

Urn models for response-adaptive randomized designs: a simulation study based on a non-adaptive randomized trial

Andrea Ghiglietti¹, Maria Giovanna Scarale^{2,3}, Rosalba Miceli⁴,
Francesca Ieva⁵, Luigi Mariani⁴, Cecilia Gavazzi⁶,
Anna Maria Paganoni⁵, and Valeria Edefonti²

February 20, 2018

¹ Dipartimento di Matematica “F. Enriques”, Università degli Studi di Milano, via Saldini, 50, 20133 Milano, Italy.

² Laboratorio di Statistica Medica, Biometria, ed Epidemiologia “G. A. Maccacaro”, Dipartimento di Scienze Cliniche e di Comunità, Università degli Studi di Milano, via A. Vanzetti, 5, 20133 Milano, Italy.

³ Unit of Biostatistics, Poliambulatorio “Giovanni Paolo II”, IRCCS Casa Sollievo della Sofferenza, Viale Padre Pio, 71013 San Giovanni Rotondo, Italy.

⁴ Struttura Semplice di Epidemiologia Clinica e Organizzazione Trials, Fondazione IRCCS Istituto Nazionale Tumori, Via G. Venezian 1, 20133 Milano, Italy.

⁵ MOX – Modellistica e Calcolo Scientifico, Dipartimento di Matematica, Politecnico di Milano, via Bonardi 9, 20133 Milano, Italy.

⁶ Struttura Semplice Dipartimentale di Terapia Nutrizionale, Fondazione IRCCS Istituto Nazionale dei Tumori, via G. Venezian, 1, 20133 Milano, Italy.

Corresponding Author: Valeria Edefonti

Laboratorio di Statistica Medica, Biometria, ed Epidemiologia “G. A. Maccacaro”

Dipartimento di Scienze Cliniche e di Comunità

Università degli Studi di Milano

via A. Vanzetti, 5, 20133 Milano, Italy

email: valeria.edefonti@unimi.it

Funding: Nothing to declare.

Conflicts of interest: The authors have no conflict of interest to declare.

Author contribution: AG, MGS, AMP, and VE conceived the study. AG, FI, and AMP carried out the simulation study and provided the R code that is presented in Section 4. RM designed and carried out the statistical analysis

of the original Home Enteral Nutrition trial. LM provided useful suggestions on the theoretical and practical issues of adaptive designs in clinical trials and revised the final manuscript. CG was the PI of the Home Enteral Nutrition trial, provided data for the current paper, and revised the final manuscript. MGS, AG, and FI wrote Subsection 2.2 and 2.3 and revised the final manuscript. AMP revised also the final manuscript. VE and RM wrote Subsection 2.1, Sections 3, 4, 5 and parts of the remaining sections. VE and RM have primary responsibility for the final content of the paper. All the authors read and approved the final version of the paper.

Abstract

Recently, response-adaptive designs have been proposed in randomized clinical trials to achieve ethical and/or cost advantages by using sequential accrual information collected during the trial to dynamically update the probabilities of treatment assignments. In this context, urn models - where the probability to assign patients to treatments is interpreted as the proportion of balls of different colors available in a virtual urn - have been used as response-adaptive randomization rules.

We propose the use of Randomly Reinforced Urn (RRU) models in a simulation study based on a published randomized clinical trial on the efficacy of home enteral nutrition in cancer patients after major gastrointestinal surgery. We compare results with the RRU design with those previously published with the non-adaptive approach. We also provide a code written with the R software to implement the RRU design in practice.

In detail, we simulate 10,000 trials based on the RRU model in three setups of different total sample sizes. We report information on the number of patients allocated to the inferior treatment and on the empirical power of the t-test for the treatment coefficient in the ANOVA model. We carry out a sensitivity analysis to assess the effect of different urn compositions. For each sample size, in approximately 75% of the simulation runs, the number of patients allocated to the inferior treatment by the RRU design is lower, as compared to the non-adaptive design. The empirical power of the t-test for the treatment effect is similar in the two designs.

Keywords: non-adaptive trial design; Randomly Reinforced Urn model; Randomized trials; Response-adaptive randomization; Simulation study.

1 Introduction

In the statistical literature, urn models have been widely studied as mathematical tools to implement randomization in the context of clinical trials (e.g. see Wei, 1978; Wen et al. 2017). These designs randomly assign those subjects that sequentially enter the trial to the treatment arms according to the color of the balls sampled from a virtual urn. Hence, the probability to assign a patient to a treatment arm is modelled by the proportion of the different types of balls in the urn. Recently, interest has been increased in the use of urn models for responses-adaptive designs, in which the probability to sample a ball of a certain type depends on the treatment performance observed on the subjects previously randomized (Durham et al., 1998; Atkinson and Biswas, 2014). These designs are, therefore, able to achieve desirable statistical properties taking into account the ethical aspects of the clinical experiment (see e.g. Hu and Rosenberger, 2006). A popular class of such designs is the Randomly Reinforced Urn (RRU) model, which has been introduced in Durham et al. (1998) for binary treatment responses and extended in Muliere et al. (2006) to handle continuous responses. The main asymptotic results on the proportion of subjects assigned to the treatment groups by a RRU design have been established in Flournoy and May (2009) and Muliere et al. (2006). For the purposes of this paper, we simply remind that a RRU design assigns patients to the superior treatment with a probability that converges to one as the sample size increases. For an overview on the RRU designs and its properties, we refer to Flournoy et al. (2012).

