Journal Pre-proof

Approximating the operating characteristics of Bayesian Uncertainty directed trial Designs

Marta Bonsaglio, Sandra Fortini, Steffen Ventz, Lorenzo Trippa

PII: S0378-3758(22)00028-3
DOI: https://doi.org/10.1016/j.jspi.2022.03.001
Reference: JSPI 6002

To appear in: *Journal of Statistical Planning and Inference*

Received date : 1 July 2021
Revised date : 10 March 2022
Accepted date : 15 March 2022

Please cite this article as: M. Bonsaglio, S. Fortini, S. Ventz et al., Approximating the operating characteristics of Bayesian Uncertainty directed trial Designs. *Journal of Statistical Planning and Inference* (2022), doi: https://doi.org/10.1016/j.jspi.2022.03.001.

Approximating the Operating Characteristics of Bayesian Uncertainty Directed Trial Designs

Marta Bonsaglio [1], Sandra Fortini [1], Steffen Ventz [2,3], Lorenzo Trippa [2,3]

[1]Department of Decision Sciences, Università Bocconi, Italy, [2]Dana-Farber Cancer Institute, US,

[3]Department of Biostatistics, Harvard T.H. Chan School of Public Health, US

March 10, 2022

Abstract

Bayesian response adaptive clinical trials are currently evaluating experimental therapies for several diseases. Adaptive decisions, such as pre-planned variations of the randomization probabilities, attempt to accelerate the development of new treatments. The design of response adaptive trials, in most cases, requires time consuming simulation studies to describe operating characteristics, such as type I/II error rates, across plausible scenarios. We investigate large sample approximations of pivotal operating characteristics in Bayesian Uncertainty directed trial Designs (BUDs). A BUD trial utilizes an explicit metric u to quantify the information accrued during the study on parameters of interest, for example the treatment effects. The randomization probabilities vary during time to minimize the uncertainty summary u at completion of the study.

We provide an asymptotic analysis (i) of the allocation of patients to treatment arms and (ii) of the randomization probabilities. For BUDs with outcome distributions belonging to the natural exponential family with quadratic variance function, we illustrate the asymptotic normality of the number of patients assigned to each arm and

of the randomization probabilities. We use these results to approximate relevant operating characteristics such as the power of the BUD. We evaluate the accuracy of the approximations through simulations under several scenarios for binary, time-to-event and continuous outcome models.

Keywords: Adaptive Designs, Almost Sure Convergence, Bayesian Uncertainty Directed Trial Designs, Central Limit Theorem, Large Sample Approximations of Operating Characteristics, Stochastic Approximation

1 Introduction

Randomized clinical trials (RCTs) are essential to demonstrate the efficacy of novel experimental therapies [13]. The landscape of clinical studies has changed during the last decades, with an increasing number of trials that utilize adaptive designs, in some cases to evaluate several experimental treatments in biomarker-defined subpopulations [3, 29]. Adaptive designs are attractive to reduce the duration of the study and to allocate efficiently limited resources [11]. Most adaptive designs use data generated during the clinical trial for interim decisions [11], for example to vary the randomization probabilities during the study [3, 8, 29, 34] or to discontinue the evaluation of an experimental treatment [29]. In multi-arm studies adaptive randomization algorithms unbalance the randomization probabilities, in most cases, towards the most promising treatments. This can increase power compared to balanced randomization, and it can reduce the overall sample size necessary to test experimental treatments [31]. Adaptive randomization procedures have been developed for several designs, including multi-arm studies [8, 12], platform and basket studies [3, 29, 34].

The decision theoretic paradigm has been used to develop trial designs [7, 9, 15]. The study aims and costs are represented by a utility function $u(\cdot)$ of the data Σ generated during the trial and the study design d. Using a Bayesian joint model for patient profiles, outcomes and other key variables, candidate designs d can be compared by computing their expected utility $E[u(\Sigma, d)]$. The optimal design maximizes $E[u(\Sigma, d)]$ among all candidate designs. Several approximations of the described optimization have been proposed. For example, [30] discussed *Bayesian Uncertainty directed trial Designs* (BUDs), a class of approximate decision theoretic designs. The utility function u in BUDs coincides with an information metric. In different words, the goal is to minimize uncertainty at completion of the study. BUDs for dose-finding and basket trials have been discussed in [16] and [28]. Previous work, related to BUDs, proposed information-based sampling schemes [10, 25, 27].

There is a rich literature on large sample analyses of adaptive designs. For instance, [1, 32] studied the behavior of sequential urn schemes. See also [2, 18, 26, 33] for a recent

summary on large sample results for urn schemes. The limiting behavior of adaptive biased coin designs have been investigated, among others, by [17, 19] and [20]. Relevant work connecting stochastic approximation with response-adaptive clinical trials include [4] and [22].

In this manuscript we focus on the asymptotic characteristics of BUDs. The design of adaptive clinical trials requires the estimation of pivotal operating characteristics, such as type I and II error rates and the distribution of patients randomized to each treatment arm. In most cases these estimates are based on time consuming Monte Carlo simulations, conducted for different candidate designs and varying key parameters, including sample sizes, enrollment rates, and outcome distributions. Approximations of the operating characteristics, beyond simulations, using asymptotic results, are crucial to compare designs across plausible scenarios.

The need for computationally efficient approximations of design-specific operating characteristics motivates our study. We show the almost sure convergence and asymptotic normality of the relative allocation of patients to treatment arms in BUDs. We first derive analytic results assuming that the treatment-specific outcome distributions belong to natural exponential family [14], and later relax this assumption. In our analysis, we represent BUD randomization procedures as stochastic approximations (SAs). We study the ordinary differential equations associated with the resulting SAs and the stability of the stationary points, following the framework developed in [5] and using results of [21, 22]. We illustrate through examples the accuracy of the asymptotic approximations by comparing asymptotic and Monte-Carlo estimates of operating characteristics of BUDs.

Our asymptotic results allow investigators to quickly approximate the distribution of the number of patients that will be assigned to each arm and the power of BUDs.

2 Trial Design

We consider a clinical experiment that assigns n patients sequentially to K arms. We use $A_t \in \mathcal{A} = \{0, \cdots, K-1\}$ to indicate the assignment of individual $t = 1, \cdots, n$ to arm A_t, $Y_t \in \mathbb{R}$ is the response to treatment A_t, and the accumulated data up to enrollment t is summarized by $\Sigma_t = \{(A_\ell, Y_\ell)\,; \ell \leq t\}$.

The BUD [30] is defined by first specifying a Bayesian model. Outcomes are conditionally independent $Y_t \mid A_t = a \sim f_{\theta_a}(y|a)$ and identically distributed within each arm a, $\theta_a \sim \pi(\theta_a)$ indicates the prior for the unknown parameter θ_a. The vector of parameters $\theta = (\theta_0, \ldots, \theta_{K-1})$ has joint distribution $\theta \sim \pi(\theta) = \prod_{a=0}^{K-1} \pi(\theta_a)$.

