

Unbiased estimation for response adaptive clinical trials

Jack Bowden^{1,2} and Lorenzo Trippa³

Statistical Methods in Medical Research
2017, Vol. 26(5) 2376–2388

© The Author(s) 2015



Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0962280215597716

journals.sagepub.com/home/smm



Abstract

Bayesian adaptive trials have the defining feature that the probability of randomization to a particular treatment arm can change as information becomes available as to its true worth. However, there is still a general reluctance to implement such designs in many clinical settings. One area of concern is that their frequentist operating characteristics are poor or, at least, poorly understood. We investigate the bias induced in the maximum likelihood estimate of a response probability parameter, p , for binary outcome by the process of adaptive randomization. We discover that it is small in magnitude and, under mild assumptions, can only be negative – causing one's estimate to be closer to zero on average than the truth. A simple unbiased estimator for p is obtained, but it is shown to have a large mean squared error. Two approaches are therefore explored to improve its precision based on inverse probability weighting and Rao–Blackwellization. We illustrate these estimation strategies using two well-known designs from the literature.

Keywords

Clinical trial, adaptive randomization, bias adjusted estimation, Horvitz-Thompson estimator, inverse probability weighting, Rao-Blackwellization

1 Introduction

From the emergence of clinical trials in medical research in the middle of the 20th century to the present day, a common paradigm has dominated the design and analysis of clinical trials as a means to improving patient care. Usually two treatments are assessed, one being the standard therapy and the other being an experimental treatment, that may offer an improvement in an appropriate patient outcome. The size of the trial is fixed in advance and analysis performed only at its conclusion. Patients entering the trial are randomized with a fixed probability (usually 1/2 for efficiency reasons) to either the experimental or control arm. The advantages of this sort of design are many: it is simple to calculate how many patients are needed to achieve a desired power for rejecting a given null hypothesis; by fixing the timing of the analysis the type I error rate can also be strictly controlled; randomization ensures (asymptotically) that patients on each of the treatment arms will be balanced with respect to all potential known and unknown confounding variables, except the treatment they took. This provides a solid basis for attributing, in a causal sense, the difference in outcomes between the two groups at the end of the trial to the treatment, and for unbiasedly estimating this difference using standard methods.

Over the last 20–30 years, modifications to the standard design have been proposed and accepted as a viable alternative within clinical research, although their use is still the exception rather than the rule. Group sequential designs^{1,2} introduced the idea of multiple analyses at fixed time points during its course, thus enabling early stopping when definitive evidence exists. Adaptive designs^{3–5} have added additional flexibility to the way in which mid-course data can be used to influence a study's future scope. However, the above approaches still share a common characteristic: whilst it is ongoing, a fixed allocation probability is used for all arms still active in the trial. We subsequently refer to this as 'fixed randomization' (FR).

¹MRC Integrative Epidemiology Unit, University of Bristol, Bristol, UK

²MRC Biostatistics Unit, Cambridge, UK

³Dana-Farber Cancer Institute, Boston, USA

Corresponding author:

Jack Bowden, MRC Integrative Epidemiology Unit, University of Bristol, Bristol, UK.

Email: jack.bowden@mrc-bsu.cam.ac.uk

By contrast, the MD Anderson cancer centre (and others) has pioneered the use of an approach to clinical trials incorporating adaptive randomization (AR) as a fundamental component.^{6,7} The basic premise of AR is to use the accumulating evidence on the performance of all treatments to decide how to allocate future patients in the trial.⁸ They are becoming an increasingly popular vehicle for identifying effective patient-specific treatments in the new era of stratified medicine.⁹

There has been a weariness among the medical community to embrace designs incorporating AR, due to concerns over their frequentist properties¹⁰ – such as bias and type I error inflation – their inferiority to a comparable FR designs,¹¹ and even their ethical validity.¹² Indeed, recent Food and Drug Administration (FDA) guidance for the pharmaceutical industry cautions against the use of AR and recommends that operating characteristics of an AR design should be explored and understood before implementation¹³. In this paper we examine the affect that data-adaptive allocation has on the estimate of the true response probability, p , parameterising a binary outcome. We show that the bias induced in the maximum likelihood estimator (MLE) for p by AR is generally small and towards zero. A simple formula is derived that explains this phenomenon and a simple unbiased estimate for p is developed that is closely related to the Horvitz–Thompson (HT) estimator from survey sampling theory.¹⁴ Unfortunately, this estimator can have a large mean squared error (MSE) and even give estimates that lie outside the parameter space. Two strategies are proposed to improve the unbiased estimator's precision. The first makes use of inverse probability of treatment weighting, as is commonly applied in the causal inference literature (see for example Hernan et al.¹⁵). The second approach involves the process of 'Rao–Blackwellization', a procedure commonly applied in sequential and adaptive trials.^{16,17}

In Section 2 we define our notation and review the process of AR from a statistical perspective, in particular, its impact on the likelihood function. In Section 3 we investigate the small sample bias in the MLE induced by AR and compare this to the bias induced by other types of design adaptations. In Section 4 bias adjusted estimation is explored and in Section 5 we apply our bias adjustment strategies to a recent Bayesian adaptive design discussed in Trippa et al.¹⁸ Concluding remarks are made in Section 6 and possible avenues of future research are also discussed.

2 Data-adaptive allocation and the MLE

We assume that a randomized clinical trial is to be conducted on n patients with K treatment arms, the treatments being assessed via a binary response. The outcome for patient i given assignment to treatment k , Y_{ik} , follows a Bernoulli distribution with parameter p_k . Let δ_{ik} be a binary variable equalling 1 if patient i is assigned to treatment k and zero otherwise, and let $\boldsymbol{\delta}_i = (\delta_{i1}, \dots, \delta_{iK})$. Define $\Delta_i \in \{1, \dots, K\}$ as the treatment patient i is actually assigned to, so that $\delta_{i\Delta_i} \equiv 1$. For simplicity, assume that patient responses are immediate, so that patient i 's response is known before patient $i+1$ is recruited. The total information in the trial before patient i is recruited is the $2 \times (i-1)$ matrix:

$$d_{i-1} = \begin{pmatrix} \Delta_1 & \dots & \Delta_{i-1} \\ y_{1,\Delta_1} & \dots & y_{i-1,\Delta_{i-1}} \end{pmatrix} \quad (1)$$

and let d_n be the complete data on treatment assignments and outcomes at the end of the trial. Further let $n(k) = \sum_{i=1}^n \delta_{ik}$ be the number of patients assigned to treatment k and $\sum_{i=1}^n \delta_{ik} y_{ik}$ be the number of responders to treatment k . Note that, whilst $n(k)$ is random, the marginal total, $\sum_{k=1}^K n(k) = n$, is fixed.

