 DOI: 10.1111/biom.12437

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 School of Public Health, Boston, Massachusetts, U.S.A.
[2]Department of Biostatistics and Computational Biology, Dana-Farber Cancer Institute, Boston, Massachusetts, U.S.A.
[*]*email:* gp@jimmy.harvard.edu

Summary. Clinical biomarkers play an important role in precision medicine and are now extensively used in clinical trials, particularly in cancer. A response adaptive trial design enables researchers to use treatment results about earlier patients to aid in treatment decisions of later patients. Optimal adaptive trial designs have been developed without consideration of biomarkers. In this article, we describe the mathematical steps for computing optimal biomarker-integrated adaptive trial designs. These designs maximize the expected trial utility given any pre-specified utility function, though we focus here on maximizing patient responses within a given patient horizon. We describe the performance of the optimal design in different scenarios. We compare it to Bayesian Adaptive Randomization (BAR), which is emerging as a practical approach to develop adaptive trials. The difference in expected utility between BAR and optimal designs is smallest when the biomarker subgroups are highly imbalanced. We also compare BAR, a frequentist play-the-winner rule with integrated biomarkers and a marker-stratified balanced randomization design (BR). We show that, in contrasting two treatments, BR achieves a nearly optimal expected utility when the patient horizon is relatively large. Our work provides novel theoretical solution, as well as an absolute benchmark for the evaluation of trial designs in personalized medicine.

Key words: Biomarkers; Bayesian adaptive designs; Dynamic programming; Optimal strategy; 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. A tremendous amount of resources is presently allocated to developing precision cancer medicine (Chin et al., 2011; La Thangue and Kerr, 2011). Since the Food and Drug Administration (FDA) approved trastuzumab in human epidermal growth factor receptor 2 (HER2) positive breast cancer patients, personalized treatments have been developed for chronic myelogenous leukemia (Druker et al., 2001), colon cancer (Allegra et al., 2009), lung cancer (Paez et al., 2004), and other malignacies (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 (Kim et al., 2011), non-small cell lung cancer (NSCLC) patients have been classified into five subgroups defined by biomarkers reflecting tumor genetics. Another biomarker-adaptive design, also allowing treatments to enter and exit the trial is the I-SPY 2 (Barker et al., 2009; Alexander et al., 2013). The goal of these trials is to identify the most beneficial treatment separately for each patient subgroup.

Adaptive designs for cancer clinical trials have been studied extensively, from both frequentist and Bayesian perspectives (Yin, 2013), although biomarkers are not often included in these methodologies. From a decision theoretic perspective, we lack results on optimal trial designs integrating biomarkers. In this article, we fill this gap by deriving an optimal Bayesian biomarker-integrated adaptive trial design. Our 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, the outcome data define optimal treatments, which vary across biomarker profiles. Subsequently, optimal treatments are individually assigned to patients for the remainder of the horizon of interest.

Having defined and studied 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), a frequentist play-the-winner rule with integrated biomarkers and a marker-stratified balanced randomization design (BR). Our comparison will illustrate scenarios in which using these designs does not lead to a significant loss in terms of the expected trial utility when compared to the theoretical optimum, as well as other where this is not the case. These comparisons are the primary practical goal of this article.

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. We use $Pr(b_k = 1)$ to denote the prevalence of biomarker mutation b_k in the population. Patients' profiles can be defined by these biomarkers indicators and summarized by the vector $G = (b_1, ...b_K)$. The patient population is divided into different biomarker groups based on their marker profiles G. The optimal trial design proposed in this article can be applied when biomarker groups are mutually exclusive. We use the index $j \in \{1, ..., J\}$ to indicate the biomarker groups in the trial.

Similarly to (Berry and Eick, 1995), we consider a patients horizon. It includes patients in the trial, as well as subsequent patients who, after completion of the experimental study, will receive treatments accordingly to the trial results. The size of the patients horizon is an important element in adaptive designs, because it controls how the design balances the two potentially competing needs of learning the treatment effects and assigning patients to the best treatment options within the trial.

We assume that the outcome of interest is binary and that it is known immediately after treatment assignment. Let $w_{i,j}$ be the probability of observing a success, or positive outcome, for treatment i in subgroup j, for $i = 1, \ldots, I$ and $j = 1, \ldots, J$.

We take a Bayesian approach to estimate the $w_{i,j}$'s and to quantify 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 $w_{i,1} = \cdots = w_{i,J}$. With prior probability π_i, treatment i has the same effect in every biomarker group. Otherwise, the effects of treatment i take a distinct value in each subgroup. These values are a priori independent with uniform marginal distributions. Also, conditional on event $w_{i,1} = \cdots = w_{iJ}$, the parameters $w_{i,j}$ are uniformly distributed. To summarize, the joint prior distribution of the unknown response rates for treatment i is

$$Pr(w_{i,1} \le a_1, ..., w_{i,J} \le a_J) = (1 - \pi_i) \prod_{j=1}^{J} a_j + \pi_i \min_{j=1}^{J}\{a_j\}, \quad a_1, \ldots, a_J \in (0, 1). \quad (1)$$

The second term captures the case where the response probability is the same in all biomarker groups. When a π_i is equal to 0, the biomarker-specific response probabilities become independent, and, during the course of the study, prediction of a patient outcome is based exclusively on previous data from individuals with the same biomarker profile. The π_i's can be selected using data from early phase trials or elicited from expert judgment (Hammitt and Zhang, 2013). The response rates of two treatments $(w_{i,1}, \ldots, w_{iJ})$ and $(w_{i'1}, \ldots, w_{i'J})$, $i \neq i'$ are a priori independently distributed.

In this article, we will use prior (1). However, our approach is more general, in that the algorithm that we describe can easily accommodate other prior distributions or under particular sets of treatment response rates. If, for example, reliable evidence on treatment effects is available from early phase trials, it can be summarized via prior distributions for the $(w_{i,j})$'s, modifying the assumption of uniformity. In the supplementary material we provide a brief description of an optimal design obtained under a different prior distribution using a probit link and an example of optimal design obtained under particular sets of treatment response rates.

Let X_n denote the treatment outcome of the n^{th} patient enrolled in the trial, and G_n the associated marker profile. The arrival of patient n also marks the n^{th} decision point. The notation $d(n|X_{n-1}, ..., X_1, G_n, ..., G_1) = i$ 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 also on the marker profiles of the previous $n-1$ patients $G_{n-1}, ..., G_1$, and on the marker profile, 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 within a biomarker group j will receive the same treatment.

