Aust. N. Z. J. Stat. **46**(4), 2004, 675–684

A RANDOMIZED LONGITUDINAL PLAY-THE-WINNER DESIGN FOR REPEATED BINARY DATA

ATANU BISWAS[1]* AND ANUP DEWANJI[1]

Indian Statistical Institute

Summary

In some clinical trials with two treatment arms, the patients enter the study at different times and are then allocated to one of two treatment groups. It is important for ethical reasons that there is greater probability of allocating a patient to the group that has displayed more favourable responses up to the patient's entry time. There are many adaptive designs in the literature to meet this ethical constraint, but most have a single binary response. Often the binary response is longitudinal in nature, being observed repeatedly over different monitoring times. This paper develops a randomized longitudinal play-the-winner design for such binary responses which meets the ethical constraint. Some performance characteristics of this design have been studied. It has been implemented in a trial of pulsed electro-magnetic field therapy with rheumatoid arthritis patients.

Key words: adaptive design; allocation probability; longitudinal observations; performance characteristics; pulsed electro-magnetic field; randomized play-the-winner rule.

1. Introduction

Adaptive allocation design in the context of sequential clinical trials involving two (or more) treatments has been a challenging research area in the recent past. If patients enter any study in a sequential manner, it is desirable that an entering patient is allocated to the treatment group that so far has yielded more favourable responses. In such adaptive allocation designs, which depend on past allocations and response histories, an entering patient has higher probability of being allocated to the treatment group 'doing better' so far, which leads to more patients in the 'better' treatment group in the long run. Zelen (1969) introduced the popular concept of the play-the-winner sampling design to compare two single (not repeated) binary treatment responses. Subsequently, for a single dichotomous response, various authors have reported different adaptive allocation designs that can be explained by urn models: initially the urn contains an equal number of balls of each kind, then an appropriate number of balls are added after each response. See, for example, the randomized play-the-winner (RPW) rule of Wei & Durham (1978) where a fixed number of balls of the same (opposite) kind are added to the urn for a success (failure), and the success-driven design of Durham, Fluornoy & Li

Received August 2002; revised July 2003; accepted November 2003.

*Author to whom correspondence should be addressed.

[1] Applied Statistics Unit, Indian Statistical Institute, 203 B. T. Road, Kolkata - 700 108, India.
e-mail: atanu@isical.ac.in

Acknowledgments. The authors thank two referees, an associate editor, the Technical Editor, Dr Amita Pal, Ms Noelle Noble, Ms Kathy Nyland and Ms Sandy Walbrek for their valuable comments, which led to an improved presentation of the paper. The authors thank the Biometry Unit of the Indian Statistical Institute, Kolkata, for providing data from the PEMF experiment conducted therein. The authors also thank Mr Shyam Sundar Sahoo for some numerical computations, and Drs K.S. Ganguly and Tapas Samanta for their interest and encouragement in the present work.

(1998) where balls are added only for a success. Some applications of the RPW rule follow Bartlett *et al.* (1985) who studied the effect of extracorporeal membrane oxygenation (ECMO) in treating newborns with respiratory failure, and Tamura *et al.* (1994) who studied the effect of fluoxetine hydrochloride in treating out-patients with depressive disorder.

Recently an adaptive randomized placebo-controlled clinical trial was designed by the Indian Statistical Institute, Kolkata, to investigate the effect of pulsed electro-magnetic field (PEMF) therapy for the treatment of rheumatoid arthritis patients. In this particular trial, the patients are scheduled to go through an 'adjustment period' of 4 weeks followed by approximately 16 weeks of treatment (by therapy or placebo). The patients undergo this treatment three times a week and their conditions are monitored once a week. Therefore, the response pattern is longitudinal. Each response is a five-component ordinal vector on pain, tenderness, swelling, joint stiffness and functional disability, with a four-category response (nil/mild/moderate/severe) on each component. For simplicity, the multivariate ordinal response at any observation time is dichotomized as 'recurrence' or 'non-recurrence' by consulting the medical expert involved in the study (see Section 5). These binary responses are, in practice, treated as response histories, which update the urn composition and, hence, the future allocation probabilities. Therefore, we have longitudinal binary responses. In this paper, we develop an allocation design in the spirit of the RPW rule to meet the ethical demand of allocating more patients to the 'better' treatment group. Although the role of prognostic factors is important in the allocation design and the final analysis, we assume homogeneity among the patients in the following development for simplicity in notation and mathematical development. We also consider a sample design with a fixed number of patients.