We quantify the information accrued by the experiment through the accumulated data Σ_t until time t by considering an utility function $u(\Sigma_t)$. Large values of $u(\Sigma_t)$ correspond to low uncertainty levels. The function $\tilde{u}$ translates the posterior distribution for the parameters of interest $\pi(\theta \mid \Sigma_t)$ into an utility. In particular, we define $u(\Sigma_t) = \tilde{u}(\pi(\theta \mid \Sigma_t))$. The information metric $\tilde{u}$ is specified by a convex functional over the convex space of distributions of the parameters: $\tilde{u}\left(\omega\pi_1 + (1-w)\pi_2\right) \leq w\tilde{u}(\pi_1) + (1-\omega)\tilde{u}(\pi_2)$ for every pair of probability measures π_1 and π_2, when $\omega \in [0, 1]$.

By Jensen's inequality, the information, on average, increases with each enrollment,

$$\Delta_t(a) := E\left(u(\Sigma_{t+1})|A_{t+1} = a, \Sigma_t\right) - u(\Sigma_t) \geq 0, \tag{1}$$

for every $a \in \mathcal{A}$. The myopic and deterministic policy $A_{t+1} = \arg\max_{a \in \mathcal{A}} \Delta_t(a)$, which is often inappropriate for clinical experiments [8] is relaxed in BUDs by a randomized version, with probabilities

$$p_{t,a} := p(A_{t+1} = a \mid \Sigma_t) \propto \Delta_t(a)^h, \tag{2}$$

where $h \geq 0$ is a tuning parameter. The randomization probabilities coincide with the myopic policy when $h \to \infty$, while with $h = 0$ the randomization probabilities become identical across arms.

2.1 Outcome distributions within the natural exponential family

We focus on outcome distributions f_{θ_a} in the natural exponential family (NEF) [6],

$$f_{\theta_a}(y|a) = f_{\psi_a}(y) \propto \exp\{y\psi_a - b(\psi_a)\}, \tag{3}$$

where $\psi_a \in \mathbb{R}$ is the canonical parameter and $b(\cdot)$ is the cumulant transform.
We indicate the mean and the variance of the responses to treatment a with $\theta_a = E_\psi(Y_t|A_t = a) = \int y f_{\psi_a}(y)dy = b'(\psi_a)$ and $\sigma_a^2 = \mathrm{Var}_\psi(Y_t|A_t = a) = \int y^2 f_{\psi_a}(y)dy - \left(\int y f_{\psi_a}(y)dy\right)^2 = b''(\psi_a)$, respectively. We use the equivalent parametrization f_{ψ_a} and f_{θ_a} interchangeably and we consider independent conjugate prior distributions [14] for ψ_a,

$$\pi(\psi_a \mid n_{0,a}, y_{0,a}) \propto \exp\{n_{0,a}\tilde{y}_{0,a}\psi_a - n_{0,a}b(\psi_a)\}, \tag{4}$$

with hyper-parameters $n_{0,a} > 0$ and $\tilde{y}_{0,a} \in \mathbb{R}$. The posterior distribution for $\psi = (\psi_0, \cdots, \psi_{K-1})$ is $\pi(\psi \mid \Sigma_t) = \prod_{a=0}^{K-1} \pi(\psi_a \mid \Sigma_t)$, where $\pi(\psi_a \mid \Sigma_t)$ has the same form as (4) with updated parameters $n_{t,a} = n_{0,a} + t\hat{p}_{t,a}$ and $\tilde{y}_{t,a} = \left(n_{0,a}\tilde{y}_{0,a} + \sum_{s=1}^{t} Y_s 1(A_s = a)\right)/n_{t,a}$. Here $\hat{p}_{t,a}$ is the proportion of patients assigned to treatment a by time t and $1(A_s = a) = 1$ if patient s received treatment a and zero otherwise. In this paper we consider

$$u(\Sigma_t) = -\sum_{a=0}^{K-1} \mathrm{Var}(\theta_a|\Sigma_t). \tag{5}$$

The expected information increment (1) becomes

$$\Delta_t(a) = \mathrm{Var}(\theta_a \mid \Sigma_t) - E\left(\mathrm{Var}(\theta_a \mid \Sigma_{t+1}) \mid A_{t+1} = a, \Sigma_t\right). \tag{6}$$

We recall a useful result from the literature on conjugate Bayesian models [6, 14], $\tilde{y}_{t,a} = E(\theta_a \mid \Sigma_t)$. Since A_{t+1} and θ_a are conditionally independent, given Σ_t, the information gain equals

$$\begin{aligned}\Delta_t(a) &= \mathrm{Var}(E(\theta_a \mid \Sigma_{t+1}) \mid A_{t+1} = a, \Sigma_t) \\ &= \mathrm{Var}\left(\frac{n_{0,a} + \tilde{y}_{0,a} + \sum_{s=1}^{t+1} Y_s 1(A_s = a)}{n_{0,a} + t\hat{p}_{t,a} + 1} \mid A_{t+1} = a, \Sigma_t\right) \\ &= \mathrm{Var}\left(\frac{Y_{t+1}}{n_{0,a} + t\hat{p}_{t,a} + 1} \mid A_{t+1} = a, \Sigma_t\right),\end{aligned}$$

where the first equality follows from the law of total variance and the second equality is a consequence of the properties of the natural exponential family. We can therefore write

$$\Delta_t(a) = \frac{\sigma^2_{t,a}}{(n_{0,a} + t\hat{p}_{t,a} + 1)^2},$$

where $\sigma^2_{t,a} = \mathrm{Var}(Y_{t+1} \mid A_{t+1} = a, \Sigma_t)$.

3 Asymptotic properties

We discuss asymptotic properties of BUDs with sum of the (negative) posterior variances of $\theta_a, a = 0, \ldots, K-1$, as information measure $u(\Sigma_t)$. In [30] a criterion is given for the allocation proportions to have a limit. Based on this result, we first prove convergence of allocation proportions and randomization probabilities under the assumption that the outcome distributions belong to the natural exponential family. We then investigate the rate of convergence of these quantities in the case $K = 2$.

Proposition 1. *Consider a two arm BUD, $K = 2$, with outcome distribution belonging to the NEF (3), conjugates prior (4) and information metric $u(\Sigma_t)$ in (5). Then, as $t \to \infty$,*

(i) the allocation of patients to treatments $a = 0, 1$ converges almost surely (a.s.),

$$\widehat{p}_{t,a} \longrightarrow \rho_a := \frac{\sigma_a^{\frac{2h}{2h+1}}}{\sigma_0^{\frac{2h}{2h+1}} + \sigma_1^{\frac{2h}{2h+1}}} \ a.s. \ as \ t \to \infty \tag{7}$$

(ii) the randomization probability converges a.s. to the same limit,

$$p_{t,a} \longrightarrow \rho_a \ a.s. \ as \ t \to \infty. \tag{8}$$

The proofs of the proposition and of all subsequent results are included in the Supplementary material. The extension of the result presented in Proposition 1 to multi-arm BUDs is formalized in Corollary 1.