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

We are interested in choosing a design d from the set D of all possible sequential biomarker-dependent allocation rules. A trial of size N is conducted to study the unknown treatment effects $\mathbf{w} = (w_{i,j})_{I \times J}$. Once a generic sequential allocation rule d is selected, the sequence of outcomes $\mathbf{X} \in \mathcal{X}$ is observed with conditional distribution $f_d(\mathbf{X}|\mathbf{w})$, where $\mathbf{X} = (X_1, \ldots, X_N)$.

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 $N_h - N \ge 0$ additional patients whose treatment will be identified by the information accrued during the trial. In what follows $d = \left[d(n|X_{n-1}, \ldots, X_1, G_n, \ldots, G_1), \;\; n = 1, \ldots, N\right]$ will denote the allocation of patients during the trial, while d^h will denote the treatment assignment rule for the remaining $N_h - N$ patients after completion of the trial. In our work d^h is allowed to depend on the individual patient's marker profile.

The optimal sequential design d and final d^h selection rules are defined with respect to a utility function $U(d, d^h, \mathbf{X})$, where $\mathbf{X} = (X_1, \ldots, X_{N_h})$, designed to capture the most important goals of conducting the trial. Comparing treatment

efficacy is helpful insofar as it allows to improve treatment outcomes. The utility function coincides with the total number of favorable treatment outcomes:

$$U(d, d^h, \mathbf{X}) = \underbrace{\sum_{n=1}^{N} X_n}_{\text{inside the trial}} + \underbrace{\sum_{n=N+1}^{N_h} X_n}_{\text{outside the trial}} \quad (2)$$

This specific form of utility function has been used by several authors (Armitage, 1985; Berry and Eick, 1995) in the design of trials without incorporating biomarkers. We will use this utility function to illustrate the behavior of the optimal design under different treatment effects across biomarker groups, and to make comparisons among trial designs.

2.2. *Optimal Solution*

The optimal choice of the adaptive strategy (d, d^h) is obtained by dynamic programming (Bellman, 2003; Parmigiani and Inoue, 2009). We begin with the second component d^h, the final decision rule, which is computed conditional on the available information $(M^{(N)}, S^{(N)})$ at completion of the trial. 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 optimal decision rule by $\tilde{d}^h$; it selects, for every subgroup j, the treatment with the highest estimated effect, i.e., $\arg\max_i \mathbb{E}(w_{i,j} \mid M^{(N)}, S^{(N)})$.

Similarly, $\tilde{d}$ is computed by maximizing the expected utility. The optimal design targets 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 during the trial.

Let $R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}]$ 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. Our goal is to select the design $\tilde{d}$ that maximizes the expected utility. We can write the expected utility of a design d as $R^{(0)}_{d,\tilde{d}^h}[M^{(0)}, S^{(0)}]$. Also, let $\mu^{(n)}_{i,j}$ be the posterior mean of $w_{i,j}$ calculated conditional on $(M^{(n)}, S^{(n)})$, that is $\mu^{(n)}_{i,j} = \mathbb{E}[w_{i,j}|X_n, ..., X_1, G_n, ..., G_1]$. 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 $\mu^{(n-1)}_{i,j}$ if this patient is assigned to treatment i. This notation can be used to compute the optimal design $\tilde{d}$ and its expected utility. At the end of a trial, based on the final decision rule $\tilde{d}^h$, the expected number of future favorable treatment outcomes is

$$R^{(N)}_{d,\tilde{d}^h}[M^{(N)}, S^{(N)}] = \sum_{n=N+1}^{N_h} \sum_j p_j \max_i(\mu^{(N)}_{i,j}),$$

where p_j is the probability of belonging to the *j-th* biomarker subgroup.

The steps for deriving the optimal design $\tilde{d}$ start with the computation of $R^{(N)}_{d,\tilde{d}^h}[M^{(N)}, S^{(N)}]$ and $\mu^{(N)}_{i,j}$ for all possible combinations of $M^{(N)}$ and $S^{(N)}$. Note that $R^{(N)}_{d,\tilde{d}^h}[M^{(N)}, S^{(N)}]$ does not vary across designs d, as it only depends on the sufficient statistics $(M^{(N)}, S^{(N)})$. In contrast $R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}]$, $n = 0, \ldots, N-1$, varies across potential designs. With backward computations, if we fix a design d, then the conditional expected value $R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}]$, $n < N$, can be calculated using the collections of $R^{(n+1)}_{d,\tilde{d}^h}[M^{(n+1)}, S^{(n+1)}]$ values at all possible $(M^{(n+1)}, S^{(n+1)})$ configurations:

$$R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}] = \mathbb{E}(R^{(n+1)}_{d,\tilde{d}^h}[M^{(n+1)}, S^{(n+1)}] \mid M^{(n)}, S^{(n)}) + \mathbb{E}(X_{n+1} \mid M^{(n)}, S^{(n)}).$$

For a fixed d, $M^{(n)}$ and $S^{(n)}$ combination, $R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}]$ is computed by integrating $R^{(n+1)}_{d,\tilde{d}^h}[M^{(n+1)}, S^{(n+1)}]$ with respect to the conditional distribution of $M^{(n+1)}$, $S^{(n+1)}$, and predicting X_{n+1}. This argument can be iterated for $n = N-1, ..., 0$.

Computation of the optimum $\tilde{d}$ is based on moving backward through the steps just described. Each adaptive strategy d assigns the $(n+1)$-th patient to a treatment conditionally on any combination $(M^{(n)}, S^{(n)})$, and on the patient's marker profile G_{n+1}. We can, therefore, write $d(M^{(n)}, S^{(n)}, G_{n+1}) \in \{1, 2, \ldots, J\}$. The optimum $\tilde{d}$ is derived through the following steps:

1) Initialize $R^{(N)}_{d,\tilde{d}^h}[M^{(N)}, S^{(N)}]$. This conditional expectation varies with the $(M^{(N)}, S^{(N)})$ values, but is invariant with respect to the design d. These computations provide us the conditional expectations $R^{(N)}_{d,\tilde{d}^h}[M^{(N)}, S^{(N)}]$, for all possible $(M^{(N)}, S^{(N)})$ values.
2) Compute $\tilde{d}(M^{(n)}, S^{(n)}, G_{n+1})$ at every $(M^{(n)}, S^{(n)}, G_{n+1})$ combination, and $R^{(n)}_{\tilde{d},\tilde{d}^h}[M^{(n)}, S^{(n)}]$, using the previously computed values of $R^{(n+1)}_{\tilde{d},\tilde{d}^h}[M^{(n+1)}, S^{(n+1)}]$ and the predicted values of X_{n+1} under each treatment $j = 1, \ldots, J$. That is, $\tilde{d}(M^{(n)}, S^{(n)}, G_{n+1})$ is computed by maximizing

$$\mathbb{E}\left(X_{(n+1)} \mid M^{(n)}, S^{(n)}, G_{n+1}\right) + \mathbb{E}\left(R^{(n+1)}_{\tilde{d},\tilde{d}^h}[M^{(n+1)}, S^{(n+1)}] \mid M^{(n)}, S^{(n)}, G_{n+1}\right). \quad (3)$$

3) Repeat the previous step for $n = N-1, ..., 0$. The expected utility of the optimal design is $\mathcal{U}(\tilde{d}) = R^{(0)}_{\tilde{d},\tilde{d}^h}[M^{(0)}, S^{(0)}]$.