The main features of this trial were staggered entry (sometimes in a batch) of the patients, missing observations and unequal numbers of observations for different patients. This complicated scenario motivated us to develop this new adaptive design. So far, not much work has appeared in the literature on the development of adaptive designs for longitudinal responses. We describe the design in Section 2. Section 3 considers two characteristics of the design: the proportions of allocations to the two competing treatments, and the proportions of both types of balls in the urn. Section 4 discusses the maximum likelihood estimation of the parameters. Section 5 describes the study at hand, adopting the design of Section 2, and illustrates the results of Sections 3 and 4. We end with some discussion in Section 6.

2. The proposed adaptive design

The patients enter the system sequentially, one or more at a time, where the interarrival times are not necessarily the same. At entry time, a patient is allocated to one of the two treatment groups (therapy or placebo) following an adaptive design, so that the patient has higher probability of being allocated to the 'better' group. The design can be illustrated by means of an urn model as follows.

We start with an urn having 2α balls of two types (T and P), α balls of each kind. Here T stands for PEMF therapy and P stands for placebo. The incoming patients are randomly allocated to either of the two treatment groups according to the urn composition at that time. We allocate the initial $2m$ patients randomly to the two treatment groups, in such a way that exactly m patients are allocated to each of the treatment groups, mimicking the initial balanced state of the urn. This is justified since, at this time, we have no information about superiority of one treatment over the other. The patients are monitored for a fixed (but possibly different) number of times and the monitoring time-points are not necessarily equispaced. At

each monitoring time-point, we observe whether recurrence (R) has occurred or not (N) for every patient under study. If recurrence (R) is observed for a particular patient, we treat it as a failure for the corresponding group (T or P) at that monitoring time, and, accordingly, add β balls of the opposite type to the urn. On the other hand, if non-recurrence (N) is observed, we treat it as a success for the corresponding group at that time, and, accordingly, add β balls of the same type. The idea behind the addition of balls is to skew the urn composition in favour of the treatment group having the 'better' record of success until that time.

When a patient enters the study, we draw a ball from the urn, note the type of the drawn ball, and return it immediately to the urn. The type of the ball drawn indicates the group to which the entering patient is allocated.

This sampling design can be viewed as a longitudinal version of the popular and well-known randomized play-the-winner rule. We call this design a randomized longitudinal play-the-winner (RLPW) rule. For any fixed (α, β), we denote it the RLPW(α, β) rule. Although m is also a design parameter, we do not mention it in the notation, by convention. From the above description, it is seen (in the next section) that the rule essentially depends on only one design parameter $\gamma = \beta/\alpha$, in which case we start with one ball of each kind, and add γ balls for every response. Conceptually, therefore, γ is allowed to be a fraction.

3. Some performance characteristics of the design

Let n denote the number of patients in the study. Let x_s and y_s denote the entry and exit times, respectively, of the sth patient entering the study, who is monitored on k_s occasions. These entry and exit times (x_s and y_s) and the number of monitorings (k_s) are assumed to be non-stochastic. Let $t_{s1}, \ldots, t_{sk_s}$ denote the monitoring times (assumed to be non-stochastic) for the sth patient. We then have $x_s \le t_{s1} < \cdots < t_{sk_s} \le y_s$. The sth patient is allocated to a group (T or P) at time x_s and monitoring starts from time t_{s1}. For the sth patient entering the study, we define the indicator variables δ_s and Z_{sj} as follows: δ_s takes the value 1 or 0 depending on whether allocation is made to the T or the P group, and Z_{sj} takes the value 1 or 0 corresponding to R or N in the jth monitoring at time t_{sj}.

For $s \ge 2m + 1$, we summarize the history of the clinical trial at time x_s- (just before the entrance of the sth patient) as follows:

$$R_{Tx_s} = \sum_{\ell=1}^{s-1} \delta_\ell \sum_{j=1}^{k_\ell} Z_{\ell j}\, \mathrm{I}\,(t_{\ell j} < x_s)\,, \qquad N_{Tx_s} = \sum_{\ell=1}^{s-1} \delta_\ell \sum_{j=1}^{k_\ell} \mathrm{I}\,(t_{\ell j} < x_s)\,, \tag{1}$$

