

Bayesian adaptive bandit-based designs using the Gittins index for multi-armed trials with normally distributed endpoints

Adam L. Smith^a and Sofía S. Villar^b

^aDepartment of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge, UK;

^bMRC Biostatistics Unit, University of Cambridge, School of Clinical Medicine, Cambridge, UK

ABSTRACT

Adaptive designs for multi-armed clinical trials have become increasingly popular recently because of their potential to shorten development times and to increase patient response. However, developing response-adaptive designs that offer patient-benefit while ensuring the resulting trial provides a statistically rigorous and unbiased comparison of the different treatments included is highly challenging. In this paper, the theory of *Multi-Armed Bandit Problems* is used to define near optimal adaptive designs in the context of a clinical trial with a normally distributed endpoint with known variance. We report the operating characteristics (type I error, power, bias) and patient-benefit of these approaches and alternative designs using simulation studies based on an ongoing trial. These results are then compared to those recently published in the context of Bernoulli endpoints. Many limitations and advantages are similar in both cases but there are also important differences, specially with respect to type I error control. This paper proposes a simulation-based testing procedure to correct for the observed type I error inflation that bandit-based and adaptive rules can induce.

ARTICLE HISTORY

Received 23 June 2016
Accepted 29 May 2017

KEYWORDS

Multi-armed bandit; Gittins index; response adaptive procedures; normally distributed endpoint; sequential sampling; patient allocation

1. Introduction

The medical and statistical communities have long held as a ‘gold standard’ for clinical trials the so-called *randomised controlled trial* (RCT), where patients are allocated to a treatment arm with a fixed probability which is equal across all arms and for all patients. This scheme ensures the trial is well balanced, eliminates possible sources of bias, and makes the results as sound as possible. However, this design makes no concession to the wellbeing of patients in the trial: in a K -arm RCT, on average $(K - 1)/K$ of the patients will be assigned to a treatment other than the most effective one (if it exists).

This creates one of the foremost ethical concerns inherent in any clinical trial: the conflict between *learning* (ensuring the selection of the best treatment) and *earning* (treating

CONTACT Sofía S. Villar  sofia.villar@mrc-bsu.cam.ac.uk  MRC Biostatistics Unit, School of Clinical Medicine
Cambridge Institute of Public Health, University of Cambridge, Robinson Way, Cambridge Biomedical Campus, Cambridge
CB2 0SR, UK

© 2017 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

most patients effectively). The scientific aim of a traditional RCT is to *learn* about new treatments and identify the most effective one. It is inevitable under this paradigm, however, that a fixed number of patients will be given an inferior treatment. The research on adaptive methods for trial designs, such as response-adaptive randomisation methods, has developed as a response to this ethical dilemma, seeking to improve the *earning* resulting from a trial while preserving its *learning*. The challenge is to find response-adaptive methods which improve patient welfare during the trial, but do not allow extreme imbalance or bias to hinder the statistical validity of the trial, and are conclusive enough truly to influence future medical practice.

The need to consider patients' wellbeing during the trial is particularly acute in the case of a treatment for a rare disease. In this situation the trial patients represent a high proportion of all those with the disease, and a trial aiming solely to identify the most effective treatment will benefit only the small number of patients remaining to be treated after the end of the trial. The ethical concerns with randomising patients onto an inferior treatment are most severe in the case of a serious or life-threatening disease. Thus, the motivation for an adaptive trial design is arguably strongest in the case of life-threatening rare diseases such as the new types of rare cancers identified by the advances of genetics. However, the challenges of maintaining statistical rigour are even more acute when recruitable patients are sparse and sample sizes are small.

The majority of the response-adaptive randomisation methods proposed in the literature use Bayesian learning and a binary endpoint, with information on the effectiveness of the treatments gained throughout the trial deployed immediately, to increase the chances of patients in the trial receiving a better performing treatment (see e.g. [22]). A limitation of these approaches is that they are myopic (they only make use of past information to alter treatment allocation probabilities) and hence they are not influenced at all by the number of patients that remain to be treated in the trial (nor by the expected number of patients outside the trial). An approach recently proposed and modified for addressing this limitation and developing 'forward looking algorithms' is to consider clinical trial design within the framework of the *Multi-Armed Bandit Problem (MABP)*. The optimal solution to the classic MABP has been known since the 1970s [8], and those responsible for its solution saw clinical trials as the 'chief practical motivation' for their work [9, p. 561]; despite this, it has never been applied to a real life clinical trial. Villar *et al.* [23, pp. 2–3]

In Villar *et al.* [23] some of the benefits and limitations of applying the MABP solution to clinical trials are explored, considering in particular the case where the trial's primary endpoint is dichotomous (i.e. the treatment arms are modelled as reward processes by Bernoulli random variables). The objective of the paper is two-fold. The first is to apply some of the considerations and techniques of Villar *et al.* [23] to define a response-adaptive Bayesian design for a clinical trial whose primary endpoint is normally distributed with known variance, a case that has been less commonly studied in the response-adaptive literature. Specifically, we investigate whether the same conclusions in terms of patient-benefit and operating characteristics hold as in the case of trials with binary endpoints and, since many trials do have normally distributed endpoints, in this way we hope further to bridge the gap between MABP theory and clinical trial practice. The second objective is to identify and address issues that may limit the use in practice of the MABP-based designs considered in this paper. Specifically, we consider in detail the level of bias and type I error rates observed under this setting and further suggest appropriate procedures to control them.

Results are illustrated by simulations in the context of a currently ongoing clinical trial: TAILoR trial, described in Wason *et al.* [25].

The structure of this paper is as follows: In Section 2 an overview of the general MABP with a continuous state variable and its solution for the special case of a normally distributed reward is provided together with an adaptive patient allocation rule based on it. Then, Section 3 presents some simulations of two-armed and multi-armed trials implementing bandit strategies for normally distributed endpoints and comparing them to alternative trial designs. Section 4 concludes with a discussion of our findings and lines of further research.

2. The classic Bayesian MABP with a continuous state variable and known variance

Let $K \in \mathbb{N}$ and consider a collection $\{X_{k,t} : k = 0, 1, \dots, K, t = 0, 1, 2, \dots\}$ of independent (real-valued) random variables, where for each fixed k the distributions of $X_{k,0}, X_{k,1}, X_{k,2}, \dots$ are identical and parametrised by some unknown $\theta_k \in \mathbb{R}^p$. At each time $t = 0, 1, 2, \dots, T - 1$ we obtain a reward by choosing some distribution $k \in \{0, \dots, K\}$ and sampling from $X_{k,t}$. In the context of a clinical trial, this corresponds to choosing the treatment allocation of the t th patient, and $X_{k,t}$ corresponds to the endpoint observation for patient t on treatment k . In order to incorporate the adaptive learning element into the model, we take a Bayesian viewpoint and assume Θ_k is a random variable taking the value θ_k . We assign Θ_k a prior distribution $\pi_k^{(0)}$, which is assumed to be a density function with respect to Lebesgue measure. By Bayes' Theorem, the posterior density of Θ_k , having observed values $x_{k,i_1}, \dots, x_{k,i_n}$ in n independent samples from $X_{k,i_1}, \dots, X_{k,i_n}$ after having treated t patients, is

$$\pi_k^{(t)}(\theta | x_{k,i_1}, \dots, x_{k,i_n}) \propto \pi_k^{(0)}(\theta) \prod_{j=1}^n f_k(x_{k,i_j} | \theta),$$

where $f_k(\cdot | \theta)$ is the density of $X_{k,t}$ (with respect to Lebesgue measure) [14]. Note that we have used the subscript i_j (for $j = 1, \dots, n$) to emphasise that the sample of n observations from distribution k is a subset of the total number of sampling observations possible at time t .

Formally, the classic Bayesian MABP within this general setting is defined by formulating a Markov decision process as follows. Consider each distribution (or arm) k as a Markov process B_k with a Borel state space (E_k, \mathcal{E}_k) , by taking the state $\xi_k(t) \in E_k$ of B_k at time t to be the value \tilde{x} of some chosen sufficient statistic \tilde{X}_k for the posterior density of Θ_k , and updating the state every time we sample from this arm. At each time $t = 0, \dots, T - 1$ a decision variable $a_{k,t} \in \{0, 1\}$ is chosen for process B_k for each $0 \leq k \leq K$, such that exactly one arm receives action 1 (is *sampled*) and all others receive action 0 (their posterior density remains *frozen*). If $a_{k,t} = 0$ then B_k is frozen (and so is its associated value for the sufficient statistics), thus $\xi_k(t + 1) = \xi_k(t)$ with probability 1. If $a_{k,t} = 1$ then ξ_k evolves according to a Markovian transition kernel \mathcal{P}_k , *i.e.* for any $A \in \mathcal{E}_k$ and $\tilde{x}_0, \tilde{x}_1, \dots, \tilde{x}_{t-1}, \tilde{x} \in E_k$ we have

$$\begin{aligned} \mathbb{P}[\xi_k(t + 1) \in A | \xi_k(t) = \tilde{x}, \xi_k(t - 1) = \tilde{x}_{t-1}, \dots, \xi_k(0) = \tilde{x}_0] \\ = \mathbb{P}[\xi_k(t + 1) \in A | \xi_k(t) = \tilde{x}] = \mathcal{P}_k(\tilde{x}, A). \end{aligned}$$

The transition kernel \mathcal{P}_k is a density p_k (with respect to Lebesgue measure on \mathbb{R}):

$$p_k(\tilde{x}, \tilde{x} \star y) = \int_{\mathbb{R}^p} f_k(y|\theta)\pi_k^{(t)}(\theta|\tilde{x}) d\theta, \tag{1}$$

where $\tilde{x} \star y$ denotes the updated value of \tilde{X}_k if y is the next value sampled.

