MOX–Report No. 60/2013

An urn model to construct an efficient test procedure for response adaptive designs

Ghiglietti, A.; Paganoni, A.M.

MOX, Dipartimento di Matematica "F. Brioschi"
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox@mate.polimi.it http://mox.polimi.it

An urn model to construct an efficient test procedure for response adaptive designs

Andrea Ghiglietti[a], Anna Maria Paganoni[a]

November 27, 2013

[a] Dipartimento di Matematica "F. Brioschi"
Politecnico di Milano
via Bonardi 9, 20133 Milano, Italy
andrea.ghiglietti@polimi.it
anna.paganoni@polimi.it

Abstract

We study statistical performance of different tests for comparing the mean effect of two treatments. Given a test $\mathcal{T}_0$, we determine which sample size and proportion allocation guarantee to a test $\mathcal{T}$ to be better than $\mathcal{T}_0$, in terms of (a) higher power and (b) fewer subjects assigned to the inferior treatment. The adoption of a response adaptive design to implement the random allocation procedure is necessary to ensure that both (a) and (b) are satisfied. In particular, we propose to use a Modified Randomly Reinforced Urn design (MRRU) and we show how to perform the model parameters selection for the purpose of this paper. The opportunity of relaxing some assumptions is examined. Results of simulation studies on the test performance are reported and a real case study is analyzed.
Keywords: Response adaptive designs; Clinical trials; Randomly Reinforced Urns; Tests based on adaptive procedures.

1 Introduction

In this paper, we focus on statistical performance of an hypothesis test for comparing the means of two populations. The procedures introduced in the paper are illustrated within a clinical trial framework, even if the generality of the mathematical setting would allow the method to be applied to a broad set of applications. So, we consider a clinical trial aiming at comparing the mean effect of two competitive treatments, say R and W. We assume the normality of the responses to the treatments. We consider a test $\mathcal{T}_0 = (p_0, n_0)$ that involves n_0 patients with a fixed proportion p_0 of subjects allocated to treatment R. We do not specific the experimental design used to implement the test $\mathcal{T}_0$. In Section 2 we consider a test $\mathcal{T}$ with different sample

size and proportion allocation and we realize a statistical study on the comparison of their performance. In particular, the analysis aims at determining which characteristics guarantee to a test $\mathcal{T}$ to perform better than $\mathcal{T}_0$, in terms of (a) higher power and (b) fewer subjects assigned to the inferior treatment. To attain both these goals we adopt a response adaptive design, i.e. an allocation procedure able to change its strategy during the trial depending on the responses collected by that moment. In a clinical setting, adaptive designs are very attractive because they aim to achieve two simultaneously goals, concerning both statistical and ethical purposes: (i) collect evidence to determine the superior treatment, and (ii) minimize the number of units allocated to the inferior treatment. For a complete literature review on response adaptive designs see [10, 12].

The adaptive procedure we propose to adopt is the *Modified Randomly Reinforced Urn* design (MRRU) introduced in [1]. A wide class of response-adaptive randomized designs is based on urn models, because it is a classical tool to guarantee a randomized device [5, 18]. Asymptotic results concerning urn models with an irreducible mean reinforcement matrix could be found in [2, 3, 5, 11, 18]. This irreducibility assumption is not satisfied by the Randomly Reinforced Urn (RRU) studied in [13, 16, 17], which has a diagonal mean replacement matrix. The RRU models were introduced by [6] for binary responses, applied to the dose-finding problems in [7, 8] and then extended to the case of continuous responses by [4, 16]. An interesting property concerning RRU models is that the probability to allocate units to the best treatment converges to one as the sample size increases, that is a very attractive feature from an ethical point of view. However, because of this asymptotic behavior, RRU models are not in the class of designs targeting a proportion in $(0, 1)$, that usually is fixed ad hoc or computed by satisfying some optimality criteria. Hence, all the asymptotic desirable properties concerning these procedures presented in literature (see for instance [14] and [15]), are not straightforwardly fulfilled by the RRU designs. Then, in [1] the Modified Randomly Reinforced Urn design (MRRU) was introduced in order to target any prespecified asymptotic allocation proportion in $(0, 1)$ with still an urn design. Section 3 focuses on describing the MRRU model. We also report some results proved in [1] and [9] concerning the almost sure convergence of the urn composition and of the proportion of balls of a specific color sampled from the urn, when the reinforcement means are different. These asymptotic results, together with Proposition 3.3 proved in [15], have been crucial to construct the theory of this paper.

Section 4 is focused on the urn parameter selection in order to use the MRRU model to construct the competitor test $\mathcal{T}$ described in Section 2. In Section 5 some assumptions on reinforcement distribution made in Section 2 are relaxed. We focus on the case of responses with unknown variances and non Gaussian response distribution (exponential and Bernoulli). These analysis has been mainly realized through

simulations.
Section 6 is focused on the simulation studies that have been computed to illustrate and support the theory presented in the paper.
In Section 7 we report the analysis of a real case study realized with the tools presented in the paper.
A short conclusion ends the paper (Section 8).
Data analysis and simulations have been carried out using R statistical software [19].

2 The *proportion - sample size* space

This section focuses on the statistical properties of the classical hypothesis test aiming at comparing the means of two Gaussian populations. Even if the mathematical framework is very general and the results shown in this section hold for many designs used in different areas, this paper is set in the context of clinical trials. However, in this paper the context will be represented by clinical trials. The goal of the study is the comparison among the response means to two competing treatments, the patients are sequentially assigned to. The allocation rule applied to the sequence of patients depends on the specific experimental design adopted in the trial. Let us fix $p_0 \in (0,1)$. Consider any procedure able to allocate a proportion of patients p_0 to treatment R, $1-p_0$ to treatment W. Let $n_0 \in \mathbb{N}$ be the total number of subjects involved in the experiment. In what follows, $n_{0,R}$ and $n_{0,W}$ indicate the number of subjects assigned to treatment R and W, respectively ($n_{0,R} + n_{0,W} = n_0$). Moreover, we denote

- $M_1, M_2, .., M_{n_{0,R}}$: the responses to treatment R, modeled as i.i.d. random variables with distribution μ_R and expected value m_R
- $N_1, N_2, .., N_{n_{0,W}}$: the responses to treatment W , modeled as i.i.d. random variables with distribution μ_W and expected value m_W

We assume the distributions to be Gaussian, i.e. $\mu_R = \mathcal{N}(m_R, \sigma_R^2)$ and $\mu_W = \mathcal{N}(m_W, \sigma_W^2)$, with known variances. Consider the classical hypothesis test

$$H_0 : \ m_R - m_W = 0 \qquad vs \qquad H_1 : \ m_R - m_W \neq 0. \tag{2.1}$$

In this context the critical region and the power curve of the test are well known. Let us first fix

- α : the significance level of the test;
- Δ_0 : the smallest difference among the means detected with high power;
- β_0 : the minimum power for a difference among the means of $\pm\Delta_0$;

Then, once fixed the proportion p_0, it is univocally determined the value of the sample size n_0 which allows the test to satisfy the proprieties required by those parameters. Moreover, we have the following

expression for critical region of level α

$$R_\alpha = \left\{ \ |\overline{M}_{n_{0,R}} - \overline{N}_{n_{0,W}}| > \sqrt{\frac{\sigma_R^2}{n_{0,R}} + \frac{\sigma_W^2}{n_{0,W}}} z_{\frac{\alpha}{2}} \ \right\} \qquad (2.2)$$

where $\overline{M}_{n_{0,R}} = \sum_{i=1}^{n_{0,R}} M_i/n_{0,R}$ and $\overline{N}_{n_{0,W}} = \sum_{i=1}^{n_{0,W}} N_i/n_{0,W}$ and $z_{\frac{\alpha}{2}}$ is the quantile of order $1-\alpha/2$ of a standard normal distribution. Furthermore, the power of the test (2.2), is a function of the real difference $\Delta = m_R - m_W$ (see Figure 1 in the case of equal variances), i.e.

$$\beta(\Delta) = P\left(Z < -z_{\frac{\alpha}{2}} - \frac{\Delta}{\sqrt{\frac{\sigma_R^2}{n_{0,R}} + \frac{\sigma_W^2}{n_{0,W}}}}\right) + P\left(Z > z_{\frac{\alpha}{2}} - \frac{\Delta}{\sqrt{\frac{\sigma_R^2}{n_{0,R}} + \frac{\sigma_W^2}{n_{0,W}}}}\right)$$

Let us call $\mathcal{T}_0$ the test defined in (2.2), with n_0 as sample size and p_0 as