Although the theoretical result of assigning most of the patients to the superior treatment is very attractive from the ethical point of view, the RRU design have rarely been implemented in clinical trials or in simulation studies based on a real set-up (e.g. see Chapter 12 in Lachin and Rosenberger, 2002). This may depend on some feasibility issues that affect the practical implementation of adaptive designs in general. In addition, although only few decision rules and

basic programming skills are needed to implement the RRU design in clinical practice, there are currently a few groups of researchers that combines knowledge of the theoretical properties of urn models and experience in planning and running clinical trials.

The substantial lack of dedicated software in standard statistical packages used in clinical practice is an additional issue that have prevented a broader use of RRU designs in this field.

The aim of the current paper is to popularize the statistical and ethical advantages of the RRU design, and of urn schemes in general, and to promote their use in clinical practice through a dedicated code written in R. In detail, we will simulate a large number of trials that follow the RRU model starting from the real-life data collected in a (previously published) Home Enteral Nutrition (HEN) randomized trial (Gavazzi et al., 2016), where a non-adaptive design was originally adopted. Comparing the performance of the RRU with that of the original non-adaptive design, we expect that the RRU design will: 1) assign fewer patients to the inferior treatment; 2) maintain similar inferential properties. This will be an advantage in terms of both statistical performance and ethical responsibility.

The paper is structured as follows. Section 2 provides some preliminary information on the HEN trial and its results (Gavazzi et al., 2016), introduces the RRU model as a form of response-adaptive design, and describes how we carried out the simulations of the RRU design based on the original HEN data. Section 3 provides a comparison of the performance of the RRU versus the non-adaptive design in the simulation study based on the HEN data. Section 4 provides some suggestions on tuning parameters and the R codes for the implementation of a RRU design in the practice of randomized clinical trials. We conclude the paper with a Discussion (Section 5).

2 Materials and Methods

2.1 A randomized controlled trial of home enteral nutrition versus nutritional counselling

The RRU model was here implemented in a simulation study based on results from a multicenter, controlled, open-label, two-parallel groups, randomized clinical trial conducted at the Fondazione IRCCS Istituto Nazionale dei Tumori (INT), Milan, Italy, and at the European Institute of Oncology, Milan, Italy, between December 2008 and June 2011 (Gavazzi et al., 2016). Malnutrition in gastrointestinal cancer patients is an independent risk factor for post-operative morbidity and mortality (Bozzetti et al., 2000) and a prognostic factor for worst long-term outcome, especially after major surgery (Martin et al. 2015). Therefore, the trial was primary aimed at investigating the effectiveness of enteral nutrition in limiting weight loss after home discharge from surgery, in comparison to nutritional counselling.

The enrolled subjects were adult (> 18 years) patients with documented upper gastrointestinal cancer (esophagus, stomach, pancreas, biliary tract) who were candidates for major elective surgery and showed a preoperative nutritional risk score that indicated a potential benefit from any nutritional intervention. A random permuted block design (stratified for referring center) randomly assigned patients before discharge to receive either HEN to cover the basal energy requirement (experimental group), or nutritional counselling by an expert dietitian, including oral supplements only when needed (Control Group - CG), in a 1:1 ratio. The protocol allowed the removal of HEN after two months from discharge if a weight gain $\geq 5\%$ was reported and oral diet was regular and adequate. Therefore, the minimum treatment period in this trial was two months. The treatment effect was defined as the difference between the mean “weight change” (weight after two months - weight at baseline) in the HEN and nutritional counselling arms (primary end-point). The total sample size required to detect a statistically significant treatment effect was of 140 patients (70 per

arm). The sample size was calculated with $\alpha = 5\%$ (two-sided) and power $1 - \beta = 80\%$ under the following assumptions derived from a previous pilot study conducted at INT:

- the baseline standard deviation of the weight distribution was equal to 10 kg;
- normality and homogeneity of weight variances was assumed across times of assessment and arms;
- 5 kg of difference was expected in the two-months mean weight change of treated versus control patients;
- a correlation coefficient of 0.5 was assumed between weights at baseline and weights after two months.

The planned efficacy analyses included one interim and one final analysis, with the interim analysis to be carried out when half of the patients had been followed for at least two months. In order not to exceed an overall type I error of 5%, the nominal significance level required by each analysis for the evaluation of efficacy was 2.94%, according to the Pocock's procedure (Pocock, 1977). The main analysis on the primary end-point was conducted with a univariate ANOVA including treatment as the main effect, after checking that standard ANOVA assumptions were satisfied.

In total, 79 patients were initially randomized; however, 11 patients had a missing weight at two months; the final analysis was, therefore, performed on 68 patients, of which 33 patients were allocated to the HEN group and 35 to the CG. The main result of the primary end-point analysis was that the mean weight loss in the patients undertaking the HEN treatment was significantly lower than that in the CG, with a treatment effect estimated by the corresponding ANOVA model coefficient (95% confidence interval) of 3.2 (1.1-5.3) and a p-value from the corresponding two-sided t-test equal to $0.31\% < 2.94\%$. For this reason, the trial was stopped at the interim analysis and results from this analysis were

published in (Gavazzi et al., 2016). So, the HEN was found to be the superior treatment in this trial.

2.2 Randomly Reinforced Urn design

We briefly introduce a RRU model for continuous responses to two treatments (Muliere et al., 2016), which has been implemented in accordance with the design characteristics of the HEN trial.

