expected utility increases when the subgroups are becoming more balanced and the maximum is achieved when the two

subgroups are of equal size. At a fixed biomarker prevalence, the largest difference in expected utility is observed when the prior expects most difference in the treatment efficacy in patient subgroups ($\pi = 0$). When the subgroups are more closely related, that is, when the biomarker status is expected to play a less important role in predicting patients' responses to treatments, the difference between the two trial designs is smaller. The difference between BAR and the theoretical optimum is increasing with diminishing π.

The second panel plots the relationship between the difference in expected utility and the prevalence of mutations. With fixed π and as $Pr(b = 1)$ is more balanced, $\mathcal{U}(\tilde{d}) - \mathcal{U}(d_{BAR})$ increases when the expected trial utility of both the optimal design and BAR decreases.

3.3 The Bayesian Adaptive Randomization and the Play the Winner Rule

There has been a growing interest in comparing Bayesian and frequentist adaptive designs. The "play-the-winner"(PW) rule [Zelen, 1969, Wei and Durham, 1978] has been used in designing several clinical trials [Bartlett et al., 1985, Tamura et al., 1994, Yao and Wei, 1996].In this section, we compare the BAR to the PW in terms of the expected trial utility. Our goal here is to quantify, from a Bayesian standpoint, the change in utility associated with using a practical approach such as PW as compared to the BAR.

Each curve in Figure 3 plots the difference in expected trial utility $\mathcal{U}(d_{BAR}) - \mathcal{U}(d_{PW})$. The three panels show the scenarios where treatment effects are generated from priors with different mixture weights (Π). Depending on the scenario chosen, PW can have a worse or better utility than BAR, though BAR achieves a higher utility in the vast majority of cases combined. The difference can be large, as many as 50 additional successfully treated patients out of a patient horizon of size 1000. The difference in expected utility decreases when the subgroups become more balanced. When the subgroups are highly imbalanced, the difference is relatively robust to the choice of π. PW has better utility when the prior rules out a biomarker effect when there is in fact a strong one, and when the subgroups are

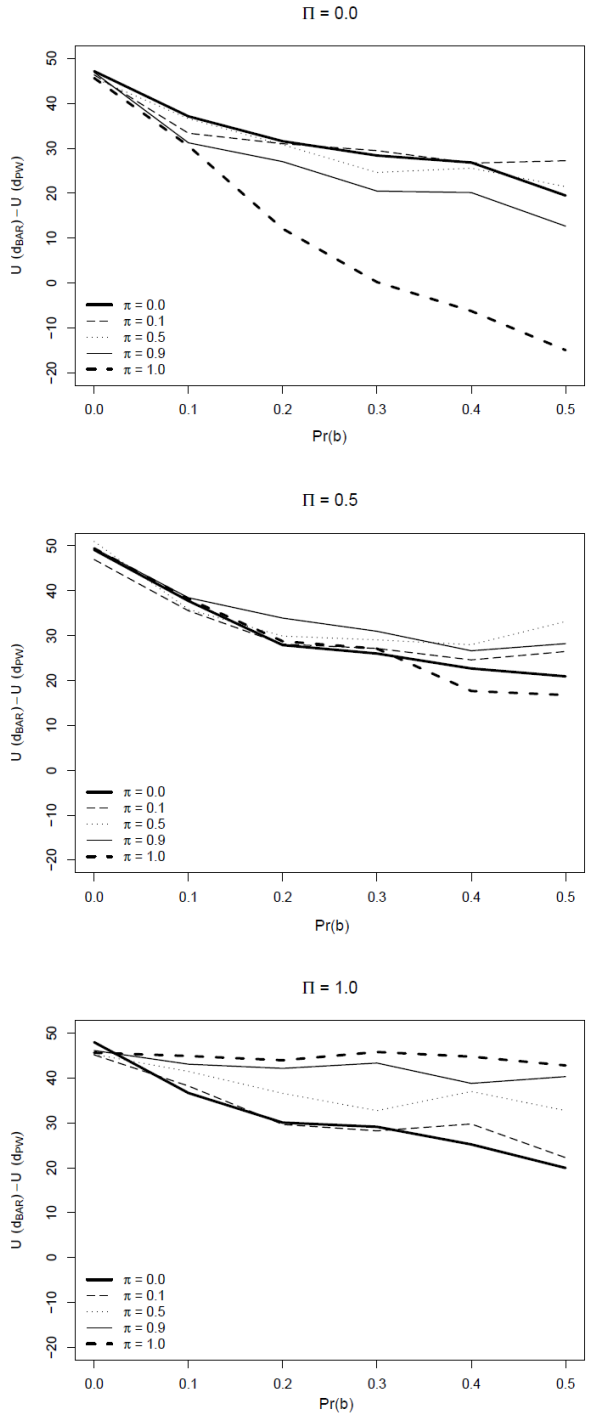

Figure 3: At different combinations of $\Pi \times \pi$: comparison between $\mathcal{U}(d_{BAR})$ and $\mathcal{U}(d_{PW})$. Patient horizon $N_h = 1000$, trial size $N = 50$

relatively balanced.