If the process B_k is sampled at time t we earn the random reward $R_k(\xi_k(t), \xi_k(t + 1))$, where $R_k : E_k^2 \rightarrow \mathbb{R}$ is the *reward function* of B_k . In the classic MABP this function is given by $R_k(\tilde{x}, \tilde{x} \star y) = y$, that is, the value of R_k is the value taken by $X_{k,t}$. We define $r_k : E_k \rightarrow \mathbb{R}$ as the *expected reward* from the process in a given state [18], given by

$$\begin{aligned} r_k(\tilde{x}) &= \mathbb{E}[R_k(\xi_k(t), \xi_k(t + 1)) | \xi_k(t) = \tilde{x}] \\ &= \mathbb{E}[X_{k,t} | \tilde{X}_k = \tilde{x}]. \end{aligned} \tag{2}$$

Let $E = E_0 \times \dots \times E_K$ be the joint state space of the MABP, and $\xi(t) = (\xi_0(t), \dots, \xi_K(t)) \in E$ the joint state vector of the MABP at time t . Let Π be the set of all *feasible* sampling policies, that is, those in which the decision at time t depends only on past information and only sample one arm (or distribution) at a time. Writing $a_{k,t}^\pi$ for the sequence of sampling decisions chosen by policy π , the value function for the classic MABP with a continuous state variable is

$$V_d^*(\xi) = \sup_{\pi \in \Pi} \mathbb{E} \left[\sum_{k=0}^K \sum_{t=0}^{T-1} d^t r_k(\xi_k(t)) a_{k,t}^\pi \mid \xi(0) = \xi \right]. \tag{3}$$

Thus, the MABP is the problem of finding a policy $\pi \in \Pi$ which maximises the value of the expected total discounted reward of the sampling process. Notice that d is a discount factor (i.e. $0 \leq d < 1$) introduced for reasons of tractability, so that the infinite horizon problem ($T = \infty$) can be considered.

One approach to solve the MABP in (3) would be via the dynamic programming equation

$$V_d^*(\xi) = \max_{i \in \{0, \dots, K\}} \left\{ r_i(\xi_i) + d \int_{E_i} \mathcal{P}_i(\xi_i, dy) R(\xi_1, \dots, \xi_{i-1}, y, \xi_{i+1}, \dots, \xi_K) \right\}. \tag{4}$$

Standard theory on Markov processes ensures that there is an optimal solution to (3), and approximations to it may be obtained using value iteration on (4) [7], but such an approach is computationally expensive, exploding with the truncation horizon T even for a small number of arms $K > 3$ [23]. For the infinite horizon MABP Gittins and Jones [8] provided a theorem by which there exists a function $v = v(B_k, \tilde{x}_k)$ such that at any time the optimal strategy is to sample the process which has the highest value of v . There is a clear computational advantage to this approach: if we can compute a grid of values of v for each bandit process, then the policy can be followed any number of times by looking up values of v for each process at each decision time. Several proofs of the Index Theorem are given in Gittins *et al.* [7]. Gittins and Jones referred to v as a *dynamic allocation index*, but this is now known widely as the *Gittins index*.

In the case where the MABP state space is discrete, as in the Bernoulli case, values of v can be looked up from a matrix. With a continuous state space it is less clear that all the necessary calculations can be performed in advance. However, in most useful cases, including the normally distributed case, the function $v(B_k, \tilde{x}_k)$ is a linear function of B_k and \tilde{x}_k under some discrete boundary conditions; the discrete part can thus be calculated in advance as a matrix of values [7].

For the finite horizon problem, which is the relevant case in the clinical trial context, [23] suggested an index-based solution to the finite horizon MABP based on the Whittle index [26]. However, the Whittle index is omitted from the studies in this paper since in most trials its performance was near identical to that of the Gittins index (calibrated through the choice of the discount factor d) which further has a lower computational cost.

2.1. The Gittins index as a Bayesian adaptive patient allocation rule for the normally distributed endpoint (with known variance)

In this paper we consider clinical trials for which the endpoint of each treatment arm k is assumed to be normally distributed with unknown mean μ_k and known variance σ_k^2 . We therefore consider the rewards from arm k to be an independent identically distributed (*iid*) sequence $X_{k,t} \stackrel{\text{iid}}{\sim} N(\mu_k; \sigma_k^2)$, and μ_k is given a prior distribution $\pi_k^{(0)}$, which we will take to be the improper uniform distribution on the whole real line. The uniform prior distribution assumption will allow us to isolate the effect on patient welfare and other relevant statistical properties of the MAB adaptive design alone, that is, without the use of prior (historical) data. Let $f(\cdot|\mu; \sigma^2)$ denote the density of a $N(\mu; \sigma^2)$ distribution. If we have observed n independent samples $x_{k,i_1}, \dots, x_{k,i_n}$ from $X_{k,i_1}, \dots, X_{k,i_n}$, then, writing $\bar{x}_{k,n} = (1/n) \sum_{j=1}^n x_{k,i_j}$ for the sample mean, by Bayes' Theorem the posterior density of μ_k at time t is $\pi_k^{(t)}(\mu_k | x_{k,i_1}, \dots, x_{k,i_n}) \sim N(\bar{x}_{k,n}; \sigma_k^2/n)$. A sufficient statistic for the posterior distribution of μ_k is $(\bar{x}_{k,n}, n)$, thus the state vector of process B_k in this case will be (\bar{x}, n) where n is the number of observations so far sampled from arm k after having treated t patients, and \bar{x} is the mean of these observations.

As explained in Gittins *et al.* [7] for the MABP with normally distributed rewards with known variance, the indices $v(\bar{x}, n; \sigma_k^2, d)$ can be written as follows:

$$v(\bar{x}, n; \sigma_k^2, d) = \bar{x} + \sigma_k v(0, n; 1, d). \quad (5)$$

Therefore, to implement the Gittins index policy at very low computational cost it suffices to calculate in advance the values of $v(0, n; 1, d)$. This can be done to a good accuracy using value iteration on (4) in the case of the two-armed bandit calibration setup. Details are given in do Amaral [1, pp. 131–162] and Gittins *et al.* [7, Chapters 7 and 8]. Computational results for this case were first computed in Jones [12]. The values of the indices $v(0, n; 1, d)$ used in this paper have been interpolated from the tables printed in Gittins *et al.* [7, pp. 261–262].

Figure 1 shows the values of the indices $v(0, n; 1, d)$ for a range of discount factors d . In Gittins and Wang [10], the *learning* component of the index is defined as the difference between the index value and the expected immediate reward, which for this MABP corresponds to the reward from sampling an arm with posterior mean \bar{x} from previous samples, that is, simply \bar{x} . Therefore, $\sigma_k v(0, n; 1, d)$ can be interpreted as a measure of the *learning*




Figure 1. Gittins Index values (normal reward process, known variance) for various discount factors d .

reward associated with continuing sampling an arm which has already been sampled n times. Figure 1 illustrates clearly that $\nu(0, n; 1, d)$ increases with d , since a larger discount factor puts greater value on future rewards and increases the value of learning. However, for any choice of d , the value of learning drops very quickly as n increases; in the limit as n tends to infinity, the value of learning tends to 0 and the sample mean converges by the Law of Large Numbers, so the index tends to the true value of the parameter μ_k .

The Gittins index solution for the case of both μ_k and σ_k being unknown exists and is similar to that in (5). The difference is that the model requires a joint prior distribution on both parameters and the known variances in (5) are replaced by sample variances.

2.2. Some considerations specific to the use of bandit strategies in a clinical trial context

In Villar *et al.* [23] simulation results comparing a number of alternative patient allocation rules to index-based solutions for trial scenarios with dichotomous endpoints were provided. The authors conclude that, alongside the clear advantages, there are a number of limitations to the use of the Gittins index as an allocation mechanism for clinical trials. Some of these disadvantages are still going to be an issue in the normally distributed case. The endpoint needs to be immediately observable so that index rules can be applied. This means that a patient in the trial cannot be treated until all previous outcomes have been observed. This is a strong limitation that affects all adaptive designs in general and not only MAB-led designs. In practice, this limits the speed at which new patients can be recruited to the trial; however, this may be less problematic in a rare disease context, where the rate of patient recruitment is likely to be slow already. Applying the adaptive algorithms in batches of patients rather than patient after patient is a way of acknowledging and partially addressing this issue [17,24].

Another limitation still present is that the allocation of treatments in a Gittins index-based design is highly deterministic, which can lead to the introduction of different sources of bias. As explained in Atkinson and Biswas [2], randomisation prevents the trial results from being influenced by ‘secular trends in the population’s health and our ability to measure it, in the quality of recruits to the trial and in the virulence of a disease’. In trials where the clinician can influence which patient receives the next treatment, so-called *selection bias* (the ability of the experimenter to predict which treatment will be allocated next) can influence the results. These extrinsic bias effects are absent from the simulations next presented and from those in Villar *et al.* [23], but could have a significant impact when deterministic rules are used on trials with real populations. Recent work [24] addresses this particular limitation proposing a simple modification of the Gittins index rule for the Bernoulli case that is randomised. Notice that the lack of randomisation of the resulting patient allocations is a limitation shared with most bandit-based algorithms, even those that introduce random terms in their definitions as, for example, [3] or [11].

For other limitations and also for the patient-benefit advantages of index-based designs reported in Villar *et al.* [23], the magnitude or even their existence requires careful consideration. This is the case for the possibility of introducing intrinsic sources of bias. Response-adaptive trials in general can result in biased estimate of a treatment’s outcomes. For example, in a two-arm trial scenario [23] found that the use of the Gittins index introduced a significant negative bias in the estimate of treatment outcomes; the magnitude of the bias is greatest for inferior treatments (since they are more likely to be dropped early in the trial) and the treatment effects are likely to be overestimated. Similar considerations apply with respect to the resulting rates of type I error (a false positive result, i.e. incorrectly rejecting the null hypothesis H_0) and of type II error (failing to detect that an experimental treatment is effective, i.e. incorrectly accepting H_0). Villar *et al.* [23] reported that the index-based designs achieved a level of statistical power that was far below the level of an RCT with the same number of patients T and also that control of the type I error rate required adjusting the statistical test to correct for its conservativeness (i.e. moderate deflation).

