Journal of Statistical Planning and
Inference 136 (2006) 1845–1852

Response-adaptive designs for continuous outcomes

Anastasia Ivanova[a,*], Atanu Biswas[b], Anna Lurie[c]

[a]*Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-7420, USA*
[b]*Applied Statistics Unit, Indian Statistical Institute, 203 B. T. Road, Kolkata 700 108, India*
[c]*Department of Mathematics, St. Mary's University, 1 Camino Santa Maria, San Antonio, TX 78228, USA*

Available online 29 August 2005

Abstract

In this paper, we propose two new response-adaptive designs to use in a trial comparing treatments with continuous outcomes. Both designs assign more subjects to the better treatment on average. The new designs are compared with existing procedures and the equal allocation. The power of the treatment comparison is assessed.

MSC: Primary 62L05; secondary 60F15

Keywords: Sequential procedure; Response-adaptive design; Urn model

1. Introduction

Consider the situation where two treatments with continuous outcomes, X_1 and X_2, are compared in a randomized trial. We assume that treatment outcomes are normally distributed with $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$. Without loss of generality, we assume that a higher value indicates a more favourable response, and $\mu_1 \geqslant \mu_2$. Response-adaptive designs are allocation procedures that change the allocation away from equal allocation according to a certain goal. One of the goals is to assign more subjects to the better treatment on average. Another goal is to optimize a certain criterion, for example, to minimize the sample size given fixed precision of estimation or fixed power of the test (Hu and Ivanova, 2004). Most of the response-adaptive designs have been developed for treatments with binary outcomes. Urn models, such as the randomized play-the-winner design (Wei and Durham , 1978; Wei, 1979), are designs that assign more subjects to the better treatment on average. The doubly adaptive coin design (Eisele, 1994) is an example of an allocation procedure that can be used to optimize a certain criterion. In this paper, we focus on response-adaptive designs for treatments with continuous outcomes. Among existing designs for continuous outcomes are the randomized design (Melfi et al., 2001) and the design of Bandyopadhyay and Biswas (2001).

A well-known Neyman allocation for treatments with continuous outcomes is the allocation that assigns $\pi = \sigma_1/(\sigma_1 + \sigma_2)$ subjects to the first treatment. This allocation minimizes the sample size when the power is fixed. The randomized design of Melfi et al. (2001) targets the Neyman allocation. If the better treatment has smaller variance then $\pi < 0.5$, hence the optimal allocation π will result in assigning fewer subjects to the better treatment. The design of Bandyopadhyay and Biswas (2001) has been developed with the goal of assigning more subjects to the better treatment on average. However, the power of the treatment comparison can be lost. In this paper, in Section 2 we introduce two response-adaptive designs for treatments with continuous outcomes. Both designs assign more subjects to the better treatment on average. Section 3 considers a response-adaptive procedure, where sequential estimation is used to tune the design parameters. Simulation results are presented in Section 4 and conclusions in Section 5.

2. Two designs for continuous outcomes

In this section, we define two new response-adaptive designs for continuous outcomes and discuss their properties.

Design 1: The urn contains balls of three types. Balls of types 1 and 2 represent the two treatments. Balls of the third type are called immigration balls. Initially, the urn contains one ball of each type. When a subject arrives a ball is drawn from the urn at random. If the ball is an immigration ball, no subject is treated, and the ball is returned to the urn together with 2 additional balls, one of each treatment type. If a ball of type i, $i = 1, 2$, is drawn the next subject is given the corresponding treatment and an outcome $X_i = x_i$, $i = 1, 2$, is observed. If $x_i > k$ for some cut-off value k, the ball is replaced and hence the urn composition remains unchanged. If $x_i \leqslant k$, the ball is not replaced.

In Design 1, a more favourable response to treatment i will maintain the existing urn composition, whereas a less favourable response will reduce the number of balls of that type. Design 1 is the drop-the-loser design (Ivanova, 2003), where the outcome is a success if $x_i > k$ and a failure otherwise. The probability of not replacing a ball of type i is $q_i = \mathrm{P}(X_i \leqslant k) = \Phi\{(k - \mu_i)/\sigma_i\}$, where Φ is the cumulative distribution function for the standard normal random variable. The choice of the cut-off value k will be discussed later.