In a trial that utilises AR, patient i 's allocation, Δ_i , can be chosen to depend on all or part of d_{i-1} , that is, on all previous patient randomization and outcome data. For illustration, Figure 1 depicts the dependence structure




Figure 1. Illustrative diagram of data on the first three patients entered into a clinical trial using adaptive randomization.

governing the randomization and outcome data for the first three patients recruited into a trial using AR. Here, p_{Δ_i} represents the true response parameter for treatment Δ_i . Since Δ_i is a random variable, so is p_{Δ_i} . From Figure 1, we see that patient two’s randomization (Δ_2) is allowed to depend on patient one’s outcome y_{1,Δ_1} and patient one’s randomization, Δ_1 – both directly and indirectly (through p_{Δ_1}). However, once $d_1 = (\Delta_1, y_{1,\Delta_1})$ has influenced the value of Δ_2 , y_{2,Δ_2} is conditionally independent of d_1 , given Δ_2 . This statement clearly generalizes to the randomization and outcome of the i th patient given d_{i-1} . We can therefore factorize the likelihood function for the parameter vector $\mathbf{p} = (p_1, \dots, p_K)$ at the end of the trial as

$$\begin{aligned} L(\mathbf{p}|d_n) &= \prod_{i=1}^n f_{\mathbf{p}}(y_{i,\Delta_i}|\Delta_i) f(\Delta_i|\Delta_1, y_{1,\Delta_1}, \Delta_2, y_{2,\Delta_2}, \dots, \Delta_{i-1}, y_{i-1,\Delta_{i-1}}) \\ &= \prod_{i=1}^n f_{\mathbf{p}}(y_{i,\Delta_i}|\Delta_i) f(\Delta_i|d_{i-1}) \\ &= \prod_{i=1}^n f_{\mathbf{p}}(y_{i,\Delta_i}|\Delta_i) \pi_{i,\Delta_i} \end{aligned} \tag{2}$$

The first component is the probability under \mathbf{p} of observing patient i ’s outcome given randomization to treatment Δ_i . The second component, π_{i,Δ_i} , is the probability that patient i was randomized to treatment Δ_i conditional on previous data $d_{i-1} = (\Delta_1, Y_{1,\Delta_1}, \Delta_2, Y_{2,\Delta_2}, \dots, \Delta_{i-1}, Y_{i-1,\Delta_{i-1}})$. Since this is a known function of the data it does not depend on the parameter vector \mathbf{p} . From standard likelihood theory (see for example Boos and Stefanski¹⁹), π_{i,Δ_i} can be removed from equation (2) when calculating the maximum likelihood estimate (MLE). This fact is obvious in a standard trial design using FR, where all n patients are randomized with equal probability to one of K arms. In this case there is clearly no information about the parameters in π_{i,Δ_i} – it takes the constant value $1/K$. However, as we have shown above, the same is true within the AR framework described above. Removing π_{i,Δ_i} from the likelihood, and using the fact that the patient outcomes are independently Bernoulli distributed, expression (2) reduces to

$$L(\mathbf{p}|d_n) \propto \prod_{k=1}^K p_k^{\sum_{i=1}^n \delta_{ik} y_{ik}} (1 - p_k)^{\sum_{i=1}^n \delta_{ik} (1 - y_{ik})}$$

and the MLE for p_k , $k = 1, \dots, K$ is simply

$$\hat{p}_k = \frac{\sum_{i=1}^n \delta_{ik} y_{ik}}{n(k)} \tag{3}$$

Melfi and Page²⁰ clarified that, in order for the MLE in equation (3) to be strongly consistent, there must exist a non-zero probability of being allocated to treatment k during the trial, as the trial size tends to infinity.

2.1 Case study: The Randomized Play the Winner rule

Wei and Durham²¹ proposed a randomized extension to Zelen’s original deterministic ‘play-the-winner’ allocation rule²² for a two arm clinical trial, which serves as a useful toy example. The basic design works as follows: a hypothetical urn is imagined with one ball labelled ‘ T_1 ’ and another labelled ‘ T_2 ’, representing treatments one and two, respectively. A ball of type T_k is picked at random from the urn, patient 1 is assigned to T_k and the ball is returned to the urn. If patient 1 subsequently experiences a success, then another ball of type T_k is added to the urn. However, if the patient experiences a failure, a ball of the alternative type is added instead. Thus, when a ball is picked at random from the urn to decide on the allocation of patient 2, they either have a $1/3$ or a $2/3$ chance of receiving treatment k . This procedure continues up until the planned end of the study and skews the treatment allocation towards the treatment that is performing the best. The probability that patient i is assigned to treatment T_1 (so that $\Delta_i = 1$) is simply the proportion of T_1 balls in the urn after $i - 1$ patients, and is given by

$$\pi_{i,1} = \frac{1 + \sum_{j=1}^{i-1} \delta_{j1} y_{j1} + \delta_{j2} (1 - y_{j2})}{i + 1}$$

The probability of randomization to T_2 , $\pi_{i,2} = 1 - \pi_{i,1}$. We will refer to this as the ‘randomized play-the-winner’ (RPW) design. When we implement the RPW design in later sections, we introduce a small modification whereby the first two patients are always split between arm 1 and 2. This guarantees that the estimates for p_1 and p_2 will always be defined but does not materially change its operating characteristics otherwise.