Corollary 1. *Under the same assumptions of Proposition 1, if $K > 2$, then, as $t \to \infty$, the allocation of patients to treatments $(\widehat{p}_{t,0}, \ldots, \widehat{p}_{t,K-1})$ and the randomization probabilities $(p_{t,0}, \ldots, p_{t,K-1})$ converge a.s. to $(\rho_0, \ldots, \rho_{K-1})$, where for $a \in \{0, \ldots, K-1\}$*

$$\rho_a = \frac{\sigma_a^{\frac{2h}{2h+1}}}{\sum_{i=0}^{K-1} \sigma_i^{\frac{2h}{2h+1}}}. \tag{9}$$

We recall that the NEFs with quadratic variance function consist of all NEFs such that $\sigma_a^2 = v_0 + v_1\theta_a + v_2\theta_a^2$ for some constants v_0, v_1, v_2. In different words, the variance is a polynomial function of order ≤ 2 of the mean [23]. This class contains models, such as the normal, Poisson, gamma, negative binomial and binomial distributions. We refer to Morris [23, 24] for a detailed study of this class of distributions.

We derive the rate of convergence and show the asymptotic normality of the randomization probabilities $p_{t,a}$ and of the allocation proportions $\widehat{p}_{t,a}$ in two-arm BUDs characterized by the utility criteria $u(\Sigma_t) = -\sum_{a=0}^{1} \mathrm{Var}(\theta_a|\Sigma_t)$ when the statistical model f_{θ_a} is a NEF with quadratic variance.

First, for $a = 0, 1$, we define

$$\begin{aligned}
\Delta\sigma^2_{t,a} &= \mathrm{Var}(Y_{t+2} \mid A_{t+2} = a, \Sigma_{t+1}) - \mathrm{Var}(Y_{t+1} \mid A_{t+1} = a, \Sigma_t), \\
v(\tilde{y}_{t,a}) &= \left(v_0 + v_1\tilde{y}_{t,a} + v_2\tilde{y}^2_{t,a}\right)^{\frac{1}{2}}, \\
W_t &= [p_{t,1}, \tilde{y}_{t,1}, \tilde{y}_{t,0}]'
\end{aligned}$$

and

$$k_a(W_t) = \left[1 + \left(\frac{p_{t,a}}{p_{t,1-a}}\right)^{\frac{1}{2h}} \frac{v(\tilde{y}_{t,1-a})}{v(\tilde{y}_{t,a})}\right],$$

where $p_{t,0}$ can be written as $1 - p_{t,1}$.

We also write $X(t) = \mathcal{O}_P(t^{-\alpha})$, where $\alpha > 0$, to intend that $\forall\, \epsilon > 0 \;\; \exists\, T, M > 0$ finite such that $P(|X(t)| > Mt^{-\alpha}) < \epsilon \, \forall t > T$.

Lemma 1 approximates, for $a \in \{0, 1\}$ the variables $\sigma_{t,a}$ and $\tilde{y}_{t+1,a}$ with functions of ($\tilde{y}_{t,a}$, $p_{t,1}$) and (Y_{t+1}, A_{t+1}).

Lemma 1. *If the outcome distributions $f_{\psi_a}, a = 0, 1$, of a two-arm BUD belong to the NEF with quadratic variance function, then*

(i) $\sigma_{t,a} = v(\tilde{y}_{t,a}) + \mathcal{O}_P(t^{-1})$

(ii) $\tilde{y}_{t+1,a} = \tilde{y}_{t,a} + (Y_{t+1} - \tilde{y}_{t,a})\dfrac{1(A_{t+1} = a)}{t} k_a(W_t) + \mathcal{O}_P(t^{-2})$

(iii) $\Delta\sigma^2_{t,a} = (v_1 + 2v_2\tilde{y}_{t,a})(Y_{t+1} - \tilde{y}_{t,a})\dfrac{1(A_{t+1} = a)}{t} k_a(W_t) + \mathcal{O}_P(t^{-2})$.

In spirit to the previous result, the next Lemma illustrates that $p_{t+1,1}$ can be approximated by a function of W_t, Y_{t+1}, A_{t+1} with an $\mathcal{O}_P(t^{-2})$ error term.

Lemma 2. *Let the outcome distributions $f_{\psi_a}, a = 0, 1$ of a two-arm BUD belong to the NEF*

with quadratic variance function. For the randomization probabilities $p_{t,1}, t \geq 1$ *it holds that*

$$p_{t+1,1} = p_{t,1} + hp_{t,1}(1-p_{t,1})\Bigg\{ \left[\frac{(v_1 + 2v_2\tilde{y}_{t,1})(Y_{t+1} - \tilde{y}_{t,1})}{v(\tilde{y}_{t,1})^2} - 2\right] \frac{A_{t+1}}{t} k_1(W_t) + \left[2 - \frac{(v_1 + 2v_2\tilde{y}_{t,0})(Y_{t+1} - \tilde{y}_{t,0})}{v(\tilde{y}_{t,0})^2}\right] \frac{(1-A_{t+1})}{t} k_0(W_t)\Bigg\} + \mathcal{O}_P(t^{-2}). \tag{10}$$

The two Lemmas above suggest how to approximate $\tilde{y}_{t+1,a} - \tilde{y}_{t,a}$ for $a \in \{0,1\}$ and $p_{t+1,1} - p_{t,1}$. For $t \geq 1$, we define the random vector $\tilde{G}_{t+1} = [G_{t+1}, G_{t+1,1}, G_{t+1,0}]'$, whose components are approximations of $t(p_{t+1,1} - p_{t,1})$ and $t(\tilde{y}_{t+1,a} - \tilde{y}_{t,a})$, respectively, where

$$G_{t+1} := hp_{t,1}(1-p_{t,1})\Bigg\{ \left[\frac{(v_1 + 2v_2\tilde{y}_{t,1})(Y_{t+1} - \tilde{y}_{t,1})}{v(\tilde{y}_{t,1})^2} - 2\right] A_{t+1} k_1(W_t) + \left[2 - \frac{(v_1 + 2v_2\tilde{y}_{t,0})(Y_{t+1} - \tilde{y}_{t,0})}{v(\tilde{y}_{t,0})^2}\right] (1-A_{t+1}) k_0(W_t)\Bigg\},$$

and $G_{t+1,a} := 1(A_{t+1} = a)(Y_{t+1} - \tilde{y}_{t,a})k_a(W_t)$ for $a = 0, 1$. By computing the conditional expectations $\tilde{g}(W_t) = -E_\psi(\tilde{G}_{t+1} \mid \Sigma_t)$ we define the map $\tilde{g}(\cdot) = [g(\cdot), g_1(\cdot), g_0(\cdot)]'$, whose components are

$$g(W_t) := -2h \frac{v(\tilde{y}_{t,1})}{v(\tilde{y}_{t,0})} \frac{(1-p_{t,1})^{\frac{2h+1}{2h}}}{p_{t,1}^{\frac{1-2h}{2h}}} k_1(W_t)^2 \left(\frac{1}{k_1(W_t)} - p_{t,1}\right) - hp_{t,1}(1-p_{t,1}) \times \left[p_{t,1} \frac{(v_1 + 2v_2\tilde{y}_{t,1})(b'(\psi_1) - \tilde{y}_{t,1})}{v(\tilde{y}_{t,1})^2} k_1(W_t) - (1-p_{t,1}) \frac{(v_1 + 2v_2\tilde{y}_{t,0})(b'(\psi_0) - \tilde{y}_{t,0})}{v(\tilde{y}_{t,0})^2} k_0(W_t)\right],$$

$$g_a(W_t) := -p_{t,a}(b'(\psi_a) - \tilde{y}_{t,a})k_a(W_t) \quad \text{for } a \in \{0,1\}.$$

In Proposition 2 we rewrite $t(W_{t+1} - W_t)$ as the sum of (i) a function of W_t, (ii) a Σ_t -martingale-difference sequence $\Delta\tilde{M}_{t+1}$ and (iii) a Σ_{t+1}-measurable sequence of remainder terms. In particular, $\Delta\tilde{M}_{t+1} = [\Delta M_{t+1}, \Delta M_{t+1,1}, \Delta M_{t+1,0}]'$ is defined by $\tilde{G}_{t+1} + \tilde{g}(W_t)$.