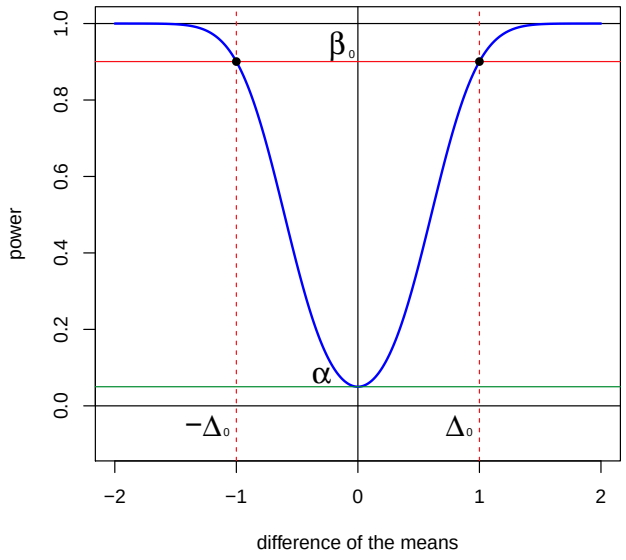

Figure 1: The picture represents the power function $\beta : \mathbb{R} \to [0,1]$ of the test defined in (2.2), in the case of $\alpha = 0.05$ and $\beta_0 = 0.9$.

proportion of patients allocated to the treatment R. To construct a test with equal parameters (α, Δ_0, β_0) and better statistical performance, the proportion of assignment or the sample size has to be conveniently modified. The test $\mathcal{T}_0$ could be represented in the space $((0,1) \times \mathbb{N})$, that we call *proportion - sample size* space, by the couple (p_0, n_0). Any other test $\mathcal{T}$ can be represented by a point (ρ, n) in the same space. The goal of this section is to point out regions of this space characterized by tests performing better than $\mathcal{T}_0$. A test $\mathcal{T}$ will be considered strictly better than $\mathcal{T}_0$ if it satisfies both the following conditions

(a) $\mathcal{T}$ has a power function uniformly higher than the power function of $\mathcal{T}_0$;

(b) $\mathcal{T}$ assigns to the worst treatment fewer patients than $\mathcal{T}_0$.

Let us call $\beta_{\mathcal{T}_0}$ and $\beta_{\mathcal{T}}$, the power functions of the tests $\mathcal{T}_0$ and $\mathcal{T}$ respectively. To achieve condition (a) we impose the following constraint

$$\beta_{\mathcal{T}}(\Delta) \geq \beta_{\mathcal{T}_0}(\Delta) \quad \forall \Delta \in \mathbb{R} \;\Leftrightarrow\; \frac{\sigma_M^2}{n\rho} + \frac{\sigma_N^2}{n(1-\rho)} \leq \frac{\sigma_M^2}{n_0 p_0} + \frac{\sigma_N^2}{n_0(1-p_0)} \tag{2.3}$$

Now, if we denote as p_{opt} the Neyman allocation proportion $\frac{\sigma_M}{\sigma_M+\sigma_N}$, we can rewrite inequality (2.3) in a more suitable form

$$\frac{{p_{opt}}^2}{n\rho} + \frac{(1-p_{opt})^2}{n(1-\rho)} \leq \frac{{p_{opt}}^2}{n_0 p_0} + \frac{(1-p_{opt})^2}{n_0(1-p_0)} \tag{2.4}$$

Inequality (2.4) divides the *proportion - sample size* space in two regions. The boundary is computed by imposing the equality in (2.4) and expressing the sample size n as a function of the proportion ρ.

$$n_\beta(\rho) \;=\; \left(\frac{p_{opt}^2}{\rho} + \frac{(1-p_{opt})^2}{1-\rho}\right)\left(\frac{p_{opt}^2}{n_0 p_0} + \frac{(1-p_{opt})^2}{n_0(1-p_0)}\right)^{-1} \tag{2.5}$$

We refer to function (2.5) as n_β, since it was computed by imposing the condition related with the power of the test β. This relationship between ρ and n is visualized in Figure 2 by a red line. Each point over this curve is a test $\mathcal{T}$ with a power uniformly higher than $\mathcal{T}_0$. Points under the red line represent tests with a power uniformly lower than $\mathcal{T}_0$. Notice that the function $n_\beta : (0,1) \to (0,\infty)$ expressed in (2.5) grows boundlessly for proportions close to zero and to one and its global minimum is reached in $\rho = p_{opt}$. This is reasonable as p_{opt} is the allocation proportion which requires the minimum number of patients to get any fixed value of power. Besides, the farther is proportion ρ from p_{opt}, the greater is the number of subjects necessary to get that power. More specifically, the minimum lies on a very interesting curve, which is univocally identified by the parameters of the classical test. Denoting with $g_{min} : (0,1) \to (0,\infty)$ the function associated with that curve, we are able to express it in an analytic form

$$g_{min}(x) \;=\; n_0\left(\frac{x^2}{p_0} + \frac{(1-x)^2}{1-p_0}\right)^{-1} \quad \forall x \in (0,1) \tag{2.6}$$

The curve is represented in Figure 2 by a red dotted line. The functions n_β and g_{min} cross in two points, in general different, that we denote M and Q. The point M is the minimum of the function n_β and it corresponds to the Neyman allocation proportion

$$M = \left(p_{opt} \;\;,\;\; n_0\left(\frac{p_{opt}^2}{p_0} + \frac{(1-p_{opt})^2}{1-p_0}\right)^{-1}\right) \tag{2.7}$$

The point Q is the maximum of the function g_{min} and it corresponds to the test $\mathcal{T}_0$: $Q = (p_0, n_0)$. The points M and Q coincide only when $p_0 = p_{opt}$. In this case, the curves n_β and g_{min} are tangents in $M \equiv Q$. Moreover, there are other relevant points highlighted by the function g_{min}. In fact, the curve starts in $X_{W,0} = (0, n_0(1-p_0))$ and ends in $X_{R,0} = (1, n_0 p_0)$. The ordinates of points $X_{W,0}$ and $X_{R,0}$ tell us how many patients have been allocated by the test $\mathcal{T}_0$ to the treatment W and R, respectively.
To satisfy (b) we have to distinguish two different cases, depending on which is the superior treatment

- if $m_R > m_W \Rightarrow$ the superior treatment is R and the condition to be imposed is

$$n(1-\rho) < n_0(1-p_0) \quad \Leftrightarrow \quad \rho > 1 - \frac{n_0}{n}(1-p_0); \tag{2.8}$$

- if $m_R < m_W \Rightarrow$ the superior treatment is W and the condition to be imposed is

$$n\rho < n_0 p_0 \quad \Leftrightarrow \quad \rho < \frac{n_0}{n} p_0. \tag{2.9}$$

Both these constraints are depicted in blue in the *proportion - sample size* plane. Below each of these lines, the first or the second condition is verified. In conclusion, we divided the *proportion - sample size* space in three regions:

- Region A :

$$A \;=\; \left\{ \; (x,y) \in (0,1) \times (0,\infty) \;:\; n_\beta(x) < y < \frac{p_0}{x} n_0 \; \right\}$$

  tests $\mathcal{T} \in A$ have a power uniformly higher and allocate to treatment R less patients than $\mathcal{T}_0$.

- Region B :

$$B \;=\; \left\{ \; (x,y) \in (0,1) \times (0,\infty) \;:\; y > \max\left\{ \frac{p_0}{x}; \frac{1-p_0}{1-x} \right\} \cdot n_0 \; \right\}$$

  tests $\mathcal{T} \in B$ have a power uniformly higher and allocate to both treatments more patients than $\mathcal{T}_0$.

- Region C :

$$C \;=\; \left\{ \; (x,y) \in (0,1) \times (0,\infty) \;:\; n_\beta(x) < y < \frac{1-p_0}{1-x} n_0 \; \right\}$$

  tests $\mathcal{T} \in C$ have a power uniformly higher and allocate to treatment W less patients than $\mathcal{T}_0$.

Hence, a test $\mathcal{T}$ with better performance than $\mathcal{T}_0$ is a point (ρ, n) in the region A if $m_R < m_W$, or in the region C if $m_R > m_W$. Unfortunately, the experimenter cannot know which is the superior treatment before conducting the trial. For this reason, it could be useful to adopt a response adaptive design to construct the test, since this method is able

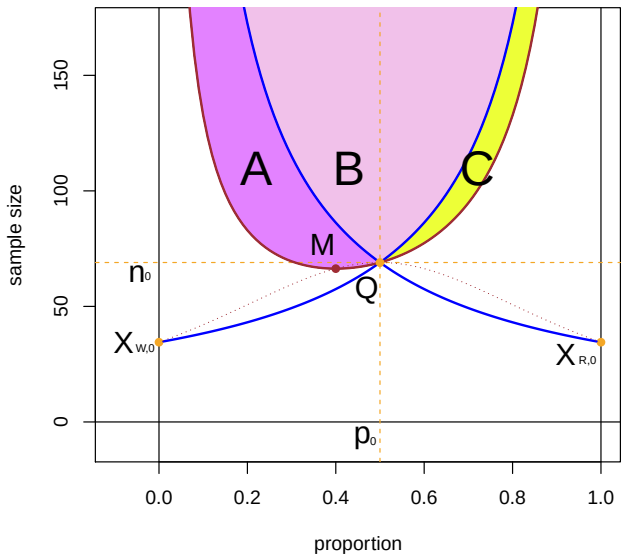

Figure 2: The picture represents the regions A, B and C, on the *proportion - sample size* plane. The red line represents the function n_β in (2.5); it separates the test $\mathcal{T}$ with power $\beta_{\mathcal{T}}(\Delta) > \beta_{\mathcal{T}_0}(\Delta)$, from the test with power $\beta_{\mathcal{T}}(\Delta) < \beta_{\mathcal{T}_0}(\Delta)$. Blue lines separates tests according on the number of patients allocated to the treatments R and W, with respect to $n_{0,R}$ and $n_{0,W}$. The dotted red line represents the function g_{min} in (2.6).

