

Statistical properties of urn designs in clinical trials

Proprietà statistiche di disegni d'urna nei test clinici

Andrea Ghiglietti and Anna Maria Paganoni

Abstract We propose a response-adaptive design, described in terms of urn model, whose allocation proportion converges to prespecified values. The main asymptotic results concerning the urn designs are presented and discussed. We adopt the urn model to implement the random allocation procedure of an experiment that aims at testing the mean effect of two treatments. We conduct a statistical analysis on the inferential performance of different tests and we show that, given a non adaptive test $\mathcal{T}_0$, the response adaptive model constructs a test $\mathcal{T}$ that is better than $\mathcal{T}_0$, in terms of (a) higher power and (b) fewer subjects assigned to the inferior treatment. A retrospective real case study is presented.

Abstract Proponiamo un disegno adattivo alla risposta, descritto in termini di modello d'urna, la cui proporzione di allocazione converge a valori prefissati. I principali risultati asintotici relativi al disegno d'urna sono presentati e discussi. Adottiamo il modello d'urna per implementare la procedura di allocazione di un esperimento che ha lo scopo di confrontare l'effetto medio di due trattamenti. Condurremo un'analisi statistica sulle prestazioni inferenziali di diversi test e mostriamo che, dato un test non adattivo $\mathcal{T}_0$, il modello adattivo alla risposta costruisce un test $\mathcal{T}$ che risulta migliore di $\mathcal{T}_0$ in termini di (a) maggiore potenza e (b) meno pazienti assegnati al trattamento inferiore. Un caso studio retrospettivo è presentato.

Key words: Clinical trials; Randomly Reinforced Urns; Tests based on adaptive procedures.

Andrea Ghiglietti

Politecnico di Milano, piazza Leonardo da Vinci 32 Milano, e-mail: andrea.ghiglietti@polimi.it

Anna Maria Paganoni

Politecnico di Milano, piazza Leonardo da Vinci 32 Milano, e-mail: anna.paganoni@polimi.it

1 Response-adaptive procedures and urn models

This work focuses on the statistical properties of urn models used as response-adaptive designs for clinical trials (see [7]). We consider urn models with random non-negative reinforcements concerning only the extracted color. These designs have been called Play the Winner or Randomly Reinforced Urn (RRU) designs (see [2, 6]). Although these designs are able to asymptotically allocate subjects to the optimal treatment, their statistical properties present some problems. At first, because there are many results for designs whose asymptotic allocation is $\rho \in (0, 1)$, that cannot be applied to RRU models since their asymptotic allocation is $\rho \in \{0, 1\}$. Moreover, these models generate groups with very different sample sizes. Then, the inferential procedures based on these designs are usually characterized by a very low power in comparing treatments effects. For these reasons, we modify the reinforcement scheme of the urn to construct a design that asymptotically targets an allocation proportion $\rho \in (0, 1)$. The term target indicates the limit of the urn proportion process. We refer to this urn model as the Modified Randomly Reinforced Urn (MRRU) design and its properties have been mainly investigated in [1, 3]. In order to assign a small proportion of subjects to the inferior treatment, the MRRU model presents two possible values for the limit of the allocation proportion: δ and η , with $0 < \delta \leq \eta < 1$. In [1, 3] we prove that both the urn proportion Z_n and the allocation proportion $N_R(n)/n$ converge a.s. to $\rho = \eta \mathbf{1}_{\{m_R > m_W\}} + \delta \mathbf{1}_{\{m_R < m_W\}}$, where m_R and m_W indicate the response means to two competing treatments, called R and W respectively. Then, this model achieves the ethical goal of assigning an arbitrarily small proportion of subject to the inferior treatment. Moreover, since the limiting proportion is within $(0, 1)$, all the results for designs with asymptotic allocation $\rho \in (0, 1)$ can be applied and the inferential performances are improved.

2 An urn procedure to construct efficient test for response-adaptive designs

In this section we conduct an analysis on the statistical performance of different tests for comparing the mean effect of two treatments ([4]). 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 the Modified Randomly Reinforced Urn design (MRRU) described in [1] and we show how to perform the model parameters selection for the purpose of this chapter.

Consider the classical hypothesis test aiming at comparing the response means to two treatment R and W

$$H_0 : m_R - m_W = 0 \quad \text{vs} \quad H_1 : m_R - m_W \neq 0. \quad (1)$$

We assume response 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. Let $p_0 \in (0, 1)$ be the proportion of patients allocated to treatment R and let n_0 be the total number of subjects involved in the experiment. Let $n_{0,R}$ 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, let us fix n_0 according on the following parameters

- 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, we have the following expression for critical region of level α

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

where $\bar{M}_{n_{0,R}} = \sum_{i=1}^{n_{0,R}} M_i / n_{0,R}$ and $\bar{N}_{n_{0,W}} = \sum_{i=1}^{n_{0,W}} N_i / n_{0,W}$ and $z_{\frac{\alpha}{2}}$ is the quantile of order $1 - \alpha/2$ of a standard normal distribution. Furthermore, the power of the test, is a function of the real difference $\Delta = m_R - m_W$, 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), with n_0 as sample size and p_0 as proportion of patients allocated to the treatment R . To construct a test with equal parameters ($\alpha, \Delta_0, \beta_0$) 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 $((0, 1) \times \mathcal{N})$, 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