where I denotes the indicator variable. These are the number of recurrences and the number of monitorings, respectively, in the T group before time x_s. The number of balls added for responses of the T patients before time x_s is βN_{Tx_s}, of which βR_{Tx_s} are balls of kind P. Similarly, we can obtain R_{Px_s} and N_{Px_s} with δ_ℓ replaced by $(1 - \delta_\ell)$ in (1). For allocating the sth patient, we use only the summarized data history up to time x_s-, that is $(R_{Tx_s}, R_{Px_s}, N_{Tx_s}, N_{Px_s})$. The conditional probability that the sth patient is allocated to the T group, given all the earlier allocations and responses, depends only on this summary history and is

$$\Pr(\delta_s = 1 \mid R_{Tx_s}, R_{Px_s}, N_{Tx_s}, N_{Px_s}) = \frac{\alpha + \beta((N_{Tx_s} - R_{Tx_s}) + R_{Px_s})}{B_{x_s}}\,, \tag{2}$$

where $B_{x_s} = 2\alpha + \beta(N_{Tx_s} + N_{Px_s})$ denotes the total number of balls in the urn just before time x_s. For our purpose, B_{x_s} is non-random.

Observing that $\Pr(\delta_s = 1)$ denotes the unconditional probability that the sth patient is allocated to the T group, we write

$$r_s = \Pr(\delta_s = 1), \quad \text{for } s \geq 1.$$

Note that $r_s = 0.5$ for $s = 1, \ldots, 2m$. We also write

$$\pi_{Tj} = \Pr(Z_{sj} = 1 \mid \delta_s = 1) \quad \text{and} \quad \pi_{Pj} = \Pr(Z_{sj} = 1 \mid \delta_s = 0),$$

which are the probabilities of R at the jth monitoring of a patient in the T and P groups, respectively.

As stated above, the primary goal of adopting such an adaptive allocation design is to skew the allocation pattern in favour of the treatment doing 'better'. Hence, the most reasonable performance characteristic is the proportion of allocations to the T group, given by $T_p = (n-2m)^{-1} \sum_{i=2m+1}^{n} \delta_i$, leaving aside the first $2m$ allocations. Our objective is to study the distribution of T_p. In particular, we obtain $\mathrm{E}(T_p)$, the expectation. Using (1), we have

$$\mathrm{E}(N_{Tx_s}) = \sum_{\ell=1}^{s-1} r_\ell \sum_{j=1}^{k_\ell} \mathrm{I}(t_{\ell_j} < x_s)$$

and $$\mathrm{E}(R_{Tx_s}) = \sum_{\ell=1}^{s-1} \mathrm{E}\Big(\sum_{j: t_{\ell_j} < x_s} Z_{\ell j} \mid \delta_\ell = 1 \Big) \Pr(\delta_\ell = 1) = \sum_{\ell=1}^{s-1} \Big(\sum_{j: t_{\ell_j} < x_s} \pi_{Tj} \Big) r_\ell .$$

A similar expression for $\mathrm{E}(R_{Px_s})$ can be obtained by replacing π_{Tj} by π_{Pj} and r_ℓ by $(1-r_\ell)$ in the expression for $\mathrm{E}(R_{Tx_s})$. Taking expectation in (2), r_s can be recursively obtained from

$$r_s = \frac{\alpha + \beta \sum_{\ell=1}^{s-1} \sum_{j: t_{\ell_j} < x_s} ((1-\pi_{Tj}) r_\ell + \pi_{Pj}(1-r_\ell))}{B_{x_s}}, \qquad \text{for } s \geq 2m+1. \quad (3)$$

Hence we have

$$\mathrm{E}(T_p) = \frac{1}{n-2m} \sum_{i=2m+1}^{n} r_i = \bar{r} \quad \text{(say)}. \quad (4)$$

Note that, from (3), the specification of the marginal probabilities, π_{Tj} and π_{Pj}, is sufficient for calculating r_s, and, hence, $\mathrm{E}(T_p)$.

As a closely related performance characteristic, we study the ultimate proportion of T balls in the urn at the end of the study, denoted by T_p^*: this is the right-hand side of (2) calculated at x_{n+1}, the time of termination of the study, and it gives us the probability of allocation to the therapy group at x_{n+1}. Similar to the derivation of (3), we see that

$$\mathrm{E}(T_p^*) = r_{n+1} = \frac{\alpha + \beta \sum_{\ell=1}^{n} \sum_{j=1}^{k_\ell} ((1-\pi_{Tj}) r_\ell + \pi_{Pj}(1-r_\ell))}{2\alpha + \beta N}, \quad (5)$$

where $N = N_{Tx_{n+1}} + N_{Px_{n+1}}$ is the total number of monitorings in the T and P groups. Later, in Table 1, we report the values of these two expectations, $\mathrm{E}(T_p)$ and $\mathrm{E}(T_p^*)$, under some simple modelling assumptions and monitoring schemes. We also have the following convergence result for T_p in a somewhat simplified scheme, where the numbers of monitorings per patient are equal. The proof is provided in the Appendix.