to target different allocation proportions according to the responses collected during the trial.
Let us introduce a vector $(X_1, X_2, ..., X_n) \in \{0; 1\}^n$ composed by the allocations to the treatments according to the adaptive design, i.e. $X_i = 1$ if the subject i receives treatment R or $X_i = 0$ if the subject i receives treatment W. Then, we define the quantities $N_R(n) = \sum_{i=1}^n X_i$ and $N_W(n) = \sum_{i=1}^n (1 - X_i)$, that represent the number of patients allocated to treatments R and W, respectively. Notice that the sample sizes $N_R(n)$ and $N_W(n)$ are random variables. Let us also define the adaptive estimators based on the observed responses until time n, i.e.

$$\overline{M}(n) = \frac{\sum_{i=1}^n X_i M_i}{N_R(n)} \qquad \text{and} \qquad \overline{N}(n) = \frac{\sum_{i=1}^n (1 - X_i) N_i}{N_W(n)}. \quad (2.10)$$

Then, the test $\mathcal{T}$ is defined by the following critical region

$$R_\alpha^{adaptive} = \left\{ \; |\overline{M}(n) - \overline{N}(n)| > \sqrt{\frac{\sigma_R^2}{N_R(n)} + \frac{\sigma_W^2}{N_W(n)}} z_{\frac{\alpha}{2}} \; \right\} \quad (2.11)$$

whose properties depend on the type of adaptive design has been applied in the trial. The authors propose to adopt the *Modified Randomly*

Reinforced Urn design (MRRU) described in [1]. The scheme of this urn model is presented in the next section.

3 The Modified Randomly Reinforced Urn Design

Let us consider the response probability laws μ_R and μ_W introduced in Section 2. In general, we can define an opportune utility function u to turn the responses into values which can be interpretable as urn reinforcements. To ease of notation, in this paper we will use the identity as utility function, i.e. we will interpret the response distributions to treatment R and W as the reinforcement distributions of red and white balls, respectively. The model requires some assumptions on the reinforcement probability laws μ_R and μ_W: both the supports are contained in $[a, b]$, where $0 \leq a \leq b < +\infty$, both the means $m_R = \int_a^b x\mu_R(dx)$ and $m_W = \int_a^b x\mu_W(dx)$ are strictly positive, and both the variances σ_R^2 and σ_W^2 are finite. So, without loss of generality, we consider the superior treatment as the one associated to the color with higher reinforcement mean.
Now, let us describe the urn model. First, let $(U_n)_n$ be a sequence of independent uniform random variables on $[0, 1]$. Then, visualize an urn initially containing r_0 balls of color R and w_0 balls of color W. Set

$$R_0 = r_0,\ W_0 = w_0,\ D_0 = R_0 + W_0,\ Z_0 = \frac{R_0}{D_0}.$$

At time $n = 1$, a ball is sampled from the urn; its color is $X_1 = \mathbf{1}_{[0,Z_0]}(U_1)$, a random variable with Bernoulli(Z_0) distribution. Let M_1 and N_1 be two independent random variables with distribution μ_R and μ_W, respectively; assume that X_1, M_1 and N_1 are independent. Next, if the sampled ball is R, it is replaced in the urn together with X_1M_1 balls of the same color if $Z_0 < \eta$, where $\eta \in (0, 1)$ is a suitable parameter, otherwise the urn composition does not change; if the sampled ball is W, it is replaced in the urn together with $(1 - X_1)N_1$ balls of the same color if $Z_0 > \delta$, where $\delta < \eta \in (0, 1)$ is a suitable parameter, otherwise the urn composition does not change. So doing we update the urn composition in the following way

$$\begin{aligned} R_1 &= R_0 + X_1M_1\mathbf{1}_{[Z_0<\eta]}, \\ W_1 &= W_0 + (1 - X_1)N_1\mathbf{1}_{[Z_0>\delta]}, \\ D_1 &= R_1 + W_1,\ Z_1 = \frac{R_1}{D_1}. \end{aligned} \tag{3.1}$$

Now iterate this sampling scheme forever. Thus, at time $n + 1$, given the sigma-field $\mathcal{F}_n$ generated by $X_1, ..., X_n, M_1, ..., M_n$ and $N_1, ..., N_n$, let $X_{n+1} = \mathbf{1}_{[0,Z_n]}(U_{n+1})$ be a Bernoulli(Z_n) random variable and, independently of $\mathcal{F}_n$ and X_{n+1}, assume that M_{n+1} and N_{n+1} are two independent random variables with distribution μ_R and μ_W, respec-

tively. Set

$$\begin{aligned} R_{n+1} &= R_n + X_{n+1} M_{n+1} \mathbf{1}_{[Z_n < \eta]}, \\ W_{n+1} &= W_n + (1 - X_{n+1}) N_{n+1} \mathbf{1}_{[Z_n > \delta]}, \\ D_{n+1} &= R_{n+1} + W_{n+1}, \\ Z_{n+1} &= \frac{R_{n+1}}{D_{n+1}}. \end{aligned} \tag{3.2}$$

$Z = (Z_n, n = 1, 2, ...)$ is a sequence of random variables that are proportions of balls of color R in the urn. The $(n+1)$-patient entering the study is allocated to treatment R according to this probability Z_n. The Modified Randomly Reinforced Urn model has been widely studied in [1] and [9]. In those works, many convergence results have been proved and the asymptotic behavior of the urn process has been discussed. The properties highlighted in those papers make the MRRU design very attractive from many point of view and for this reason we decided to adopt this model as the adaptive design to be applied in the framework presented in Section 2. In particular, in [1] the following result is proved

Theorem 3.1 *The sequence of proportions $(Z_n)_{n\in\mathbb{N}}$ of the urn process converges almost surely and*

$$lim_{n\to\infty} Z_n = \begin{cases} \eta & \text{if } \int_a^b x\mu_R(dx) > \int_a^b x\mu_W(dx), \\ \delta & \text{if } \int_a^b x\mu_R(dx) < \int_a^b x\mu_W(dx). \end{cases} \tag{3.3}$$

The urn proportion process $(Z_n)_{n\in\mathbb{N}}$ converges to a value which depends on the unknown means of the reinforcement distributions. This aspect characterizes the adaptive nature of the design based on the urn model. In particular, this urn model generates a process $(Z_n)_{n\in\mathbb{N}}$ that converges almost surely to one of the values $\{\delta, \eta\}$, according to which reinforcement presents the distribution with the greatest mean. When $m_R = m_W$ we do not have the explicit form of the asymptotic distribution of the urn proportion Z_n. Nevertheless, we know that $(Z_n)_{n\in\mathbb{N}}$ converges to a random variable Z_∞ whose distribution has no atoms and its support is $S_\infty = [\delta, \eta]$. In [9] the asymptotic behavior of different quantities related to the MRRU model has been studied. One of those results concerns the number of red balls sampled from the urn divided by the total number of draws. This quantity converges almost surely to the same limit of the urn proportion of red balls Z_n.

Proposition 3.2 *The sequence $(N_R(n)/n)_{n\in\mathbb{N}}$ converges almost surely and*

$$lim_{n\to\infty} \frac{N_R(n)}{n} = \begin{cases} \eta & \text{if } \int_a^b x\mu_R(dx) > \int_a^b x\mu_W(dx), \\ \delta & \text{if } \int_a^b x\mu_R(dx) < \int_a^b x\mu_W(dx). \end{cases}.$$

Notice that this result is very useful because the number of red (white) balls sampled from the urn, divided by the number of draws, represents

exactly the proportion of patients allocated to treatment R (W).

Consider an estimation problem for the means m_R and m_W of the responses to treatments. The a.s. limit of the urn process $(Z_n)_n$ in the MRRU model is in the open interval $(0,1)$, so both the sequences $N_R(n) = \sum_{i=1}^{n} X_i$ and $N_W(n) = \sum_{i=1}^{n}(1 - X_i)$ diverge to infinity almost surely. This allows us to apply the results proved in [15] concerning the adaptive estimators (2.10). The main asymptotic result is the following:

Proposition 3.3 *The estimators $\overline{M}(n)$ and $\overline{N}(n)$ are consistent estimators of m_R and m_W, respectively. Moreover as $n \to \infty$,*

$$\left(\sqrt{N_R(n)} \frac{(\overline{M}(n) - m_R)}{\sigma_R}, \sqrt{N_W(n)} \frac{(\overline{N}(n) - m_W)}{\sigma_W} \right) \to (Z_1, Z_2)$$

in distribution, where (Z_1, Z_2) are independent standard normal random variables.

This result gives us the asymptotic normality of the adaptive estimators. This property is not trivial since the sample sizes $N_R(n)$ and $N_W(n)$ are random quantities that also depend on the treatment responses. This result is very useful in an inferential setting, when a statistics based on the adaptive estimator is used. In fact, in that context the Proposition 3.3 provides the probability model which allows us to build the critical region or to compute the p-value of a test to compare m_R and m_W based on adaptive design driven by MRRU model.