An important contribution of this paper is to assess the extent to which further considerations different from the ones mentioned above apply to the normally distributed endpoint. In particular, assessing how important the bias and statistical error levels are in the normally distributed case, and suggesting how to control for the type I error rate at a desired level, are two of the main contributions of this work.

3. Simulation studies

In this section we evaluate the performance of a range of patient allocation rules in a clinical trial context, including bandit-based solutions using the Gittins index. As a case study for simulations we shall use a generalisation of the currently ongoing TelmisArtan and Insulin Resistance in HIV trial (*TAILoR trial*), which is described and also used as a case study in Wason *et al.* [25]. See also [15] for discussion of the design of the TAILoR trial.

The TAILoR trial is a one-sided test of K experimental treatments against a control treatment (i.e. testing for superiority). Treatment k is assumed to have endpoint outcomes $X_{k,t} \stackrel{\text{iid}}{\sim} N(\mu_k; \sigma^2)$, for $k = 0, 1, \dots, K$ (where $k=0$ is the control treatment), and σ^2 is

known and common to all treatments. Setting $\delta_k = \mu_k - \mu_0$, the global null hypothesis is $H_{0,G} : \delta_1, \dots, \delta_K \leq 0$ and the alternative hypotheses are $H_{1,k} : \delta_k > 0, 1 \leq k \leq K$.

We focus on the following: statistical power $(1 - \beta)$; type I error rate (α) ; expected proportion of patients in the trial assigned to the best treatment $(\mathbb{E}p^*)$; the *Expected Outcome* $(\mathbb{E}O)$ defined as the mean patient outcome across the trial realisations; and, for the two-arm case, bias in the maximum likelihood estimate of treatment effect associated with each decision rule.

For testing these hypotheses we shall use the following test statistics:

$$Z_k = \frac{\bar{X}_k - \bar{X}_0}{\sigma \sqrt{\frac{1}{n_k} + \frac{1}{n_0}}}, \quad k = 1, \dots, K,$$

where n_k is the number of sample observations taken from arm k and \bar{X}_k is the sample mean of arm k . Under the assumption that the n_k 's are independent and identically distributed samples, these k test statistics will follow a normal distribution with mean $\delta_k / (\sigma \sqrt{1/n_k + 1/n_0})$ and variance 1. In the case of a two-arm trial with one experimental treatment to be tested against a control, this simplifies to the case of a standard z -test using a univariate normal distribution. For the multi-armed case we will consider the joint distribution of Z_1, \dots, Z_k and use a critical value C_α that controls the *Family-Wise type I Error Rate (FWER)*, defined as $\mathbb{P}[\{\text{reject } H_{0,G}\} | H_{0,G}]$, within a specified level $\alpha \in (0, 1)$.

For each scenario we set the size of the trial T to ensure that an RCT with equal randomisation achieves a specified power $(1 - \beta)$ to detect a specified effective treatment difference $\delta^{(1)}$ between each arm and the control, while controlling the FWER within α . Because we are interested in the marginal type II error rate in a single test, rather than a family-wise error rate, we consider the marginal distributions rather than a joint distribution to determine a required sample size per arm for an RCT. Following this rationale it can be computed that the total required size of the RCT trial (i.e. across all arms) is

$$T = \sum_{k=0}^K n_k = (K + 1) \left(\frac{2\sigma^2(C_\alpha + z_\beta)^2}{(\delta^{(1)})^2} \right), \tag{6}$$

where z_β is the $100(1 - \beta)$ th-percentile of a standard $N(0, 1)$ distribution. See Appendix 1 for details of how C_α is determined and, for example, [28] for a review of sample size calculation in RCTs.

Following [25], we shall assume the variance in the outcomes is $\sigma^2 = 1$, and specify the treatment difference to be detected as $\delta^{(1)} = 0.545$ (chosen such that the probability of a patient given a treatment k with $\delta_k = \delta^{(1)}$ having a better outcome than a patient on the control treatment is 0.65). We will consider the usual error rates of $\alpha = 0.05$ and $\beta = 0.10$.

In every scenario we consider the following patient allocation procedures:

- *Fixed Randomised (FR)*: For each patient, treatments are allocated randomly with fixed probability $1/(K + 1)$ across all treatments;
- *Thompson Sampling (TS)* [21]: For each patient, treatments are allocated randomly, where the probability $\pi_{k,t}$ of allocating treatment k to patient t is proportional to the

posterior probability that treatment k is the best, i.e.

$$\pi_{k,t} = \frac{\mathbb{P} \left[\max_i \mu_i = \mu_k \mid X_{k,t} \right]^c}{\sum_{l=0}^K \mathbb{P} \left[\max_i \mu_i = \mu_l \mid X_{l,t} \right]^c},$$

where c is a tuning parameter defined as $\frac{t}{2T}$ introduced to stabilise the resulting allocation probabilities [20]. The probabilities in the fraction are estimated by simulation at each t ;

- *Upper Confidence Bound (UCB)*: For the first $K+1$ patients, patient t is allocated treatment $k = t - 1$; each patient $t > K+1$ is allocated the treatment k with the highest value of the index $\bar{x}_k + \sigma \sqrt{2 \ln t / n_k}$, as proposed in [3,13].
- *Kullback-LeiblerUCB (KLU)*: For the first $K+1$ patients, patient t is allocated treatment $k = t - 1$; each patient $t > K+1$ is allocated the treatment k with the highest value of the index $\bar{x}_k + \sigma \sqrt{2(\ln t + 3(\ln \ln t)) / n_k}$. This variant of UCB was shown in [6] to have improved asymptotic regret bounds compared to UCB.
- *Current Belief (CB)*: The next patient is allocated the treatment with the highest posterior mean \bar{x}_k .
- *Gittins Index (GI)*: The next patient is allocated the treatment with the highest value of the Gittins Index $v(\bar{x}_k, n_k; \sigma^2, d)$, where d is the value of the discount factor;
- *Randomised Gittins Index (RGI)*: as first suggested in Glazebrook [11], the next patient is allocated the treatment with the highest value of the semi-randomised index $v(\bar{x}_k, n_k; \sigma^2, d) + ((K + 1) / n_k) Y_t$, where Y_t is a random variable sampled from the exponential distribution with mean $1 / (K + 1)$ (this choice of randomisation element is the same as that used by Villar *et al.* [23]).
- *Randomised Belief Index (RBI)*: As first suggested in Bather [4], the next patient is allocated the treatment with the highest value of the semi-randomised index $\bar{x}_k + ((K + 1) / n_k) Y_t$, where Y_t is a random variable sampled from the exponential distribution with mean $1 / (K + 1)$.

For the multi-armed scenarios we have additionally considered the following rules:

- *Trippa et al. Procedure (TP)*: For each patient, treatments are allocated randomly, where the probability $\pi_{k,t}$ of allocating treatment k to patient t is defined by

$$\pi_{k,t} = \frac{\bar{\pi}_{k,t}}{\sum_{l=0}^K \bar{\pi}_{l,t}},$$

where

$$\bar{\pi}_{k,t} = \begin{cases} \frac{\mathbb{P}[\mu_k > \mu_0 \mid \bar{x}_{k,t-1}]^{\gamma_t}}{\sum_{l \geq 1} \mathbb{P}[\mu_l > \mu_0 \mid \bar{x}_{l,t-1}]^{\gamma_t}}, & k \geq 1 \\ \frac{1}{K} \exp \left[\max_{k=1, \dots, K} [n_{k,t} - n_{0,t}] \eta_t \right], & k = 0, \end{cases}$$

in these simulations we have considered $\gamma_t = 3(t/T)^{1.75}$ and $\eta_t = 0.25(t/T)$, as in Trippa *et al.* [22]. Note that this procedure is only considered for multi-arm trials because by design its allocation to the control arm will closely follow that of the best experimental arm.

- *Controlled Gittins (CG)*: Each patient is randomly assigned the control treatment with (fixed) probability $1/K$. If the patient is not randomly assigned to the control group in this way, then she is assigned to the treatment with the greatest Gittins index. Although CG deviates from the optimality of GI, it was still found in Villar *et al.* [23] to offer a significant improvement in patient welfare over FR; moreover, it largely combatted the issue of reduced power. In fact, when there existed a clear superior treatment among the K arms it was found to achieve even higher power than FR.
- *Controlled UCB (CUC)*: A variant of UCB with the control allocation protected as in CG above.

In all scenarios we also include ‘batched’ versions of the Bayesian rules in which the allocation probabilities are updated after a block of b patients are treated instead of after every patient. This idea was implemented in [17,24] as a means of overcoming the practical limitations imposed by the assumption of immediate outcome observability of these algorithms. Their inclusion is intended to more closely replicate the constraints of a real life clinical trial without fully sequential design. Specifically, we consider:

- *Batched Thompson Sampling (TSB)* and *Batched Trippa et al. Procedure (TPB)*: As per TS and TP above, but the allocation probabilities $\pi_{k,t}$ are only updated after every 20 patients. The size of the batch ($b=20$) was chosen to illustrate the effects of a moderate delay in relation to the first two trial sizes considered in this paper and a more severe delay for the trial size reflecting a rare disease scenario. Note that none of the trial sizes assumed are exactly divided by $b=20$. This leaves a few remaining patients which are allocated using the allocation probabilities resulting after observing the outcomes of last block.

In every scenario considered and for every procedure we assumed that the prior for the μ_k parameters is the improper uniform distribution on the whole real line. Notice that a fully Bayesian approach to the design could make use of historical data existing before the trial through appropriate choice of the prior distribution. In this paper we have chosen to use an uninformative prior to make results comparable to the case study in [25] and to isolate the effects of the adaptive designs in the different performance measures.