Consider patients that sequentially enter a trial and have to be randomly assigned to either treatment \mathcal{R} or treatment \mathcal{W} . To model this, we assume that, before subject $i \geq 1$ enters the trial, we have a virtual urn with $R_{i-1} > 0$ red balls and $W_{i-1} > 0$ white balls. We indicate with (R_{i-1}, W_{i-1}) the urn composition before subject $i \geq 1$ enters the trial. We also set the initial urn composition balanced (i.e., $R_0 = W_0$), to reflect the 1:1 randomization.

When subject i enters the trial, a ball is sampled from the virtual urn and he/she is assigned to treatment \mathcal{R} if the sampled color is red ($X_i = 1$) or to treatment \mathcal{W} if the sampled color is white ($X_i = 0$). When his/her response to the assigned treatment is ascertained, we indicate it by ξ_{Ri} if the assigned treatment is \mathcal{R} or by ξ_{Wi} if the assigned treatment is \mathcal{W} . The responses conditional on treatment are assumed independent and identically distributed.

The urn is then updated by adding balls of the same color as the sampled one; in detail, the number of balls added to the urn is represented by the utility function u , which is a suitable positive monotone increasing function of the response observed on subject i . Formally, the urn composition is updated as follows:

$$\begin{cases} R_i = R_{i-1} + X_i u(\xi_{Ri}) \\ W_i = W_{i-1} + (1 - X_i) u(\xi_{Wi}), \end{cases} \quad (1)$$

where we called 'reinforcement' the quantities $u(\xi_{Ri})$ and $u(\xi_{Wi})$.

The updating rule in (1) implies the single responses are available before the next patient enters the trial. In the case of 'delayed responses', we propose here

a variant of the previous design in the same spirit of Bai et al. (2002): the urn updating is based only on those responses that were available during the time interval between the arrivals of subject i and $i + 1$. Formally, for any $i \geq 1$, let us denote by A_i the set of patients whose responses to treatments are available before subject i arrives. Then, the urn composition is updated as follows:

$$\begin{cases} R_i = R_{i-1} + \sum_{k \in (A_{i+1} \setminus A_i)} X_k u(\xi_{Rk}) \\ W_i = W_{i-1} + \sum_{k \in (A_{i+1} \setminus A_i)} (1 - X_k) u(\xi_{Wk}), \end{cases} \quad (2)$$

where $(A_{i+1} \setminus A_i)$ refers to those subjects whose responses are available during the time interval between the arrivals of subject i and $i + 1$. In case of no delayed responses, $(A_{i+1} \setminus A_i) = i$, and hence (1) and (2) are equivalent.

It follows from the RRU design definition that the probability to assign a subject i to the treatment \mathcal{R} is the proportion of red balls in the urn at the moment of his/her entrance in the trial:

$$\mathbf{P}(X_i = 1 | R_{i-1}, W_{i-1}) = \frac{R_{i-1}}{R_{i-1} + W_{i-1}}, \quad (3)$$

where the right hand side of the formula indicates the urn proportion at time $i - 1$. Hence, the sequence $\{X_i; i \geq 1\}$ of the subject assignment indicators is composed by conditionally Bernoulli random variables. In addition, it is worth noting that the urn proportion:

- needs not to be a rational number in general, as reinforcement values may be any real positive numbers;
- changes as far as a new response is made available; as a consequence, the probability to assign any new subject to one treatment or to the other depends on the treatment performance, in accordance with other response-adaptive designs.

Now, define $N_R(n) = \sum_{i=1}^n X_i$ as the number of subjects assigned to treat-

ment \mathcal{R} among the first n patients enrolled in the trial and $N_W(n) = n - N_R(n)$ as the number of subjects assigned to \mathcal{W} .

The main asymptotic result of the RRU design is that the proportion of subjects assigned to the superior treatment converges to one, as the sample size increases to infinity. Formally, denoting by $m_R := \mathbf{E}[u(\xi_{R1})]$ and $m_W = \mathbf{E}[u(\xi_{W1})]$, from Muliere et al. (2016) we have that

$$\frac{N_R(n)}{n} \xrightarrow{a.s.} \begin{cases} 1 & \text{if } m_R > m_W, \\ 0 & \text{if } m_R < m_W. \end{cases} \quad (4)$$

Hence, the RRU design asymptotically targets the superior treatment \mathcal{R} . As a consequence, we expect that, as the sample size increases, a RRU design assigns a lower number of subjects to the inferior treatment with a higher probability, as compared to a non-adaptive design.

2.3 Simulations of Randomly Reinforced Urn designs

In this subsection, we describe how we simulated the RRU design starting from the HEN trial data and how we derived the results for comparing the RRU design with the non-adaptive one. We considered the following main steps:

- (i) using the HEN trial dataset (Gavazzi et al. 2016) described in Subsection 2.1:
 - (1) we estimated the parameters of the Gaussian distribution of the responses to the HEN group;
 - (2) we estimated the parameters of the Gaussian distribution of the responses in the CG;
 - (3) we computed the empirical distribution of the difference between arrival times of consecutive subjects;
- (ii) we simulated N independent trial samples based on the RRU model; for

each sample, responses to both treatments and intervals between arrival times were randomly generated from distributions introduced in point (i);

(iii) we computed from these N trials:

- (1) the empirical distribution of the number of subjects assigned to the inferior treatment W ;
- (2) the empirical power of the corresponding test.

Previous literature (Flournoy and May, 2009; Ghiglietti and Paganoni, 2014) demonstrates the asymptotic normality of the distribution of the test statistic under the null hypothesis of equal response means. However, since this is an approximated test and the variances of the treatment responses are estimated from the data, in this paper we opt for a Student's t-test.

The previous steps are detailed in the following.