3 The bias of the MLE

Whilst AR leaves the asymptotic properties of the MLE intact, this does not give any indication as to its finite sample properties in a real trial context, in particular its bias. Without loss of generality, consider estimation of a single treatment’s parameter p_k . We start by noting that

$$E[n(k)\hat{p}_k + (n - n(k))p_k] = E[n(k)\hat{p}_k] + E[(n - n(k))p_k] \quad (4)$$

$$= E[n(k)(\hat{p}_k - p_k)] + np_k \quad (5)$$

$$= np_k$$

since

$$E[n(k)(\hat{p}_k - p_k)] = E\left[n(k)\left(\frac{\sum_{i=1}^n \delta_{ik}Y_{ik}}{n(k)} - p_k\right)\right] = 0$$

From the standard definition for the covariance of two random variables, and making use of equations (4) and (5), we can write

$$\text{Cov}[n(k), \hat{p}_k] + E[n(k)]E[\hat{p}_k] + E[(n - n(k))p_k] = np_k$$

Further rearrangement and cancellation yield

$$\frac{\text{Cov}[n(k), \hat{p}_k]}{E[n(k)]} = p_k - E[\hat{p}_k] = -\text{Bias}(\hat{p}_k) \quad (6)$$

Equation (6) gives a simple characterisation of the bias. It is zero when $n(k)$ is independent of \hat{p}_k and increases in magnitude as they become more dependent. The bias is a decreasing function of $E[n(k)]$ and will be opposite to the sign of their covariance. FR guarantees a zero covariance (and hence zero bias) since $E[n(k)] = 1/K$ is a constant. However, in the common case where AR is used to direct more patients towards treatments that appear to work well, $n(k)$ will be non-constant and positively correlated with \hat{p}_k .

The bias induced by AR can be explained heuristically. Suppose that, in the early part of a clinical trial using AR, a treatment’s true effect is overestimated because an unusually high number of patients experience a positive response. In this case, AR will assign more patients to receive the treatment in the latter stages of the trial. The ‘random high’ in the treatment effect estimate then has the scope to regress back down towards its true value, the final MLE. However, suppose instead that in the early part of a trial a treatment’s effect is underestimated by chance. In this case fewer patients are assigned to receive the treatment in the latter stages and the treatment effect estimate does not have the same ability to regress back up towards its true value. This asymmetry creates a positive covariance between the MLE and its sample size and hence a negative bias.

The black dots in Figure 2 show the bias induced in the response probability estimates for \hat{p}_k as a function of $\text{Cov}(n(k), \hat{p}_k)$ under the RPW design with $n = 25$ patients. Each dot represents a different parameter constellation in the region $(p_1, p_2) \in (0.2, 0.8)$. As equation (6) suggests, positive covariance implies negative bias in the MLE, and vice versa.

3.1 Further examples

Bauer et al.²³ noted that a similar mechanism acts to induce a negative bias in a traditional sequential trial with early stopping. For example, consider a Simon’s two-stage single arm trial measuring a treatment’s response with respect to a binary outcome.²⁴ Assume that 13 patients are initially allocated a treatment, with the trial continuing to full enrolment (34 patients) only if four or more patients are ‘responders’ at stage 1. If there are nine or more

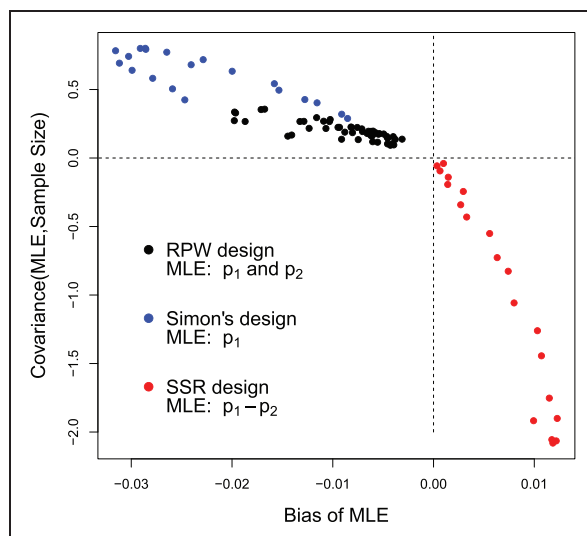


Figure 2. Bias of the MLE versus its covariance with the overall sample size for under three different trial designs. RPW: randomized play the winner, SSR: sample size re-estimation.

responders at the end of the trial the null hypothesis $H_0 : p \leq 0.2$ can be rejected with a type I error rate of 10%.²⁵ The number of patients allocated to treatment 1, $n(1)$, is therefore a random variable:

$$n(1) = \begin{cases} 13 & \text{if } \sum_{i=1}^{13} y_{i1} \leq 3 \\ 34 & \text{Otherwise} \end{cases}$$

The blue dots in Figure 2 show the bias in the estimates for the treatment’s response probability, p_1 , as the true value is varied between 0.2 and 0.4. Again, when large estimates go hand-in-hand with large sample sizes then a negative bias necessarily follows.

A much discussed adaptive design strategy exists that can induce exactly the opposite effect – that is a negative covariance between the MLE and the sample size it is based on (and therefore a positive bias). Consider a two-arm, two-stage trial that incorporates unblinded sample size re-estimation.^{4,26,27} In stage 1, 100 patients are initially allocated to each arm (using FR) and the MLE for the response parameters are estimated as \hat{p}_1 and \hat{p}_2 . From Li et al.,²⁶ we calculate the estimated standardised mean difference $\hat{\theta}$ and stage 1 test statistic z_1 via the variance stabilising transformation:

$$\hat{\theta} = 2[\sin^{-1}(\sqrt{\hat{p}_1}) - \sin^{-1}(\sqrt{\hat{p}_2})], \quad z_1 = \hat{\theta} \sqrt{\frac{100}{2}}$$

Suppose that if z_1 is less than 1 or greater than 2.74 then the trial is stopped for efficacy or futility. However, if $z_1 \in (1, 2.74)$ then further patients are recruited in order to give an 80% probability of rejecting the null hypothesis $H_0 : \theta \leq 0$ at the alternative $\theta = \hat{\theta}$ with an unconditional type I error rate of 5%. This is similar to design 1 explored in Bowden and Mander.²⁷ Making use of the Ceiling function ‘[.]’, the number of patients allocated to each arm ($n(1)$ and $n(2)$) is again a random variable:

$$n(1) = n(2) = \begin{cases} 100 & \text{if } z_1 \notin (1, 2.74) \\ \left\lceil \left[\frac{(1.92 + 0.84)^2}{z_1^2} - 1 \right] 100 \right\rceil & \text{Otherwise} \end{cases}$$

The red dots in Figure 2 show the bias induced in the MLE for $p_1 - p_2$, $\hat{p}_1 - \hat{p}_2$, versus its covariance with the second stage sample size when $p_1 \in (0.45, 0.65)$ and p_2 is fixed at 0.3. Since small interim estimates of $p_1 - p_2$ (with a corresponding z_1 close to 1) are associated with large overall sample sizes, and large interim estimates (with a

corresponding z_1 close to 2.74) are associated with small sample sizes, the aforementioned covariance is negative and the bias is therefore positive.