Theorem 1. *For k monitorings per patient, as $n \to \infty$, $T_p \xrightarrow{p} \pi$, where*

$$\pi = \frac{\sum_{j=1}^{k} \pi_{Pj}}{\sum_{j=1}^{k} (\pi_{Tj} + \pi_{Pj})}. \tag{6}$$

In the proof, $k_s = k$ is taken for convenience. The convergence in probability still holds for unequal but finite k_s values. Considering the Cesaro limit, the expectations in (4) and (5) have the same limit. Hence, $E(T_p^*)$ also has the limit π. The sequence $\{r_s, s \geq 2m+1\}$ is seen, in our numerical work, to be monotonically increasing or decreasing while converging to π. In the Appendix, we prove and characterize this result in a special scenario.

Both (4) and (5) depend solely on the marginal probabilities π_{Pj} and π_{Tj}. Since specification of these probabilities seems somewhat artificial, we consider the following simple model for numerical illustration. We write $\pi_{T1} = q_T$ and $\pi_{P1} = q_P$. For equispaced monitoring times, we assume that $\Pr(Z_{sj} = 1 \mid Z_{s1} = z_{s1}, \ldots, Z_{s,j-1} = z_{s,j-1}, \delta_s)$ depends, besides δ_s, only on the time since the last recurrence (that is, on $j - \ell_0$, where ℓ_0 is the maximum $\ell (< j)$ such that $z_{s\ell} = 1$). For example, if $z_{s\ell} = 1$ and $z_{s,\ell+1} = \cdots = z_{s,j-1} = 0$, then the above probability depends only on $j - \ell$ and is denoted by $q_{T,j-\ell}$ or $q_{P,j-\ell}$, for $\delta_s = 1$ or 0, respectively. In particular, if $z_{s1} = \cdots = z_{s,j-1} = 0$, then the above probability is q_{Tj} or q_{Pj}. This modelling is similar to that of Bonney (1987) for correlated binary data with some natural ordering in the observations. We obtain the recursive relation

$$\pi_{uj} = \sum_{\ell=1}^{j-1} q_{u,j-\ell} \pi_{u\ell} \left(\prod_{i=1}^{j-1-\ell} (1 - q_{ui}) \right) + q_{uj} \prod_{i=1}^{j-1} (1 - q_{ui}), \quad \text{for } u = T \text{ and } P. \tag{7}$$

We can then obtain r_s, using (3) and (7), in terms of q_{uj}. We now consider further modelling of the q_{uj} as

$$q_{uj} = 1 - (1 - q_u)^j, \quad \text{for } u = T \text{ and } P, \tag{8}$$

so that q_{uj} increases with j. As j increases, this model is found to have an increasing odds-ratio and decreasing correlation between $Z_{s\ell}$ and $Z_{s,\ell+j}$ for fixed ℓ — desirable features in this context.

We assume equal numbers ($k = 10$) of monitoring times for each patient. Patients enter the study one at a time with a gap of five monitoring times between successive patients. We take $m = 2$, $\alpha = 2$ and $\beta = 1$. Table 1 presents values of the two expectations (4) and (5) for different values of q_T and q_P using the model given by (7) and (8). Table 1 also provides the limiting value π, given by (6). For $q_T < q_P$, the expectations are greater than 0.5. The expectations for $q_T > q_P$ are less than 0.5 and, as expected, symmetric (in the case $q_T < q_P$) around 0.5, and therefore not reported. When $q_T = q_P$, from the symmetry of the design, and also using (3), (7) and (8), $r_s = 0.5$ for all s. The two expectations are very similar for all values of n. However, when the values of q_T and q_P are very small, the expectations are somewhat different from the limiting value since the convergence is slower. The difference decreases as n becomes large.