To start, we considered the following three different choices of trial sample sizes:

- (a) $n = 58$;
- (b) $n = 68$;
- (c) $n = 78$,

where the total sample size 68 of the HEN trial (Section 2.1) was used as the reference set-up and we moved $\pm 15\%$ from that to get other two reasonable sample sizes.

For each set-up, we performed $N = 10,000$ simulations of independent trials based on the RRU design: in each run we have a virtual urn to be sampled and reinforced as described in Subsection 2.2. Formally, we denote by (R_i^j, W_i^j) the urn composition and by $R_i^j / (R_i^j + W_i^j)$ the urn proportion in simulation $j = \{1, \dots, N\}$ at time $i \in \{1, \dots, n\}$.

All the urns start with the same (fixed) initial composition, i.e. $(R_0^j, W_0^j) = (R_0, W_0)$ for any $j = \{1, \dots, N\}$. Then, the urn composition (R_i^j, W_i^j) is updated

as in (2):

$$\begin{cases} R_i^j = R_{i-1}^j + \sum_{k \in (A_{i+1}^j \setminus A_i^j)} X_k^j u(\xi_{Rk}^j) \\ W_i^j = W_{i-1}^j + \sum_{k \in (A_{i+1}^j \setminus A_i^j)} (1 - X_k^j) u(\xi_{Wk}^j), \end{cases}$$

where X_k^j is a Bernoulli random variable with parameter $R_{k-1}^j / (R_{k-1}^j + W_{k-1}^j)$ and the set A_i^j here includes all the patients who arrived two months earlier than subject i . Indeed, in the HEN trial, responses were available only two months after treatment administration.

In addition, as normality assumptions in the original data were not rejected (see Subsection 2.1), responses to both treatments were generated as independent Gaussian random variables with arm-specific means and variances computed using the HEN dataset and given by: $m_R = -0.315$ and $\sigma_R = 3.868$ for treatment \mathcal{R} (HEN group), $m_W = -3.571$ and $\sigma_W = 4.789$ for treatment \mathcal{W} (CG). Formally, we generated the following quantities:

- (1) $\xi_{R1}^j, \dots, \xi_{Rn}^j \sim \mathcal{N}(m_R, \sigma_R^2)$ potential responses to treatment \mathcal{R} (HEN group);
- (2) $\xi_{W1}^j, \dots, \xi_{Wn}^j \sim \mathcal{N}(m_W, \sigma_W^2)$ potential responses to treatment \mathcal{W} (CG),

where either ξ_{Ri}^j or ξ_{Wi}^j is observed, as each subject just receives one treatment. We also randomly generated the potential arrival times from the corresponding empirical distribution in the HEN dataset.

For any sample size n (cases (a)-(b)-(c)) and any simulation $j = \{1, \dots, N\}$, we finally reported:

- (1) the number of patients $N_W^j(n) = \sum (1 - X_i^j)$ assigned to the CG, known to be the inferior treatment in the HEN trial (Subsection 2.1);
- (2) the result $I_n^j \in \{0, 1\}$ of the t-test for equal mean changes at level $\alpha = 0.05$ (corresponding to the treatment coefficient in the ANOVA model): $I_n^j = 0$ if the test does not reject H_0 , while $I_n^j = 1$ if the test rejects H_0 .

It is worth noting that $N_W^j(n)$ (and consequently $N_R^j(n)$) typically differs across simulations, because the urn processes are independent and the subjects are

allocated to the treatments depending on the urn-specific path of colors of the sampled balls. We also estimated the power of the t-test from the N simulated trials referring to the empirical power $1 - \hat{\beta} = N^{-1} \sum_{j=1}^N I_n^j$.

Without loss of generality, we set the u function as: $u(x) = (x + 20)/40$. Since in the HEN trial the response values, x , range in the interval $(-20, 20)$, this function was chosen to map linearly our simulated responses, x , in $(0, 1)$. We also assumed the initial urn composition to be $R_0 = W_0 = 1$ (i.e. one ball of each color initially put into the urn). However, we carried out a sensitivity analysis to assess the effect of different initial urn compositions for the different total sample sizes available. In detail, we considered the cases: $R_0 = W_0 = 5$ or $R_0 = W_0 = 10$.

To carry out the comparison with the non-adaptive design, we calculated the number of subjects allocated to the inferior treatment when the non-adaptive design was assumed. Let us denote this by n_W . In case (b) (reference setup: $n = 68$), n_W was known to be equal to 35, as in the HEN trial 35 out of 68 subjects were allocated to the inferior treatment. In addition, we have to estimate n_W in cases (a) and (c). In case (a) ($n = 58$), we built several ($\simeq 10,000$) subsamples of size 58 from the original HEN sample of total size 68; we estimated n_W as the mean number of subjects allocated to the CG across the available samples of size 58. To estimate n_W in case (c) ($n = 78$), we applied a proportion similar to that found in (b) on the 78 available subjects of this case. The corresponding n_W were equal to 29 for case (a) and 38 for case (c). The empirical power of the adaptive design was compared with the theoretical power of the non-adaptive t-test which was computed assuming that the true difference of the mean weight changes between the two arms is equal to the value obtained in the HEN trial.