Proposition 2. *Let the outcome distributions* $f_{\psi_a}, a = 0, 1$*, of a two-arm BUD belong to*

the NEF with quadratic variance function. Then,

$$W_{t+1} = W_t - \frac{1}{t}\tilde{g}(W_t) + \frac{1}{t}(\Delta\tilde{M}_{t+1} + \tilde{r}_{t+1}), \tag{11}$$

where the reminder terms $\tilde{r}_{t+1} := [r_{t+1}, r_{t+1,1}, r_{t+1,0}]$ *are three* $\mathcal{O}_P(t^{-1})$ *sequences.*

Using Proposition 2 we use the theory of stochastic approximation [5, 21] to derive the asymptotic distribution of W_t. Following the stochastic approximation framework, we consider equation (11) together with the ordinary differential equation (ODE)

$$\frac{dW_t}{dt} = -\tilde{g}(W_t), \tag{12}$$

where $t \in (0, +\infty)$ denotes continuous time. The ODE has arbitrary initial conditions. Note that if we ignore the residual term $\tilde{r}_{t+1}$, the difference $W_{t+1} - W_t$ in (11) is equal to $-\frac{1}{t}\tilde{g}(W_t)$ plus a Σ_t-martingale-difference sequence.

We describe the distribution of W_t, for large t, using on an asymptotic analysis of the ODE (12). By identifying the stationary point $[\rho_1, b'(\psi_1), b'(\psi_0)]$ of the ODE, assessing its stability, and using some regularity conditions on $\Delta\tilde{M}_{t+1}$ and $\tilde{r}_{t+1}$, we prove a central limit type result for W_t. In particular, Theorem 1 indicates the asymptotic normality of the randomization probability $p_{t,1}$.

Theorem 1. *Under the assumptions of Proposition 2,* $t^{1/2}(p_{t,1} - \rho_1) \to \mathcal{N}\Big(0, \dfrac{\Gamma}{1+4h}\Big)$, *as* $t \to +\infty$, *where*

$$\Gamma = h^2\rho_1^2(1-\rho_1)^2\left[\frac{(v_1 + 2v_2b'(\psi_1))^2}{\rho_1\sigma_1^2} + \frac{(v_1 + 2v_2b'(\psi_0))^2}{(1-\rho_1)\sigma_0^2} + \frac{4}{\rho_1(1-\rho_1)}\right]. \tag{13}$$

The following corollary verifies the asymptotic normality of the relative allocation $\hat{p}_{t,1}$ by applying the Delta method and Slutsky's Theorem.

Corollary 2. *Under the assumptions of Theorem 1, as* $t \to +\infty$,

$$t^{1/2}(\hat{p}_{t,1} - \rho_1) \longrightarrow \mathcal{N}\Big(0, \Psi\Big),$$

$$\textit{with } \Psi = \frac{2h^2\rho_1^2(1-\rho_1)^2}{(1+4h)(1+2h)}\left[\frac{(v_1+2v_2b'(\psi_1))^2}{\rho_1\sigma_1^2} + \frac{(v_1+2v_2b'(\psi_0))^2}{(1-\rho_1)\sigma_0^2}\right] + \frac{\rho_1(1-\rho_1)}{1+4h}.$$

4 Applications and examples

We apply the results in the previous section to the design of clinical trials. We consider three common outcomes, binary, time to event and continuous outcomes.

Binary outcomes. For $Y_t \in \{0,1\}$, we use the Bernoulli model $f_{\psi_a}(1) = 1 - f_{\psi_a}(0) = \theta_a$, $\theta_a = 1/(1+e^{-\psi_a})$, and conjugated prior $\theta_a \sim \text{Beta}(\alpha,\beta)$. The outcome variance σ_a^2 in expression (7) is $\theta_a(1-\theta_a)$, and the parameters of the quadratic variance function in (13) are $v_1 = 1$ and $v_2 = -1$. Therefore, $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ converges in distribution to a mean zero Gaussian variable with variance

$$\frac{2h^2\rho_1^2(1-\rho_1)^2}{(1+4h)(1+2h)}\left[\frac{(1-2\theta_1)^2}{\rho_1\theta_1(1-\theta_1)} + \frac{(1-2\theta_0)^2}{(1-\rho_1)\theta_0(1-\theta_0)}\right] + \frac{\rho_1(1-\rho_1)}{1+4h}.$$

The top panel of the second column of Figure 1 shows a trajectory $\hat{p}_{t,1}$ for a single simulated two-arm BUD trial (black curve). The response probabilities (θ_0, θ_1) are set equal to 0.2 and 0.4. We used $\alpha = \beta = 2$ and $h = 5$. The shaded area shows (point-wise at each t) upper and lower 2.5% quantiles of the distribution of $\hat{p}_{t,1}$ across 10000 simulations. The second row illustrates the distribution of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ across 10000 simulations of the two-arm BUD trial. The empirical distribution of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ has been smoothed with a kernel density estimator. The panel compares the $\mathcal{N}(0, 0.097)$ density (asymptotic approximation) to the empirical distribution of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ across simulations, when $t = 100, 1000$ and 10000. The last row compares the empirical distribution distribution of $t^{1/2}(p_{t,1} - \rho_1)$ to the

$\mathcal{N}\Big(0, \dfrac{\Gamma}{1+4h}\Big)$ density.

Time-to-event outcomes. We consider an exponential model $f_{\psi_a}(y) = \exp\{-y\psi_a\}\psi_a, y \geq 0$ with mean $\theta_0 = 1/\psi_a$, and we use the conjugated gamma prior $\psi_a \sim \text{Gamma}(\alpha, \beta)$. The outcome variance σ_a^2 in expression (7) is $1/\psi_a^2$, the parameters of the quadratic variance function in (13) are $(v_1, v_2) = (0, 1)$ and $h = 5$. Therefore, the asymptotic variance of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ is

$$\rho_1^2(1-\rho_1)^2\left(\frac{1}{\rho_1} + \frac{1}{1-\rho_1}\right)\left(1 + \frac{2}{1+4h} - \frac{2}{1+2h}\right).$$

The third column of Figure 1 compares the asymptotic and empirical distributions of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$ and $t^{1/2}(p_{t,1} - \rho_1)$, based on 10000 simulations of the BUD trial. In this example $(\theta_0, \theta_1) = (5, 7)$ and $(\alpha, \beta) = (3, 3)$.