The algorithm uses iterative computations of the $R^{(n)}_{d,\tilde{d}^h}[M^{(n)}, S^{(n)}]$ conditional expectation when the biomarker subgroup j of the n-th patient together with $(M^{(n-1)}, S^{(n-1)})$ are known, and $d(M^{(n-1)}, S^{(n-1)}, G_n) = i$. The conditional expectation, for every $n = 1, \ldots, N$, has a simple closed form expression:

$$\mu^{(n-1)}_{i,j} \times R^{(n)}_{d,\tilde{d}^h}[M^{(n-1)} + \mathbf{1}_{i,j}, S^{(n-1)} + \mathbf{1}_{i,j}] + \left(1 - \mu^{(n-1)}_{i,j}\right) \times R^{(n)}_{d,\tilde{d}^h}[M^{(n-1)} + \mathbf{1}_{i,j}, S^{(n-1)}], \quad (4)$$

here $\mathbf{1}_{i,j}$ is a matrix of zeros with the exception of the (i, j)-th entry equal to one. At each iteration the optimization algorithm sets $\tilde{d}(M^{(n-1)}, S^{(n-1)}, G_n)$ equal to i if the i-th treatment maximizes the conditional expectation of $X_n + R^{(n)}_{\tilde{d},\tilde{d}^h}[M^{(n)}, S^{(n)}]$. In summary, the conditional expectation of the final utility, when the n-patient belong to the j-th group, under $\tilde{d}$ is

$$\max_i \left[\mu^{(n-1)}_{i,j}\left(1 + R^{(n)}_{\tilde{d},\tilde{d}^h}[M^{(n-1)} + \mathbf{1}_{i,j}, S^{(n-1)} + \mathbf{1}_{i,j}]\right) + \left(1 - \mu^{(n-1)}_{i,j}\right) R^{(n)}_{\tilde{d},\tilde{d}^h}[M^{(n-1)} + \mathbf{1}_{i,j}, S^{(n-1)})]\right] + \sum_{i,j} s^{(n)}_{i,j}.$$

When we compute the optimal design, at most configurations of the matrices (M^n, S^n), there is a single treatment that maximizes the conditional expectation in (3). The only exception are configuration with predicted utilities in (3) identical across multiple arms. In these cases the optimal design is allowed to arbitrarily select one of the treatments, randomly or deterministically. If, for example, the prior models for arms $i = 1, 2$ are identical, then the optimal design can assign the first patient to either of the two treatments.

The algorithm to compute the optimal design can be used with any prior distribution on the biomarker-specific response probabilities, not just that in Section 2.1. In the Supplementary Material we briefly describe an optimal design obtained using a different prior distribution, defined, similarly to the BATTLE trial (Kim et al., 2011), using a probit link. An important advantage of the prior that we use in this article is the possibility of computing the predictive probabilities of positive outcomes, during the course of the study, without recurring to Markov chain Monte Carlo simulations or other approximation methods. The same advantage is shared by all finite mixtures of Beta distributions. These constitute a large class of prior distributions, often referred to as Bernerstein priors, that contains our specification as a special case. These distributions have been extensively discussed in the literature because they combine flexibility and analytic tractability.

3. Numerical Results and Comparison

3.1. *Properties of the Theoretical Optimum in Two-Treatments Comparisons*

We consider the comparison of two cancer treatments. Assume that early phase studies showed that patients with and without a specific mutation may respond differently to these treatments, though there is no certainty about the difference between effects across biomarker subgroups. In other words, patients are divided into two groups depending on whether they present with a specific genetic abnormality, which may determine whether they respond to the experimental treatments.

The treatment response rate $w_{i,j}$ 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, that is the response rates $(w_{1,1}, w_{1,2})$ under treatment 1 are independent of the response rates $(w_{2,1}, w_{2,2})$ under treatment 2. The prior distribution of the response rates is a special case of (1):

$$Pr(w_{i,1} \le a_1, w_{i,2} \le a_2) = (1 - \pi_i)a_1 a_2 + \pi_i \min\{a_1, a_2\}, \quad \text{for every } a_1, a_2 \in (0, 1). \tag{5}$$

Marginally, $w_{i,j} \sim Beta(1, 1)$, and $(w_{1,1}, w_{1,2}) \perp (w_{2,1}, w_{2,2})$. We refer to Web Appendix 1 for the computation of the posterior distribution of the response rates.

Two key quantities describe a biomarker: $Pr(b = 1)$, the probability of drawing a biomarker positive patients during the patient horizon, and π, as defined in (5). Our simulations illustrate that, under the model chosen here, prior distributions with stronger dependence between response probabilities in the two population subgroups are associated to higher expected utilities. A more detailed description of the absolute performance of the optimal design can be found in Web Appendix 2. Web Appendix 3 shows a similar trend under a different prior distribution that also allows to modulate the degree of dependence of the response probabilities.

In the following sections, we will compare the optimal design with other adaptive designs. In preparation we investigate sensitivity to the patient horizon. Reliance on the patient horizon is an important difference between the optimal design obtained with dynamic programming, and those obtained from BAR or balanced designs. In some contexts, such as in the study of treatments for rare diseases, it is valuable to explicitly consider the likely size of the future patient population. In other contexts it might be difficult to select a realistic horizon. Nonetheless, comparisons for different N_h values can help in the evaluation of practical randomized schemes like BAR. We used sensitivity analyses to illustrate how the parameter N_h influence adaptation. First, we generate treatment effects $w_{i,j}$ from the same prior that is also used for designing the trial and for treatment assignment. Next we generate treatment effects $w_{i,j}$ from a mixture distribution that is different from the Bayesian model used for treatment assignment. More precisely we replace the parameter π in the prior with a different value Π and obtain the distribution of the treatment assignments in each subgroup. We use this set of simulations to explore robustness to the choice of the prior. In both cases we average across a collection of values of $\mathbf{w}$'s, where $\mathbf{w} = (w_{i,j},\ i = 1, 2,\ j = 1, 2)$. The sensitivity analysis results are reported in Table 1 and Web Appendix 5. These results show that the expected trial utility of the optimal design is roughly proportional to N_h when fixing the trial size. For example, if biomarker subgroups are balanced, $\Pi = 0.1$ and $\pi = 0.1$, then $\mathcal{U}(\tilde{d}) = 160.28$ when $N_h = 250$, that is about half of the utility $\mathcal{U}(\tilde{d}) = 319.87$ obtained when $N_h = 500$. Similar comparisons can be made by fixing other hyperparameter values and varying the size of the patient horizon N_h.