All the analyses have been performed using a specialized code (available upon request from the authors) within the framework of the open-source statistical software R.

n	$N_W(n)$				n_W	$1 - \beta$	$1 - \hat{\beta}$
	1 st quartile	Mean	Median	3 rd quartile			
(a) 58	19	25.6	25	31	29	0.88	0.83
(b) 68	22	29.6	29	36	35	0.92	0.88
(c) 78	25	33.6	33	41	38	0.94	0.92

Table 1: Summary statistics (1st and 3rd quartiles, mean, and median) of the empirical distribution of the number of subjects assigned to the inferior treatment, $N_W(n)$, and empirical power, $1 - \hat{\beta}$, of the t-test for equal mean weight changes (corresponding to the treatment coefficient in the ANOVA model) for the different sample sizes n in the Randomly Reinforced Urn design, in comparison with the corresponding results for the non-adaptive design, n_W and $1 - \beta$. We reported in bold typeface the results obtained with the same sample size of the original Home Enteral Nutrition trial. The initial composition of the urns in all simulations was set at: $R_0 = W_0 = 1$.

3 Results

In this section we show the performance of the simulated RRU trials based on the HEN data. Table 1 shows some descriptive statistics of the empirical distribution of the number of subjects assigned to the inferior treatment, $N_W(n)$, and the empirical power of the t-test, $1 - \hat{\beta}$, for the different sample sizes n and a fixed initial urn composition $R_0 = W_0 = 1$, in comparison with the corresponding results for the non-adaptive design, n_W and $1 - \beta$.

For all sample sizes under consideration [cases (a)-(b)-(c)], the mean and the median of $N_W(n)$ were smaller than n_W , the number of subjects assigned to the inferior treatment by the non-adaptive design. It follows that the RRU design presented the 50% of probability (or more) to assign fewer subjects to the inferior treatment, as compared to the non-adaptive design. Although higher than n_W for all the sample sizes considered, the third quartile of $N_W(n)$ in the RRU design was very close to n_W for any n under consideration. In addition, the obtained values for the t-test's empirical power under the RRU design were close, but slightly smaller than, the corresponding power values derived in the non-adaptive design.

Further information on the distribution of $N_W(n)$ is provided by the boxplots reported in Figure 1. For any sample size, the median of $N_W(n)$ was below the dashed line indicating the number of subjects assigned to the inferior treatment

by the non-adaptive design. Similarly, we confirmed that, although higher, the third quartile was closer than the median to the dashed line for the three cases under consideration. In addition, the probability that $N_W(n)$ was less than n_W was close to 75% for any sample size under consideration. Finally, although mostly symmetric, the empirical distributions of the number of subjects assigned to the inferior treatment showed a high level of variability. This variability increases, as the total sample size increases.

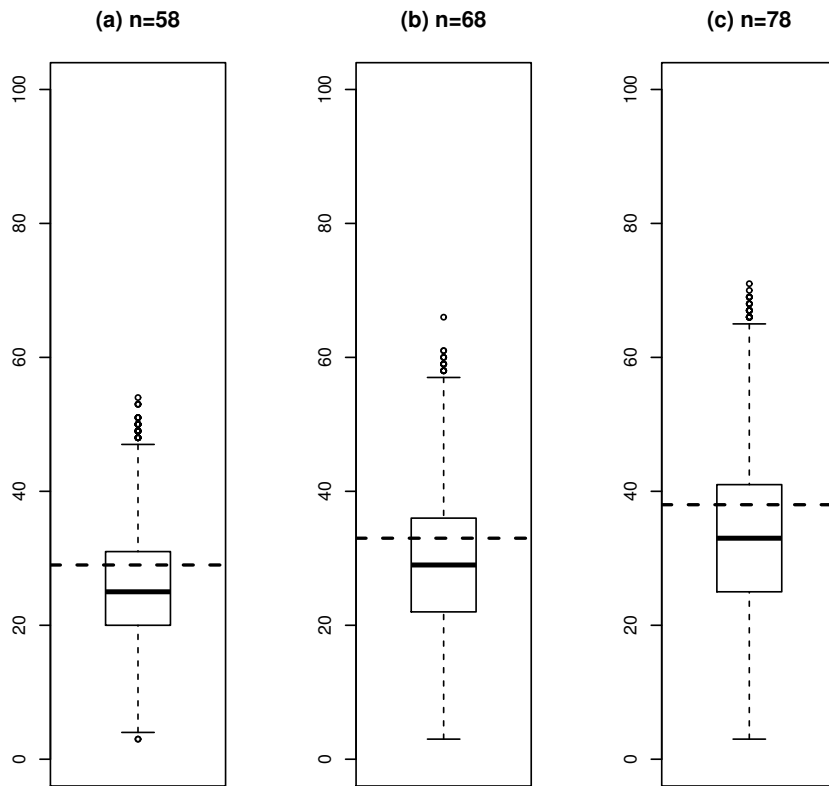


Figure 1: Boxplots of the number of subjects assigned to the inferior treatment (Control Group) in the three cases reported above each picture: (a) $n = 58$, (b) $n = 68$, (c) $n = 78$. The dashed line indicated the number of subjects assigned to the control group in the non-adaptive trial in the three cases.