For the rules that are based on the Gittins index values there is an obvious ethical concern around the choice of a discounting factor when calculating the indices: clearly current and future patients’ wellbeing should be valued equally. So for scenarios with large sample sizes we will take d close to 1, usually $d=0.995$. In the case of a rare disease, if it is known that not more than N patients will ever be treated, the value of d could be chosen so that $d^N \approx 0$, to ensure that the possibility of treating patients beyond N has little impact on current decisions. By doing this the current estimation of the patient population could be used to indirectly affect the choice of trial design. In all trial designs and in all simulations, ties among index values are broken randomly.

3.1. Two-arm trial

We first simulate the TAILoR trial with one experimental arm to be compared with a control treatment, that is, $K=1$. The trial is implemented under H_0 with $\delta_1 = \mu_0 = \mu_1 = 0$ and under H_1 with $\mu_0 = 0$, $\mu_1 = 0.545$ (i.e. $\delta_1 = 0.545$). In both scenarios the common variance is $\sigma^2 = 1$ and $d=0.995$. The sample size considered is of $T=116$ patients. This

size ensures a 5% type I error rate (using $C_\alpha = 1.645$ as a critical value) and $(1 - \beta) = 90\%$ power to detect a difference of $\delta^{(1)} = 0.545$ through a FR design.

3.1.1. Type I error control for adaptive designs

For the adaptive allocation mechanisms, including the Gittins index-based, the use of a critical value of $C_\alpha = 1.645$ is found in simulations to generate a type I error rate inflated above 5%. This is in stark difference to the type I error deflation reported in [23,24]. We will next explain this phenomenon in detail in terms of the GI rule but a similar logic applies for other adaptive rules.

In each realisation of the trial, if one arm performs badly early on and is dropped (or allocated with a very low probability), then the sample mean from this arm will not have a chance to regress upwards to the mean (or do so more slowly), being therefore negatively biased. Figure 2 illustrates this in a typical GI trial run under H_0 , displaying the posterior mean $\bar{x}_{k,t}$ of the outcomes for each treatment arm k after t patients have been allocated that arm. In this example, the control arm $k=0$ performs badly early on in the trial, so is dropped with just $n_0 = 19$ patients, leaving the trial’s estimate of this treatment’s outcomes negatively biased. The experimental arm performs better early on, so is continued and regresses to its mean; thus the trial’s final estimate of this treatment’s effect is close to the true value of 0. The result is that the test statistic takes the value $1.81 > 1.645$, so a hypothesis test using the normal cut-off value of 1.645 would generate a type I error, incorrectly concluding the superiority of the experimental arm.

In order to choose a more suitable critical value for the hypothesis tests when using adaptive designs, we estimate the distribution of the test statistic Z under each trial design by a Monte Carlo simulation with 10^4 repeats of the trial under H_0 . Figure 3 shows the observed empirical distributions of Z for the GI trials, implemented under H_0 and under H_1 . In each case, as well as a histogram of the observed empirical distribution, also displayed is a curve of the standard normal distribution which the test statistic is expected to follow in a FR trial, for comparison.

In Figure 3(a) we see that, in the GI trial under H_0 , the distribution of the test statistic is starkly different from a normal distribution. The sample standard deviation of 1.37 is




Figure 2. The posterior mean $\bar{x}_{k,t}$ of each treatment arm’s outcomes after each patient in a typical GI trial under H_0 .




Figure 3. Histograms of empirical distributions of the test statistic Z in GI trials, implemented under each hypothesis. Also marked is the standard normal distribution which Z should follow in the FR trial (red). The sample mean \bar{Z} , standard deviation S_Z and an empirical 95th-percentile $C_{0.05}$ have been calculated under H_0 . The empirical 95th-percentile under H_0 will correspond to the critical value for hypothesis testing, and is marked by a vertical dotted line on the histograms. (a) GI trial under H_0 (b) GI trial under H_1

much greater than the standard deviation of 1.00 in the FR case, and the heavier tails than a normal distribution correspond to an inflated type I error rate when hypothesis testing is carried out with the normal critical value of 1.645. Notice that because both left and right tails are heavier than the normal tails, the inflation of type I error rate when testing hypotheses at the normal cut-off value would be even greater in a two-tailed test.

The empirical cumulative distribution function evaluated at 1.645 is $\hat{F}_{GI,H_0}(1.645) = 0.89$, indicating that we might expect a type I error rate of 11% if hypothesis testing was carried out with this critical value. Instead, the empirical 95th-percentile of the distribution is $C_{0.05} = \hat{F}_{GI,H_0}^{-1}(0.95) = 1.951$, marked on the histogram by a vertical dotted line. We will therefore use this as the critical value for hypothesis testing in the GI trials to control the type I error rate within 5%.

Notice that the two peaks in the frequency density arise from the two situations in which the estimate of one arm’s outcomes is negatively biased, and the other is unbiased: the right-hand peak corresponds to an incorrect conclusion that the experimental treatment is superior to the control treatment (the situation illustrated in Figure 2), and the left-hand peak corresponds to an incorrect conclusion of the opposite.

Figure 3(b) illustrates that if the GI trial is implemented under H_1 the bimodality of the distribution of the test statistic is greatly reduced, but still present to some extent. $\hat{F}_{GI,H_1}(1.951) = 0.77$, i.e. 77% of the distribution still lies to the left of the empirical critical value of 1.951, marked by a vertical dotted line; thus we expect to observe greatly reduced power of around 23% in the GI trials. The (small) left-hand peak has a weight of $\hat{F}_{GI,H_1}(-0.5) = 5\%$, indicating that in 5% of trials the superior arm is dropped early on due to poor initial performance, and the trial has ended up favouring the wrong arm.

Following the same procedure, 95th-percentiles of the test statistic distribution are estimated for the other adaptive trial designs. Histograms for the distributions of the test statistics in the other trial designs are displayed in Appendix II in Figure A1. Notably, TS,

RGI, UCB and KLU are the only ones of the adaptive designs which appear unimodal in both scenarios. The unimodality of TS, RGI, UCB and KLU under H_1 (Figure A1(b)) indicates that, in almost all realisations of these trials under H_1 , the trial is correctly favouring the superior experimental arm by the end of the trial.

3.1.2. Results and discussion

We now present results of 10^4 repetitions of each trial design using the estimated values described before as *a priori* critical values. The results of the simulation are displayed in Table 1. $\mathbb{E}p^*$ under H_0 is computed as the proportion of patients receiving the control treatment, and under H_1 as the proportion of patients receiving the experimental treatment. The (s.d) values are the standard deviations associated with each measurement. The *Upper Bound (UB)* row displays a theoretical optimum for each measurement based on a design which assigns every patient to the best treatment (i.e. $p^* = 1$) in every trial.

All the adaptive rules achieve better patient welfare than the FR design under H_1 . In this scenario RBI, RGI, UCB and GI all perform similarly well in patient welfare, with $\mathbb{E}O$ values between 0.47 and 0.48, the closest values to the theoretical UB of 0.545. Note that this contrasts with the findings in Villar *et al.* [23] for the Bernoulli case, where GI was found to achieve much better patient welfare than either of the semi-randomised designs. However, the results for these rules are in line with their poorer performance in terms of power when compared to the findings in Villar *et al.* [23]. The TS trial is outperformed by the other adaptive designs in terms of patient welfare; this is explained by the tuning parameter c in the TS mechanism which stabilises the randomisation probabilities.

The high standard deviations in p^* for all the adaptive designs under H_0 indicate that p^* has a broad distribution across the realisations of the trial, so the trials are not consistent and are frequently unbalanced. The standard deviation of 0.48 for CB is close to the limiting case where, in each trial, $p^* \sim \text{Bernoulli}(\frac{1}{2})$, that is, all patients within a trial are assigned to the same treatment, which would give a standard deviation of $\sqrt{\frac{1}{2}(1 - \frac{1}{2})} = 0.50$ in p^* across the trials. This indicates that most trials under CB (and to some extent also GI

Table 1. Comparison in 10^4 trial replicates of operating characteristics of different two-arm trial designs of size $T = 116$, under both hypotheses.

	C_α	$H_0 : \mu_0 = \mu_1 = 0$						$H_1 : \mu_0 = 0, \mu_1 = 0.545$				
		α	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)	$(1 - \beta)$	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)	
(Adaptive random)	FR	1.645	0.0510	0.4997	(0.05)	-0.0001	(0.09)	0.8996	0.4997	(0.05)	0.2718	(0.10)
	TS	1.701	0.0528	0.5006	(0.11)	0.0003	(0.09)	0.8723	0.7317	(0.10)	0.3997	(0.11)
	TSB	1.676	0.0519	0.4994	(0.10)	-0.0001	(0.09)	0.8824	0.6962	(0.09)	0.3816	(0.11)
(Semi-random index-based)	RBI	1.998	0.0509	0.5041	(0.37)	-0.0001	(0.09)	0.3493	0.8891	(0.17)	0.4845	(0.13)
	RGI	1.941	0.0487	0.5005	(0.27)	0.0000	(0.09)	0.5494	0.8764	(0.09)	0.4765	(0.10)
(Index-based)	UCB	2.068	0.0508	0.5050	(0.24)	0.0012	(0.09)	0.5575	0.8697	(0.10)	0.4734	(0.11)
	KLU	1.867	0.0481	0.5021	(0.17)	-0.0001	(0.09)	0.7777	0.8225	(0.08)	0.4489	(0.10)
	CB	1.782	0.0420	0.4918	(0.48)	0.0007	(0.09)	0.1724	0.7624	(0.40)	0.4139	(0.24)
	GI	1.951	0.0437	0.5006	(0.38)	-0.0010	(0.09)	0.2373	0.8786	(0.23)	0.4796	(0.16)
	UB					0.0000	(0.09)		1.0000	(0.00)	0.5450	(0.09)

Note: C_α : critical value used in hypothesis testing; α : type I error rate; $\mathbb{E}p^*$: mean proportion of trial patients assigned the best treatment; (s.d.): standard deviation for each measurement; $\mathbb{E}O$: mean patient outcome; $(1 - \beta)$: statistical power. UB: theoretical UB from assigning all patients best treatment.

and RBI) are highly unbalanced, with one arm being dropped early on and most patients receiving the same treatment).