Design 2: The second response-adaptive design requires specification of the design parameters c and $T > 0$. Design 2 is similar to Design 1, except that when an outcome $X_i = x_i$ is observed the ball is replaced with probability $\Phi\{(x_i - c)/T\}$ and not replaced with probability $1 - \Phi\{(x_i - c)/T\}$.

Design 2 is equivalent to the following procedure. When an outcome $X_i = x_i$ is observed, random variable $V \sim \mathrm{N}(c, T^2)$ is generated independently of X_i. The ball is then replaced if $x_i > V$ and not replaced if $x_i \leqslant V$. Hence, the probability of not replacing the ball of type i, $i = 1, 2$, in Design 2 is equal to

$$\mathrm{P}(X_i \leqslant V) = \Phi\left(\frac{c - \mu_i}{\sqrt{\sigma_i^2 + T^2}}\right) = \mathrm{P}\left\{X_i \leqslant \frac{(c - \mu_i)\sigma_i}{\sqrt{\sigma_i^2 + T^2}} + \mu_i\right\}.$$

That is, Design 2 yields the same marginal probability of replacing a ball as does Design 1 with treatment specific cut-offs $k_i = \mu_i + (c - \mu_i)\sigma_i / \sqrt{\sigma_i^2 + T^2}, i = 1, 2$. Design 2 is the drop-the-loser rule with $q_i = \Phi\left\{(c - \mu_i)/\sqrt{\sigma_i^2 + T^2}\right\}$. Following Ivanova (2003), the limiting allocation proportion to the first treatment for the two designs is equal to $q_2/(q_1 + q_2)$, with corresponding q_1 and q_2.

3. The choice of design parameters

3.1. Specifying the design parameters in terms of unknown $\mu_1, \mu_2, \sigma_1, \sigma_2$

Consider the choice of the design parameter k in Design 1. Let $w \in (0, 1)$, then define $k = w\mu_1 + (1 - w)\mu_2$, so that $k \in (\mu_1, \mu_2)$. For this choice of k, if $\mu_1 > \mu_2$, $q_1 = \Phi\{-(1 - w)(\mu_1 - \mu_2)/\sigma_1\} < 0.5$ and $q_2 = \Phi\{w(\mu_1 - \mu_2)/\sigma_2\} > 0.5$. Hence, in the limit, more subjects are allocated to the treatment with larger mean. We recommend using $(\mu_1 + \mu_2)/2$ as k in Design 1. Similarly, we recommend setting $c = (\mu_1 + \mu_2)/2$ in Design 2. Let $q_i^{(1)}$ be the probability of not returning the ball of type i, $i = 1, 2$, in Design 1 and let $q_i^{(2)}$ be corresponding probabilities in Design 2. If $k = c \in (\mu_1, \mu_2)$, $q_1^{(1)} < q_1^{(2)}$ and $q_2^{(1)} > q_2^{(2)}$, hence $q_2^{(1)}/(q_1^{(1)} + q_2^{(1)}) > q_2^{(2)}/(q_1^{(2)} + q_2^{(2)}) > 0.5$. That is, the limiting allocation proportion to the first treatment in Design 1 is further away from 0.5 than the limiting allocation proportion in Design 2. As we will see in Section 4 allocation proportions that are too extreme negatively affect the efficiency of treatment comparison. Fig. 1 shows the limiting allocation proportion to the first treatment for the two designs for a wide range of cut-off values k and c, $k = c$.