Table 2 shows the results of the sensitivity analysis to different initial urn compositions. Our analysis was robust with respect to the initial urn composi-

Scenarios		Results				$1 - \hat{\beta}$
n	$R_0 = W_0$	$N_W(n)$				
		1 st quartile	Mean	Median	3 rd quartile	
58	1	19	25.6	25	31	0.83
	5	23	27.4	27	31	0.86
	10	24	27.9	28	31	0.87
68	1	22	29.6	29	36	0.88
	5	27	31.7	32	36	0.91
	10	29	32.6	32	37	0.91
78	1	25	33.6	33	41	0.92
	5	31	36.1	36	41	0.94
	10	33	37.3	37	42	0.94

Table 2: Sensitivity analysis to different urn initial compositions with $R_0 = W_0$: summary statistics (1st and 3rd quartiles, mean, and median) of the empirical distribution of the number of subjects assigned to the inferior treatment, $N_W(n)$, and empirical power, $1 - \hat{\beta}$, of the t-test for the combination of different available sample sizes n and urn initial compositions. We reported the reference scenario in bold typeface.

tion chosen. Indeed, the mean and median number of subjects allocated to the inferior treatment in the RRU design was still below the corresponding number of subjects in the non-adaptive design, for any n and fixed urn composition under consideration. In addition, as far as the number of balls initially inserted into the urn increases, for fixed n , the medians increase and, with $R_0 = W_0 = 10$, they almost reached the number of patients assigned to the inferior treatment in the non-adaptive design, n_W . The empirical power of the t-test was correspondingly higher than in the reference scenario of $R_0 = W_0 = 1$ for any n under consideration, thus making it almost identical to the empirical power in the non-adaptive design (see column $1 - \beta$ in Table 1). Similarly, as far as the number of balls initially inserted into the urn increases, for fixed n , the variability of $N_W(n)$ decreases and the adaptive design becomes closer and closer to the non-adaptive one.

4 Practical implementation of the RRU design

In the following, we give some technical details on how to implement a RRU design in the practice of clinical trials. The initial set-up at the trial start involves:

- total sample size n ;
- initial urn composition (R_0, W_0) ;
- utility function u .

We highlight that the implementation of the RRU design does not require any theoretical support or add-on code for sample size calculation. We just suppose that the trial investigators have calculated a total sample size n according to some approach, including traditional non-adaptive techniques.

There is no standard approach to choose the initial urn composition. However, extreme choices in the initial urn composition or in reinforcement values may lead to more and more extreme urn compositions, and, on this way, to swamp the urn and to compromise the adaptive mechanism of the trial. Such extreme situations may be avoided by setting the initial urn composition at the same order of magnitude of the reinforcement values. To this purpose, a general rule of thumb could be to set R_0 and W_0 such that: (i) their sum $(R_0 + W_0)$ is similar to the mean number of balls added to the urn at any time a new response is available, and (ii) the initial proportion of red balls in the urn, $Z_0 = R_0 / (R_0 + W_0)$, may reflect the *a priori* belief on which treatment is superior: the better the treatment \mathcal{R} , the higher is Z_0 . In our simulation study, we always set $R_0 = W_0$ and therefore: $Z_0 = 0.5$, meaning that we have no reason to believe *a priori* that one treatment is superior. In accordance with the equipoise principle, this proportion is typically set to $Z_0 = 0.5$ in the clinical practice.

The utility function, u , is, in principle, any positive monotone increasing function that maps the range of continuous responses into a positive bounded interval. For instance, in our simulation of the HEN trial, since the response values,

x , range in the interval $(-20, 20)$, we set $u(x) := (x + 20)/40$, in order to obtain reinforcements in $(0, 1)$.

The RRU design is practically implemented as follows:

- information storing:

The minimal set of information for implementing the RRU design may be collected in two databases. In the former one, we store for each subject (in rows) the following variables (in columns):

- subject ID;
- date of entry in the study;
- treatment assignment;
- date of response;
- response value.

In the latter one, we store for each date of subject response (in rows) the updated urn composition (R, W) (in columns). In the first row, we have the randomization date of the first patient entered in the study and (R_0, W_0) .

- subject randomization:

Equation (3) is implemented in the following R function:

```
new_subject <- function(R,W){
    ## (R,W): urn composition
    ## before the arrival of the new subject
    Z <- R/(R+W)          ## compute the urn proportion
    U <- runif(1,0,1)     ## sample a ball from the urn
    if (U<=Z) {T <- 'R'}  ## if the sampled ball is red,
                          ## assign the subject to treatment 'R'
    else if (U>Z) {T <- 'W'} ## if the sampled ball is white,
```

```

                                ## assign the subject to treatment 'W'
return(T)    ## return the treatment assigned to the subject
}

```

At the arrival of the i^{th} subject, we assign him/her to treatment calling the function with the current urn composition:

```
> new_treatment<-new_subject(R,W)
```

For instance, if the current urn composition is $(R, W) = (20, 25)$, we may obtain:

```
> new_treatment<-new_subject(20,25)
> new_treatment
[1] "R"
```

- urn updating:

Equation (1) is implemented in the following R function:

```

new_response <- function(R,W,xi,T,u){
    ## (R,W): the urn composition
    ## at the time the new response is available,
    ## xi: value of the new response,
    ## T: treatment associated to the new response ('R' or 'W')
    ## u: utility function
    if (T=='R') {R <- R+u(xi)}
        ## if xi is a response to treatment 'R',
        ## we increase the number of red balls in the urn
    else if (T=='W') {W <- W+u(xi)}
        ### if xi is a response to treatment 'W',
        ## we increase the number of white balls in the urn
    return(data.frame(R,W))
        ### return the urn composition after the update
}

```

The utility u has to be specified in a suitable R function. For instance, in our simulation the utility function is implemented in R as follows:

```
utility<-function(xi){(xi+20)/40}
```

When the subject gives his/her response to treatment, we call the `new_response()` function, with the updated information on treatment and response:

```
> new_urn_comp<-new_response(R,W,xi,new_treatment,utility)
```

and we obtain the updated urn composition. For instance, if the response to treatment is $xi=10$, we obtain:

```
> new_urn_comp<-new_response(20,25,10,'R',utility)
```

```
> new_urn_comp
```

```
      R  W
```

```
1 20.75 25
```

to be inserted as the new urn composition in the corresponding database, together with the new date of response.