4 Bias adjusted estimation

4.1 Simple bias adjustment

A simple bias-corrected estimate for p_k exists that utilises π_{i,Δ_i} (when $\Delta_i = k$), namely

$$\hat{p}_{k,HT} = \frac{1}{n} \sum_{i=1}^n \frac{\delta_{ik} Y_{ik}}{\pi_{i,k}} \quad (7)$$

$\hat{p}_{k,HT}$ bears strong resemblance to the HT estimator used extensively in survey methodology to correct for stratified sampling.¹⁴ It is not strictly identical because the $\pi_{i,k}$'s are not independent in equation (7). However, for ease we will refer to it as the HT estimator. We can see that it is indeed unbiased, because

$$E\left(\frac{1}{n} \sum_{i=1}^n \frac{\delta_{ik} Y_{ik}}{\pi_{i,k}}\right) = \frac{1}{n} \sum_{i=1}^n E\left[E\left(\frac{\delta_{ik} Y_{ik}}{\pi_{i,k}} \mid d_{i-1}\right)\right]$$

Since δ_{ik} is independent of Y_{ik} given d_{i-1} this can be written as

$$\frac{1}{n} \sum_{i=1}^n E\left[Y_{ik} \left(\frac{E[\delta_{ik} \mid d_{i-1}]}{\pi_{i,k}}\right)\right] = \frac{1}{n} \sum_{i=1}^n E[Y_{ik}] = p_k$$

Although unbiased, $\hat{p}_{k,HT}$ can have a large variance, in part because it can be larger than 1, which is clearly nonsensical. Of course, if we crudely constrain it to be less than 1 its unbiasedness is not maintained. We are therefore interested in alternative estimation strategies that improve upon $\hat{p}_{k,HT}$, by naturally shrinking it to be within the unit interval and reducing its variance. A simple estimator that achieves this is a 'normalised' version of $\hat{p}_{k,HT}$:

$$\hat{p}_{k,IPW} = \frac{n \hat{p}_{k,HT}}{\sum_{i=1}^n \frac{\delta_{ik}}{\pi_{i,k}}} \quad (8)$$

This can be thought of as the inverse probability (of treatment) weighted (IPW) estimate that is commonly used for causal inference of observational data¹⁵ and is the solution to the estimating equation

$$\sum_{i=1}^n \frac{\delta_{ik}}{\pi_{i,k}} (y_{ik} - p_k) = 0$$

It is not surprising that the IPW estimate appears a perfect fit for this context, given the similarity of the dependence structure present in trial data d_n (and shown in Figure 1) to the phenomenon of time-varying confounding. The IPW estimate (8) is not unbiased for finite sized trials but is at least constrained to be in (0,1). Since the π_{i,Δ_i} terms are known exactly (one could not implement AR without them) the HT and IPW estimators are as trivial to calculate as the MLE.

Figure 3 (left) shows the bias of the MLE and IPW estimators for the RPW design with $n = 25$ patients. The bias of the HT estimate is not shown because it is zero (our simulations confirm this). Six different parameter combinations for $\mathbf{p} = (p_1, p_2)$ are considered – see the vertical columns in Figure 3 (left) – where the values of p_1 and p_2 and the bias in their estimates are denoted by circles and triangles, respectively. The results are the average of 50,000 simulations. Under RPW allocation, the MLE for \hat{p}_k is always negatively biased. The bias is largest for the treatment with the smallest true effect size, and grows as the difference between the best and worst treatment increases. Although only the HT estimator is unbiased, the bias of the IPW estimator is essentially negligible for scenarios 1–4. However, for scenarios 5 and 6, a small negative bias is present. Figure 3 (right) compares the MSE of the MLE, IPW and HT estimators across the six scenarios (the RBHT estimator is discussed in the next section). Across all scenarios, the IPW estimator is shown to have a comparable MSE to the MLE. Predictably, the HT estimator has the largest MSE.




Figure 3. Left: bias of the MLE and IPW estimators. Right: MSE of the MLE, IPW, HT and RBHT estimators.

4.2 Rao–Blackwellization

A second option for improving the HT estimator is to use Rao–Blackwellization. Define the sufficient statistic, S_n , for the parameter vector \mathbf{p} to be the number of patients assigned to each arm, $n(k)$, and the number of responders in each arm, $\sum_{i=1}^n \delta_{ik}y_{ik}$, for $k = 1, \dots, K$. Let the value of the sufficient statistic for the observed data be denoted by s_n . Now define the set Ω^* containing all possible trial realisations with sufficient statistics equal to s_n , that is:

$$\Omega^* = \{d_n^* : n^*(k) = n(k), \sum_{i=1}^n \delta_{ik}^* y_{ik}^* = \sum_{i=1}^n \delta_{ik} y_{ik}, k = 1, \dots, K\}$$