Estimates of expectations of T_p and T_p^*, corresponding to the monitoring scheme of our study with $n = 22$, are obtained by 10 000 simulations and presented in the last two columns of Table 1. As expected, for $q_T = q_P$, these estimated expectations are around 0.5; but for $q_T < q_P$, they are greater than 0.5. We see remarkable agreement of these estimates

TABLE 1

Values of $\bar{r}$ and r_{n+1} and the corresponding π for equal numbers ($k = 10$) of monitorings and equal (five monitoring times) interarrival times for each patient. Here $m = 2$, $\alpha = 2$, $\beta = 1$.

q_P	q_T	$n = 50$		$n = 100$		π	PEMF study	
		$\bar{r}$	r_{n+1}	$\bar{r}$	r_{n+1}		$\bar{r}$	r_{n+1}
0.002	0.001	0.506	0.508	0.507	0.509	0.663	0.515	0.524
0.004	0.001	0.517	0.522	0.521	0.527	0.792	0.540	0.562
0.005	0.001	0.522	0.529	0.527	0.535	0.824	0.550	0.578
0.01	0.005	0.523	0.530	0.528	0.536	0.649	0.537	0.557
0.02	0.01	0.535	0.545	0.543	0.553	0.636	0.549	0.576
0.05	0.01	0.601	0.627	0.620	0.644	0.765	0.642	0.706
0.05	0.02	0.566	0.581	0.577	0.592	0.652	0.591	0.633
0.1	0.01	0.666	0.701	0.691	0.724	0.832	0.713	0.799
0.1	0.05	0.565	0.576	0.573	0.582	0.604	0.584	0.614
0.2	0.05	0.644	0.663	0.656	0.672	0.696	0.659	0.708
0.2	0.1	0.580	0.588	0.585	0.592	0.600	0.582	0.604
0.5	0.1	0.705	0.714	0.711	0.717	0.722	0.655	0.682
0.5	0.2	0.628	0.631	0.630	0.632	0.633	0.577	0.586
0.8	0.1	0.779	0.780	0.780	0.781	0.781	0.671	0.693
0.8	0.2	0.707	0.705	0.705	0.704	0.704	0.592	0.603
0.8	0.5	0.583	0.580	0.581	0.579	0.579	0.518	0.518

with the corresponding exact values, although the monitoring schemes are very different. The corresponding standard deviations have also been estimated by simulation (not reported here), and seem to decrease with the increase in q_T and q_P values.

4. Maximum likelihood estimation

To describe maximum likelihood estimation of the model parameters, note that the allocation probabilities (of the δ_ℓ) involve only the design parameters α and β, and, given the allocations, the probabilities of responses at different monitoring times involve only the model parameters (see also Ware, 1989). Therefore, the relevant likelihood contribution, in terms of the q_{Tj} and q_{Pj} described in the previous section, from the sth patient is given by

$$L_s = \left(\prod_{i=1}^{z_{s\cdot}+1} \prod_{\ell=1}^{\Delta s_i - 1} (1 - q_{u\ell}) \right) \left(\prod_{i=1}^{z_{s\cdot}} q_{u,\Delta s_i} \right),$$

where $z_{s\cdot} = \sum_{j=1}^{k_s} z_{sj}$ is the number of recurrences in the sth patient with recurrence times, say, $u_{s1} < \cdots < u_{sz_{s\cdot}}$, $\Delta s_i = u_{si} - u_{s,i-1}$, for $i = 1, \ldots, z_{s\cdot} + 1$, with $u_{s0} = 0$ and $u_{s,z_{s\cdot}+1} = t_{sk_s} + 1$, and $u = T$ or P depending on whether $\delta_s = 1$ or 0. The total likelihood is then the product of L_s over $s = 1, \cdots, n$, which, for the model (8), simplifies to

$$L = \left(p_T^{f_T} \prod_{s:\,\delta_s=1} \prod_{i=1}^{z_{s\cdot}} (1 - p_T^{\Delta s_i}) \right) \left(p_P^{f_P} \prod_{s:\,\delta_s=0} \prod_{i=1}^{z_{s\cdot}} (1 - p_P^{\Delta s_i}) \right), \tag{9}$$

where $p_u = 1 - q_u$, for $u = T, P$, $f_T = \frac{1}{2} \sum_{s:\,\delta_s=1} \sum_{i=1}^{z_{s\cdot}+1} (\Delta s_i - 1)\Delta s_i$, and f_P is similar to f_T with the first sum being over all s with $\delta_s = 0$.

The likelihood (9) can be used to obtain maximum likelihood estimates of p_T and p_P by solving

$$f_u = \sum_{s:\,\delta_s = I_{u=T}} \sum_{i=1}^{z_{s\bullet}} \frac{\Delta s_i}{p_u^{-\Delta s_i} - 1}, \quad \text{for } u = T \text{ and } P. \tag{10}$$

The likelihood can also be used for testing a hypothesis involving q_T and q_P. For example, the null hypothesis H_0: $q_T = q_P$, can be of interest. Under H_0, the common parameter $p_T = p_P = p$, say, can be estimated from an equation similar to (10) with f_u replaced by $f_\cdot = f_T + f_P$ and the first sum being over all s.