4 The parameters selection to construct the test $\mathcal{T}$

Consider the situation presented in Section 2. Initially the problem is faced with a classical no-adaptive test. Let us denote this test as $\mathcal{T}_0$. Assume a sample size n higher than the one of the test $\mathcal{T}_0$ (i.e., $n = c \cdot n_0$ with $c > 1$). For any $n \geq n_0$, we can individuate the following intervals

- $I_n^A = \{x \in (0,1) : (x,n) \in A\}$
- $I_n^B = \{x \in (0,1) : (x,n) \in B\}$
- $I_n^C = \{x \in (0,1) : (x,n) \in C\}$

Notice that

- $I_n^A \bigcup I_n^B \bigcup I_n^C \subset (0,1)$
- $I_n^A \bigcap I_n^B = \emptyset$, $I_n^B \bigcap I_n^C = \emptyset$, $I_n^A \bigcap I_n^C = \emptyset$,

The aim is to point out an adaptive test $\mathcal{T}$ represented in the *proportion - sample size* space by a point in region A when R is the inferior

treatment, or in the I_n^C when W the inferior one. This goal is achieved when

$$\begin{cases} \frac{N_R(n)}{n} \in I_n^C & \text{if } \int_a^b x\mu_R(dx) > \int_a^b x\mu_W(dx), \\ \frac{N_R(n)}{n} \in I_n^A & \text{if } \int_a^b x\mu_R(dx) < \int_a^b x\mu_W(dx). \end{cases}.$$

Inspired by Proposition 3.2, we set $\delta \in I_n^A$ and $\eta \in I_n^C$, so that $\lim_{k\to\infty} \frac{N_R(k)}{k} \in I_n^A$ if $m_R < m_W$ and $\lim_{k\to\infty} \frac{N_R(k)}{k} \in I_n^C$ if $m_R > m_W$. This choice implies that the test $\mathcal{T}$ is in the right region, where both condition (a) and (b) are satisfied. In Figure 3 we show how the urn process Z_n converges towards the right region.

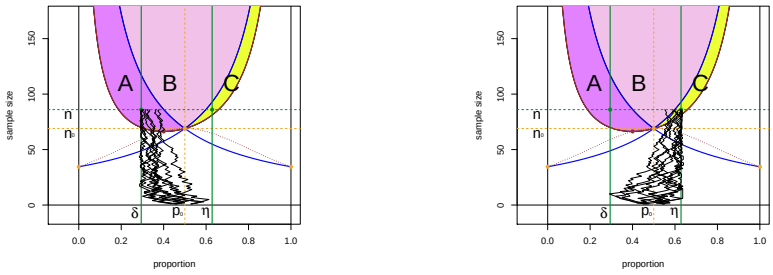

Figure 3: The pictures represents the regions A, B and C, for a particular choice of α, β_0, Δ_0 and p_0. For each fixed sample size n, the parameters of the urn model $\delta, \eta \in (0,1)$ are chosen such that $(\delta, n) \in A$ and $(\eta, n) \in C$. On the left: simulations with $m_R < m_W$. On the right: simulations with $m_R > m_W$. In both pictures, the black lines represent 10 replications of the urn process $(Z_k)_k$.

The speed of convergence of the urn model is a key point for the success of this procedure. In general, the asymptotic behavior of the urn process $(Z_n)_{n\in\mathbb{N}}$ depends on the reinforcement distributions (μ_R, μ_W) and on the parameters (δ, η). Once the assumptions on the reinforcement probability laws are made and the statistical parameters are fixed, the regions A, B, C can be determined and the rate of convergence depends only on the unknown means m_R and m_W; in particular, the speed of convergence is an increasing function of the mean distance $|\mu_R - \mu_W|$. Moreover, since the value of the sample size n has been computed as a decreasing function of Δ_0, the closeness of the urn proportion Z_n to its limit (η or δ) after n draws, depends mainly on the size of the normalized distance $\frac{|\mu_R - \mu_W|}{\Delta_0}$. If this ratio is large it means that the treatments' performance are very different with respect to the minimum relevant distance $|\Delta_0|$. In this case, the quantity $\frac{N_R(n)}{n}$ will be quickly closed to the limit of the urn process and so the procedure will actually design a test $\mathcal{T}$ which lies in the right region. At the contrary, if $\frac{|\mu_R - \mu_W|}{\Delta_0}$ is small, it means that the difference between the treatments becomes less relevant. The urn proportion $(Z_n)_{n\in\mathbb{N}}$ will be a process which slowly converges to its limit. Therefore, in this situation, the

assumption that $\frac{N_R(n)}{n}$ is a good approximation of its limit is less reasonable and so the test $\mathcal{T}$ will be easily found outside the right region. Naturally, it is useless to choose an excessively little value of Δ_0 just to increase the ratio $\frac{|\mu_R - \mu_W|}{\Delta_0}$; in fact, this change would heavy increase the sample size n_0, in order to fulfill the level and power constraint of $\mathcal{T}_0$. As a consequence, the power evaluated at the real difference of the means $\beta(\Delta)$ would be so high that there would be no need to maximize it.

There are other factors which influence the speed of convergence of the process $(Z_n)_{n\in\mathbb{N}}$, like the values of the parameters η and δ. In fact, it is known that the closer to a border point of the interval $(0,1)$ the limit is, the slower the process converges. This fact is relevant when we propose to improve the approximation of $\frac{N_R(n)}{n}$ with its limit (δ or η) by increasing the sample size n, i.e. using $\widetilde{n} = \widetilde{c} \cdot n_0$ (with $\widetilde{c} >> c$) instead of n. Naturally, since we are using more subjects here, it will be more likely that the urn proportion $Z_{\widetilde{n}}$ will be closer to the limit η (or δ), which was previously fixed in the interval I_n^A (or I_n^C). The problem is that the points $(\delta, \widetilde{n})$ and $(\eta, \widetilde{n})$ could be not in the regions A and C anymore. In fact, when we use the sample size $\widetilde{n}$ instead of n, we should locate the parameters η and δ in the intervals $I_{\widetilde{n}}^A$ and $I_{\widetilde{n}}^C$ instead of I_n^A and I_n^C; so doing, we can be sure that the points $(\delta, \widetilde{n})$ and $(\eta, \widetilde{n})$ are in the right regions. Moreover, as the sample size n grows the intervals I_n^A and I_n^C become smaller and move towards the border points 0 and 1. This slows down the convergence of the process $(Z_n)_{n\in\mathbb{N}}$ and makes negligible the initial gain obtained by increasing the sample size.

Remark 4.1 *The main inferential problem of this paper is a two-sided hypothesis test for comparing the mean effect of two treatments (2.1). It's worth to notice that nothing changes if we consider an one-sided test, where the alternative hypothesis states that one treatment is better than the other one, for instance* $H_0 : m_R \leq m_W$ *and* $H_1 : m_R > m_W$. *In this case the goal (b) reduces to assign more patients to treatment* W*, so we can fix the parameter* δ *arbitrarily in the interval* $(0, \eta)$*. In Figure 4 we show the partition of the plane proportion - sample size and the choice of the parameters* δ *and* η *with an one-sided test.*

5 Different response distributions

In this section we relax some assumptions on reinforcement distributions. First, we consider the situation with Gaussian laws but unknown variances, then, we discuss the case of non-Gaussian response distributions (exponential and Bernoulli).

In Section 2 we made the assumption that the variances of the responses' distributions σ_R^2 and σ_W^2 are known. This hypothesis is very strong and in many cases unrealistic, since the variability of a new phenomenon is typically unknown and the variance usually has to be esti-

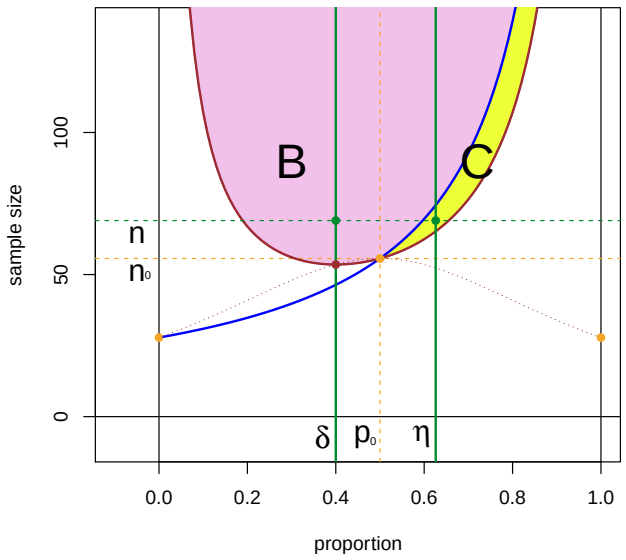

Figure 4: The picture shows the case of an one-sided test. The regions B and C are defined for a fixed level α and a test $\mathcal{T}_0$ characterized by (p_0, n_0). Once fixed a new sample size $n > n_0$, the parameters of the urn model $\delta, \eta \in (0,1)$ are chosen such that $(\delta, n) \in B$ and $(\eta, n) \in C$

mated through the same observations used to realize the test. Then, a good design should incorporate the possibility of estimating variances, updating them at each step of the procedure and maintaining the good properties obtained with known variances.
First, fix $\delta = \eta = p_0$. Then, we denote as $S_R^2(n)$ and $S_W^2(n)$ the adaptive estimators for the responses' variances, expressed as follows

$$S_R^2(n) = \frac{\sum_{i=1}^n X_i (M_i - \overline{M}(n))^2}{N_R(n) - 1}, \text{ and } S_W^2(n) = \frac{\sum_{i=1}^n (1 - X_i)(N_i - \overline{N}(n))^2}{N_W(n) - 1}. \tag{5.1}$$