To generate elements of Ω^* we refer back to the data matrix in formula (1) and note that any permutation of the order of columns constituting matrix d_n leaves s_n unchanged and will therefore be in Ω^* . However, at the same time, column permutations are the only transformation of d_n allowed, since materially changing any individual element of d_n will alter s_n . In order to compute the Rao–Blackwellized Horvitz–Thompson (RBHT) estimator $E(\hat{\mathbf{p}}_{HT|s_n})$ we need to integrate the HT estimate over Ω^* as follows:

$$\hat{\mathbf{p}}_{RBHT} = E(\hat{\mathbf{p}}_{HT|s_n}) = \frac{\sum_{\Omega^*} \hat{\mathbf{p}}_{HT}(d_n^*) L(d_n^*; \mathbf{p})}{\sum_{\Omega^*} L(d_n^*; \mathbf{p})} \tag{9}$$

where $L(d_n^*; \mathbf{p})$ is the likelihood of observing data d_n^* (a permutation of d_n), under the true unknown parameters \mathbf{p} , as given in equation (2). The change of notation is deliberate, and for the reasons explained below. Since the cardinality of Ω^* increases rapidly with the overall sample size, direct integration becomes unfeasible in most realistic settings and a Monte Carlo approximation is necessary. In order to achieve this end, we first note that ratio $\frac{L(d_n^*; \mathbf{p})}{\sum_{\Omega^*} L(d_n^*; \mathbf{p})}$ in equation (9) remains constant across all values of \mathbf{p} . To see this, we rewrite $L(d_n^*; \mathbf{p})$ as

$$\prod_{i=1}^n f_p(y_{i,\Delta_i} | \Delta_i) \prod_{i=1}^n \pi_{i,\Delta_i}^* \tag{10}$$

The first term is the likelihood of the data given treatment assignment; this product is constant for any permutation of d_n , d_n^* , in Ω^* (it does not matter what order the n terms are multiplied together). However, the second component of the likelihood does change under permutation, and so we write π_{i,Δ_i}^* to denote the probability that the i th patient (in the permuted data set d_n^* , not the original data set d_n) would have been randomized to treatment Δ_i in d_n^* . We therefore have that

$$L(d_n^*; \mathbf{p}) \propto \prod_{i=1}^n \pi_{i,\Delta_i}^* \tag{11}$$

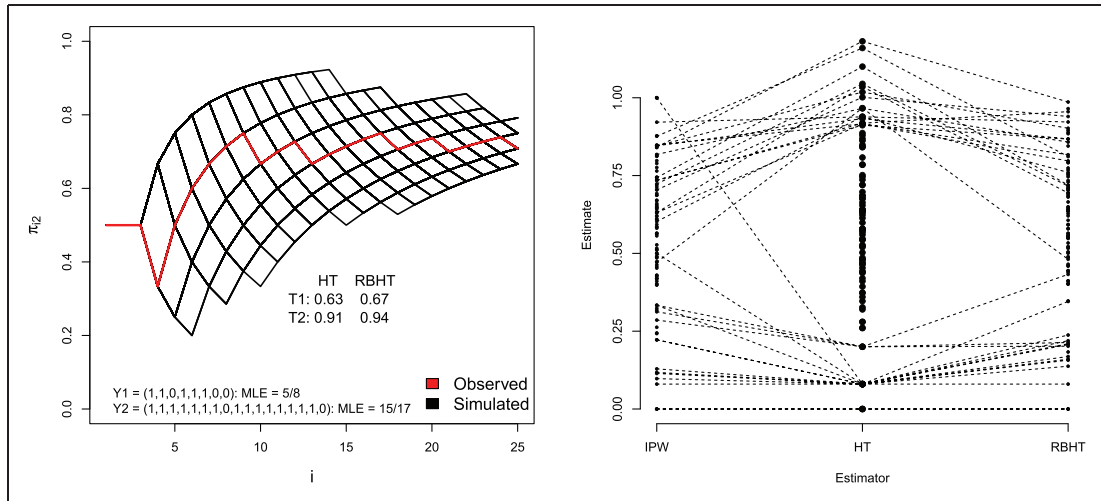


Figure 4. Left: π_{12} for single trial realisation under Scenario 6. Right: IPW, HT and RBHT estimates for p_1 for 100 trial realisations under Scenario 6. Trial simulations giving rise to HT estimates above 0.8 and below 0.2 are linked to their respective IPW and RBHT estimates by dotted lines, in order to highlight the additional benefit that inverse probability weighting and Rao–Blackwellization provides.

Given uniformly sampled trial realisations $d_n^*(j); j = 1, \dots, M$, we can therefore approximate expectation (9) as:

$$\hat{p}_{RBHT} \approx \frac{\sum_{j=1}^M \hat{p}_{HT}(d_n^*(j)) \Pi^*(j)}{\sum_{j=1}^M \Pi^*(j)} \tag{12}$$

where $\Pi^*(j) = \prod_{i=1}^n \pi_{i,\Delta_i}^*$ under $d_n^*(j)$. Although this looks a simple and attractive procedure, most random permutations of the data will have an extremely small likelihood of occurrence, and equation (9) will generally be dominated by a small number of trial realisations clustered around the observed value, d_n . An alternative Monte-Carlo approach consistent with this fact is to sample trial realisations (d_n^* 's) from the conditional distribution of d_n given s_n and obtain an approximation to equation (9) via an unweighted average:

$$\hat{p}_{RBHT} \approx \frac{\sum_{j=1}^M \hat{p}_{HT}(d_n^*(j))}{M} \tag{13}$$

For the RPW case study in Section 2.1, which contains only two arms and 25 patients, one can simply generate trial realisations from $(d_{25}|s_{25})$ by drawing d_{25} 's unconditionally, and saving those with sufficient statistics equal to s_{25} . A simple way to do this is to specify a value for the parameter vector \mathbf{p} and, whilst any choice is valid, it is both natural and efficient to use the MLE. For illustration we consider a single trial realisation of the RPW design, simulated under scenario 6. In this case eight patients are randomized to treatment 1 (with 5 responders) and 17 to treatment 2 (with 15 responders), which defines s_{25} . The red line in Figure 4 (left) shows the value of π_{12} for $i = 1, \dots, 25$. It starts at 1/2 (by design) and after 10 patients or so it stabilises at around 70%.

The MLE $\hat{\mathbf{p}} = (5/8, 15/17)$. The grid of black lines in Figure 4 (left) shows all realised values of $\pi_{i,2}$ over 500 simulated trials, $d_{25}^*(1), \dots, d_{25}^*(500)$. The HT and RBHT estimates for treatment 1 and 2 are also shown, the latter being the average of the HT estimates obtained from the 500 simulated trials. Figure 4 (right) shows the IPW, HT and RBHT estimates for p_1 , obtained across 100 trial realisations, again under scenario 6 ($p_1 = 0.6$). Values above 0.8 and below 0.2 are linked to highlight that the IPW and RBHT estimators both constrain the HT estimate to be within (0,1) and shrink its variance. Figure 3 (right) in Section 4.1 shows RBHT estimator's MSE compared to the other estimators for the RPW design. Its MSE is normally half way between that of the HT and IPW estimators. In summary, for the RPW design, the simple IPW estimator performs perfectly adequately, and there appears to be nothing gained in applying a more complex form of bias correction.