The observation above indicates that the ability of the optimal design to identify the best treatment is relatively robust to the choice of N_h. With a small trial size, in our sensitivity analysis, the optimal design selects the best treatment options with comparable probabilities at different values of the

Table 1

Sensitivity analysis to illustrate the influence of the parameters π and N_h on the optimal design $\tilde{d}$ characteristics. This table reports the expected trial utility and, in brackets, the utility standard deviation of four different designs: optimal design (optimal), Bayesian adaptive randomization(BAR), play-the- winner(PW) and marker-stratified balanced randomization(BR).

	$N = 30, N_h = 250, \Pi = 0.1$											
	$\pi = 0.1$				$\pi = 0.5$				$\pi = 0.9$			
	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR
$Pr(b = 1) = 0.1$	162.86 (56.40)	149.95 (58.68)	145.36 (62.82)	156.72 (52.11)	162.46 (56.45)	149.88 (58.50)	145.36 (62.82)	156.72 (52.11)	160.60 (57.29)	148.31 (61.63)	145.36 (62.82)	156.72 (58.50)
$Pr(b = 1) = 0.5$	160.28 (46.38)	147.82 (50.05)	149.16 (48.05)	155.89 (44.94)	163.69 (44.68)	149.95 (59.58)	149.16 (48.05)	155.89 (44.94)	158.67 (47.39)	145.60 (51.34)	149.16 (48.05)	155.89 (44.94)
	$N = 30, N_h = 500, \Pi = 0.1$											
	$\pi = 0.1$				$\pi = 0.5$				$\pi = 0.9$			
	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR
$Pr(b = 1) = 0.1$	325.03 (105.77)	302.90 (119.80)	298.40 (125.58)	319.48 (110.74)	322.57 (111.14)	299.83 (121.68)	298.40 (125.58)	319.48 (110.74)	324.92 (108.83)	305.09 (122.35)	298.40 (125.58)	319.48 (110.74)
$Pr(b = 1) = 0.5$	319.87 (92.09)	291 (101.52)	295.15 (101.59)	319.66 (87.46)	320.92 (94.31)	292.89 (103.44)	295.15 (101.59)	319.66 (87.46)	315.62 (93.47)	295.74 (98.36)	295.15 (101.59)	319.66 (87.46)
	$N = 30, N_h = 1000, \Pi = 0.1$											
	$\pi = 0.1$				$\pi = 0.5$				$\pi = 0.9$			
	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR
$Pr(b = 1) = 0.1$	653.67 (223.91)	615.29 (243.87)	594.58 (256.47)	636.84 (212.01)	653.33 (225.15)	608.67 (240.02)	594.58 (256.47)	636.84 (212.01)	649.77 (222.86)	608.99 (241.00)	594.58 (256.47)	636.84 (212.01)
$Pr(b = 1) = 0.5$	644.54 (185.61)	594.86 (202.36)	596.59 (197.07)	638.09 (177.26)	638.25 (181.72)	590.98 (201.22)	596.59 (197.07)	638.09 (177.26)	639.53 (180.72)	586.70 (202.62)	596.59 (197.07)	638.09 (177.26)
	$N = 30, N_h = 1500, \Pi = 0.1$											
	$\pi = 0.1$				$\pi = 0.5$				$\pi = 0.9$			
	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR	Optimal	BAR	PW	BR
$Pr(b = 1)$=0.1	988.19 (320.12)	910.60 (362.80)	912.87 (383.98)	955.68 (333.30)	989.32 (317.72)	911.64 (375.87)	912.87 (383.98)	955.68 (333.30)	965.52 (327.54)	900.58 (353.18)	912.87 (383.98)	955.68 (333.30)
$Pr(b = 1) = 0.5$	975.48 (273.80)	887.87 (302.83)	887.53 (318.77)	963.05 (281.88)	967.47 (282.36)	905.76 (301.05)	887.53 (318.77)	963.05 (281.88)	970.64 (285.14)	894.05 (302.35)	887.53 (318.77)	963.05 (281.88)

horizon N_h. However, we still recommend using the best estimate of N_h whenever possible. For example when comparing designs, a better estimate of N_h can provide more accurate information about the difference in expected trial utilities.

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

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

where $Pr^{(n)}(w_{1,j} < w_{2,j}) = Pr(w_{1,j} < w_{2,j}|X_n, \ldots, X_1, G_n, \ldots, G_1)$ is the posterior probability that treatment 2 is the best treatment for patients in subgroup j. With the same priors defined in Section 2, Supplementary Material Section 2 shows the calculation of $r_j^{(n)}$. We follow (Thall and Wathen, 2007) for the choice of the tuning parameter c.

To facilitate the graphical display of the results, we assumed that $\Pi_1 = \Pi_2 = \Pi$ and $\pi_1 = \pi_2 = \pi$. Figure 1 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