So we can replace the true variances σ_R^2 and σ_W^2 with their estimators $S_R^2(i)$ and $S_W^2(i)$; then, in the critical region (2.11) the quantile of the t-student substitutes the quantile of the Gaussian distribution. Moreover, the function $n_\beta(\cdot)$ introduced in (2.5) has to be redefined as follows

$$n_\beta(\rho; i) := \left(\frac{\widehat{p}_{opt}^2(i)}{\rho} + \frac{(1 - \widehat{p}_{opt}(i))^2}{1 - \rho} \right) \left(\frac{\widehat{p}_{opt}^2(i)}{n_0 p_0} + \frac{(1 - \widehat{p}_{opt}(i))^2}{n_0 (1 - p_0)} \right)^{-1}$$

where $\widehat{p}_{opt}(i) = \frac{S_R(i)}{S_R(i) + S_W(i)}$. This procedure has to be done at every step $i \leq n$, after that a new response is collected and one of the two estimates can be updated. Notice that the function $n_\beta(\cdot; i)$ is random

and changes for any $i \leq n$, because now it depends on the observations. As a consequence, also the intervals I_i^A, I_i^B, I_i^C will be random too and we have to recompute them for any $i \leq n$. This leads to two sequences $(\delta_i)_i, (\eta_i)_i$ instead of two parameters δ, η, since we need to maintain the property that the parameters of the urn model are chosen in the corresponding intervals: $\delta_i \in I_i^A$ and $\eta_i \in I_i^C$.
In [15] it has been proved that when the sequences $N_R(n)$ and $N_W(n)$ are divergent, adaptive estimators like $S_R^2(n)$ and $S_W^2(n)$ are strongly consistent. This result implies the $n_\beta(t;i) \to_i n_\beta(t)$ almost surely for any $t \in (0,1)$. This fact ensures that it's always possible to create two convergent sequences $(\delta_i)_i \to \delta, (\eta_i)_i \to \eta$ such that $\delta \in I^A$ and $\eta \in I^C$.

When we relax the normality assumption on the reinforcements distribution it is difficult to write the power function of the test in an analytic form. It is not always possible to solve the condition $\beta_{\mathcal{T}}(\Delta) \geq \beta_{\mathcal{T}_0}(\Delta)$ and then to compute the function n_β. Anyway, this task can be realized in simulation and so we will show that the *proportion - sample size* plane can be partitioned again in the regions $A-B-C$ also with non-Gaussian reinforcements. In particular, we focus on two situations: exponential and Bernoulli responses.

Exponential responses:
Let us make the following assumptions on patients' responses

- $M_1, M_2, .., M_{n_{0,R}}$: the responses to treatment R, modeled as i.i.d. random variables with distribution $\mu_R = \mathcal{E}(\lambda_R)$
- $N_1, N_2, .., N_{n_{0,W}}$: the responses to treatment W, modeled as i.i.d. random variables with distribution $\mu_W = \mathcal{E}(\lambda_W)$

Our aim is to perform the following hypothesis test

$$H_0 : \ \lambda_R = \lambda_W \qquad vs \qquad H_1 : \ \lambda_R \neq \lambda_W. \tag{5.2}$$

We will keep the notation of Section 2. We use the likelihood ratio test to compute the critical region. The likelihood function of the whole sample is

$$\begin{aligned} L(\lambda_R, \lambda_W, data) &= \lambda_R^{n_{0,R}} \lambda_W^{n_{0,W}} \exp\left(-\lambda_R \sum_{i=1}^{n_{0,R}} M_i - \lambda_W \sum_{i=1}^{n_{0,W}} N_i \right) \\ &= \left(\lambda_R^{p_0} \lambda_W^{1-p_0} \exp\left(-\lambda_R \overline{M}_{n_{0,R}} p_0 - \lambda_W \overline{N}_{n_{0,W}} (1-p_0) \right) \right)^n \end{aligned}$$

where $\overline{M}_{n_{0,R}} = \sum_{i=1}^{n_{0,R}} M_i / n_{0,R}$ and $\overline{N}_{n_{0,W}} = \sum_{i=1}^{n_{0,W}} N_i / n_{0,W}$. Then, the likelihood ratio test gives us the following critical region

$$\left\{ \frac{\sup_{\lambda_R = \lambda_W \in (0,\infty)} L(\lambda_R, \lambda_W, data)}{\sup_{(\lambda_R, \lambda_W) \in (0,\infty)^2} L(\lambda_R, \lambda_W, data)} \ < \ c_\alpha \right\} = \left\{ \frac{\overline{M}_{n_{0,R}}^{p_0} \cdot \overline{N}_{n_{0,W}}^{1-p_0}}{\overline{M}_{n_{0,R}} \cdot p_0 + \overline{N}_{n_{0,W}} \cdot (1-p_0)} \ < \ \sqrt[n]{c_\alpha} \right\}$$

where $c_\alpha \in (0,1)$ can be determined setting the significance level of this critical region to α.

Bernoulli responses:
Let us make the following assumptions on patients' responses

- $M_1, M_2, .., M_{n_{0,R}}$: the sequence of the responses to treatment R, modeled as i.i.d. random variables with distribution $\mu_R = \mathcal{B}(p_R)$
- $N_1, N_2, .., N_{n_{0,W}}$: the sequence of the responses to treatment W, modeled as i.i.d. random variables with distribution $\mu_W = \mathcal{B}(p_W)$

Let us consider now the following hypothesis test

$$H_0 : \ p_R = p_W \qquad vs \qquad H_1 : \ p_R \neq p_W. \tag{5.3}$$

The likelihood function for two samples of Bernoulli variables is

$$L(p_R, p_W, data) \ = \\ \left(p_R^{\overline{M}_{n_{0,R}} p_0} (1-p_R)^{(1-\overline{M}_{n_{0,R}})p_0} p_W^{\overline{N}_{n_{0,W}}(1-p_0)} (1-p_W)^{(1-\overline{N}_{n_{0,W}})(1-p_0)} \right)^n$$

Then, the likelihood ratio test gives us the following critical region

$$\left\{ \ \frac{\sup_{p_R = p_W \in (0,1)} L(p_R, p_W, data)}{\sup_{(p_R, p_W) \in (0,1)^2} L(p_R, p_W, data)} \ < \ c_\alpha \ \right\} \ = \\ \left\{ \ \frac{\overline{P}^{\overline{P}} (1-\overline{P})^{1-\overline{P}}}{\overline{M}_{n_{0,R}}^{\overline{M}_{n_{0,R}} p_0} (1-\overline{M}_{n_{0,R}})^{(1-\overline{M}_{n_{0,R}})p_0} \overline{N}_{n_{0,W}}^{\overline{N}_{n_{0,W}}(1-p_0)} (1-\overline{N}_{n_{0,W}})^{(1-\overline{N}_{n_{0,W}})(1-p_0)}} \ < \ \sqrt[n]{c_\alpha} \ \right\}$$

where

$$\overline{P} = \frac{\sum_{i=1}^{n_{0,R}} M_i + \sum_{i=1}^{n_{0,W}} N_i}{n} = \overline{M}_{n_{0,R}} p_0 + \overline{N}_{n_{0,W}} (1-p_0).$$

Also in this case $c_\alpha \in (0,1)$ can be determined setting the significance level of this critical region to α.

The power function $(\widehat{\beta}_{(p_0,n_0)})$ in both cases (5.2) and (5.3) can be computed through simulations and so we can empirically compute function $n_\beta(\cdot)$ in this way: for any $\rho \in (0,1)$

$$n_\beta(\rho) \ := \ \min \left\{ \ n \geq 1 \ : \ \widehat{\beta}_{(\rho,n)} \geq \widehat{\beta}_{(p_0,n_0)} \ \right\}$$

Now that we have defined the function $n_\beta(\cdot)$, we can partition the *proportion - sample size* plane, introduce the intervals I_n^C and I_n^A and after that fix the parameters η and δ within them. When the urn model is used to allocate the patients the design becomes adaptive and the critical region should be written in a different form, replacing $\overline{M}_{n_{0,R}}$, $\overline{N}_{n_{0,W}}$ and ρ with $\overline{M}(n)$, $\overline{N}(n)$ and $\rho(n)$. As we can see from Figures 5 and 6, the structure of the regions is the same of those computed in the Gaussian response case.