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

$$\beta_{\mathcal{T}}(\Delta) \geq \beta_{\mathcal{T}_0}(\Delta) \quad \forall \Delta \in \mathcal{R} \Leftrightarrow \frac{\sigma_R^2}{n\rho} + \frac{\sigma_W^2}{n(1-\rho)} \leq \frac{\sigma_R^2}{n_0 p_0} + \frac{\sigma_W^2}{n_0(1-p_0)} \quad (3)$$

Now, if we denote as p_{opt} the Neyman allocation proportion $\frac{\sigma_R}{\sigma_R + \sigma_W}$, we can rewrite inequality (3) in a more suitable form

$$\frac{p_{opt}^2}{n\rho} + \frac{(1-p_{opt})^2}{n(1-\rho)} \leq \frac{p_{opt}^2}{n_0 p_0} + \frac{(1-p_{opt})^2}{n_0(1-p_0)} \quad (4)$$

Inequality (4) divides the *proportion - sample size* space in two regions. The boundary is computed by imposing the equality in (4) and expressing the sample size n as a function of the proportion ρ .

$$n_\beta(\rho) = \left(\frac{p_{opt}^2}{\rho} + \frac{(1-p_{opt})^2}{1-\rho} \right) \left(\frac{p_{opt}^2}{n_0 p_0} + \frac{(1-p_{opt})^2}{n_0(1-p_0)} \right)^{-1} \quad (5)$$

We refer to function (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 1 by a red line. Each point over this curve is a test $\mathcal{T}$ with a power uniformly higher than $\mathcal{T}_0$.

To satisfy condition (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

$$n(1-\rho) < n_0(1-p_0) \Leftrightarrow \rho > 1 - \frac{n_0}{n}(1-p_0); \quad (6)$$

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

$$n\rho < n_0 p_0 \Leftrightarrow \rho < \frac{n_0}{n} p_0. \quad (7)$$

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 $\mathcal{T} \in \mathcal{C}$ 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 to target different allocation proportions according to the responses collected during the trial. The asymptotic behavior of statistics based on adaptive estimators have been deeply studied in literature (see for instance [5]). The authors propose to adopt the *Modified Randomly Reinforced Urn* design (MRRU).

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 the intervals are mutually disjoint and contained in $(0, 1)$. 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

$$\begin{cases} \frac{N_R(n)}{n} \in I_n^C \text{ if } m_R > m_W, \\ \frac{N_R(n)}{n} \in I_n^A \text{ if } m_R < m_W. \end{cases}$$

Since the sequence $\frac{N_R(n)}{n}$ converge a.s. to $\rho = \eta \mathbf{1}_{\{m_R > m_W\}} + \delta \mathbf{1}_{\{m_R < m_W\}}$, we set $\delta \in I_n^A$ and $\eta \in I_n^C$, so that $\lim_{k \rightarrow \infty} \frac{N_R(k)}{k} \in I_n^A$ if $m_R < m_W$ and $\lim_{k \rightarrow \infty} \frac{N_R(k)}{k} \in I_n^C$ 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 1 we show how the urn process Z_n converges towards the right region.

The speed of convergence of the urn model is a key point for the success of this procedure. A complete discussion over the quantities that influence the asymptotic behavior of the urn process is reported in [4].

In [4] we analyzed a real case study, where the application of the methodology presented in this 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 Balloon) must be reduced as much as possible. So the Door to Balloon time (DB) is our treatment's response. We




Fig. 1 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 $(Z_k)_k$.

have two different treatments: the patients managed by the 118 (free-toll 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. The results of the study are presented in [4] and they show the goodness of the procedure presented in the paper.

References

1. Aletti, G., Ghiglietti, A., Paganoni, A.M.: A modified randomly reinforced urn design. *Journal of Applied Probability*, 50, (2), 486–498 (2013)
2. Durham, S.C., Flournoy, N., Li, W.: A sequential design for maximizing the probability of a response. *The Canadian Journal of Statistics*, 26 (3), 479–495 (1998)
3. Ghiglietti, A., Paganoni, A. M.: Statistical properties of two-color randomly reinforced urn design targeting fixed allocations. Available via Mox-report 48/2012, Politecnico di Milano (2012)
<http://mox.polimi.it/it/progetti/pubblicazioni/quaderni/48-2012.pdf>
4. Ghiglietti, A., Paganoni, A. M.: An urn procedure for efficiently detecting the parameters of an hypothesis test. Available via Mox-report 60/2013, Politecnico di Milano (2013)
<http://mox.polimi.it/it/progetti/pubblicazioni/quaderni/60-2013.pdf>
5. Melfi, F., Page, C., Gerdal, M.: An adaptive randomized design with application to estimation. *The Canadian Journal of Statistics*, 29, 107–116 (2001).
6. Muliere, P., Paganoni, A.M., Secchi, P.: A randomly reinforced urn. *Journal of Statistical Planning and Inference*, 136, 1853–1874 (2006).
7. Rosenberger, W. F. Zhang, L.: Response-adaptive randomization for clinical trials with continuous responses. *Biometrics*, 62, 562–569 (2006).