5. An example: the PEMF trial

In this section, we briefly describe the trial and the results of the placebo-controlled PEMF therapy for which the proposed RLPW sampling design was initially developed. This is an implementation of the proposed design with $\alpha = 2$, $\beta = 1$ and $m = 2$. The study was conducted in the Indian Statistical Institute, Kolkata, from January 1999 to March 2000. Twenty-two patients participated in the study. The first four patients were randomly assigned in such a manner that two patients were treated by the PEMF therapy and the other two by the placebo. The remaining 18 patients were randomly allocated to either T or P group using the updated state of the urn at their entry times. The number of patient monitorings differed according to the condition of the patients, varying from 7 to 62. Of the 798 total monitorings, 16 recurrences were observed in the 22 patients of which four were in the T group and 12 in the P group. The study was a double-blind trial in the sense that neither the patients nor the medical expert were aware of the group-identification of the patients. During a monitoring, a patient was considered to have a recurrence if pain or joint stiffness was severe.

Data on the number of patients in the two groups clearly exhibits the superiority of the PEMF therapy over the placebo. Excluding the initial four patients, 14 out of 18 have been treated with the PEMF therapy by our adaptive design. Therefore, the observed value of T_p is $14/18 \approx 0.778$, which is well above the 50% mark. The urn proportions at different entry times exhibit a skewed pattern in favour of the PEMF therapy, with values of 0.630 at the entry time of the fifth patient and 0.731 at the entry of the 22nd patient. The ultimate proportion $T_p^* = 0.733$, showing benefit of therapy over the placebo. The observed value of T_p^* is quite close to that of T_p, showing a large sample characteristic in this small sample case also.

The maximum likelihood estimates of q_T and q_P are 0.00033 and 0.00409 and, under H_0, the maximum likelihood estimate of the common parameter $q_T = q_P = q$ is 0.00106. The likelihood ratio test for H_0: $q_T = q_P$ leads to the observed value of the statistic as 23.03 which, by comparison with the χ_1^2 distribution, is strong evidence for the superiority of the PEMF therapy.

6. Discussion

The choice of m, depending on the total number of patients, can be looked upon as a compromise between fully adaptive and fully balanced allocation. It ensures at least m patients to each treatment, though subsequently the adaptive procedure may allocate all the remaining patients to one particular treatment. In our study, the choice of $m = 2$ was driven by some prior knowledge of the flow of patient accrual.

The choice of design parameters α and β is crucial, although the limiting proportion of allocation π in (6) is free from these parameters. A large value of α (as compared to β) assigns considerable weight to the initial balanced allocation and, accordingly, the first few allocations are little influenced by the few response-driven added balls, as the urn is dominated by the initial balls. Therefore, a very large α forces the allocation pattern to stay near the 50% proportion (the allocation becomes almost a balanced random allocation scheme), which is a loss in the sense that a large number of patients get the inferior treatment. If m is moderately large, the effects of these parameters diminish quickly, as the urn soon becomes dominated by the response-driven added balls. If $\alpha = 0$, the urn composition is totally controlled by the response-driven added balls indicating a total belief in data for the allocation purpose. Thus, a decrease in α makes the allocation sensitive to responses and so increases the variability of the number of allocations in a group (see Rosenberger (1999) for such result in the context of the RPW rule). There is, therefore, a trade-off which needs to be addressed appropriately. An alternative and possibly more realistic design is to start with α balls of each kind and to add β_1 balls for non-recurrence and β_2 balls for recurrence. If recurrence is much less frequent than non-recurrence, as in our study, one may choose $\beta_1 < \beta_2$. Then, the denominator of the conditional probability (2) is random, resulting in a more complicated mathematical formulation. However, there is no problem in interpretation, understanding and implementation of such a design in practice.