Continuous outcomes. We consider a normal outcome model $\mathcal{N}(\theta_a, \sigma_a^2)$ with known variance σ_a^2. We use a conjugated prior $\theta_a \sim \mathcal{N}(0, v_{0,a}^2)$. In this case $v_1 = v_2 = 0$, $h = 5$, and

$$t^{1/2}(\hat{p}_{t,1} - \rho_1) \xrightarrow[t\to\infty]{} \mathcal{N}(0, \frac{\rho_1(1-\rho_1)}{1+4h}). \tag{14}$$

Column 1 of Figure 1 illustrates the empirical distribution of $t^{1/2}(\hat{p}_{t,1} - \rho_1)$, $t = 100, 1000$ or 10000, and the normal approximation.

Figure 1 illustrate that the normal approximations of the sampling distribution of $\sqrt{t}\widehat{p}_{t,a}$ and $\sqrt{t}\widehat{p}_{t,a}$ are accurate for $t = 100$ and, as expected, the approximation improves with the size t. To quantify the accuracy of the approximation of the distribution $\widehat{F}$ of $\sqrt{t}\widehat{p}_{t,a}$ (and $\sqrt{t}p_{t,a}$) across simulations by the normal approximation F_N we used the index $O_p(\widehat{F}, F_N) = \frac{|A_{p,\widehat{F}} \cap A_{p,F_N}|}{|A_{p,\widehat{F}} \cup A_{p,F_N}|} \in [0, 1]$. Here $A_{p,F} = [F^{-1}(p), F^{-1}(1-p)]$ indicates the interval between the

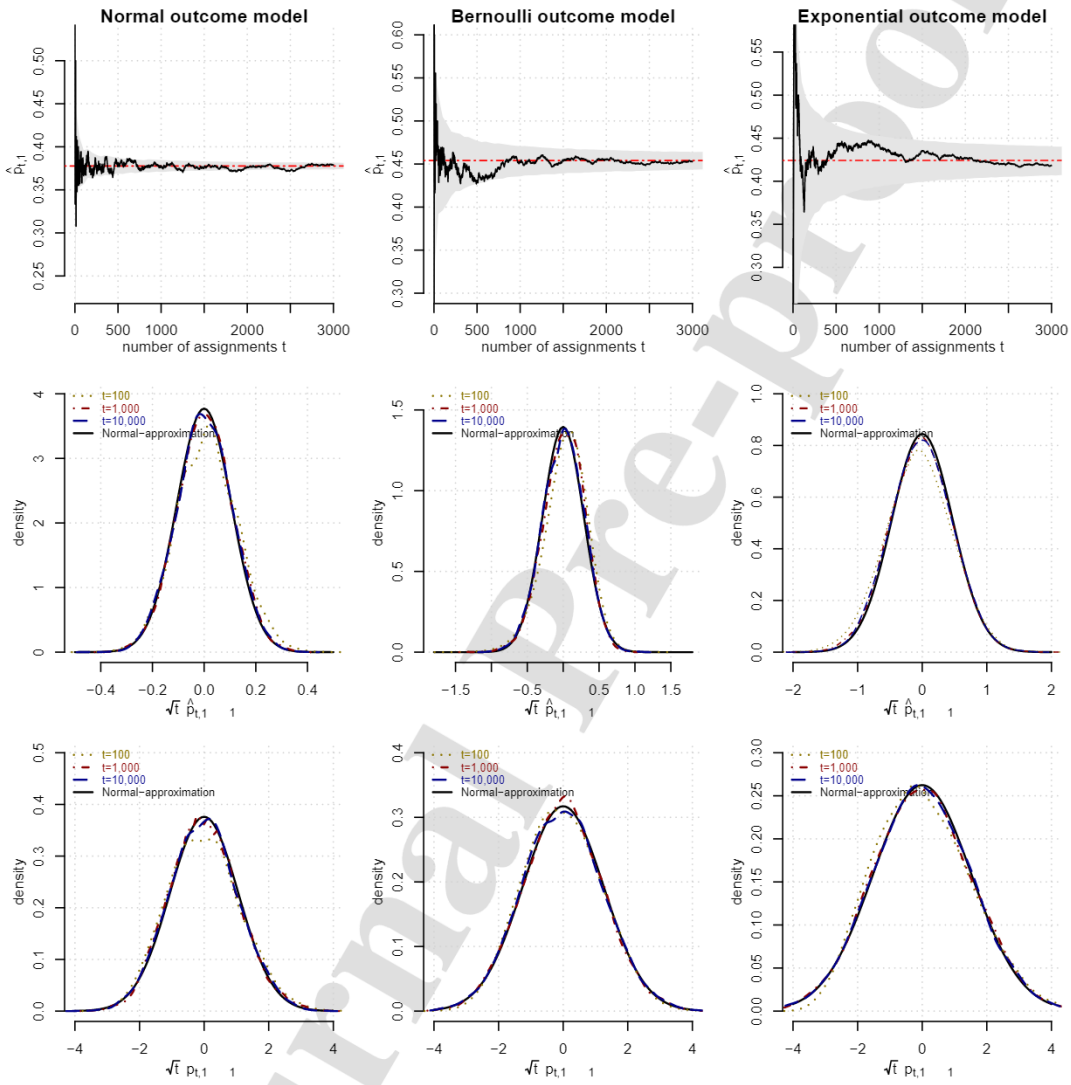

Figure 1: The panels in the first row compare $\hat{p}_{t,1}$ with the limit ρ_1 (red line) in each of the three examples (binary, continuous and time to event outcomes). The other panels compare asymptotic and empirical distributions of randomization probabilities $p_{t,1}$ and allocation proportions $\hat{p}_{t,1}$. The empirical distributions in each of the three examples (binary, continuous and time to event outcomes) are based on 10000 simulations of the two-arm BUD trial.

p-th and $(1-p)$-th percentile of $F = \widehat{F}$ and F_N, and $|[x, y]| = y - x$, for $y > x$, denotes the length of the interval. $O_p = 1$ if and only if $A_{p,\widehat{F}} = A_{p,F_N}$. For instance, for $\sqrt{t}\widehat{p}_{t,a}$ and the normal outcome model, $O_{0.05}(\widehat{F}, F_N)$ was 0.87, 0.97 and 0.97 when $t = 100, 1000, 10000$, whereas 0.96, 0.96 and 0.98 for the binary outcome model, and 0.90, 0.98 and 0.99 for the exponential outcome model. For $\sqrt{t}p_{t,a}$ the $O_{0.05}(\widehat{F}, F_N)$ was 0.82, 0.99 and 0.99 when $t = 100, 1000, 10000$ for the normal outcome model, 0.91, 0.99 and 1.00 for the binary outcome model, and 0.94, 0.99 and 0.99 for the exponential outcome model.

Power analysis and sample size selection. We use the results in Section 3 to select the sample size of BUD studies accordingly to the targeted type I and II error rates. We approximate the power function of the BUD under fixed scenarios using Corollary 2.