4 Discussion

In this article, we derive a theoretical optimal adaptive design, where treatment assignment is allowed to depend on a binary biomarker. The optimal design, maximizes the expected trial utility given by the expected number of favorable outcomes with a given patient horizon. When treatment efficacy depends on biomarkers, our analysis shows how this relationship affects the best achievable results in terms of the expected trial utility in trials where treatment assignments are adaptive on early treatment outcomes. Our work provides absolute benchmark for the evaluation of trial designs in personalized medicine.

Our analysis is the first to consider optimal designs when treatment efficacy depends on biomarkers. We hope it will provide the basis to consider more complex situations as well. Although here we focus on maximizing the number of favorable outcomes, our approach can be easily modified to handle other utility functions.

The optimal design is an application of dynamic programming, which records every possible outcome path when conducting a trial. This feature of the algorithm imposes heavy computational burden, even for a modest trial size. In this context we wrote simulation programs that reduce the dimension of the data structure and free up machine memory during the implementation of trials. However, computational demands remain a challenge for this type of approach.

In practice it is uncommon to implement trials when the treatment assignment is deterministic for each patient. To address this concern, we also consider an adaptive design where the treatment assignment is randomized. By definition, the largest utility gain is achieved when the treatment assignment is optimal and deterministic, but our work allows one to benchmark a proposed suboptimal randomized design against the optimum. Our comparison between the optimal design and other designs quantifies the difference in expected trial

utility under different scenarios. When the number of patients to be cured is of primary concern, the optimal design may be preferable. There could also be situations where a small portion of the expected trial utility is sacrificed for a design that has easier interpretation or includes a randomized treatment assignment.

Acknowledgements

The computations in this paper were conducted on the machines of Dana-Farber Cancer Institute and the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University.

References

C. J. Allegra, J. M. Jessup, M. R. Somerfield, S. R. Hamilton, E. H. Hammond, D. F. Hayes, P. K. McAllister, R. F. Morton, and R. L. Schilsky. American society of clinical oncology provisional clinical opinion: testing for kras gene mutations in patients with metastatic colorectal carcinoma to predict response to anti–epidermal growth factor receptor monoclonal antibody therapy. *Journal of Clinical Oncology*, 27(12):2091–2096, 2009.

P. Armitage. The search for optimality in clinical trials. *International Statistical Review/Revue Internationale de Statistique*, pages 15–24, 1985.

A. C. Atkinson and A. Biswas. Bayesian adaptive biased-coin designs for clinical trials with normal responses. *Biometrics*, 61(1):118–125, 2005.

A. Barker, C. Sigman, G. Kelloff, N. Hylton, D. Berry, and L. Esserman. I-spy 2: an adaptive breast cancer trial design in the setting of neoadjuvant chemotherapy. *Clinical Pharmacology & Therapeutics*, 86(1):97–100, 2009.

R. H. Bartlett, D. W. Roloff, R. G. Cornell, A. F. Andrews, P. W. Dillon, and J. B. Zwischenberger. Extracorporeal circulation in neonatal respiratory failure: a prospective randomized study. *Pediatrics*, 76(4):479–487, 1985.

J. Baselga, L. Norton, J. Albanell, Y.-M. Kim, and J. Mendelsohn. Recombinant humanized anti-her2 antibody (herceptin) enhances the antitumor activity of paclitaxel and doxorubicin against her2/neu overexpressing human breast cancer xenografts. *Cancer research*, 58(13):2825–2831, 1998.

R. E. Bellman. *Dynamic programming*. Courier Dover Publications, 2003.

D. A. Berry and S. G. Eick. Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. *Statistics in medicine*, 14(3):231–246, 1995.

P. L. Canner. Selecting one of two treatments when the responses are dichotomous. *Journal of the American Statistical Association*, 65(329):293–306, 1970.

L. Chin, J. N. Andersen, and P. A. Futreal. Cancer genomics: from discovery science to personalized medicine. *Nature medicine*, 17(3):297–303, 2011.

T. Colton. A model for selecting one of two medical treatments. *Journal of the American Statistical Association*, 58(302):388–400, 1963.

C. Compton. Getting to personalized cancer medicine. *Cancer*, 110(8):1641–1643, 2007.

B. J. Druker, M. Talpaz, D. J. Resta, B. Peng, E. Buchdunger, J. M. Ford, N. B. Lydon, H. Kantarjian, R. Capdeville, S. Ohno-Jones, et al. Efficacy and safety of a specific inhibitor of the bcr-abl tyrosine kinase in chronic myeloid leukemia. *New England Journal of Medicine*, 344(14):1031–1037, 2001.

J. R. Eisele. The doubly adaptive biased coin design for sequential clinical trials. *Journal of Statistical Planning and Inference*, 38(2):249–261, 1994.