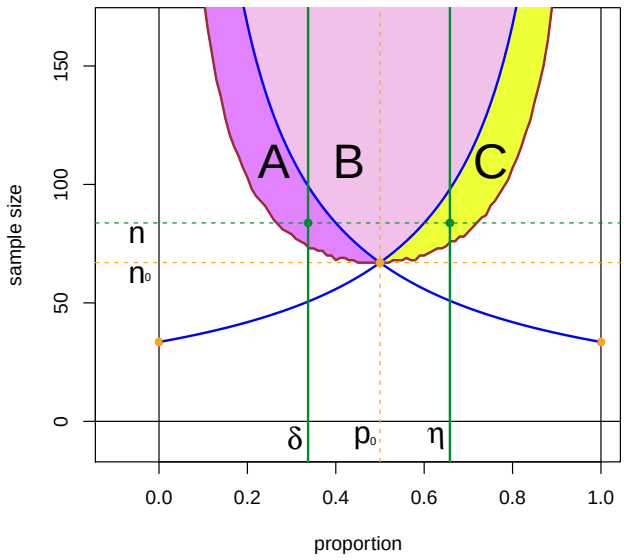

Figure 5: This is an example with exponential distributed responses ($\lambda_R = 2$ and $\lambda_W = 1$). The parameters are: $\alpha = 0.05$, $1 - \beta_0 = 0.2$, $\Delta_0 = \Delta = 1/2$. The test $\mathcal{T}_0$ uses an allocation proportion $p_0 = 1/2$ and needs a sample size of $n_0 = 67$. The red line represents the function $n_\beta(\cdot)$ computed by simulation.

6 Simulation Studies

We realized some simulation studies aiming at illustrating the theory presented in the paper. In this section, we are going to show some of those simulations; in particular, we want to highlight the good properties provided by the use of an adaptive design in the framework of Section 2.

Let us consider the two-sided hypothesis test (2.1), for comparing the mean effect of two treatments R and W. We simulated the responses to treatments R and W from two sequences of i.i.d. random variables, with probability laws μ_R and μ_W Gaussian with means m_R and m_W and variances σ_R^2 and σ_W^2, respectively. In all the simulations, $m_W = 10$ and m_R ranges from 5 to 15; we analyze separately the situation of equal variances ($\sigma_R^2 = 1.5^2, \sigma_W^2 = 1.5^2$) and different variances ($\sigma_R^2 = 1, \sigma_W^2 = 4$). We set the significance level $\alpha = 0.05$ and the minimum power $\beta_0 = 0.9$ for a difference of $\Delta_0 = 1$. We assume to have a balanced non adaptive design $p_0 = 0.5$. Then, we compute the right value for the sample size n_0 to fulfill the conditions of significance level and power set in advance, which is $n_0 = 96$ when the variances are equal and $n_0 = 106$ when the variances are different.

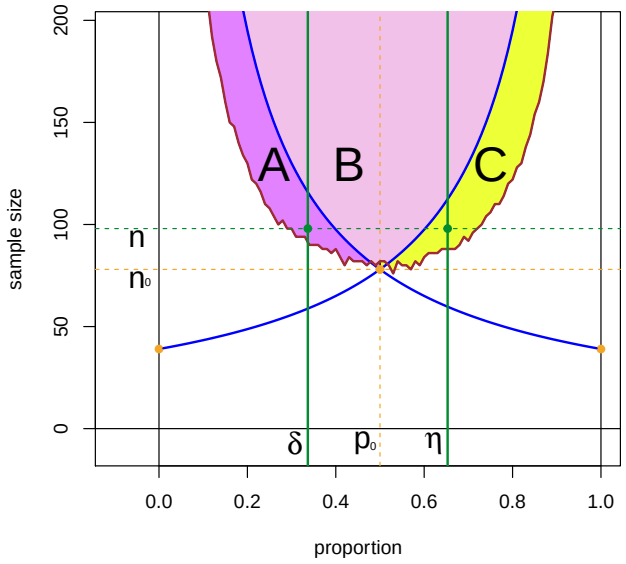

Figure 6: This is an example with Bernoulli distributed responses ($p_R = 0.2$ and $p_W = 0.5$). The parameters are: $\alpha = 0.05$, $1 - \beta_0 = 0.2$, $\Delta_0 = \Delta = 0.3$. The test $\mathcal{T}_0$ uses an allocation proportion $p_0 = 1/2$ and needs a sample size of $n_0 = 76$. The red line represents the function $n_\beta(\cdot)$ computed by simulation.

At this point, we apply the procedure described in Section 2 to get a new adaptive test $\mathcal{T}$ performing better than $\mathcal{T}_0$. The sample size of $\mathcal{T}$ has been increased of a 25% ($n = 1.25 \cdot n_0$), obtaining $n = 120$ in the case of equal variances and $n = 132$ with different variances. In both cases, we can design the regions A, B and C and the corresponding intervals I_n^A, I_n^B and I_n^C; we set δ in the center of I_n^A and η in the center of I_n^A. In particular, we have

- $\sigma_R^2 = 1.5^2$, $\sigma_W^2 = 1.5^2 \quad \Rightarrow \quad I_n^A = (0.127, 0.402)$, $I_n^C = (0.598, 0.632)$.
- $\sigma_R^2 = 1$, $\sigma_W^2 = 4 \quad \Rightarrow \; I_n^A = (0.279, 0.403)$, $I_n^C = (0.597, 0.721)$

In all simulations, the urn has been initialized with a total number of balls $d_0 = (m_R + m_W)/2$; the initial urn proportion z_0 has been set at the center of the interval (δ, η). Then, for each value of $m_R \in \{5, 7, 9, 9.5, 10.5, 11, 13, 15\}$, we have run 1000 urn processes $(Z_k)_k$ stopped at time n, following the algorithm described in Section 3. The results are reported in Table 1 (equal variances) and 2 (different variances).

The proportion of simulation runs the test $\mathcal{T}$ has a power higher than $\mathcal{T}_0$ is very high. In other words, it means that most of the simulations yields an allocation proportion after n step such that $(N_R(n)/n, n) \in$

m_R	Δ	$\#\{\beta_{\mathcal{T}} \geq \beta_{\mathcal{T}_0}\}$	$\#\{N_R(n) < n_{0,R}\}$	$\#\{N_W(n) < n_{0,W}\}$
5	-5	0.954	(0.766)	0.011
7	-3	0.967	(0.573)	0.057
9	-1	0.970	(0.320)	0.178
9.5	-0.5	0.973	(0.301)	0.201
10.5	0.5	0.969	0.210	(0.283)
11	1	0.976	0.182	(0.319)
13	3	0.961	0.083	(0.486)
15	5	0.962	0.040	(0.608)

Table 1: The table represents the proportion of simulation runs $\mathcal{T}$ performs better feature than $\mathcal{T}_0$. The parenthesis indicate the column of the inferior treatment. For every choice of m_R, 1000 simulations have been realized. Here, the case of equal variances has been reported: $\sigma_R^2 = \sigma_W^2 = 1.5^2$.

$\{A \bigcup B \bigcup C\}$. Moreover, this result has been found for any values of Δ, that is remarkable since the means are unknown before doing the test. The second goal of this design was minimizing the number of subjects assigned to the inferior treatment. In Table 1 we report the proportion of runs $\mathcal{T}$ allocates to each treatment less subjects than $\mathcal{T}_0$. To better understand this aspect of the performance of the MRRU model, we report in Figure 7 the flanked boxplots of the number of subjects allocated to the inferior treatment in the 1000 replications of the urn design. The red line indicate the number of subject allocated to the inferior treatment by $\mathcal{T}_0$. Then, the goal is to maximize the number of cases below the red line. The numbers within parenthesis in Table 1 represent the proportion of simulation runs that are below the red line in Figure 7.
Notice from Figure 7 that, the greater is the mean distance $|\Delta| = |m_R - m_W|$, the smaller is the number of subjects allocated to the inferior treatment.

In the case of different variances (Table 2), in most of the runs $\mathcal{T}$ has a power greater than $\mathcal{T}_0$. Nevertheless, it seems that the larger is the value of m_R the less is the proportion of times the power of $\mathcal{T}$ is greater than $\mathcal{T}_0$. The reason of this fact is due to the asymmetry of variances: with these values of σ_R^2 and σ_W^2 the length of the interval I_n^C is very small. Then, when the urn process $(Z_k)_k$ overcomes η can occur more often that Z_n goes out from the interval I_n^C, and so does the allocation proportion $N_R(n)/n$. When this happens, we have that $(N_R(n)/n, n) \notin \{A \bigcup B \bigcup C\}$ and so the power of $\mathcal{T}$ will be smaller than the power of $\mathcal{T}_0$.