We assume that the primary aim of the clinical trial is to test the one-sided null hypothesis $H_0 : \theta_{0,1} = \theta_{0,0}$ against the alternative $H_1 : \theta_{0,1} > \theta_{0,0}$. We verified (Supplementary material) that, under the sequential BUD design, the maximum-likelihood estimates (MLE) $\widehat{\theta}_{t,a}$ of the unknown true mean response $\theta_{0,a}$ to treatment $a = 0, 1$ within the NEF of outcome models have the same limiting distribution as the MLE of a study design with fixed arm-specific sample sizes, i.e.

$$t^{1/2} \begin{bmatrix} \hat{\theta}_{t,0} - \theta_{0,0} \\ \hat{\theta}_{t,1} - \theta_{0,1} \end{bmatrix} \xrightarrow[t\to\infty]{} \mathcal{N}\left(\mathbf{0}, Diag(\eta_{0,0}, \eta_{0,1})\right), \tag{15}$$

where $\eta_{0,a} := \left(\rho_a I_{\theta_{0,a}}\right)^{-1}$ and $I_{\theta_{0,a}}$ is the Fisher information of $f_{\theta_{0,a}}$.

We use a standard Wald-statistics, $Z_a = \dfrac{\sqrt{t} \times (\hat{\theta}_{t,1} - \hat{\theta}_{t,0})}{\sqrt{\widehat{\eta}_{t,a} + \widehat{\eta}_{t,1}}}$, where $\widehat{\eta}_{t,a} = 1/(\widehat{\rho}_a \times I_{\widehat{\theta}_{t,a}})$, and the MLE $\widehat{\sigma}_a^2$ for $\widehat{\rho}_a = \rho_a(\widehat{\sigma}_a^2)$ in (7) to test H_0. The power function of the BUD design is approximated by $\Phi\Big(z_{1-\alpha} - \dfrac{\sqrt{t}(\theta_{0,1} - \theta_{0,0})}{\sqrt{\eta_{0,1} + \eta_{0,0}}}\Big)$ where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal random distribution and $\Phi(z_{1-\alpha}) = 1-\alpha$. Therefore $\widehat{t}_{1-\alpha,1-\beta} = \dfrac{(z_{1-\alpha} - z_{1-\beta})^2(\eta_{0,0} + \eta_{0,1})}{(\theta_{0,1} - \theta_{0,0})^2}$ approximates the sample size of the BUD study to achieve a power equal to 1-β and type I error rate α.

Figure 2 compares, for three BUD designs (binary, time-to-event and continuous out-

comes), power estimates based on asymptotic approximations (blue dotted lines) and on Monte Carlo simulations (1000 simulated trials, blue solid lines). The computational time for the simulation-based calculations is orders of magnitude larger than the normal approximation. We also show the empirical estimates of the type I error rates (brown solid lines) for the outlined asymptotic testing procedure with target type I error rate of $\alpha = 0.05$ (brown dotted lines). For the normal outcome model, $\sigma_0^2 = 1, \sigma_1^2 = 3$, and $\theta_0 = \theta_1 = 0$ (null scenario, brown lines) or $(\theta_1, \theta_2) = (0, 1)$ (positive treatment effect, blue lines). Similarly, for the Bernoulli and Exponential models the parameter values θ that defined null (brown lines) and alternative scenarios (blue lines) are indicated in the panels of Fig 2.

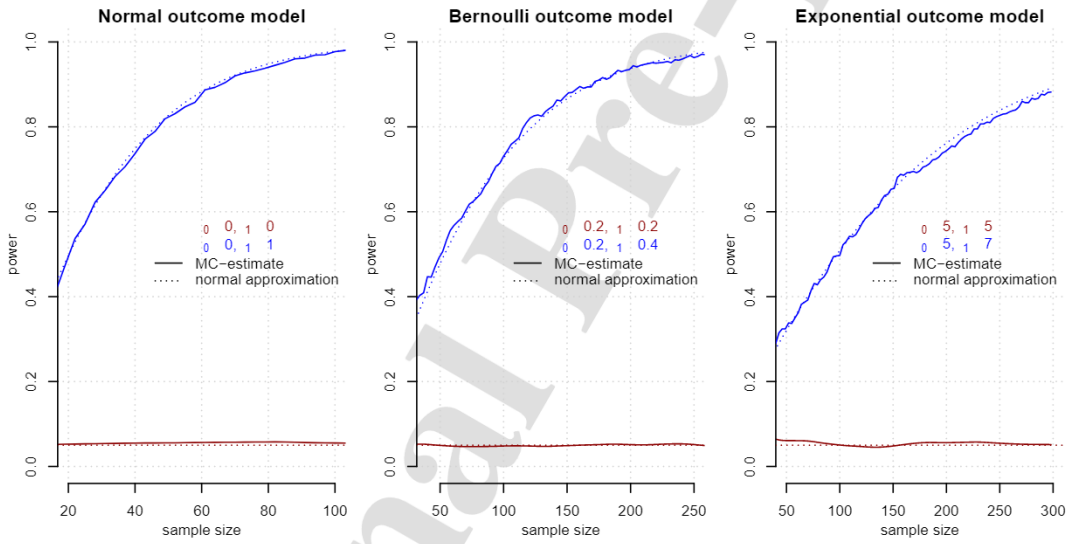

Figure 2: Power (blue lines) and type I error (red lines): comparison of estimates based on asymptotic approximations (dotted lines) and standard Monte Carlo simulations (1000 simulated trials, solid lines), for binary, continuous and time-to-event outcomes.

5 Convergence results beyond the NEF

We extend the almost sure convergence of the allocation proportion and randomization probability of a BUD (Proposition 1) to outcome distributions beyond the NEF. The following Lemma introduces approximations of the information increment (1) of BUDs with

utility (5), where $K = 2$ and θ_a is not required to be the mean of the outcomes as in Section 3. We use $X(t) = o_P(a(t))$ to indicate that $X(t)/a(t)$ converges to zero in probability.

Lemma 3. *Consider two-arm BUDs with information metric $u(\Sigma_t)$ in (5). The parameter space $\Theta \subset \mathbb{R}$ is a bounded open interval, the parameter $\theta_{0,a}$ is an interior point of Θ for $a \in \{0,1\}$ and the prior is the uniform distribution on Θ. If (i) $\inf_{y,\theta_a} f_{\theta_a}(y) > 0$, (ii) $\sup_{y,\theta_a} f_{\theta_a}(y) < \infty$, and (iii) $\sup_{y,\theta_a} \left| \frac{\partial^k f_{\theta_a}(y)}{\partial \theta_a^k} \right| < \infty$ for $k = 1, 2, 3$, then*

$$\Delta_t(a) = I_{\theta_{0,a}}^{-1} \times (t\hat{p}_{t,a})^{-2} + o_P((t\hat{p}_{t,a})^{-2}). \tag{16}$$

The following proposition states that, under the assumptions of Lemma 3, the asymptotic convergence (7) and (8) also hold outside the NEF.