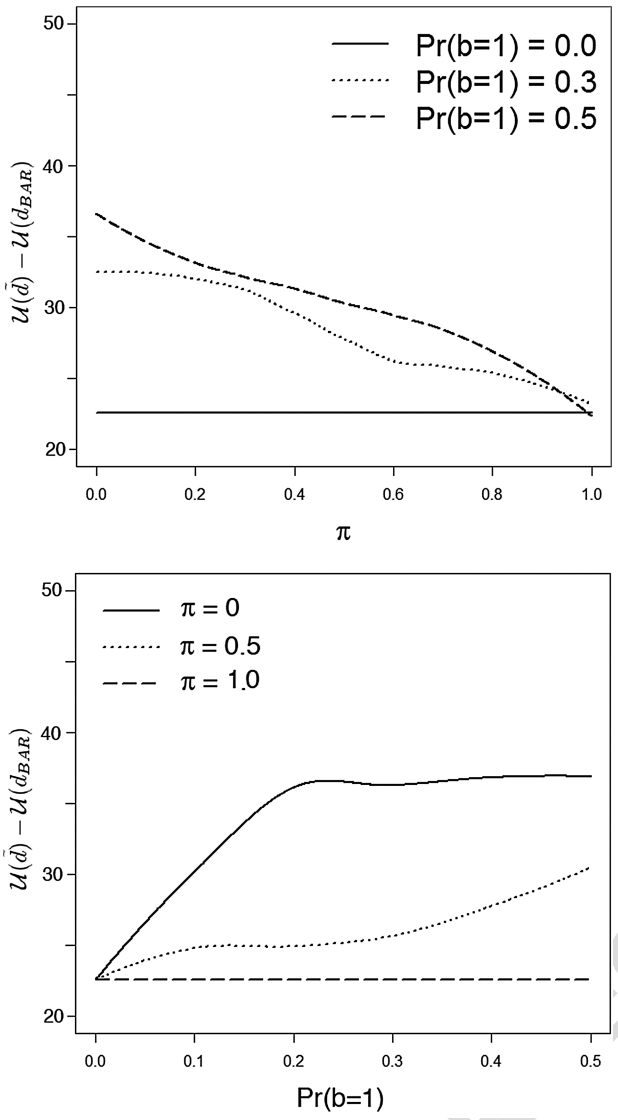

Figure 1. Comparison between the optimal design and Bayesian Adaptive Randomization (BAR). Here, we consider $\pi = \Pi$. Each curve plots the differences between the expected utilities $\mathcal{U}(\tilde{d})$ and $\mathcal{U}(d_{BAR})$. The patients' horizon is $N_h = 1000$ and the trial size is $N = 50$.

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 $Pr(b = 1)$, we use the designs $\tilde{d}$ and d_{BAR} to simulate trials. We then count the numbers of favorable treatment outcomes, during the entire patients' horizon N_h, under each design and compute the mean difference. Each curve in the two panels of Figure 1 is the expected utility lost by using BAR compared to the best achievable result $\mathcal{U}(\tilde{d})$.

In the top panel, the difference between the optimal design and BAR is minimized when the patients' population is homogeneous, that is all patients have the same biomarker profile. The difference in expected utility increases when the subgroups becomes 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 $\pi = 0$, so the prior assigns probability 1 to different efficacy levels across subgroups.

The bottom panel of Figure 1 shows the relationship between the difference in expected utility and the prevalence of mutations. With fixed π the difference $\mathcal{U}(\tilde{d}) - \mathcal{U}(d_{BAR})$ increases when biomarker subgroups become balanced.

3.3. *The Bayesian Adaptive Randomization and the ~~Play the Winner~~ Rule*

We are also interested 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; Yao and Wei, 1996). In this section, we compare the BAR design to the PW in terms of the expected trial utility. Our goal here is to quantify, from a Bayesian standpoint, the change in expected utility associated with using a practical approach such as the PW design compared to BAR.

Each curve in Figure 2 plots the difference in expected trial utility $\mathcal{U}(d_{BAR}) - \mathcal{U}(d_{PW})$. The three panels show the scenarios where treatment effects $w_{i,j}$ are generated from a two components mixture with weights Π and $(1 - \Pi)$. In this simulation study we use $\Pi \neq \pi$, that is the simulation model differs from the Bayesian model used to determine patients' allocations. 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 the cases that we explored. The difference in expected utility ranges from 20 to 50 additional successfully treated patients out of a patient horizon of size 1000. When the subgroups are highly imbalanced, the difference is relatively robust to the choice of π. PW has better utility than BAR when the prior rules out a biomarker effect ($\pi = 1$), but $\Pi = 0$, the outcomes are simulated from varying response probabilities across subgroups, and the subgroups are relatively balanced.

3.4. *Optimal Design and Balanced Randomization*

In closing, we address the comparison between the optimal adaptive design and a design that does not adapt: balanced randomization (BR), with balanced allocation within each biomarker subgroup. In Table 1 and Web Appendix (Section 5), we present the expected utility of marker-stratified balanced randomized trials, and contrast balanced randomization with the optimal design. These results show that the relative performances of these two designs vary with the size of the patient horizon and the study size. The utility of the balanced design is nearly optimal with relatively large horizon sizes. The optimal design strives to reach a compromise between exploration, or learning, and successful treatment. With a large patient horizon, the emphasis is on the learning rather than treatment success during the trial. Thus, patients assignment during the trials is mostly driven by the need of identifying the best available treatment. In contrast, with a short horizon, the emphasis shifts towards treatment success within the trial. More generally, when the sizes of the patient horizon and the study sample sizes are comparable, there is a gain in expected trial utility from using the optimal design compared to BR.

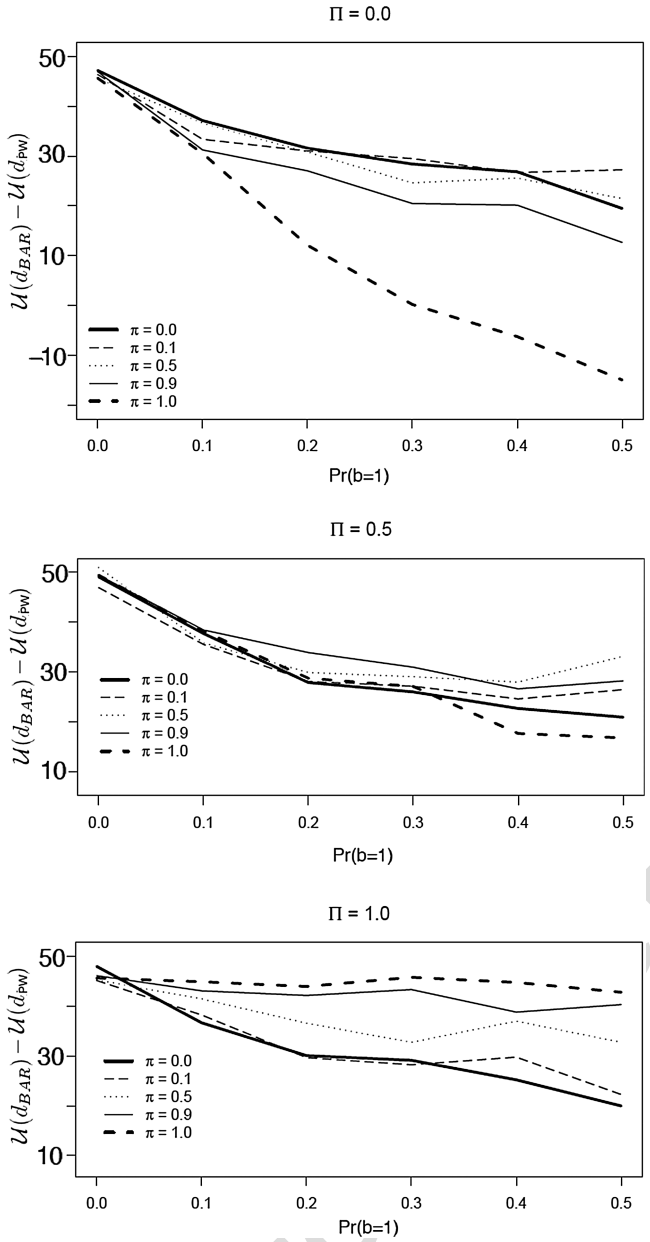

Figure 2. Comparison between $\mathcal{U}(d_{BAR})$ and $\mathcal{U}(d_{PW})$ at different combinations of Π and π. Patient horizon $N_h = 1000$, trial size $N = 50$

.