The root of the pooled variance $\sqrt{(\sigma_1^2 + \sigma_2^2)/2}$ can be used as T in Design 2. A value of T much larger than the pooled variance will result in probabilities of not returning the ball to the urn that are close to 0.5 making the adaptive mechanism very weak. Values of T much smaller than the pooled variance will result in assigning too many subjects to the better treatment and consequently might result in loss of the power of treatment comparison. If $c = (\mu_1 + \mu_2)/2$ and $T = \sqrt{(\sigma_1^2 + \sigma_2^2)/2}$ the marginal probabilities of returning the ball to the urn are $q_1 = \Phi\left\{-0.5(\mu_1 - \mu_2)/\sqrt{1.5\sigma_1^2 + 0.5\sigma_2^2}\right\}$ and $q_2 = \Phi\left\{0.5(\mu_1 - \mu_2)/\sqrt{0.5\sigma_1^2 + 1.5\sigma_2^2}\right\}$.

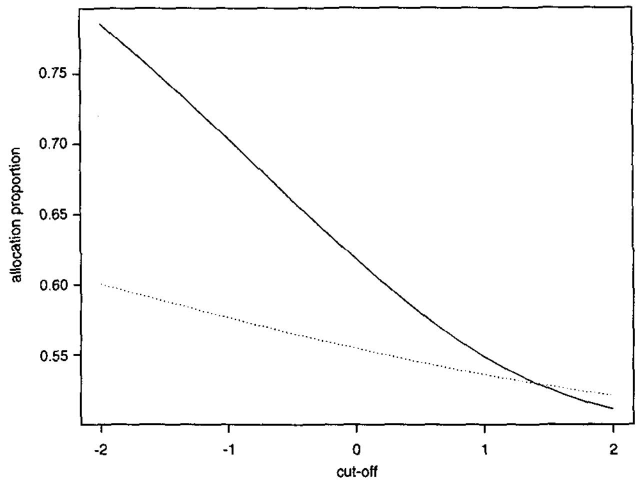

Fig. 1. Allocation proportion for Design 1 (solid line) and Design 2 (dashed line) for different values of cut-offs $k = c$ with $\mu_1 = 0.5$, $\mu_2 = 0$, $\sigma_1 = \sigma_2 = 1$, $T = 1$.

3.2. Estimating the design parameters from the data

Design parameters $k=(\mu_1+\mu_2)/2$ in Design 1, and $c=(\mu_1+\mu_2)/2$ and $T=\sqrt{(\sigma_1^2+\sigma_2^2)/2}$ in Design 2 can be sequentially estimated from the data already collected. Consider independent identically distributed sequences from the two populations $\{X_{1,j}, j \geqslant 1\}$ and $\{X_{2,j}, j \geqslant 1\}$. Let δ_j be 1 or 0 depending on whether the jth subject was allocated to treatment 1 or 2. Let $N_i(m)$ be the number of subjects assigned to treatment $i = 1, 2$, by the time m subjects are assigned, $N_1(m) = \sum_{j=1}^{m} \delta_j$ and $N_2(m) = \sum_{j=1}^{m}(1-\delta_j)$. Define sample means as $\bar{X}_1(m) = \sum_{j=1}^{m} \delta_j X_{1,j}/N_1(m)$ and $\bar{X}_2(m) = \sum_{j=1}^{m}(1-\delta_j)X_{2,j}/N_2(m)$, and sample variances as

$$s_1^2(m) = \frac{\sum_{j=1}^{m}(\delta_j X_{1,j} - \bar{X}_1)^2}{N_1(m)-1} \quad \text{and} \quad s_2^2(m) = \frac{\sum_{j=1}^{m}\{(1-\delta_j)X_{2,j} - \bar{X}_2\}^2}{N_2(m)-1}.$$

Consider the following modification of Design 2. The first $m_0 = 2m_0'$ subjects are assigned to the two treatments, m_0' to each. Assume that $m \geqslant m_0$ subjects have been assigned and their outcomes observed. Then subject $m+1$ is assigned according to Design 2 with $c_m = \{\bar{X}_1(m) + \bar{X}_2(m)\}/2$ and $T_m = \sqrt{\{s_1^2(m) + s_2^2(m)\}/2}$. The following theorem gives the limit of the allocation proportion to the first treatment, $N_1(m)/m$ as m tends to infinity.