Under H_1 , the high standard deviation in p^* under GI arises from the bimodality observed in Figure 3(b): in a small proportion of realisations of the trial, the control arm is incorrectly favoured and $p^* < \frac{1}{2}$. The lower standard deviation in p^* for RGI confirms that RGI trials are more consistent in correctly favouring the superior treatment arm. As expected, the GI trial has greatly reduced statistical power (just 24%) compared to the value of 90% achieved by the FR trial. Reduced power is also evident in the other trial designs (c.f. Figure A1); CB has the lowest power (17%).

Note that UCB outperforms KLU in patient welfare, but KLU offers significantly higher power (78%) than UCB (56%). Interestingly, despite the improved regret bounds for KLU proved in Cappé *et al.* [6] KLU only begins to dominate UCB under both power and patient-benefit when the number of patients is very large. Nevertheless, KLU seems to achieve the best compromise between patient welfare and statistical inference out of other modifications to the UCB algorithm designed to improve regret bounds reviewed for this paper, with power of 78% and $\mathbb{E}O$ of 0.45 (only slightly below the 0.48 achieved by GI). The low standard deviation of 0.10 in expected outcome indicates that the welfare benefit is more consistent than in the GI trial.

The results for the batched TS (TSB) illustrate the effects of a blocked implementation of the algorithm to deal with a moderate delay: a marginal increase in power and a considerable decrease of the patient welfare benefits. However, the patient-benefit advantages of TSB over FR are considerably large even if assuming a moderate delay in patient recruitment.

3.1.3. Bias in treatment effect estimates

Figure 4 shows the mean (across the trial realisations) of the bias ($\bar{x}_k^{(t)} - \mu_k$) in the estimated outcome of each treatment after a total of t patients have been treated across both arms in the trial, under each scenario. Figure 4(a,b) shows the GI design introducing a negative bias into estimates of both treatment's effects; within each trial realisation this bias will be restricted to one of the two arms, corresponding to the two modes of the test statistic distribution in Figure 3(a). In all scenarios, the deterministic designs GI and CB exhibit larger bias than the semi-randomised designs RBI and RGI.

3.2. Four-arm trial scenario

This scenario uses the TAILoR trial but now considers $K = 3$ experimental treatments to be compared with a control treatment. To achieve a type I error rate of 5%, the critical value is $C_\alpha = 2.0621$ for the FR trials. Once again we take $\sigma^2 = 1$, and we assume a trial size is of $T = 302$ patients since this is the total required trial size for FR to achieve $(1 - \beta) = 90\%$ power to detect a difference of $\delta^{(1)} = 0.545$ in treatment outcome. The trial is implemented under H_0 with $\mu_0 = \mu_1 = \mu_2 = \mu_3 = 0$ and under H_1 with $\mu_0 = 0$, $\mu_1 = \mu_2 = 0.178$, $\mu_3 = 0.545$. These values are chosen to give the *Least Favourable Configuration (LFC)* for the trial, with $\mu_1 = \mu_2 = \delta^{(0)}$ and $\mu_3 = \delta^{(1)}$, where, as [25] explain: ' $\delta^{(1)}$ is a prespecified clinically relevant effect, and $\delta^{(0)}$ is some threshold below which a treatment is considered uninteresting. The configuration is called least favourable as it minimises the probability of recommending a treatment with effect greater than or equal to $\delta^{(1)}$ amongst

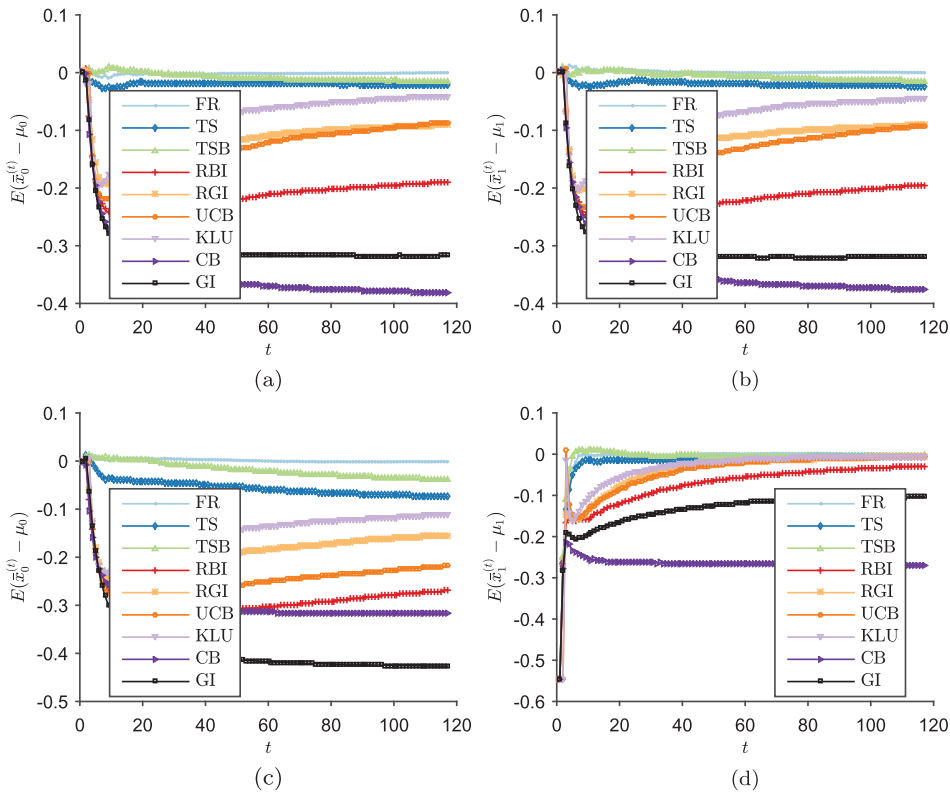


Figure 4. $\mathbb{E}(\bar{x}_k^{(t)} - \mu_k)$, the mean (across the trial realisations) of the bias in the estimated outcome of each treatment after a total of t patients have been treated across both arms in the trial, under each scenario (two-arm trial simulations). (a) H_0 , control arm $k = 0$, (b) H_0 , experimental arm $k = 1$ (c) H_1 , control arm $k = 0$ and (d) H_1 , experimental arm $k = 1$.

all configurations where at least one treatment has a treatment effect of $\delta^{(1)}$ or higher and no treatment effects lie in the interval $(\delta^{(0)}, \delta^{(1)})$.’ Following [25], $\delta^{(0)} = 0.178$ is chosen so that the probability of a patient on a treatment with this treatment effect achieving a better outcome than a patient on the control treatment is 0.55 and the corresponding probability for $\delta^{(1)} = 0.545$ is 0.65.

We will compare all the trial designs, including now the *Controlled Gittins (CG)* design, in which each patient is allocated the control treatment with probability $1/(K + 1) = 0.25$, and otherwise allocated the drug with the highest value of the Gittins Index. We compare CG design against similar procedures: the *Trippa Procedure (TP)* and *Controlled UCB (CUC)* designs. We also include the *Batched Trippa Procedure (TPB)* to assess the effects of delays in outcome observability. The Gittins Indices used are again based on discount factor $d = 0.995$. To calculate critical values for the trial designs other than FR, Monte Carlo simulations were run as explained in Section 3.1.1. Critical values are found by calculating the empirical 95th-percentile of the distribution of $Z_{\max} := \max_{j=1,2,3} Z_j$, in order to control the FWER. Trial simulations are then run using the computed quantiles as critical values; for each design the trial is run 10^4 times. Results are displayed in Table 2.

Table 2. Comparison in 10^4 trial replicates of operating characteristics of different four-arm trial designs of size $T = 302$, under both hypotheses.

	C_α	$H_0 : \mu_0 = \mu_1 = \mu_2 = \mu_4 = 0$					$H_1 : \mu_0 = 0, \mu_1 = \mu_2 = 0.178, \mu_3 = 0.545$				
		α	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)	$(1 - \beta)$	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)
FR	2.062	0.0534	0.2493	(0.02)	0.0006	(0.06)	0.8982	0.2502	(0.02)	0.2252	(0.06)
TS	2.198	0.0476	0.2504	(0.08)	-0.0001	(0.06)	0.8751	0.4997	(0.10)	0.3394	(0.07)
TSB	2.103	0.0514	0.2497	(0.07)	-0.0001	(0.06)	0.8989	0.4794	(0.10)	0.3314	(0.07)
RBI	2.041	0.0519	0.2469	(0.27)	0.0002	(0.06)	0.3608	0.7917	(0.22)	0.4624	(0.10)
RGI	2.070	0.0499	0.2479	(0.18)	-0.0006	(0.06)	0.6309	0.7603	(0.12)	0.4462	(0.07)
UCB	2.223	0.0500	0.2507	(0.13)	0.0001	(0.06)	0.7333	0.7028	(0.12)	0.4238	(0.08)
KLU	2.154	0.0434	0.2502	(0.09)	-0.0002	(0.06)	0.8718	0.6068	(0.10)	0.3848	(0.07)
CB	1.691	0.0524	0.2468	(0.41)	0.0008	(0.06)	0.1075	0.4941	(0.49)	0.3438	(0.20)
GI	1.955	0.0486	0.2457	(0.28)	-0.0008	(0.06)	0.2264	0.7743	(0.29)	0.4552	(0.13)
CG	1.923	0.0405	0.4577	(0.21)	-0.0006	(0.06)	0.8667	0.5681	(0.22)	0.3392	(0.10)
CUC	1.934	0.0572	0.3362	(0.09)	0.0007	(0.06)	0.9599	0.5357	(0.10)	0.3277	(0.07)
TP	2.027	0.0498	0.2593	(0.02)	-0.0010	(0.06)	0.9418	0.3095	(0.06)	0.2462	(0.06)
TPB	2.027	0.0479	0.2488	(0.02)	0.0003	(0.06)	0.9342	0.3082	(0.06)	0.2476	(0.06)
UB					0.0000	(0.06)		1.0000	(0.00)	0.5450	(0.06)

Note: C_α : critical value used in hypothesis testing; α : type I error rate; $\mathbb{E}p^*$: mean proportion of patients in a trial assigned the best treatment; (s.d.): standard deviation for each measurement; $\mathbb{E}O$: mean patient outcome; $(1 - \beta)$: statistical power. UB: theoretical UB from assigning all patients best treatment.