4.2.1 A Metropolis–Hastings algorithm for the RBHT estimator

As the number of patients and number of arms in a trial increase, generating Monte Carlo draws from the conditional distribution $(d_n|s_n)$ by specifying a value for \mathbf{p} also becomes infeasible, since the probability a random draw satisfying the condition is too small. We therefore propose a Monte Carlo Markov Chain approach to calculate the RBHT estimate. Now we need to construct an irreducible Markov Chain $d_n^*(j); j = 1, \dots, M$, over the space of the d_n^* sequences with stationary distribution $(d_n|s_n)$. If possible, standard ergodicity arguments guarantee that the above Monte Carlo average of the $d_n^*(j)$'s in equation (13) converges with probability 1 to $\widehat{\mathbf{p}}_{RBHT}$ as M increases. Our construction of the Markov Chain is a direct application of the Metropolis–Hastings algorithm. We start from a $d_n^*(1)$ consistent with s_n (for example $d_n^*(1) = d_n$) then at every $j = 1, \dots, M - 1$, a transition $d_n^*(j) \rightarrow d_n^*(j + 1)$ is proposed. $d_n^*(j + 1)$ is created by selecting randomly two integers (i_1, i_2) in $\{1, \dots, n\}$, and swapping the positions of the i_1 th and i_2 th columns of the $2 \times n$ matrix $d_n^*(j)$. At each step the usual Metropolis–Hastings rule accepts or rejects $d_n^*(j + 1)$. The application of the Metropolis–Hastings rule requires one only to compute the probabilities on the right side of expression (11) and generates a Markov Chain with the desired stationary distribution. In the next section we apply the Metropolis–Hastings implementation of the RBHT estimator to a recent trial example.

5 An adaptive design for glioblastoma

Trippa et al.¹⁸ have recently proposed a Bayesian adaptive design for testing multiple experimental treatments in a controlled trial setting for patients with recurrent glioblastoma. Their motivation was to find a design that required fewer patients to identify an effective treatment compared to a standard multi-arm trial with equal randomization. Although the original cancer setting determined the use of a time-to-event outcome, the design approach can easily be transferred to the binary data setting. Indeed this has been the modus operandi for others to evaluate its operating characteristics^{28–30} and in the case of the latter, to offer a strong critique against its use. Let $Pr(p_k - p_0 > 0 | d_{i-1})$ be the posterior probability that treatment k ($k = 1, \dots, K$) provides a higher chance of success than the control (treatment $k = 0$), given all of the available data from patients $1, \dots, i - 1$. The probability that patient i is randomized to treatment K under this scheme is¹⁸

$$\pi_{i,\Delta_i} = \begin{cases} \frac{Pr(p_k - p_0 > 0 | d_{i-1})^{\gamma_i}}{\sum_{j=1}^K Pr(p_j - p_0 > 0 | d_{i-1})^{\gamma_j}} & \text{for } k \text{ in } 1, \dots, K \\ \frac{1}{K} \exp\{\max(n_{i-1,k}, k = 1, \dots, K) - n_{i-1,0}\}^{\eta_i} & \text{for } k = 0 \end{cases}$$

The term $n_{i-1,k}$ represents the number of patients who have been assigned to treatment arm k after $i - 1$ patients have been recruited into the trial. The γ_i and η_i terms are positive increasing functions of i which can be ‘tuned’ to fix the desired characteristics of the trial. The total sample size of the trial, n , is fixed as before. We refer the reader to Trippa et al.^{18,31,32} for the interpretation of the tuning parameters (γ_i, η_i) and extensions with composite outcomes models, including progression free survival and overall survival times. This AR procedure achieves two goals: To allocate more patients to the experimental arm which is performing the best (for the benefit of patients in the trial); to allocate a similar number of patients to the control arm as the best experimental arm (to protect the trial’s overall power).

Following Friedlin and Korn³⁰ and Trippa et al.,²⁹ 5000 trials of $n = 140$ patients are simulated under this design with three active arms and one control arm. Uniform beta(1,1) priors were used to initialise the Bayesian analysis, so that the posterior distributions for each p_k were also beta distributed. The allocation probabilities can then be calculated using a simple Monte Carlo approximation. The function’s γ_i and η_i were chosen to be $10(i/n)^{1.75}$ and $0.25(i/n)$, respectively. Figure 5 shows the bias and MSE of the MLE and IPW estimators under the five parameter constellations considered in Friedlin and Korn³⁰ and Trippa et al.²⁹ The RBHT estimate was calculated using the Metropolis–Hastings algorithm described in Section 4.2.1. We show its MSE only since it is unbiased (again, simulations confirm this) and plotting its estimated bias at zero obscures the other results.

By design, there is very little bias in the standard MLE for the control group across all scenarios. The MLE for the experimental treatments is, however, consistently negatively biased. The IPW estimate reduces the bias in the MLE, but its performance is not uniformly as impressive as for the RPW(1,1) design, especially for treatments with relatively small effect sizes. We suspect this is because the probability of randomization to a poorly performing treatment may approach zero, which makes the IPW estimator naturally unstable. Figure 6




Figure 5. Bias (left) and MSE (right) for the glioblastoma case study.