With a large horizon BR is often a near optimal choice given the ease of implementation, and the power is nearly maximized by BR designs. On the other hand, it is also worth noting a trend towards trials involving relatively small patient horizons. For example, trials involving rare diseases can have relatively short horizons because of the low prevalence. The size of the patient horizon may also decrease as a result of the fragmentations of populations in small subgroups, and redefinition of diseases into sub-categories, which often requires different treatments, based on new diagnostic tests, technologies, and therapies.

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 within 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 benchmarks for the evaluation of trial designs in precision 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 modified to handle other utility functions. We extended previous comparisons of Bayesian adaptive designs and balanced designs (Trippa et al., 2012; Wason and Trippa, 2014) to biomarker-integrated trials. While in this article we focused on two-arm studies, the relative advantages of adaptive and balanced designs will vary with the number of experimental arms in multi-arm studies (Berry, 2011). A larger number of arms generally emphasizes the gains from adaptivity.

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

In practice it is uncommon to implement trials wherein 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 successfully treated 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.

5. Supplementary Materials

Web Appendices , Tables, and Figures referenced in the Sections 2 and 3 are available with this article at the *Biometrics* website on Wiley Online Library along with the R code to implement the optimal design.

Acknowledgements

Computing was performed on the Dana-Farber Cancer Institute high performance cluster and on the Odyssey cluster supported by the FAS Division of Science, Research Com-

puting Group at Harvard University. This project has been supported by grant 5P30 CA006516-50, grant R01 AI024643 and the Claudia Adams Barr Program in Innovative Cancer Research.

References

Alexander, B. M., Wen, P. Y., Trippa, L., Reardon, D. A., Yung, W.-K. A., Parmigiani, G., et al. (2013). Biomarker-based adaptive trials for patients with glioblastoma-lessons from I-SPY 2. *Neuro-Oncology* **15**, 972–978.

Allegra, C. J., Jessup, J. M., Somerfield, M. R., Hamilton, S. R., Hammond, E. H., Hayes, D. F., et al. (2009). 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**, 2091–2096.

Armitage, P. (1985). The search for optimality in clinical trials. *International Statistical Review/Revue Internationale de Statistique* **53**, 15–24.

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

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

Bellman, R. E. (2003). *Dynamic Programming*. ~~Courier~~ Dover Publications.[Q3]

Berry, D. A. (2011). Adaptive clinical trials: The promise and the caution. *Journal of Clinical Oncology* **29**, 606–609.

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

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

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

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

Kim, E. S., Herbst, R. S., Wistuba, I. I., Lee, J. J., Blumenschein, G. R., Tsao, A., et al. (2011). The BATTLE trial: Personalizing therapy for lung cancer. *Cancer Discovery* **1**, 44–53.

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

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

Paez, J. G., Jänne, P. A., Lee, J. C., Tracy, S., Greulich, H., Gabriel, S., et al. (2004). EGFR mutations in lung cancer: Correlation with clinical response to gefitinib therapy. *Science* **304**, 1497–1500.

Parmigiani, G. and Inoue, L. (2009). *Decision Theory: Principles and Approaches*, volume 812. John Wiley & Sons.[Q4]

Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, volume 703. John Wiley & Sons.[Q5]

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

Trippa, L., Lee, E. Q., Wen, P. Y., Batchelor, T. T., Cloughesy, T., Parmigiani, G., et al. (2012). Bayesian adaptive randomized trial design for patients with recurrent glioblastoma. *Journal of Clinical Oncology* **30**, 3258–3263.

Wason, J. and Trippa, L. (2014). A comparison of bayesian adaptive randomization and multi-stage designs for multi-arm clinical trials. *Statistics in Medicine* **33**, 2206–2221.

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

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

Yao, Q. and Wei, L. (1996). Play the winner for phase II/III clinical trials. *Statistics in Medicine* **15**, 2413–2423.

Yin, G. (2013). *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*, volume 876. John Wiley & Sons.[Q6]

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

Received October 2014. *Revised September* 2015. *Accepted September* 2015.

AUTHOR QUERY FORM

JOURNAL: BIOMETRICS

Article: BIOM12437

Dear Author,
During the copyediting of your paper, the following queries arose. Please respond to these by annotating your proofs with the necessary changes/additions using the E-annotation guidelines attached after the last page of this article.

We recommend that you provide additional clarification of answers to queries by entering your answers on the query sheet, in addition to the text mark-up.

Query No.	*Query*	*Remak*
Q1:	Please confirm that given names (red) and surnames/family names (green) have been identified correctly.	confirmed
Q2:	Please check the font used in inline and display equations and confirm that its appearance is correct.	checked and confirmed
Q3:	Please provide publisher location.	Mineola, New York
Q4:	Please provide publisher location.	Chichester, UK
Q5:	Please provide publisher location.	Hoboken, New Jersey
Q6:	Please provide publisher location.	Hobeken, New Jersey

USING e-ANNOTATION TOOLS FOR ELECTRONIC PROOF CORRECTION

WILEY

Required software to e-Annotate PDFs: Adobe Acrobat Professional or Adobe Reader (version 8.0 or above). (Note that this document uses screenshots from Adobe Reader X)
The latest version of Acrobat Reader can be downloaded for free at: http://get.adobe.com/reader/

Once you have Acrobat Reader open on your computer, click on the Comment tab at the right of the toolbar:

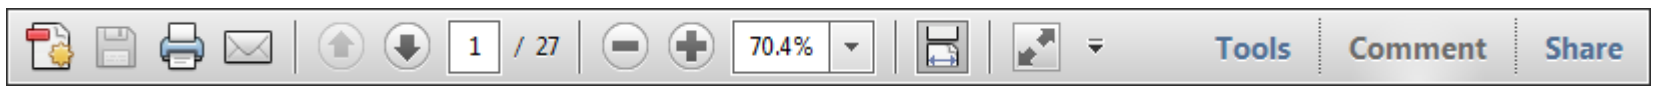

This will open up a panel down the right side of the document. The majority of tools you will use for annotating your proof will be in the Annotations section, pictured opposite. We've picked out some of these tools below:

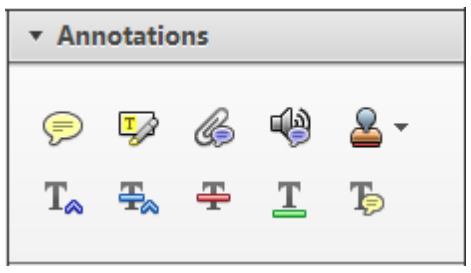

1. Replace (Ins) Tool – for replacing text.

Strikes a line through text and opens up a text box where replacement text can be entered.