As in the two-arm scenario, all the adaptive rules outperform the FR design under H_1 in terms of patient welfare, although TP only improves marginally over FR in this case. The greatest $\mathbb{E}O$ values are achieved by RGI, RBI, UCB and GI, but these designs and CB exhibit a greatly reduced power level compared with FR, rendering them less useful as trial designs from a frequentist point of view. In particular, CB, which is essentially the simplest myopic approach, exhibits the worst performance in terms of power and variability. As in the two-arm trial, KLU achieves considerably greater power than UCB and the welfare benefit is only slightly reduced, offering a very good compromise between the two conflictive objectives.

As in the two-armed case TSB (90%) achieves marginally higher power than TS (88%) in return for slightly lower patient welfare ($\mathbb{E}O$ of 0.33 compared to 0.34). Conversely, TPB results in a slightly reduced power than TP while the patient welfare is practically identical. In both cases, the difference caused by the moderate ‘batching’ of patients’ outcomes is small, indicating that these adaptive designs could offer patient-benefit advantages even if applied without a fully sequential design.

As expected, the family of ‘protected control’ designs: TP, CG and CUC, offer a compromise between learning (power) and earning (patient welfare). Whilst CG does not perform as well as GI in patient welfare, its $\mathbb{E}O$ value of 0.3392 is still a significant improvement on FR’s 0.2252, and the 87% power attained by CG is greater than that of many of the other adaptive designs, and only marginally lower than FR’s 90%. CUC compares similarly to UCB and dominates over TP by offering a significantly increased patient welfare with a slight increase in power over TP. Just as found in Villar *et al.* [23] for the Bernoulli case, fixing the control allocation in this way is a simple heuristic modification of adaptive allocation rules that results in good trial designs in terms of both patient welfare and frequentist operating characteristics.

Figure 5 shows the bias in the estimates of treatment outcomes in this scenario, for the control treatment and the best experimental treatment (arm 3). For the designs which were




Figure 5. $\mathbb{E}(\bar{x}_k^{(t)} - \mu_k)$, the mean (across the trial repeats) of the bias in the estimated treatment outcome of each drug under each scenario in the four-arm trial (large sample size). (a) H_0 , control arm $k = 0$, (b) H_0 , experimental arm $k = 3$, (c) H_1 , control arm $k = 0$, (d) H_1 , experimental arm $k = 3$.

included in the two-arm simulation, the results here are similar. CG significantly lowers the bias in the estimates of control treatment outcomes, but it does not improve the issue of negatively biased estimates of unselected experimental treatment outcomes, where it performs almost identically to the original GI.

3.3. Four-arm rare disease trial scenario

The final simulation scenario is the same as in Section 3.2 but with the trial size reduced to $T = 64$ to imitate a rare disease setting where the number of patients who can be recruited is limited. Notice that for $T = 64$ the FR trial will achieve a power of 30% while controlling the FWER within 5%. The same critical values are used for hypothesis testing as in the large trials in Section 3.2, based on the assumption that (especially in a trial where patients are recruited sequentially) the experimenter might not know at the start of the trial the total number of patients she will be able to recruit, so more appropriate critical values cannot be estimated *a priori*. Based on the same reasoning we continue to use the original choice of $d = 0.995$.

Table 3 shows the full results of the simulations. Due to the greatly reduced sample sizes, all designs now achieve much lower power, a common situation in drug development for rare diseases. In a situation where $N \gg T$, statistical power is important, and CUC and CG

Table 3. Comparison in 10^4 trial replicates of operating characteristics of different four-arm trial designs of size $T = 64$, under both hypotheses.

	C_α	$H_0 : \mu_0 = \mu_1 = \mu_2 = \mu_4 = 0$					$H_1 : \mu_0 = 0, \mu_1 = \mu_2 = 0.178, \mu_3 = 0.545$				
		α	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)	$(1 - \beta)$	$\mathbb{E}p^*$	(s.d.)	$\mathbb{E}O$	(s.d.)
FR	2.062	0.0490	0.2497	(0.05)	0.0009	(0.13)	0.2975	0.2490	(0.05)	0.2260	(0.13)
TS	2.198	0.0510	0.2491	(0.09)	-0.0001	(0.13)	0.2592	0.3594	(0.11)	0.2779	(0.13)
TSB	2.103	0.0499	0.2489	(0.08)	-0.0010	(0.13)	0.2901	0.3245	(0.09)	0.2625	(0.13)
RBI	2.041	0.0471	0.2497	(0.22)	0.0000	(0.13)	0.1619	0.5351	(0.26)	0.3529	(0.16)
RGI	2.070	0.0565	0.2489	(0.14)	0.0013	(0.12)	0.2344	0.4725	(0.18)	0.3258	(0.15)
UCB	2.223	0.0444	0.2475	(0.14)	-0.0028	(0.13)	0.1730	0.4772	(0.18)	0.3296	(0.15)
KLU	2.154	0.0503	0.2492	(0.10)	-0.0023	(0.13)	0.2452	0.4194	(0.13)	0.3043	(0.13)
CB	1.691	0.0515	0.2522	(0.37)	-0.0006	(0.12)	0.0775	0.4569	(0.46)	0.3239	(0.22)
GI	1.955	0.0477	0.2529	(0.24)	0.0015	(0.13)	0.1226	0.5445	(0.29)	0.3585	(0.17)
CG	1.923	0.0523	0.4100	(0.17)	-0.0002	(0.12)	0.3806	0.4061	(0.22)	0.2742	(0.15)
CUC	1.934	0.0573	0.3445	(0.10)	-0.0011	(0.13)	0.3851	0.3821	(0.15)	0.2670	(0.14)
TP	2.027	0.0462	0.2275	(0.04)	0.0011	(0.12)	0.3174	0.3256	(0.10)	0.2534	(0.13)
TPB	2.027	0.0472	0.1751	(0.04)	-0.0019	(0.12)	0.2674	0.3141	(0.07)	0.2620	(0.13)
UB					0.0000	(0.13)		1.0000	(0.00)	0.5450	(0.13)

Note: α : type I error rate; $\mathbb{E}p^*$: mean proportion of patients in a trial assigned the best treatment; (s.d.): standard deviation for each measurement; $\mathbb{E}O$: mean patient outcome; $(1 - \beta)$: statistical power. UB: theoretical UB from assigning all patients best treatment.

offer the best compromise. Both perform similarly well, achieving higher power than FR, and offering a marked improvement in patient welfare compared with FR. However, if the trial subjects comprise most of the total population to be treated ($T/N \approx 1$), then GI and RBI provide the best patient outcome throughout the trial.

The results in the table for the batched approaches show that, as expected, as the delay in recruitment is more severe the advantages of TSB and TPB over FR are significantly reduced (though both designs still offer important patient welfare advantages). Noticeably, the effect on power and patient welfare of a severe delay in the controlled version (i.e. TPB) differs to that of the *uncontrolled* variant (TSB). The controlled version has its power levels reduced as the delay increases (while the opposite happens to TSB). TP improves power over FR by matching the allocation of the control arm to that of the best performing arm, therefore increasing the allocation to these two arms over the other arms. With a larger delay TBP will allocate larger number of patients to all arms which therefore reduces its marginal power levels compared to TP. For TSB the power improvement is explained because the design cannot skew allocation to the best arm as fast as with TS, thus allocating more patients to all arms when compared to TS.

One distinctive feature of the results is that the Type I error rate α in the UCB trial is lower than the expected 5%, at just 4.4%. As explained above, the same critical values for hypothesis testing have been used as in Section 3.2, since the experimenter might not known in advance the total number of patients to be recruited. Figure 6 shows how the appropriate critical value $C_{0.05}$ for hypothesis testing with a Type I error rate $\alpha = 5\%$ varies according to the size T of the trial. For most trial designs, there is little variation in $C_{0.05}$ as T increases. However, for the UCB trial, $C_{0.05}$ increases significantly with T ; as a result, the appropriate critical value to ensure a 5% Type I error rate is lower for the smaller 64 person trial, at $C_{0.05} = 2.10$, compared to $C_{0.05} = 2.22$ for the 302 person trial. Therefore, the 64 person trial conducted at the higher critical value of 2.22 generates a low Type I error rate, and the power is even lower than it could be if the test was relaxed by lowering the critical




Figure 6. Empirical critical values $C_{0.05}$ for one-tailed testing to maintain 5% FWER in the four-arm trial design, against number T of patients in the trial.

value to 2.10. As a result, the UCB mechanism may be unsuitable for trials where the total number of patients to be recruited is not known in advance. This effect is less pronounced in the KLU variant making it more suitable in that case.

Since the trial size was much smaller than expected, there is a motivation to consider if using a smaller value for d would affect results, as a smaller discounting factor corresponds to putting less value on learning for the future. Note that, when varying the discount factor, we might expect the distribution of the test statistic \mathbf{Z} to change, and so critical values for the hypothesis tests would have to be recalculated for each discount factor for the Gittins index designs, via a Monte Carlo simulation as in Section 3.2. In simulations not included here we found that for this trial setting in all of GI, RGI and CG there is no significant variation in patient outcome between discount factors in $\{0.9, 0.95, 0.995, 0.99\}$.