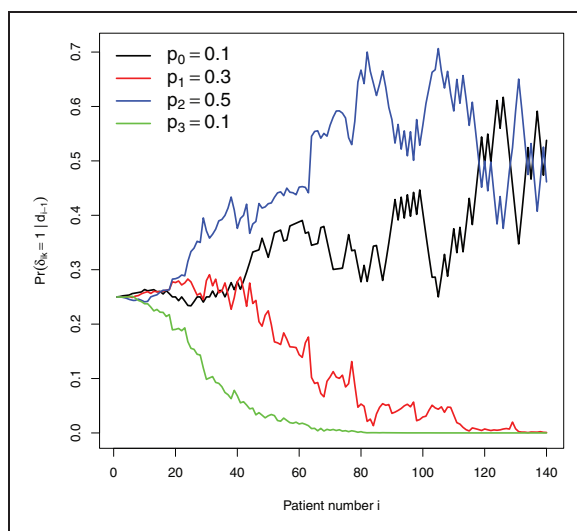


Figure 6. Randomization probabilities for a single trial realisation under parameter constellation 5.

illustrates this point for single trial realisation under scenario 5, where treatment 3 is significantly worse than the other experimental treatments. After 60 patients have been recruited the probability of allocation to arm three is essentially zero. Consequently, the IPW estimator generally has a larger MSE than the MLE, as seen in Figure 5 (right). The RBHT estimator perfectly removes the bias across all scenarios (results not shown). More importantly, it also has a smaller MSE than either the MLE or IPW estimators, making it attractive for this design.

6 Discussion

In this paper we have considered the effect that AR has on the estimation of a treatment’s true response rate within the context of a clinical trial. We have clarified that, whilst the simple proportion of total responders is still the MLE under AR, it will be biased in small samples (but not seriously so). We restricted our attention to the bias in each treatment parameter p_k and not, for example, on treatment control response differences $p_k - p_0$. This was to explain the issue of bias in its most clear and general terms. Furthermore, we do not consider the issue of hypothesis testing, for which such differences would naturally form a basis. In the glioblastoma design, since the control group’s allocation probability was not dependent on its patient response rate, the MLE \hat{p}_0 was essentially unbiased and therefore all of the bias in the difference $\hat{p}_k - \hat{p}_0$ is driven by \hat{p}_k . However, when the

control group is not treated differently and all arms in a trial are subject to the same AR scheme (as in the RPW example) the bias of the difference $\hat{p}_k - \hat{p}_j$ can be either positive or negative, but it will generally be even smaller than the bias for \hat{p}_k or \hat{p}_j . This is because some of the negative bias will cancel out.

The HT estimator was shown to perfectly adjust for the bias in p_k , but was not attractive to use in practice. Inverse probability weighting – commonly used in the causal inference and missing data literature – offered a simple means to improve the HT estimator. Its performance was poor, however, if the allocation probability weights were allowed to approach zero. IPW estimators can be improved via the use of ‘stabilised weights’,¹⁵ and it would be interesting to explore whether their use in this context.

Rao–Blackwellization was shown to offer the most comprehensive (but computationally intensive) means to improve the HT estimator. If S_n is a complete and sufficient statistic for p then \hat{p}_{RBHT} could claim to be the uniform minimum variance unbiased estimate (UMVUE) for p . This technique has been suggested in the analysis of group sequential trials by Emerson and Fleming.¹⁶ In their case, the MLE for p_k at the point of the first interim analysis (i.e. the first chance to stop the trial) is Rao–Blackwellized, since only information collected up to this point is unbiased. In our case, we are using a weighted average of the entire data. However, due to concerns about verifying the completeness of S_n , we simply refer to it as a ‘Rao–Blackwellized’ estimator. By the Rao–Blackwell theorem, the RBHT estimate will be unbiased with a smaller variance than the HT estimator. The treatment control difference, $p_k - p_0$, is likely to be the primary outcome measure in a clinical trial. Since the RBHT estimator is calculated for the complete vector of response parameters, any linear combination of the RBHT parameter estimates (such as a treatment control difference) can be subsequently calculated (with unbiasedness maintained). This makes the RBHT estimator more widely applicable in practice than showcased here.

Another strategy for bias adjusted inference in group sequential trials, due to Whitehead³³ and termed the bias-corrected MLE, is to find the vector \check{p} satisfying: $E(\hat{p}|\check{p}) = \hat{p}$. In words, \check{p} is the parameter constellation for which the expected value of the MLE (given trial data generated under \check{p}) is equal to the observed MLE. In practice, the vector \check{p} is estimated via an iterative process. It has been shown to perform well in group-sequential trials using FR^{16,25} and one could in theory apply it to this context. However, if one adopts a Monte-Carlo approach here too, then it is not attractive, because of the inherent difficulty in assessing the convergence of an iterative process containing a stochastic element. In contrast, whilst calculation of the RBHT estimate involves the generation of trial realisations, no such iteration is necessary.

Besides statistical bias, other types of biases have been highlighted as particularly problematic for trials utilising AR. These biases are very important but are beyond the scope of this paper. However, we now briefly allude to some obvious examples. The un-blinding of patient data mid-trial in order to implement AR can lead to so called ‘operational’ bias, if this information leaks to an individual who has the power and authority to misuse it.³⁴ Another example occurs when the patients recruited into the trial systematically differ over time in their characteristics, and those characteristics are not independent of the treatment’s effect. When so called ‘patient drift’ exists, AR trials are again susceptible to bias and type I error rate inflation when FR trials are not.³⁵ In this second case, however, bias adjusted estimation could potentially be extended to this setting by defining the randomization probability weights conditional on any measured time varying covariates.

A complete discussion as to the relative merits of AR versus FR is also beyond the scope of this paper. We refer the reader to the interesting discussions in Berry³⁶ and Friedlin and Korn,³⁰ that reach opposite conclusions. In recent editorial on the subject Thall et al.³⁷ conclude that in early phase trials, where the goal is often to select the most promising dose from many candidates, AR holds sufficient promise to be considered as a design choice. However, they also discuss that in other trial settings, any benefit provided by an AR scheme is likely to be negated by the logistical complications of implementation. In a recent commentary article, Hey and Kimmelman¹² appear to rule out the use of AR in two-arm trials, but echo the sentiments of Thall et al. in indicating its potential utility in multi-arm trial settings. In our investigation we also found the performance of the various bias adjusted estimators to be better in the multi-arm glioma trial context than the two arm RPW trial context. This is because multi-arm trials inherently put a stronger emphasis on treatment selection, which leads to more pronounced biases and hence the need for bias adjustment.