In Table 2 we also report the proportion of simulation runs $\mathcal{T}$ allocates to each treatment fewer subjects than $\mathcal{T}_0$. Figure 8 shows the boxplots of the number of subjects allocated to the inferior treatment with the

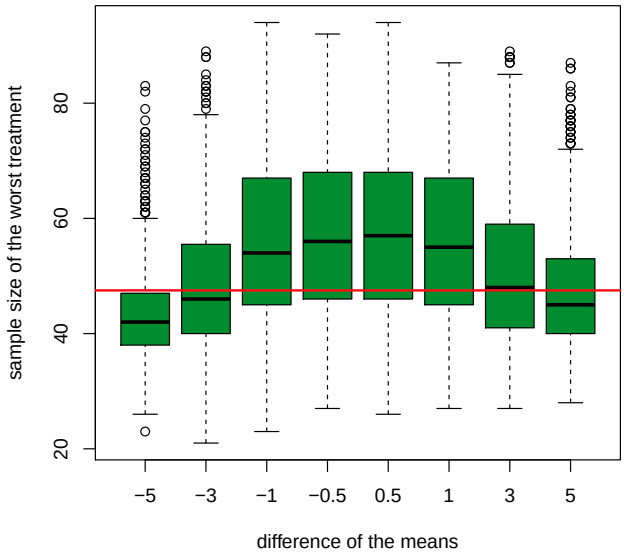

Figure 7: The picture shows, for any $\Delta \in \{-5, -3, -1, -0.5, 0.5, 1, 3, 5\}$, the flanked boxplots of the number of subjects allocated to the inferior treatment by $\mathcal{T}$. In order to compute the boxpots, 1000 replications of the urn process $(Z_k)_k$ have been used. The red line represent the number of subject allocated to the inferior treatment by $\mathcal{T}_0$, that in both cases is $n_0 p_0 = n_0(1 - p_0) = 48$. Here, the case of equal variances has been reported: $\sigma_R^2 = \sigma_W^2 = 1.5^2$.

1000 replications of the urn process.
It is easy to note from Figure 8 that, even when the variances are different, the greater the mean distance $|\Delta| = |m_R - m_W|$, the smaller the number of subjects allocated to the inferior treatment. In this case, the design performs better when the worst treatment is W. As explained before, this occurs because with these values of σ_R^2 and σ_W^2 the interval I_n^C is very short.

7 Real Case Study

In this section we show a real case study, where the application of the methodology presented in the paper would have improved the performance of a classical test, from both the statistical and ethical point of view. We consider data concerning treatment times of patients affected by ST- Elevation Myocardial. The main rescue procedure for these patients is the Primary Angioplasty. It is well known that to improve the outcome of patients and reduce the in-hospital mortality the time between the arrival at ER (called Door) and the time of intervention (called Baloon) must be reduced as much as possible. So

m_R	Δ	$\#\{\beta_{\mathcal{T}} \geq \beta_{\mathcal{T}_0}\}$	$\#\{N_R(n) < n_{0,R}\}$	$\#\{N_W(n) < n_{0,W}\}$
5	-5	1.000	(0.895)	0.003
7	-3	0.98	(0.636)	0.042
9	-1	0.928	(0.364)	0.131
9.5	-0.5	0.930	(0.345)	0.136
10.5	0.5	0.887	0.222	(0.232)
11	1	0.876	0.205	(0.265)
13	3	0.847	0.092	(0.361)
15	5	0.799	0.064	(0.447)

Table 2: The table represents the proportion of times the new test $\mathcal{T}$ presented a different feature with respect to the classical test $\mathcal{T}_0$: having higher power of assigning fewer patients to one of the two treatment. The parenthesis indicate the column of the worst treatment. For every choice of m_R, 1000 simulations have been realized. Here, the case of different variances has been considered: $\sigma_R = 1$ and $\sigma_W = 2$.

the Door to Baloon time (DB) is our treatment's response. We have two different treatments: the patients managed by the 118 (free-tall number for emergency in Italy) and the self presented ones. We design our experiment to allocate the majority of patients to treatment performing better, and simultaneously collect evidence in comparing the time distributions of DB times.

We have at our disposal the values of the door-to-baloon time (DB) in minutes of 1179 patients. Among them, 657 subjects have been managed by 118, while the others 522 subjects reached the hospital by themselves. We denote the choice of calling 118 as treatment W and the choice of going to the hospital by themselves as treatment R. In this case, since the lower are the responses (DB time) the better is the treatment, a decreasing utility function is necessary. Moreover, the urn model presented in Section 3 requires the reinforcements distributions to be positive. Then, we choose the monotonic utility function $u(x) = 6 - log(x)$ to transform responses (DB time) into reinforcement values, in order to satisfy those assumptions. To ease notation, from now on we refer to the responses transformed by the utility function as the responses collected directly from the patients. In this situation, the means and variances computed using all the data at our disposal are taken as the true means and variances of the populations R and W: $m_R = 1.503$, $m_W = 1.996$, $\sigma_R = 0.518$, $\sigma_W = 0.760$. Notice that, since the true difference of the means $\Delta = m_R - m_W = -0.493$ is negative, W is the best treatment. We want to conduct a non-adaptive test and a response adaptive one that aim at determining the best treatment, in order to compare their performance.

Initially, we imagine to conduct a non-adaptive test $\mathcal{T}_0$ to compare

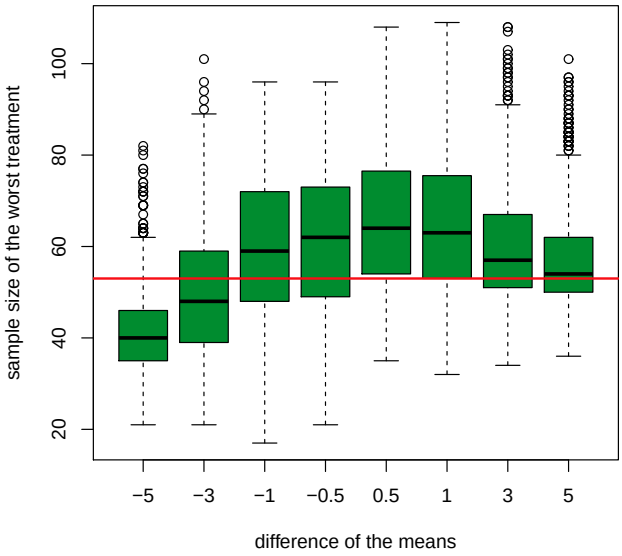

Figure 8: The picture shows, for any $\Delta \in \{-5, -3, -1, -0.5, 0.5, 1, 3, 5\}$, the boxplots of the number of subjects allocated to the inferior treatment by $\mathcal{T}$. In order to compute the boxplots, 1000 replications of the urn process $(Z_k)_k$ have been used. The red line represent the number of subject allocated to the inferior treatment by $\mathcal{T}_0$, that in both cases is $n_0 p_0 = n_0(1 - p_0) = 53$. Here, the case of different variances has been reported: $\sigma_R^2 = 1$ and $\sigma_W^2 = 4$.

the mean effects of treatments R and W. We fix a significance level $\alpha = 0.01$, a minimum power $\beta_0 = 0.95$ for a standard difference of the means $\Delta_0 = 0.5$. Then, we assume responses to treatments R and W are i.i.d random variables with distributions μ_R and μ_W, respectively. Moreover, we assume the laws are Gaussian: $\mu_R = \mathcal{N}(m_R, \sigma_R^2)$ and $\mu_W = \mathcal{N}(m_W, \sigma_W^2)$ (verified by empirical tools). The allocation proportion is set to $p_0 = 0.468$, the empirical one. With these parameters we can conduct a two-sided t-test that requires a total of $n_0 = 119$ subjects, $n_0 p_0 = 56$ allocated to treatment R and $n_0(1 - p_0) = 63$ allocated to treatment W. To compute n_0 we have assumed known variances. The power of this test computed in correspondence to the true difference of the means is $\beta_{\mathcal{T}_0}(\Delta) = 0.945$.

Now, consider the urn model presented in Section 3 to construct the adaptive test $\mathcal{T}$. $\mathcal{T}$ involves more subject in the experiment than $\mathcal{T}_0$, in particular $n = 1.25 \cdot n_0 = 148$. Nevertheless, since in practice variances are unknown, n_0 and n should be computed from the estimates of the variances. As a consequence, the total number of subjects needed for $\mathcal{T}$ is random, because it depends on the variance estimation. For this reason, we may have replications with different sample size n.

We realize 500 replications of the urn procedure. Since the data at our disposal are much more than the amount of data we need for each trial, by permutating the responses we can take at random different data with a different order in each replication. In Figure 9, we represent 10 simulations of the urn proportion process $(Z_n)_n$.

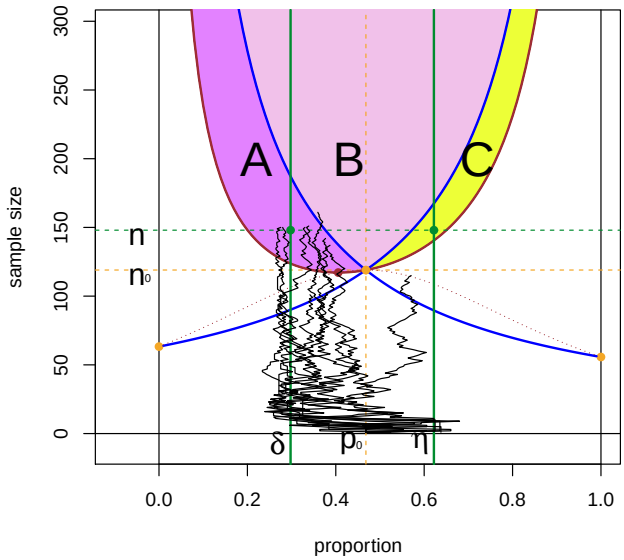

Figure 9: Black lines represent 10 replications of the urn proportion process $(Z_n)_n$. Each replication uses responses taken at random from the data at our disposal. The *proportion* - *sample size* space has been partitioned assuming the variances known.