How to use it

- Highlight a word or sentence.
- Click on the Replace (Ins) icon in the Annotations section.
- Type the replacement text into the blue box that appears.

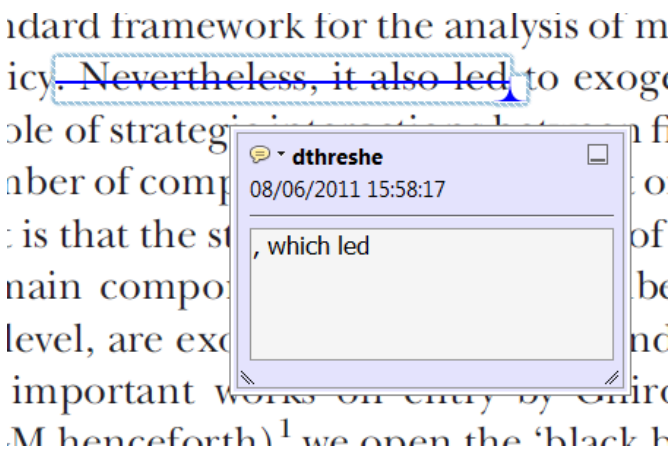

2. Strikethrough (Del) Tool – for deleting text.

Strikes a red line through text that is to be deleted.

How to use it

- Highlight a word or sentence.
- Click on the Strikethrough (Del) icon in the Annotations section.

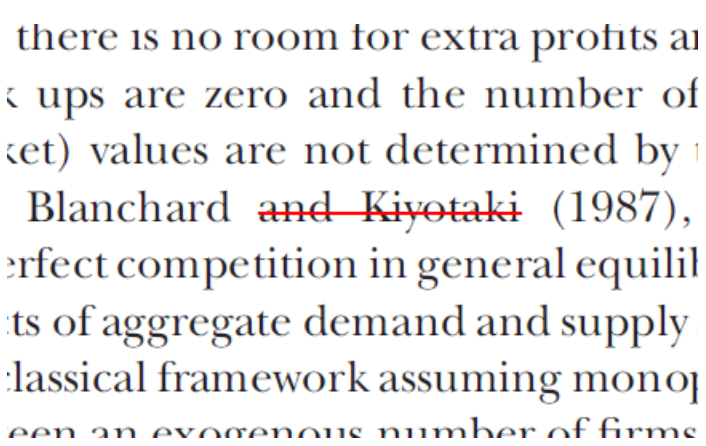

3. Add note to text Tool – for highlighting a section to be changed to bold or italic.

Highlights text in yellow and opens up a text box where comments can be entered.

How to use it

- Highlight the relevant section of text.
- Click on the Add note to text icon in the Annotations section.
- Type instruction on what should be changed regarding the text into the yellow box that appears.

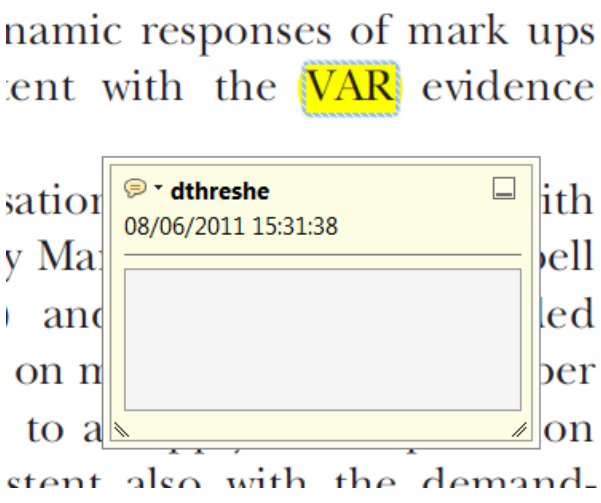

4. Add sticky note Tool – for making notes at specific points in the text.

Marks a point in the proof where a comment needs to be highlighted.

How to use it

- Click on the Add sticky note icon in the Annotations section.
- Click at the point in the proof where the comment should be inserted.
- Type the comment into the yellow box that appears.

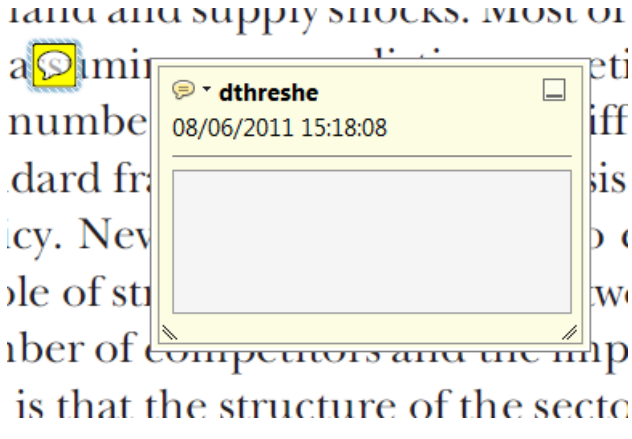

USING e-ANNOTATION TOOLS FOR ELECTRONIC PROOF CORRECTION

5. Attach File Tool – for inserting large amounts of text or replacement figures.

Inserts an icon linking to the attached file in the appropriate pace in the text.

How to use it

- Click on the Attach File icon in the Annotations section.
- Click on the proof to where you'd like the attached file to be linked.
- Select the file to be attached from your computer or network.
- Select the colour and type of icon that will appear in the proof. Click OK.

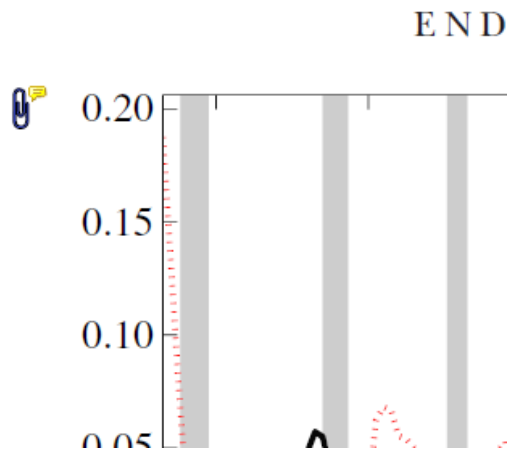

6. Add stamp Tool – for approving a proof if no corrections are required.

Inserts a selected stamp onto an appropriate place in the proof.

How to use it

- Click on the Add stamp icon in the Annotations section.
- Select the stamp you want to use. (The Approved stamp is usually available directly in the menu that appears).
- Click on the proof where you'd like the stamp to appear. (Where a proof is to be approved as it is, this would normally be on the first page).