Acknowledgements

The authors would like to thank Shaun Seaman, Stijn Vansteelandt and the two anonymous reviewers for their comments which greatly improved the paper. The first author thanks Harvard University school of Public Health and the Dana-Farber Cancer Institute for hosting him during a research visit in 2013 to undertake this work.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Dr Jack Bowden is supported by an MRC Methodology research fellowship (grant number MR/L012286/1). Lorenzo Trippa is supported by the Claudia Adams Barr Program in Innovative Cancer Research.

References

1. Pocock SJ. Group sequential methods in the design and analysis of clinical trials. *Biometrika* 1976; **64**: 191–199.
2. Jennison C and Turnbull BW. *Group sequential methods with applications to clinical trials*. London: Chapman and Hall, 1999.
3. Bauer P and Kohne K. Evaluations of experiments with adaptive interim analyses. *Biometrics* 1994; **50**: 1029–1041.
4. Proschan MA and Hunsberger SA. Designed extension of studies based on conditional power. *Biometrics* 1995; **51**: 1315–1324.
5. Cui L, Hung HMJ and Wang SJ. Modification of sample size in group sequential clinical trials. *Biometrics* 1999; **55**: 321–324.
6. Lee JJ, Gu X and Liu S. Bayesian adaptive randomization designs for targeted agent development. *Clin Trials* 2010; **7**: 584–596.
7. Alexander BM, Wen PY, Trippa L, et al. Biomarker-based adaptive trials for patients with glioblastoma? lessons from i-spy 2. *Neuro-oncology* 2013; **15**: 972–978.
8. Berry DA and Eick SG. Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. *Stat Med* 1995; **14**: 231–246.
9. Barker AD, Sigman CC, Kelloff GJ, et al. I-spy 2: an adaptive breast cancer trial design in the setting of neoadjuvant chemotherapy. *Clin Pharmacol Ther* 2009; **86**: 97–100.
10. Ventz S and Trippa L. Bayesian designs and the control of frequentist characteristics: a practical solution. *Biometrics* 2015; **71**: 218–226.
11. Korn EL and Friedlin B. Outcome-adaptive randomization: Is it useful? *J Clin Oncol* 2011; **29**: 771–776.
12. Hey SP and Kimmelman J. Are outcome adaptive allocation trials ethical? *Clin Trials* 2015; **12**: 102–106.
13. FDA. Guidance for industry: adaptive design clinical trials for drugs and biologics. Technical report, US Food and Drug Administration, 2010.
14. Horvitz DG and Thompson DJ. A generalization of sampling without replacement from a finite universe. *J Am Stat Assoc* 1952; **47**: 663–685.
15. Hernan MA, Brumback BB and Robins JM. Estimating the causal effect of zidovudine on cd4 count with a marginal structural model for repeated measures. *Stat Med* 2002; **21**: 1689–1709.
16. Emerson SS and Fleming TR. Parameter estimation following group sequential hypothesis testing. *Biometrika* 1990; **77**: 875–892.
17. Bowden J and Glimm E. Conditionally unbiased and near unbiased estimation of the selected treatment mean for multi-stage drop-the-losers trials. *Biometrical J* 2014; **56**: 332–349.
18. Trippa L, Lee EQ, Yen PY, et al. Alexander. Bayesian adaptive trial design for patients with recurrent glioblastoma. *J Clin Oncol* 2012; **30**: 3258–3263.
19. Boos D and Stefanski LA. *Essential statistical inference*. New York: Springer, 2013.
20. Melfi VE and Page C. Estimation after adaptive allocation. *J Stat Plann Inference* 2000; **87**: 353–363.
21. Wei LJ and Durham S. The randomized play-the-winner rule in medical trials. *J Am Stat Assoc* 1978; **85**: 840–843.
22. Zelen M. Play-the-winner rule and the controlled clinical trial. *J Am Stat Assoc* 1969; **64**: 131–146.
23. Bauer P, Koenig F, Brannath W, et al. Selection and bias – two hostile brothers. *Stat Med* 2010; **30**: 1–11.
24. Simon R. Optimal two-stage designs for phase II clinical trials. *Controll Clin Trials* 1989; **10**: 110.
25. Bowden J and Wason JMS. Identifying combined design and analysis procedures in two stage trials with a binary endpoint. *Stat Med* 2012; **31**: 3874–3884.
26. Li G, Shih WJ and Xie T. A sample size adjustment procedure for clinical trials based on conditional power. *Biostatistics* 2002; **3**: 277–287.
27. Bowden J and Mander A. A review and re-interpretation of a group sequential approach to sample size re-estimation in two stage trials. *Pharm Stat* 2014; **13**: 163–172.
28. Wason JMS and Trippa L. A comparison of Bayesian adaptive randomisation and multi-stage designs for multi-arm clinical trials. *Stat Med* 2014; **33**: 2206–2221.
29. Trippa L, Lee EQ, Yen PY, et al. Reply to b. Friedlin et al. *J Clin Oncol* 2013; **31**: 970–971.

30. Friedlin B and Korn EL. Letter to the editor: adaptive randomisation versus interim monitoring. *J Clin Oncol* 2013; **31**: 969–970.
31. Trippa L, Wen PY, Parmigiani G, et al. Combining progression-free survival and overall survival as a novel composite endpoint for glioblastoma trials. *Neuro oncology* 2015; **17**: 1106–1113.
32. Alexander BM and Trippa L. Progression-free survival: too much risk, not enough reward? *Neuro-oncology* 2014; **16**: 615–616.
33. Whitehead J. On the bias of the maximum likelihood estimate following a group sequential test. *Biometrika* 1986; **73**: 573–581.
34. Chow SC and Chang M. Adaptive design methods in clinical trials a review. *Orphanet J Rare Dis* 2008; **3**: 11.
35. Simon R and Simon N. Using randomization tests to preserve type I error with response-adaptive and covariate-adaptive randomization. *Stat Probab Lett* 2011; **81**: 767–772.
36. Berry DA. Adaptive clinical trials: the promise and the caution. *J Clin Oncol* 2011; **29**: 606–609.
37. Thall PF, Fox P and Wathen JK. Some caveats for outcome adaptive randomization in clinical trials. In: Sverdlov O (ed.) *Modern adaptive randomized clinical trials: statistical, operational, and regulatory aspects. Randomization in clinical trials.* Oxford: Taylor & Francis, 2015.