A. M. Gonzalez-Angulo, B. T. Hennessy, and G. B. Mills. Future of personalized medicine in oncology: a systems biology approach. *Journal of Clinical Oncology*, 28(16):2777–2783, 2010.

M. A. Hamburg and F. S. Collins. The path to personalized medicine. *New England Journal of Medicine*, 363(4):301–304, 2010.

J. K. Hammitt and Y. Zhang. Combining experts judgments: Comparison of algorithmic methods using synthetic data. *Risk Analysis*, 33(1):109–120, 2013.

R. Herbst, G. Blumenschein Jr, E. Kim, J. Lee, A. Tsao, C. Alden, S. Liu, D. Stewart, I. Wistuba, and W. Hong. Sorafenib treatment efficacy and kras biomarker status in the biomarker-integrated approaches of targeted therapy for lung cancer elimination (battle) trial. *J Clin Oncol*, 28(15 Suppl):7609, 2010.

A. Ivanova. A play-the-winner-type urn design with reduced variability. *Metrika*, 58(1):1–13, 2003.

E. S. Kim, R. S. Herbst, I. I. Wistuba, J. J. Lee, G. R. Blumenschein, A. Tsao, D. J. Stewart, M. E. Hicks, J. Erasmus, S. Gupta, et al. The battle trial: personalizing therapy for lung cancer. *Cancer Discovery*, 1(1):44–53, 2011.

N. B. La Thangue and D. J. Kerr. Predictive biomarkers: a paradigm shift towards personalized cancer medicine. *Nature Reviews Clinical Oncology*, 8(10):587–596, 2011.

U. McDermott and J. Settleman. Personalized cancer therapy with selective kinase inhibitors: an emerging paradigm in medical oncology. *Journal of Clinical Oncology*, 27(33):5650–5659, 2009.

J. G. Paez, P. A. Jänne, J. C. Lee, S. Tracy, H. Greulich, S. Gabriel, P. Herman, F. J. Kaye, N. Lindeman, T. J. Boggon, et al. Egfr mutations in lung cancer: correlation with clinical response to gefitinib therapy. *Science*, 304(5676):1497–1500, 2004.

G. Parmigiani and L. Inoue. *Decision theory: principles and approaches*, volume 812. John Wiley & Sons, 2009.

W. B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

G. L. Rosner and B. N. Bekele. Bayesian designs in clinical trials. In W. K. Kelly, editor, *Oncology Clinical Trials*. Demos Medical Publishing, 2010.

E. H. Rubin, K. M. Anderson, and C. K. Gause. The battle trial: a bold step toward improving the efficiency of biomarker-based drug development. *Cancer discovery*, 1(1): 17–20, 2011.

J. Savard. Personalised medicine: a critique on the future of health care. *Journal of bioethical inquiry*, 10(2):197–203, 2013.

R. L. Schilsky. Personalized medicine in oncology: the future is now. *Nature reviews Drug discovery*, 9(5):363–366, 2010.

R. N. Tamura, D. E. Faries, J. S. Andersen, and J. H. Heiligenstein. A case study of an adaptive clinical trial in the treatment of out-patients with depressive disorder. *Journal of the American Statistical Association*, 89(427):768–776, 1994.

P. F. Thall and J. K. Wathen. Practical bayesian adaptive randomisation in clinical trials. *European Journal of Cancer*, 43(5):859–866, 2007.

G. Upton and R. Lee. The importance of the patient horizon in the sequential analysis of binomial clinical trials. *Biometrika*, 63(2):335–342, 1976.

L. J. van't Veer and R. Bernards. Enabling personalized cancer medicine through analysis of gene-expression patterns. *Nature*, 452(7187):564–570, 2008.

A. Walther, E. Johnstone, C. Swanton, R. Midgley, I. Tomlinson, and D. Kerr. Genetic prognostic and predictive markers in colorectal cancer. *Nature Reviews Cancer*, 9(7): 489–499, 2009.

L. Wei and S. Durham. The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, 73(364):840–843, 1978.

L. D. Wood, D. W. Parsons, S. Jones, J. Lin, T. Sjöblom, R. J. Leary, D. Shen, S. M. Boca, T. Barber, J. Ptak, et al. The genomic landscapes of human breast and colorectal cancers. *Science*, 318(5853):1108–1113, 2007.

Q. Yao and L. Wei. Play the winner for phase ii/iii clinical trials. *Statistics in medicine*, 15 (22):2413–2423, 1996.

M. Zelen. Play the winner rule and the controlled clinical trial. *Journal of the American Statistical Association*, 64(325):131–146, 1969.

X. Zhou, S. Liu, E. S. Kim, R. S. Herbst, and J. J. Lee. Bayesian adaptive design for targeted therapy development in lung cancera step toward personalized medicine. *Clinical Trials*, 5(3):181–193, 2008.