5 Discussion

The current paper provides scientific and practical support to a broader use of the RRU response-adaptive design in randomized clinical trial planning, to alter the randomization ratio and favor patients allocation to the most effective treatment. We applied the RRU design in a simulation study based on real-life data from a randomized trial planned with a non-adaptive design, to compare the RRU performance with the original results from this trial. The simulation study gave scientific evidence of the possible advantages achieved with such a design, in that, as compared to the corresponding non-adaptive design, fewer subjects were allocated to the inferior treatment with a higher probability and

with a limited loss of power in hypothesis testing. We also showed the possibility of the practical use of the RRU design by providing the R code for its implementation.

The above results were derived by implementing a modification of the RRU design proposed in Muliere et al. (2006) that takes into account the more realistic case of delayed treatment responses. The adaptive randomization probabilities are determined based on the data observed thus far. This is a step forward in the modeling of real-life problems with urn models. Indeed, in the vast majority of medical research, including oncology, end-points are accessible for the evaluation of treatment efficacy during a relatively long follow-up period. In all these cases, our approach does not require suspending patient accrual and thus avoids wasting resources, adding administrative inconvenience and, in the end, having an infeasible trial. In addition, we referred to previous literature on asymptotic normality of the distribution of the test statistic under the null hypothesis of equal response means (Flournoy and May, 2009; Ghiglietti and Paganoni, 2014). In this paper, we opted for a Student's t-test. This choice is reasonable considering that the Student's t-test is more conservative than the z-test.

We acknowledge that sample size, test power and treatment effect size are three issues to be addressed within the RRU design and related one to the other. The current application considers the total sample size as a fixed aspect of the problem. Our choices were in line with the real set-up of the Italian efficacy study of HEN for gastrointestinal cancer patients. In particular, we hypothesized the following alternatives: 68 (total sample size of the HEN trial; reference set-up), together with 58 and 78, obtained by moving $\pm 15\%$ from 68. For all the three settings, our results showed that fewer subjects were allocated to the inferior treatment CG with a higher probability, as compared to the corresponding non-adaptive design. In detail, in approximately 75% of the 10,000 simulation runs the number of patients allocated to the inferior treatment by the RRU design is lower than the corresponding one in the non-adaptive de-

sign. For instance, in the reference case of 68 subjects collected in the original trial, the RRU design provided a median number of patients allocated to the inferior treatment equal to 29, versus a corresponding number as high as 35 in the non-adaptive design.

The total sample size is related to the effect size and power in a way that goes beyond the scope of this paper. Therefore, we did not include in the present work any theoretical development or add-on code for sample size calculation. One strategy could be to calculate a reference upper bound for the sample size by following standard approaches for non-adaptive designs and apply an adjustment that increases this value and protects from the loss of power expected with the RRU design. A more refined solution is to consider a trial simulation investigating the impact of the tuning parameters on total sample size and corresponding power for the target effect size.

In our application, the differences in the power of the t test for equal mean changes in the non-adaptive versus adaptive design range from 1% to 6% in absolute value. Power loss is an expected result, due to the unbalanced randomization ratio that is typical of adaptive designs. However, such loss seems to be limited and compensated by the expected gain in the number of subjects assigned to the superior treatment by the RRU design.

The amount of gain in the allocation procedure is related to the estimated effect size, too. Indeed, the convergence rate to the optimal allocation depends on the means and variances of the distributions of the responses to treatment, as proved in Flournoy and May, 2009.

Our simulation study was based on a trial with a differential effect between competitive treatments. For a high positive difference between the observed and the target effect size, the adaptive allocation will be strongly unbalanced in favor of the superior treatment, generating higher and higher loss of power. However, such loss should be compensated by the gain due to the larger-than-expected effect size. In the case of no observed effect between treatments, as the sample size increases to infinity, the stochastic process of the proportions of subjects

assigned to the superior treatment oscillates and does not convergence to 1 anymore. The distribution of this proportion is, indeed, the unique continuous solution of a functional equation involving unknown probability distributions on $[0, 1]$, as shown in Aletti et al. (2007) and Aletti et al. (2009). This is still a setting where a RRU design can be applied because there will be a sort of random selection of the arm “receiving” more subjects. In addition, although we do not know the exact functional form, we can still take advantage of the fact that the distribution of the proportion has no point mass and has lower and upper bounds. For instance, we can deal with potential outliers by setting sensible thresholds to allow updating of the urn composition only when the proportion falls within the specified range (see Aletti et al. 2013; Ghiglietti and Paganoni, 2016).

Randomly reinforced urn models share some general issues with sequential designs that might discourage their broader application in clinical trials. Patients’ characteristics may change during recruitment because time trends in the target population and changes in investigator’s conduct may happen. Indeed, since later-entered patients have a better chance of receiving the better treatment, clinicians might modify their behavior in selecting patients and assessing their characteristics. In both cases, allocation and assessment biases could distort the treatment effect estimate. All these issues are less relevant in short-running studies and when blinding procedures are applicable. Among feasibility issues, adaptive designs generally require that information on response to treatment be quickly ascertainable and used as far as it is available, but real-time data collection may be difficult in practice. In some applications, the response assessment implies a long follow-up time and, therefore, a new subject is ready for treatment assignment while responses to previously enrolled patients have not been observed yet. However, our variant of the RRU scheme is able to manage urn updating when information is available from small blocks of subjects. In addition, in the practical implementation of a RRU design, we had to select appropriate values for several parameters, including total sample

size, initial urn composition, and utility function for the urn reinforcement. Our choices were informed by the real example under consideration, as suggested in Section 4. However, the R code provided is very general and can be easily modified to cope with different real-life situations. If we assumed a different form for the utility function, we could simply insert it in the right hand side of the corresponding R function `utility`, with no modifications in the remaining R functions. Similarly, if we dealt with a binary outcome, the utility function could simply be equal to the identity function, as the response values lie in $[0,1]$. At the same time, we chose our real-life example to adhere, as far as possible, to some extra feasibility requirements. In detail, both the time interval between consecutive patient arrivals and the delay of the responses have to be similar in magnitude to the total study period divided by the total number of patients enrolled. For instance, in our example, the mean time intervals between consecutive patients was about 20 days, the mean delay of the responses was approximately 2 months, and the study lasts for about 4 years.