Proposition 3. *Under the assumptions of Lemma 3, it holds that*

$$\hat{p}_{t,a} \longrightarrow \rho_a := \frac{I_{\theta_{0,a}}^{-\frac{h}{2h+1}}}{I_{\theta_{0,0}}^{-\frac{h}{2h+1}} + I_{\theta_{0,1}}^{-\frac{h}{2h+1}}} \ a.s. \ \text{as } t \to \infty. \tag{17}$$

and

$$p_{t,a} \longrightarrow \rho_a \ a.s. \ \text{as } t \to \infty. \tag{18}$$

for $a \in \{0,1\}$.

Note that assumption (i) of Lemma 3 implies that the support of the outcome distribution f_{θ_a} is bounded. The regularity conditions of Lemma 3 can be modified, for example to cover settings where the outcome support is unbounded. A list of alternative assumptions is specified in the Supplementary material.

To illustrate the result we consider as an outcome model f_{θ_a} beyond the NEF the truncated Weibull model $f_{\theta_a}(y) = \frac{e^{-(yr)^{\theta_a}} (ry)^{\theta_a - 1} r\theta_a}{1 - e^{-(t_0 r)^{\theta_a}}}$, $y \in (0, t_0)$, with unknown shape parameter θ_a and known rate r parameter.

The top row of Figure 3 shows a trajectory of $p_{t,1}$ (Column A) and $\hat{p}_{t,1}$ (Column B) $t = 1, \ldots, 6000$ for a single simulation of a two-arm BUD trial with $r = 1, \theta_0 = 1, \theta_1 = 1.5$

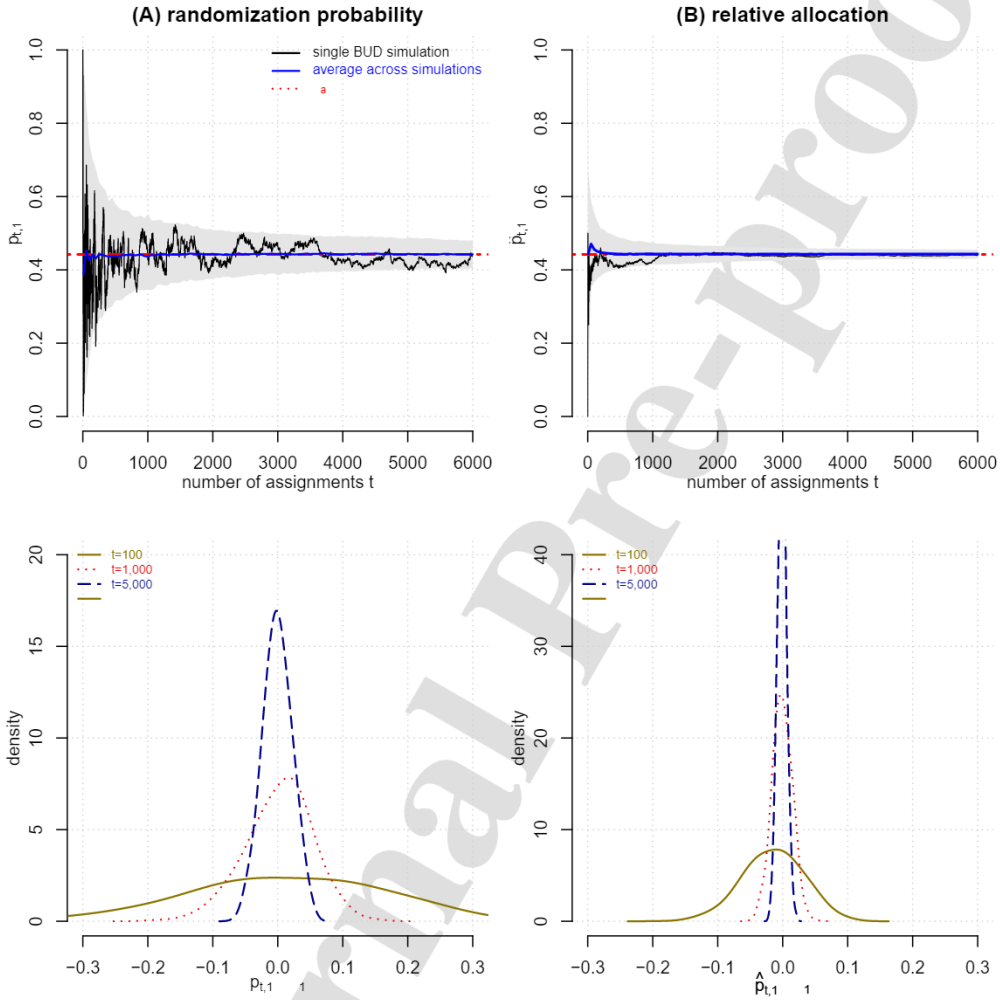

Figure 3: Allocation proportions and randomization probabilities of a two-arm BUD design. The primary outcomes are modeled with a truncated Weibull distribution. The panels in the first row compare the average allocation proportion and randomization probability across simulations (blue lines) with their limit ($t \to \infty$, red lines). The panels in the second row show kernel-density estimates of the distributions of $p_{t,1} - \rho_1$ and $\hat{p}_{t,1} - \rho_1$ across simuilations.
.

(black curve). The shaded area shows (point-wise at each t) upper and lower 2.5% quantiles of the distribution of $p_{t,1}$ and $\hat{p}_{t,1}$ across simulations. The blue lines indicate the means of $p_{t,1}$ and $\hat{p}_{t,1}$ across simulations, while the red horizontal lines indicate their limit ρ_a. The bottom row shows kernel-density estimates of the distribution of $p_{t,1}$ and $\hat{p}_{t,1}$ across simulations. Figure 3 shows that the empirical distribution of the relative allocation $\hat{p}_{t,1}$ concentrates around ρ_1 with increasing sample size t. As expected from Proposition 3, also $\hat{p}_{t,1}$ converges to ρ_1, but we observe a higher variability of $\hat{p}_{t,1}$ around ρ_1 than for $p_{t,1}$.

6 Discussion

Asymptotic analyses of Bayesian adaptive procedures simplify the design of clinical trials and reduce the need for time-consuming simulations to evaluate operating characteristics across potential trial scenarios. We derived asymptotic results for the randomization probabilities and the allocation proportions of BUDs using stochastic approximation techniques. BUD's randomization procedure was expressed as a sequence of recursive equations which allowed the application of techniques from classical stochastic approximation theory. This allowed us to derive a central limit theorem for the relative allocation of patients to treatments and for the randomization probabilities. Potential applications of stochastic approximation theory in the analysis of trial designs have been previously discussed by [22]. In our work we showed that they allow to evaluate major operating characteristics of BUDs. We considered for example the variability of the allocation proportions during the trials and the power of the BUD design with a fixed sample size under a parameter θ of interest. These approximation can reduce the required computing time to design a clinical study. For instance, generating 10,000 trials under a BUD design with $t = 10,000$ enrollments and binary outcomes took about 50 minutes on a (single core) laptop. In comparison, the computing time for the large sample approximations and the power function for $50 \leq t \leq 10,000$ was less than 0.01 seconds.