Although the responses are dichotomized, it may be of interest to extend the present design for the original type of responses (i.e. multivariate ordinal). Suppose we have an M-component response vector $\boldsymbol{Z}_{sj} = (Z_{sj1}, \ldots, Z_{sjM})$ at the jth monitoring time of the sth patient, where each of Z_{sjw}, $w = 1, \ldots, M$, can take values $0, 1, \ldots, K$ (K may be allowed to vary with w). Without loss of generality, we assume that higher values of Z_{sjw} indicate worse health conditions. We again start with a balanced urn having α balls of each of the two kinds. We go on updating the urn by adding appropriate balls at every stage. In this case, we find, for $w = 1, \ldots, M$,

$$R_{Tx_s w} = \sum_{\ell=1}^{s-1} \delta_\ell \sum_{j=1}^{k_\ell} Z_{\ell j w}\, \mathrm{I}(t_{\ell j} < x_s) \quad \text{and} \quad N_{Tx_s} = K \sum_{\ell=1}^{s-1} \delta_\ell \sum_{j=1}^{k_\ell} \mathrm{I}(t_{\ell j} < x_s),$$

and similarly $R_{Px_s w}$ and N_{Px_s}. The number of balls added to the urn for the wth component of the response vectors in the T group before time x_s is $\beta_w N_{Tx_s}$, of which $\beta_w R_{Tx_s w}$ are of kind P. Thus, using the summarized data,

$$\mathcal{H}(x_s-) = (R_{Tx_s w}, R_{Px_s w}, w = 1, \ldots, M; N_{Tx_s}, N_{Px_s})$$

up to time x_s-, the conditional probability that the sth patient is allocated to the T group is given by

$$\Pr(\delta_s = 1 \mid \mathcal{H}(x_s-)) = \frac{\alpha + \sum_{w=1}^{M} \beta_w ((N_{Tx_s} - R_{Tx_s w}) + R_{Px_s w})}{2\alpha + (N_{Tx_s} + N_{Px_s}) \sum_{w=1}^{M} \beta_w}.$$

The subsequent derivations are along similar lines. Also, we may wish to incorporate any covariate information in the process of updating the urn. In that case, when the patient's condition is favourable in terms of the covariates, we add a larger number of balls of the other kind for R and add smaller number of balls of the same kind for N.

Appendix

Proof of Theorem 1. As $\{r_s, s \geq 2m+1\}$ is a bounded sequence, there exists a subsequence $\{r_{u(s)}\}$ which is convergent. Suppose it converges to some π_0. From (3), we find that

$$\frac{r_{u(s)+1}}{r_{u(s)}} \to 1 \quad \text{as} \quad u(s) \to \infty .$$

Then, for some $\epsilon > 0$,

$$\pi_0(1-\epsilon) \leq \liminf r_{u(s)+1} \leq \limsup r_{u(s)+1} \leq \pi_0(1+\epsilon) ,$$

and hence $$\limsup r_{u(s)+1} - \liminf r_{u(s)+1} \leq 2\pi_0\epsilon .$$

Since ϵ is arbitrary, the sequence $\{r_{u(s)+1}\}$ is convergent and it converges to π_0. Proceeding in this fashion the convergence of $\{r_s\}$ can be established.

Since k is fixed and finite, for large s,

$$\frac{r_s}{r_s^*} \approx 1 , \quad \text{where} \quad r_s^* = \frac{\alpha + \beta \sum_{\ell=1}^{s-1} \sum_{j=1}^{k} (\pi_{Pj} - (\pi_{Tj} + \pi_{Pj}) r_\ell^* + r_\ell^*)}{2\alpha + \beta(s-1)k} , \tag{11}$$

which is the unconditional probability of allocating the sth patient to group T, if all the k responses of each of the previous $(s-1)$ patients are available before the entrance of the sth patient into the study. Since k is finite, r_s^* also converges to π_0 as s goes to infinity. Hence, for the sake of simplicity, we work with r_s^* to find π_0. Letting $s \to \infty$ in (11) (taking the Cesaro limit on the right hand side), we get

$$\pi_0 = \frac{1}{k} \sum_{j=1}^{k} \left(\pi_{Pj} - (\pi_{Tj} + \pi_{Pj})\pi_0 + \pi_0\right) ,$$

which yields $$\pi_0 = \frac{\sum_{j=1}^{k} \pi_{Pj}}{\sum_{j=1}^{k} (\pi_{Tj} + \pi_{Pj})} = \pi .$$

Hence, $$\mathrm{E}(T_{\text{prop}}) = \frac{1}{n-2m} \sum_{i=2m+1}^{n} r_i \to \pi \quad \text{as} \quad n \to \infty ,$$

taking the Cesaro limit once again.