4. Conclusions and discussion

The simulation results provided by this paper illustrate how the index-based response-adaptive design derived from the MABP can lead to significant improvements in patient welfare also with a normally distributed endpoint. In all situations, designs based on the Gittins index achieved the largest patient welfare gain over FR trials or myopic designs currently in use in drug development such as TP. However, there are a number of limitations to the effectiveness of the purely deterministic Gittins index design that still prevail. As in the binary case, the Gittins index rule exhibits considerably lower power than FR, and whilst the loss of power can be alleviated to some extent by the introduction of random perturbations to the indices (RGI), in the two-arm trial the power achieved is still not sufficient for most clinical trials unless the exploration term is correctly calibrated.

In a multi-armed case, the patient welfare advantages of adaptive designs, and GI-based particularly, over FR are the largest. Moreover, there are adaptive designs that can offer more power than FR together with a patient-benefit advantage, making them suitable for drug development for common conditions. In the four-arm case based on a real trial we studied, a small deviation from optimality by protecting the allocation of the control treatment (CG and CUC) offers a power close to (or even above) FR's while still providing considerable patient-benefit. In contexts where power is relatively less important (if there are very few disease sufferers outside the trial), GI, RGI or UCB offer even better patient welfare at the expense of a power reduction.

There are designs that increase power levels of the UCB algorithm by introducing modifications to improve its asymptotic regret bounds, as shown for KLU in Cappé *et al.* [6]. However, such power gains require a very large number of patients in the trial to be also accompanied by similar patient welfare advantages. For example, KLU dominates over UCB under both criteria only in scenarios where trials had more than thousands of patients. For smaller (and more realistic) trial sizes, as the ones considered in this paper, UCB had better patient welfare and less power than KLU. Nevertheless, rules like KLU offer a good trade-off between the two objectives and can be suitable designs for common diseases.

An important observation drawn from the simulations provided by this paper is that the type I error deflation of the GI observed for the Bernoulli case does not hold in the normally distributed case. Actually, if no correction is introduced using a standard test will result in an important type I error inflation. In this work we have outlined a simulation-based procedure that can be used to prevent this inflation.

As pointed out in Berry [5], trying to shoehorn trials employing an adaptive design from a Bayesian viewpoint into traditional frequentist hypothesis tests may not be the most appropriate method of inference. Alongside the statistical community's faith in randomisation is a trust in frequentist inference, so this is generally used even in Bayesian trials to make the results as persuasive as possible. But, the inferential power and the potential patient-benefit from adaptive trials could be improved by applying Bayesian inference methods combined with the use of prior data. Further research could seek an appropriate method of Bayesian inference based on index-based adaptive trials, for example, by considering which arm the adaptive design is favouring most at the end of the trial, or by incorporating information derived from historical data.

None of the Bayesian allocation mechanisms considered here manages to completely eliminate the statistical bias phenomenon; further research is needed to seek an alternative mechanism or a means of accounting for the bias introduced. Moreover, they all carry a level of selection bias which, while not studied in the simulations included in this paper, could lead to much greater bias in clinical trials on a real population. Further research is needed to investigate whether significant practical problems will arise from selection bias in real trials, and whether random perturbations to the indices are sufficient to eliminate these problems. Alternatively, to overcome this limitation further research could repeat the idea introduced in Villar *et al.* [24] to randomise group of patients based on probabilities determined by the Gittins Indices for trials with continuous endpoints. For the procedures that protect allocation to the control arm we recommend a randomised implementation (where a patient is randomised to control or experimental arms with probabilities $1/(K + 1)$, $1 - 1/(K + 1)$ respectively and then allocated to experimental arms according to the

index rule). A systematic allocation to control arm (i.e., 1 in every $K+1$ patient is allocated to control) while in theory is equivalent to its randomised counterpart in practice is subject to a very high degree of selection bias.

Some adaptive trials are designed to take account of covariates in the trial population (e.g. age, weight, blood pressure) which might affect the treatment response, by ensuring the allocations are balanced across the covariate factors [2]. Other trials incorporating covariate information combined with response-adaptive procedures with the aim of identifying superior treatments more quickly, mainly treatments that work better within subgroups, is an essential requirement to make *personalised medicine* possible. Some work has been done on incorporating covariates into the one-armed bandit problem, yet further research is needed to extend the approach to multi-armed bandits used in this work to clinical trials with biomarkers. See [16,19,27].

None of the adaptive designs considered formally accounts for the estimated population size. The index-based approaches indirectly can consider that by appropriate selection of the discount factor. However, the results in this paper suggest that the choice between implementing a traditional FR design or an adaptive design should depend on the current belief of how large the population of patients outside the trial is.

Finally, the results presented in this paper have highlighted that further analogous research is needed to extend these results and address potential specific issues to trials with other endpoints, such as continuous endpoints that are not normally distributed.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was funded by the Biometrika Trust and the UK Medical Research Council (grant number MC_UP_1302/2).

References

- [1] J.F.P. do Amaral, *Aspects of optimal sequential resource allocation*, D.Phil. thesis, University of Oxford, 1985.
- [2] A.C. Atkinson and A. Biswas, *Randomised Response-Adaptive Designs in Clinical Trials*, CRC Press, Boca Raton, FL, 2014.
- [3] P. Auer, N. Cesa-Bianchi, and P. Fischer, *Finite-time analysis of the multiarmed bandit problem*, *Mach. Learn.* 47 (2002), pp. 235–256.
- [4] J. Bather, *Randomized allocation of treatments in sequential trials*, *Adv. Appl. Probab.* 12 (1980), pp. 174–182.
- [5] D.A. Berry, *[Investigating therapies of potentially great benefit: ECMO]: comment: ethics and ECMO*, *Stat. Sci.* 4 (1989), pp. 306–310.
- [6] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz, *Kullback-Leibler upper confidence bounds for optimal sequential allocation*, *Ann. Stat.* 41 (2013), pp. 1516–1541.
- [7] J. Gittins, K. Glazebrook, and R. Weber, *Multi-Armed Bandit Allocation Indices*, 2nd ed., Wiley, Chichester, 2011.
- [8] J.C. Gittins and D.M. Jones, *A Dynamic Allocation Index for the sequential design of experiments*, in *Progress in Statistics (European Meeting of Statisticians, Budapest, 1972)*, J. Gani, K. Sarkadi, and I. Vincze, eds., North-Holland, Amsterdam, 1974, pp. 241–266.

- [9] J.C. Gittins and D.M. Jones, *A dynamic allocation index for the discounted multiarmed bandit problem*, *Biometrika* 66 (1979), pp. 561–565.
- [10] J. Gittins and Y. Wang, *The learning component of dynamic allocation indices*, *Ann. Stat.* 20 (1992), pp. 1625–1636.
- [11] K.D. Glazebrook, *On randomized dynamic allocation indices for the sequential design of experiments*, *J. R. Stat. Soc. Ser. B (Methodol.)* 42 (1980), pp. 342–346.
- [12] D. Jones, *Search procedures for industrial chemical research*. Master's thesis, U.C.W. Aberystwyth, 1970.
- [13] M.N. Katehakis and H. Robbins, *Sequential choice from several populations*, *Proc. Natl. Acad. Sci. U.S.A.* 92 (1995), pp. 8584–8585.
- [14] H. Kobayashi, B.L. Mark, and W. Turin, *Probability, Random Processes, and Statistical Analysis*, CUP, Cambridge, 2012.
- [15] D. Magirr, T. Jaki, and J. Whitehead, *A generalized Dunnett test for multi-arm multi-stage clinical studies with treatment selection*, *Biometrika* 99 (2012), pp. 494–501.
- [16] V. Perchet and P. Rigollet, *The multi-armed bandit problem with covariates*, *Ann. Stat.* 41 (2013), pp. 693–721.
- [17] V. Perchet, P. Rigollet, S. Chassang, and E. Snowberg, *Batched bandit problems*, *Ann. Stat.* 44 (2016), pp. 660–681.
- [18] M.L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, Wiley-Interscience, Hoboken, NJ, 2005.
- [19] J. Sarkar, *One-armed bandit problems with covariates*, *Ann. Stat.* 19 (1991), pp. 1978–2002.
- [20] P.F. Thall and J.K. Wathen, *Practical Bayesian adaptive randomisation in clinical trials*, *Eur. J. Cancer* 43 (2007), pp. 859–866.
- [21] W. Thompson, *On the likelihood that one unknown probability exceeds another in view of the evidence of two samples*, *Biometrika*, 25 (1933), pp. 285–294.
- [22] L. Trippa, E.Q. Lee, P.Y. Wen, T.T. Batchelor, T. Cloughesy, G. Parmigiani, and B.M. Alexander, *Bayesian adaptive randomized trial design for patients with recurrent Glioblastoma*, *J. Clin. Oncol.* 30 (2012), pp. 3258–3263.
- [23] S. Villar, J. Bowden, and J. Wason, *Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges*, *Stat. Sci.* 30 (2015), pp. 199–215.
- [24] S. Villar, J. Bowden, and J. Wason, *Response-adaptive randomization for multi-arm clinical trials using the forward-looking Gittins index rule*, *Biometrics* 71 (4) (2015), pp. 969–978.
- [25] J. Wason, N. Stallard, J. Bowden, and C. Jennison, *A multi-stage drop-the-losers design for multi-arm clinical trials*, *Statistical Methods in Medical Research* (to appear). Advance online publication. doi:10.1177/0962280214550759, 2014.
- [26] P. Whittle, *Restless bandits: Activity allocation in a changing world*, *J. Appl. Probab.* 25 (1988), pp. 287–298.
- [27] M.B. Woodroofe, *A one-armed bandit problem with a concomitant variable*, *J. Am. Stat. Assoc.* 74 (1979), pp. 799–806.
- [28] B. Zhong, *How to calculate sample size in randomized controlled trial?*, *J. Thorac. Dis.* 1 (2011), pp. 51–54.