In conclusion, we provide support to the theoretical and practical use of RRU designs in randomized clinical trials: significant ethical and cost advantages are obtained over equal randomization, with fewer subjects assigned to the inferior treatment with a higher probability. Although we are aware of the many feasibility issues related to adaptive designs in general, we still believe that the RRU designs represent a valid attempt to develop an adaptive design in randomized clinical trials.

References

- [1] WEI, L.J. (1978). An application of an urn model to the design of sequential controlled clinical trials. *JASA*, 73, 559-563.
- [2] WEN, S., NING, J., COLLINS, S. and BERRY, D. (2017). A response-adaptive design of initial therapy for emergency department patients with heart failure. *Contemp. Clin. Trials*, 52, 46-53.

- [3] DURHAM, S.C., FLOURNOY, N. and LI, W. (1998). A sequential design for maximizing the probability of a response. *The Canad. J. Stat.*, 26 (3), 479-495.
- [4] ATKINSON, A.C. and BISWAS, A. (2014). Randomised response-adaptive designs in clinical trials. Chapman and Hall/CRC.
- [5] HU, F. and ROSENBERGER, W.F. (2006). The theory of response-adaptive randomization in clinical trials. Wiley, Hoboken, NJ.
- [6] MULIERE, P., PAGANONI, A.M. and SECCHI, P. (2006). A randomly reinforced urn. *J. Stat. Plann. Inference*, 136, 1853-1874.
- [7] FLOURNOY, N. and MAY, C. (2009). Asymptotics in response-adaptive designs generated by a two-color, randomly reinforced urn. *Ann. Stat.*, 37, 1058-1078.
- [8] FLOURNOY, N., MAY, C. and SECCHI, P. (2012). Asymptotically optimal response-adaptive designs for allocating the best treatment: an overview. *Int. Statist. Rev.*, 80, (2), 293-305.
- [9] LACHIN, J.M. and ROSENBERGER, W.F. (2002). Randomization in clinical trials. Wiley, New York.
- [10] R DEVELOPMENT CORE TEAM (2016). R: 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/>.
- [11] GAVAZZI, C., COLATRUGLIO, S., VALORIANI, F., MAZZAFERRO, V., SABBATINI, A., BIFFI, R., MARIANI, L. and MICELI, R. (2016). Impact of home enteral nutrition in malnourished patients with upper gastrointestinal cancer: a multicenter randomized clinical trial. *Europ. J. Cancer*, 64, 107-112.
- [12] BOZZETTI, F., GAVAZZI, C., MICELI, R., ROSSI, N., MARIANI, L., COZZAGLIO, L., BONFANTI, G. and PIACENZA, S. (2000). Perioperative total

- parenteral nutrition in malnourished, gastrointestinal cancer patients: a randomized, clinical trial. *JPEN J Parenter Enteral Nutr*, 24, 7-14.
- [13] MARTIN, L., SENESSE, P., GIOULBASANIS, I., ANTOUN, S., BOZZETTI, F., DEANS, C., STRASSER, F., THORESEN, L., JAGOE, R.T., CHASEN, M., LUNDHOLM, K., BOSAEUS, I., FEARON, K.H. and BARACOS, V.E. (2015). Diagnostic criteria for the classification of cancer-associated weight loss. *J. Clin. Oncol.*, 33, 90-99.
- [14] POCOCK, S.J. (1977). Group sequential methods in the design and analysis of clinical trials. *Biometrika*, 64, 191-199.
- [15] BAI, Z.D., HU, F. and ROSENBERGER, W.F. (2002). Asymptotic properties of adaptive designs for clinical trials with delayed response. *Ann. Stat.*, 30 (1), 122-139.
- [16] GHIGLIETTI, A. and PAGANONI, A.M. (2014). Statistical properties of two-color randomly reinforced urn design targeting fixed allocations. *Electr. J. Stat.*, 8, 708-737.
- [17] ALETTI, G., MAY, C. and SECCHI, P. (2007). On the distribution of the limit proportion for a two-color, randomly reinforced urn with equal reinforcement distributions. *Adv. Appl. Probab.*, 39, 3, 690-707.
- [18] ALETTI, G., MAY, C. and SECCHI, P. (2009). A central limit theorem, and related results, for a two-color randomly reinforced urn. *Adv. Appl. Probab.*, 41, 3, 829-844.
- [19] ALETTI, G., GHIGLIETTI, A. and PAGANONI, A.M. (2013). Randomly reinforced urn designs with prespecified allocations. *J. Appl. Probab.*, 50, 2, 486-498.
- [20] GHIGLIETTI, A. and PAGANONI, A.M. (2016). An urn model to construct an efficient test procedure for response adaptive designs. *Statist. Meth. Appl.*, 25, 2, 211-226.