Now, as in Bandyopadhyay & Biswas (1997), it can be shown that

$$\Pr(\delta_s = 1 \mid \delta_i = 1) - \Pr(\delta_s = 1) \to 0 \quad \text{as} \quad s \to \infty ,$$

for any fixed i $(< s)$, which immediately proves that $\mathrm{cov}(\delta_i, \delta_s) \to 0$ as $s \to \infty$ (since the δ_j are binary). Hence, as in Bandyopadhyay & Biswas (1997), we have

$$\frac{1}{n^2} \mathrm{var}\left(\sum_{i=1}^{n} \delta_i\right) = \frac{1}{n^2}\left(\sum_{i=1}^{n} \mathrm{var}(\delta_i) + 2\sum_{i<j} \mathrm{cov}(\delta_i, \delta_j)\right) \to 0 , \quad \text{as} \quad n \to \infty .$$

Hence the convergence in probability is established.

Monotonic convergence of $\{r_s^*, s \geq 2m+1\}$

From (11), we have

$$r_{s+1}^* = \left(\frac{2\alpha + \beta(s-1)k}{2\alpha + \beta sk}\right) r_s^* + \frac{\beta}{2\alpha + \beta sk} \sum_{j=1}^{k} \left(\pi_{Pj} - (\pi_{Tj} + \pi_{Pj}) r_s^* + r_s^*\right). \qquad (12)$$

Therefore, using the right-hand side of (12) for r_{s+1}^*, we have

$$r_{s+1}^* \geq r_s^* \;\Leftrightarrow\; r_s^* \leq \frac{\sum_{j=1}^{k} \pi_{Pj}}{\sum_{j=1}^{k} (\pi_{Tj} + \pi_{Pj})} = \pi\,. \qquad (13)$$

Multiplying both sides of (13) by $((2\alpha + \beta sk) - \beta \sum_{j=1}^{k} (\pi_{Tj} + \pi_{Pj}))$ and rearranging, we get

$$(2\alpha + \beta(s-1)k) r_s^* + \beta \sum_{j=1}^{k} \left(\pi_{Pj} - (\pi_{Tj} + \pi_{Pj}) r_s^* + r_s^*\right) \leq (2\alpha + \beta sk) \frac{\sum_{j=1}^{k} \pi_{Pj}}{\sum_{j=1}^{k} (\pi_{Tj} + \pi_{Pj})},$$

and hence

$$r_{s+1}^* \leq \pi\,. \qquad (14)$$

Hence, from (13) and (14), if $r_{2m+1}^* \leq \pi$, the sequence $\{r_s^*, s \geq 2m+1\}$ is increasing and bounded by π. Also, from the proof of Theorem 1, the sequence converges to π. A similar result holds for $r_{2m+1}^* > \pi$, in which case the sequence $\{r_s^*, s \geq 2m+1\}$ is decreasing while converging to π. This result provides theoretical support for the numerical finding.

References

BANDYOPADHYAY, U. & BISWAS, A. (1997). Sequential comparison of two treatments in clinical trials: a decision theoretic approach based on randomized play-the-winner rule. *Sequential Anal.* **16**, 65–91.

BARTLETT, R.H., ROLOFF, D.W., CORNELL, R.G., ANDREWS, A.F., DILLON, P.W. & ZWISCHENBERGER, J.B. (1985). Extracorporeal circulation in neonatal respiratory failure: a prospective randomized trial. *Pediatrics* **76**, 479–487.

BONNEY, G.E. (1987). Logistic regression for dependent binary observations. *Biometrics* **43**, 951–973.

DURHAM, S., FLUORNOY, N. & LI, W. (1998). A sequential design for maximizing the probability of a favourable response. *Canad. J. Statist.* **26**, 479–495.

ROSENBERGER, W.F. (1999). Randomized play-the-winner clinical trials: review and recommendations. *Controlled Clinical Trials* **20**, 328–342.

TAMURA, R.N., FARIES, D.E., ANDERSEN, J.S. & HEILIGENSTEIN, J.H. (1994). A case study of an adaptive clinical trial in the treatment of out-patients with depressive disorder. *J. Amer. Statist. Assoc.* **89**, 768–776.

WARE, J.H. (1989). Investigating therapies of potentially great benefit: ECMO. *Statist. Sci.* **4**, 298–340.

WEI, L.J. & DURHAM, S. (1978). The randomized play-the-winner rule in medical trials. *J. Amer. Statist. Assoc.* **73**, 838–843.

ZELEN, M. (1969). Play-the-winner rule and the controlled clinical trial. *J. Amer. Statist. Assoc.* **64**, 131–146.