As we can see from Figure 9, the urn process seems to target region A, where parameter δ is set. This is because R is the worst treatment in this case. Test $\mathcal{T}$ has higher power and assigns to treatment R less patients than $\mathcal{T}_0$. This is our goal, since we know that R is the worst treatment ($m_R < m_W$).
For each one of the 500 replications we compute analytically the power at the true difference of the means Δ. In general, the power will be different for any simulation because different is the number of subjects assigned to the treatments (N_R and N_W). In Figure 10 we show a boxplot with the 500 values of the power computed using the urn model, to be compared with the power obtained with $\mathcal{T}_0$. Moreover, we show for each simulation the number of subjects assigned to treatment R, to be compared with the number of subjects assigned to R by $\mathcal{T}_0$.

From Figure 10, we notice that the urn design described in Section 3

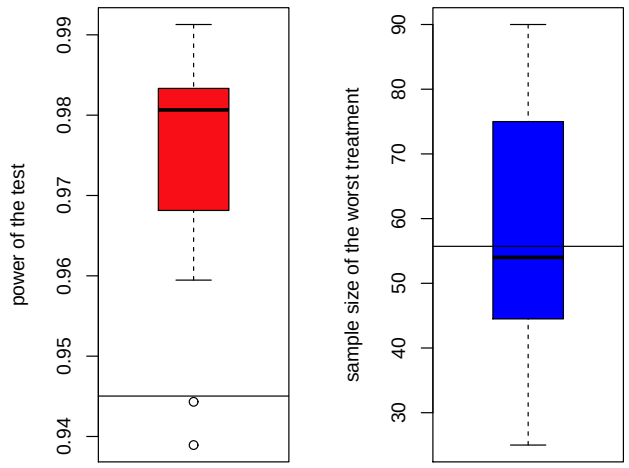

Figure 10: On the left: boxplot representing 500 values of power evaluated at the true difference of the means $\Delta = -0.493$ using $\mathcal{T}$: $\beta_{\mathcal{T}}(\Delta)$. The red line represents the power obtained with $\mathcal{T}_0$: $\beta_{\mathcal{T}_0}(\Delta) = 0.945$. On the right: boxplot representing 500 values of the number of subjects assigned to treatment R by $\mathcal{T}$: N_R. The red line represents the number of subjects assigned to treatment R by $\mathcal{T}_0$: $n_0 \cdot p_0 = 56$.

allows us to construct a test $\mathcal{T}$ with higher power than $\mathcal{T}_0$. This occurs for more than 99% of the replications, and the mean of the power computed overall the runs is

$$\frac{1}{500}\sum_{i=1}^{500}\beta_{\mathcal{T}i}(\Delta) \;=\; 0.975 \;>\; 0.945 \;=\; \beta_{\mathcal{T}_0}(\Delta).$$

Even if $\mathcal{T}$ needs a sample size n larger than $\mathcal{T}_0$, the number of subjects allocated to the inferior treatment R is less for $\mathcal{T}$ for the 52.6% of the runs. Besides, the mean of the number of units assigned to treatment R in all the runs is almost the same of the number computed with $\mathcal{T}_0$

$$\frac{1}{500}\sum_{i=1}^{500}N_{Ri} \;=\; 56.43 \;\simeq\; 56 \;=\; n_0 \cdot p_0.$$

8 Conclusions

In this paper we have conducted an analysis on the statistical properties of tests that aim at comparing the means of the responses to two

treatments. Starting from any non-adaptive test $\mathcal{T}_0$, we pointed out the features of an adaptive test $\mathcal{T}$ performing better than $\mathcal{T}_0$. Since the framework here is represented by clinical trials, this goal is achieved when $\mathcal{T}$ has (a) higher power and (b) assigns to the inferior treatment less subjects than $\mathcal{T}_0$. We investigated this task by individuating in the *proportion - sample size* space the subregions associated to tests $\mathcal{T}$ performing better than $\mathcal{T}_0$.
The test $\mathcal{T}$ can be implemented by adopting a response adaptive design. We propose an urn procedure (MRRU) that is able to target a fixed proportion allocation in (0,1). Thanks to this property, the urn model can individuate the test $\mathcal{T}$ in different regions depending on which is the inferior treatment, and both goals (a)-(b) can be accomplished. We showed that the assumption of normal responses and known variances can be relaxed and the procedure to partition the *proportion - sample size* space and to detect the test $\mathcal{T}$ still holds. We reported simulations and a case study that highlight the goodness of the procedure.

References

[1] Aletti, G., Ghiglietti, A. and Paganoni, A.M. (2013). A modified randomly reinforced urn design. *Journal of Applied Probability*, vol. 50 no. 2 486-498

[2] Bai, Z. D., Hu, F. and Zhang, L. X. (2002). Gaussian approximation theorems for urn models and their applications. *Annals of Applied Probabability*, 12 1149–1173.

[3] Bai, Z. D. and Hu, F. (2005). Asymptotics in randomized urn models. *Annals of Applied Probabability*, 15 914–940.

[4] Beggs, A.W. (2005). On the convergence of reinforcement learning. *Journal of Economic Theory*, 122, 1–36.

[5] Cheung, Z. D., Hu, F. and Zhang, L. X. (2006). Asymptotic theorems of sequential estimation-adjusted urn models. *Annals of Applied Probabability*, 16 340–369.

[6] Durham, S.C. and Yu, K.F. (1990). Randomized play-the leader rules for sequential sampling from two populations. *Probability in Engineering and Information Science*, 26(4), 355–367.

[7] Durham, S.C., Flournoy, N. and Li, W. (1996). Randomized Plya urn designs. Proc. Biometric Section of the American Statistical Association 166–170. Amer. Statist. Assoc., Alexandria, VA.

[8] Durham, S.C., Flournoy, N. and Li, W. (1998). A sequential design for maximizing the probability of a response. *The Canadian Journal of Statistics*, 26(3), 479–495.

[9] Ghiglietti, A. and Paganoni, A. M. (2012). Statistical properties of two-color randomly reinforced urn design targeting fixed allocations. Available as Mox-report 48/2012, Politecnico di Milano. In: *http://mox.polimi.it/it/progetti/pubblicazioni/quaderni/48-2012.pdf*

[10] Hu, F. and Rosenberger, W. F. (2006). The Theory of Response-Adaptive Randomization in Clinical Trials. Wiley, Hoboken, NJ.

[11] Janson, S.(2004). Functional limit theorems for multitype branching processes and generalized Pólya urns. *Stochastic Process. Appl.*, 110 177–245

[12] Lachin, J. M. and Rosenberger, W. F. (2006). Randomization in Clinical Trials. Wiley, New York.

[13] May, C., Paganoni, A.M. and Secchi, P. (2005). On a two-color generalized Polya urn. *Metron*, Vol. LXIII, 1

[14] Melfi, F. and Page, C. (2000). Estimation after adaptive allocation. *J. Statist. Plann. Inference*, 87 353–363.

[15] Melfi, F., Page, C. and Geraldes, M. (2001). An adaptive randomized design with application to estimation. *The Canadian Journal of Statistics*, 29, 107–116.

[16] Muliere, P., Paganoni, A.M. and Secchi, P. (2006). A randomly reinforced urn. *Journal of Statistical Planning and Inference*, 136, 1853–1874.

[17] Paganoni, A.M. and Secchi, P. (2007). A numerical study for comparing two response-adaptive designs for continuous treatment effects. *Statistical Methods and Applications*, 16, 321–346.

[18] Rosenberger, W. F. (2002). Randomized urn models and sequential design. *Sequential Analysis*, 21(1&2), 1–28.

[19] R Development Core Team (2011). A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org/.

MOX Technical Reports, last issues

Dipartimento di Matematica "F. Brioschi",
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

60/2013 Ghiglietti, A.; Paganoni, A.M.
An urn model to construct an efficient test procedure for response adaptive designs

59/2013 Aletti, M.; Bortolossi, A.; Perotto, S.; Veneziani, A.
One-dimensional surrogate models for advection-diffusion problems

58/2013 Artina, M.; Fornasier, M.; Micheletti, S.; Perotto, S.
Anisotropic adaptive meshes for brittle fractures: parameter sensitivity

57/2013 Antonietti, P.F.; Perugia, I.; Zaliani, D.
Schwarz domain decomposition preconditioners for plane wave discontinuous Galerkin methods

56/2013 Antonietti, P.F.; Ayuso de Dios, B.; Mazzieri, I.; Quarteroni, A.
Stability analysis for Discontinuous Galerkin approximations of the elastodynamics problem

55/2013 Laadhari, A.; Ruiz-Baier, R.; Quarteroni, A.
Fully Eulerian finite element approximation of a fluid-structure interaction problem in cardiac cells

54/2013 Biasi, R.; Ieva, F.; Paganoni, A.M.; Tarabelloni, N.
Use of depth measure for multivariate functional data in disease prediction: an application to electrocardiographic signals

53/2013 Micheletti, S.
A continuum variational approach based on optimal control to adaptive moving mesh methods

51/2013 Chen, P.; Quarteroni, A.
Weighted reduced basis method for stochastic optimal control problems with elliptic PDE constraint

52/2013 Chen, P.; Quarteroni, A.; Rozza, G.
Multilevel and weighted reduced basis method for stochastic optimal control problems constrained by Stokes equations