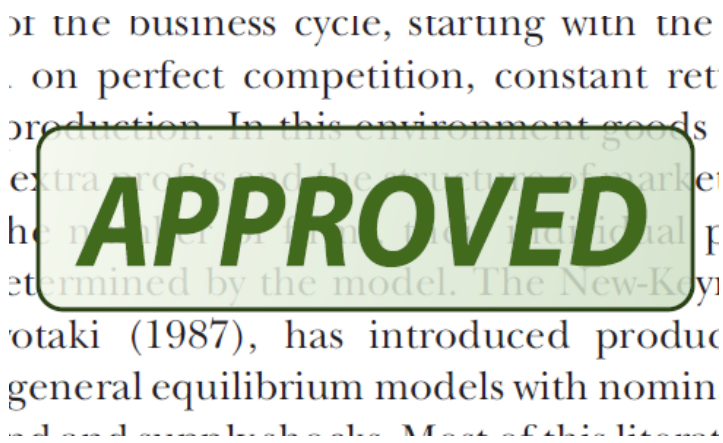

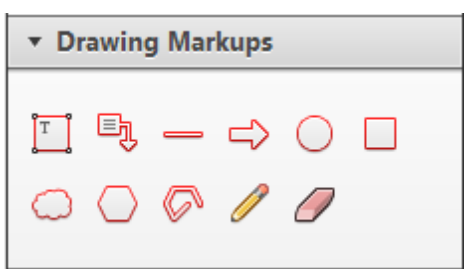

7. Drawing Markups Tools – for drawing shapes, lines and freeform annotations on proofs and commenting on these marks.

Allows shapes, lines and freeform annotations to be drawn on proofs and for comment to be made on these marks..

How to use it

- Click on one of the shapes in the Drawing Markups section.
- Click on the proof at the relevant point and draw the selected shape with the cursor.
- To add a comment to the drawn shape, move the cursor over the shape until an arrowhead appears.
- Double click on the shape and type any text in the red box that appears.

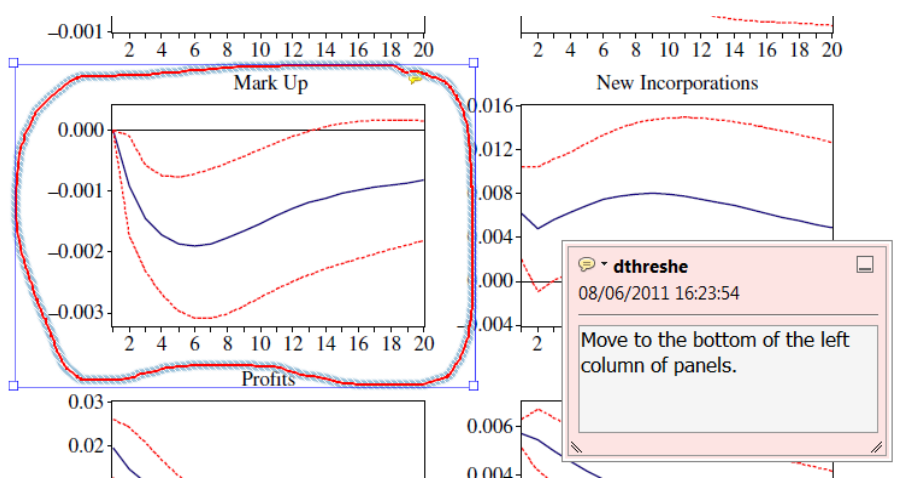

For further information on how to annotate proofs, click on the Help menu to reveal a list of further options:

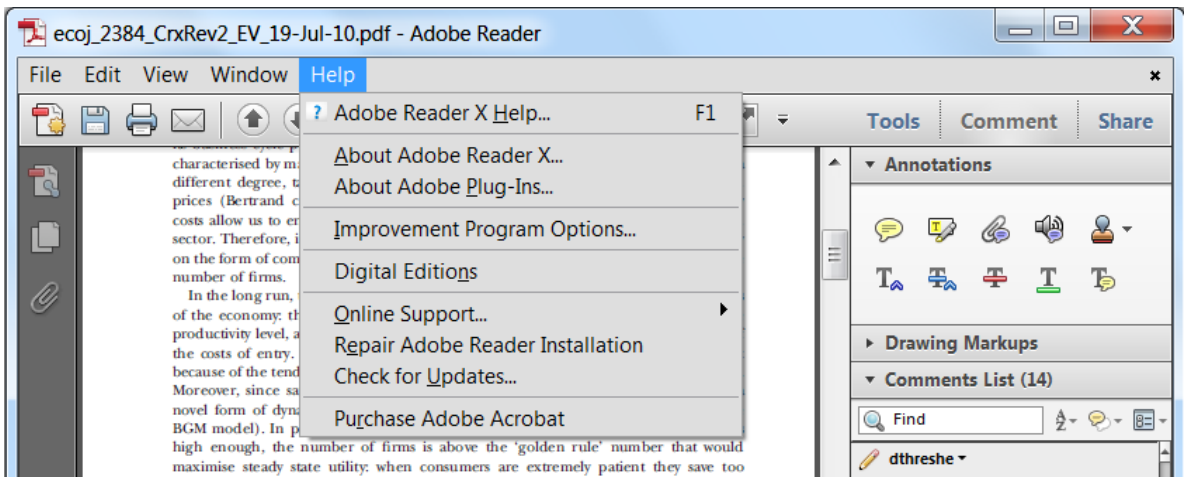

Proof Correction Marks

Please correct and return your proofs using the proof correction marks below. For a more detailed look at using these marks please reference the most recent edition of The Chicago Manual of Style and visit them on the Web at: http://www.chicagomanualofstyle.org/home.html

Instruction to typesetter	*Textual mark*	*Marginal mark*
Leave unchanged	under matter to remain	stet
Insert in text the matter indicated in the margin	∧	∧ followed by new matter
Delete	through single character, rule or underline or through all characters to be deleted	
Substitute character or substitute part of one or more word(s)	through letter or through characters	new character or new characters
Change to italics	under matter to be changed	ital
Change to capitals	under matter to be changed	caps
Change to small capitals	under matter to be changed	sc
Change to bold type	under matter to be changed	bf
Change to bold italic	under matter to be changed	bf+ital
Change to lower case		lc
Insert superscript	∨	∨ under character e.g. ∨
Insert subscript	∧	∧ over character e.g. ∧
Insert full stop	⊙	⊙
Insert comma		
Insert single quotation marks		
Insert double quotation marks		
Insert hyphen		
Start new paragraph		
Transpose		
Close up	linking characters	
Insert or substitute space between characters or words	#	#
Reduce space between characters or words		