Appendix 1. Controlling the family-wise type I error rate

In order to control the FWER when carrying out multiple testing, we need to consider the joint distribution of Z_1, \dots, Z_k . We have, for $1 \leq j \neq k \leq K$,

$$\begin{aligned} \text{cov}(Z_j, Z_k) &= \text{cov} \left(\frac{\bar{X}_j - \bar{X}_0}{\sigma \sqrt{\frac{1}{n_j} + \frac{1}{n_0}}}, \frac{\bar{X}_k - \bar{X}_0}{\sigma \sqrt{\frac{1}{n_k} + \frac{1}{n_0}}} \right) \\ &= \frac{1}{\sigma^2} \frac{1}{\sqrt{\frac{1}{n_j} + \frac{1}{n_0}}} \frac{1}{\sqrt{\frac{1}{n_k} + \frac{1}{n_0}}} \text{Var}(\bar{X}_0), \end{aligned}$$

by the independence of Z_j and Z_k . Using the fact that $\text{Var}(\bar{X}_0) = \sigma^2/n_0$,

$$\text{cov}(Z_j, Z_k) = \left[\left(1 + \frac{n_0}{n_j}\right) \left(1 + \frac{n_0}{n_k}\right) \right]^{-1/2}.$$

Error rates are lowest when the variance of the sample means is minimised, which corresponds to the trial being well balanced: in a RCT trial with fixed equal randomisation all the sample sizes are (asymptotically) equal, so we will have a good approximation for $n_0 \approx n_1 \approx \dots \approx n_K$ and $\text{cov}(Z_j, Z_k) \approx \frac{1}{2}$. Hence, under $H_{0,G}$, $\delta_1 = \dots = \delta_K = 0$ and

$$\mathbf{Z} = (Z_1, \dots, Z_K) \sim N_K(0, \Sigma_K),$$

where Σ_K is the $K \times K$ matrix given by

$$(\Sigma_K)_{ij} = \begin{cases} 1, & i = j, \\ \frac{1}{2}, & i \neq j. \end{cases}$$

So we would expect a RCT trial to control the FWER at level α by using critical value C_α satisfying

$$\int_{-\infty}^{C_\alpha} \dots \int_{-\infty}^{C_\alpha} \phi_{\Sigma_K}(x_1, \dots, x_K) dx_1 \dots dx_K = 1 - \alpha, \tag{A1}$$

where ϕ_{Σ_K} is the probability density function of a multi-variate normal $N_K(0, \Sigma_K)$ distribution, i.e. ensuring that

$$\mathbb{P} \left[\max_{k=1, \dots, K} Z_k \leq C_\alpha \right] = 1 - \alpha. \tag{A2}$$

Appendix 2: Calculating empirical cut-off values to control the type I error rate

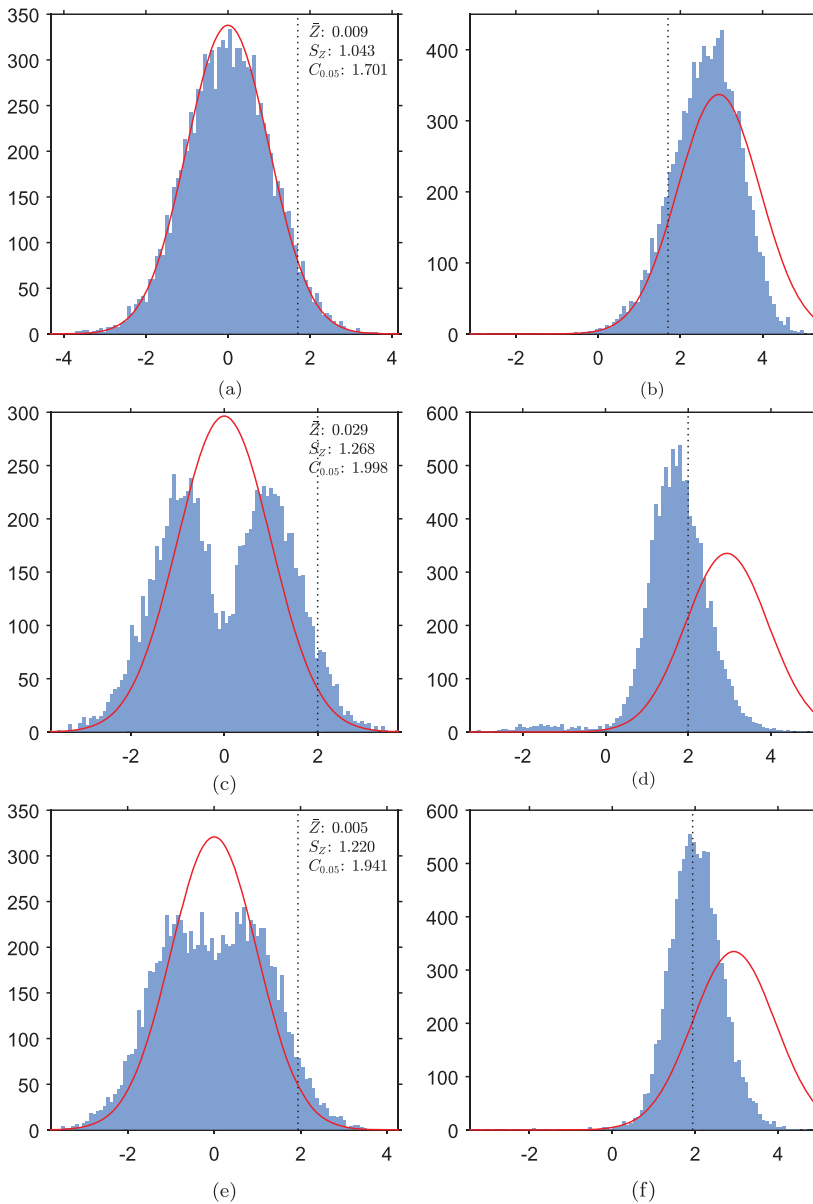


Figure A1. Histograms of empirical distributions of the test statistic $Z_{0,1}$ in TS, RBI, RGI, UCB, KLU and CB two-arm trials, implemented under each hypothesis (as in Figure 3). Also marked is the standard normal distribution which $Z_{0,1}$ should follow in the FR trial (red). For each design, the sample mean $\bar{Z}_{0,1}$, standard deviation $S_{Z_{0,1}}$ and an empirical 95th-percentile $C_{0.05}$ have been calculated under H_0 . The empirical 95th-percentile under H_0 will correspond to the critical value for hypothesis testing, and is marked by a vertical dotted line on the histograms. (a) TS trial under H_0 , (b) TS trial under H_1 , (c) RBI trial under H_0 , (d) RBI trial under H_1 , (e) RGI trial under H_0 , (f) RGI trial under H_1 , (g) UCB trial under H_0 , (h) UCB trial under H_1 , (i) KLU trial under H_0 , (j) KLU trial under H_1 , (k) CB trial under H_0 , (l) CB trial under H_1 .

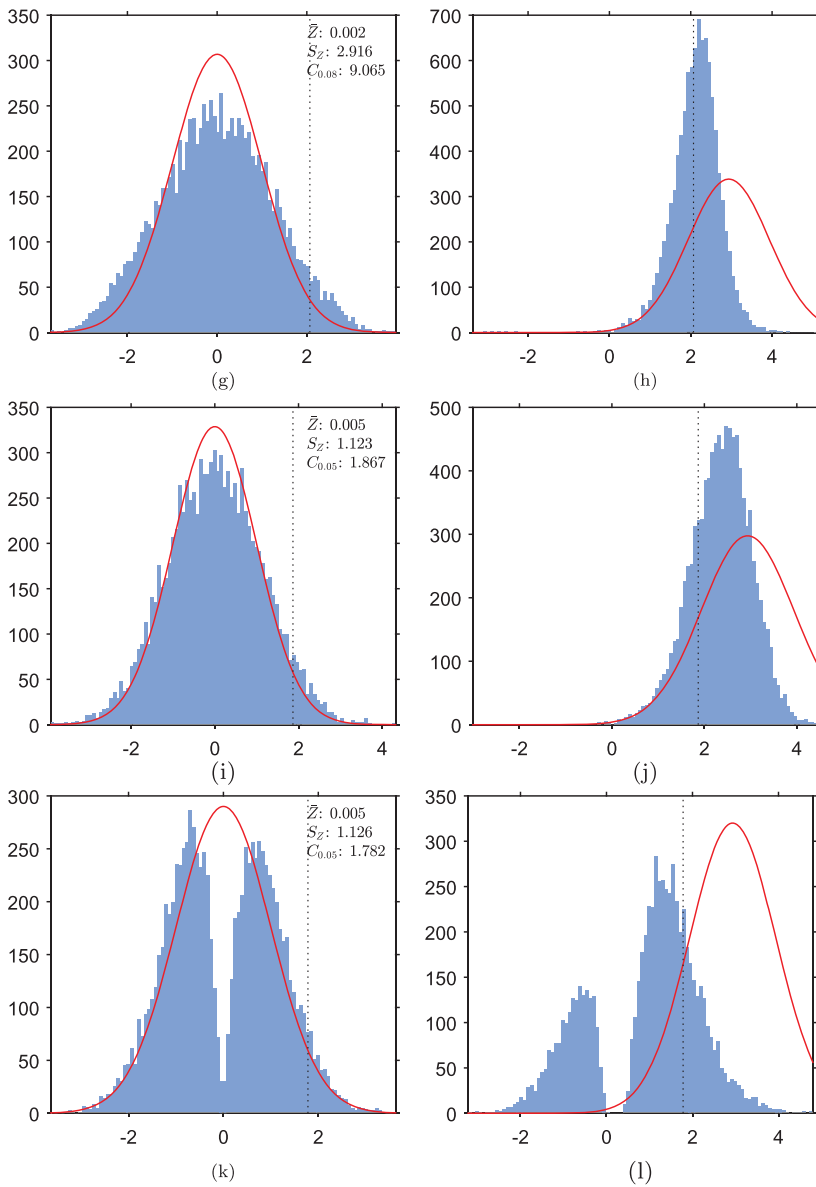


Figure A1. (Continued).