For the design of a BUD clinical trial one can apply the presented results (Proposition 1, Corollary 2 and the normal approximation of the power function) as follows:
(i) First specify (a) the outcome model (4) within the NEF, (b) an estimate of the response $\theta_{0,0}$ to therapy $a = 0$, (c) a target treatment effect $\theta_{0,1} - \theta_{0,0} > 0$ (d) the randomization parameter $h > 0$, and (e) the target type I and II error rates α and β.
(ii) Proposition 1 and Corollary 2 then provide point estimates (Proposition 1, ρ_a) of the average number of allocations to arm $a = 0, 1$ and the variability (Corollary 2, $\sqrt{\Psi/t}$) of number allocations $\widehat{p}_{t,a}$ around ρ_a during the trial.
(iii) To determine the power of the BUD design under $H_1 : \theta_{0,1} > \theta_{0,0}$ for different sample size t, we first compute the Fisher information $I_{\theta_{0,a}}$, $a = 0, 1$ for the NEL, and set $\eta_{0,a} := \left(\rho_a I_{\theta_{0,a}}\right)^{-1}$. We then use the approximation of the power function outlined in Section 4, and determine a sample size t of the BUD trial with power $1 - \beta$ and type I error rate tα.

The stochastic approximation framework, as we showed in our examples, enables useful approximations of the patients' assignment variability and other characteristics.

References

[1] Zhi-Dong Bai and Feifang Hu. Asymptotic theorems for urn models with nonhomogeneous generating matrices. *Stochastic Processes and Their Applications*, 80(1):87–101, 1999.

[2] Zhi-Dong Bai and Feifang Hu. Asymptotics in randomized urn models. *The Annals of Applied Probability*, 15(1B):914–940, Feb 2005. ISSN 1050-5164. doi: 10.1214/105051604000000774. URL http://dx.doi.org/10.1214/105051604000000774.

[3] AD Barker, CC Sigman, GJ Kelloff, NM Hylton, DA Berry, and LJs Esserman. I-spy 2: an adaptive breast cancer trial design in the setting of neoadjuvant chemotherapy. *Clinical Pharmacology & Therapeutics*, 86(1):97–100, 2009.

[4] Jay Bartroff, Tze Leung Lai, et al. Approximate dynamic programming and its applications to the design of phase i cancer trials. *Statistical Science*, 25(2):245–257, 2010.

[5] Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive algorithms and stochastic approximations*, volume 22. Springer Science & Business Media, 2012.

[6] José M Bernardo and Adrian FM Smith. *Bayesian theory*, volume 405. John Wiley & Sons, 2009.

[7] Donald A Berry. Modified two-armed bandit strategies for certain clinical trials. *Journal of the American Statistical Association*, 73(362):339–345, 1978.

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

[9] Donald A Berry and Bert Fristedt. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). *London: Chapman and Hall*, 5:71–87, 1985.

[10] Donald A Berry, Peter Mueller, Andy P Grieve, Michael Smith, Tom Parke, Richard Blazek, Neil Mitchard, and Michael Krams. Adaptive bayesian designs for dose-ranging drug trials. In *Case studies in Bayesian statistics*, pages 99–181. Springer, 2002.

[11] Donald A Berry et al. Bayesian statistics and the efficiency and ethics of clinical trials. *Statistical Science*, 19(1):175–187, 2004.

[12] Scott M Berry, Bradley P Carlin, J Jack Lee, and Peter Muller. *Bayesian adaptive methods for clinical trials*. CRC press, 2010.

[13] Medical Research Council et al. Streptomycin treatment of pulmonary tuberculosis. *British Medical Journal*, 2:769–782, 1948.

[14] Persi Diaconis and Donald Ylvisaker. Conjugate priors for exponential families. *The Annals of Statistics*, 7(2):269–281, 1979.

[15] Meichun Ding, Gary L Rosner, and Peter Müller. Bayesian optimal design for phase ii screening trials. *Biometrics*, 64(3):886–894, 2008.

[16] I Domenicano, S Ventz, M Cellamare, RH Mak, and L Trippa. Bayesian uncertainty-directed dose finding designs. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 68(5):1393–1410, 2019.

[17] Jeffrey R Eisele and Michael B Woodroofe. Central limit theorems for doubly adaptive biased coin designs. *The Annals of Statistics*, pages 234–254, 1995.

[18] Andrea Ghiglietti, Anand N Vidyashankar, William F Rosenberger, et al. Central limit theorem for an adaptive randomly reinforced urn model. *The Annals of Applied Probability*, 27(5):2956–3003, 2017.

[19] Feifang Hu, Li-Xin Zhang, et al. Asymptotic properties of doubly adaptive biased coin designs for multitreatment clinical trials. *The Annals of Statistics*, 32(1):268–301, 2004.

[20] Feifang Hu, Li-Xin Zhang, Xuming He, et al. Efficient randomized-adaptive designs. *The Annals of Statistics*, 37(5A):2543–2560, 2009.

[21] Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.

[22] Sophie Laruelle and Gilles Pagès. Randomized urn models revisited using stochastic approximation. *The Annals of Applied Probability*, 23(4):1409–1436, 2013.

[23] Carl N. Morris. Natural exponential families with quadratic variance functions. *The Annals of Statistics*, 10(1):65–80, 1982.

[24] Carl N. Morris. Natural exponential families with quadratic variance functions: Statistical theory. *The Annals of Statistics*, 11(2):515–529, 1983.

[25] Peter Müller, Don A Berry, Andrew P Grieve, and Michael Krams. A bayesian decision-theoretic dose-finding trial. *Decision analysis*, 3(4):197–207, 2006.

[26] William F Rosenberger. Randomized urn models and sequential design. *Sequential Analysis*, 21(1-2):1–28, 2002.

[27] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Operations Research*, 66(1):230–252, 2017.

[28] Lorenzo Trippa and Brian Michael Alexander. Bayesian baskets: a novel design for biomarker-based clinical trials. *Journal of Clinical Oncology*, pages JCO–2016, 2016.

[29] Steffen Ventz, William T Barry, Giovanni Parmigiani, and Lorenzo Trippa. Bayesian response-adaptive designs for basket trials. *Biometrics*, 73(3):905–915, 2017.

[30] Steffen Ventz, Matteo Cellamare, Sergio Bacallado, and Lorenzo Trippa. Bayesian uncertainty directed trial designs. *Journal of the American Statistical Association*, pages 1–13, 2018.

[31] James MS Wason and Lorenzo Trippa. A comparison of bayesian adaptive randomization and multi-stage designs for multi-arm clinical trials. *Statistics in medicine*, 33 (13):2206–2221, 2014.

[32] LJ Wei et al. The generalized polya's urn design for sequential medical trials. *The Annals of Statistics*, 7(2):291–296, 1979.

[33] Li-Xin Zhang. Central limit theorems of a recursive stochastic algorithm with applications to adaptive designs. *The Annals of Applied Probability*, 26(6):3630–3658, 2016.

[34] Xian Zhou, Suyu Liu, Edward S Kim, Roy S Herbst, and J Jack Lee. Bayesian adaptive design for targeted therapy development in lung cancer?a step toward personalized medicine. *Clinical Trials*, 5(3):181–193, 2008.