Theorem. *The limiting allocation proportion $N_1(m)/m$ for the design described above tends to $q_2/(q_1+q_2)$ with $q_1 = \Phi\left\{-0.5(\mu_1-\mu_2)/\sqrt{1.5\sigma_1^2+0.5\sigma_2^2}\right\}$ and $q_2 = \Phi\left\{0.5(\mu_1-\mu_2)/\sqrt{0.5\sigma_1^2+1.5\sigma_2^2}\right\}$.*

Proof. Consider the sequence of designs, where mth design is Design 2 with c_m and T_m as design parameters, $m \geqslant m_0$. Let $q_{1,m}$ and $q_{2,m}$ be the probability of not replacing a ball of type 1 and 2 correspondingly in Design 2 with c_m and T_m. From Theorem 1 of Ivanova (2003) the mth design of the sequence can be embedded into a family of two continuous-time linear immigration–death processes having a common immigration process with an immigration rate 1 and independent death processes with death rates $q_{1,m}$ and $q_{2,m}$. Hence for every m, $m \geqslant m_0$, the probability of selecting treatment i, $i = 1, 2$, is strictly away from 0, and $N_i(m) \to \infty$ as $m \to \infty$. Consequently, $\bar{X}_i(m)$ and $s_i^2(m)$ are consistent estimates of μ_i and σ_i^2, $i = 1, 2$. Since $q_{i,m}$ is a continuous function of μ_i and σ_i, $i = 1, 2$, $q_{1,m}$ converges almost surely to $q_1 = \Phi\left\{-0.5(\mu_1 - \mu_2)/\sqrt{1.5\sigma_1^2 + 0.5\sigma_2^2}\right\}$ and $q_{2,m}$ to $q_2 = \Phi\left\{0.5(\mu_1 - \mu_2)/\sqrt{0.5\sigma_1^2 + 1.5\sigma_2^2}\right\}$. □

Modification of Design 1 can also be considered, where the estimate of the average of the means, $k_m = \{\bar{X}_1(m) + \bar{X}_2(m)\}/2$, is used in place of the design parameter k. The limiting allocation proportion for this design is equal to $q_2/(q_1+q_2)$ with $q_1 = \Phi\{-(\mu_1 - \mu_2)/(2\sigma_1)\}$ and $q_2 = \Phi\{(\mu_1 - \mu_2)/(2\sigma_2)\}$. The proof of the result is similar to the proof of the theorem.

4. Design comparisons

In this section, we compare the two new designs with the design of Bandyopadhyay and Biswas (2001) (henceforth B&B design) and the equal allocation. The B&B design is described as follows. The first $m_0 = 2m_0'$ subjects are assigned to the two treatments, m_0' to each. Assume that $m \geqslant m_0$ subjects have been assigned and their outcomes observed. Subject $m + 1$ is allocated to treatment 1 with probability $\Phi[\{\bar{X}_1(m) - \bar{X}_2(m)\}/M]$ and to treatment 2 with probability $1 - \Phi[\{\bar{X}_1(m) - \bar{X}_2(m)\}/M]$. The B&B design was used with $M = 1$. We also simulated the B&B design with the value of M chosen such that the limiting allocation for the B&B design is the same as for Design 2. Designs 1 and 2 were used with $k = c = (\mu_1 + \mu_2)/2$, and $T = \sqrt{(\sigma_1^2 + \sigma_2^2)/2}$. We also considered Design 2, where both c and T are estimated from the data sequentially (Design 2E). We used $m_0 = 6$ in the simulations with both Design 2E and the B&B procedure. Parameters c and T were estimated after subjects 10, 20, 40 and then after every 40th subject. Each design is run 5000 times.

We used a statistic $\{\bar{X}_1(n) - \bar{X}_2(n)\}/\sqrt{s_1^2(n)/N_1(n) + s_2^2(n)/N_2(n)}$ to test the hypothesis $\mu_1 = \mu_2$. Here n is the total number of subjects in the trial. Satterthwaite's approximation was used to construct the reference distribution (Rosner, 1995, p. 272). The sample size was chosen to yield the power of 0.8 for two-sided size 0.05 test when equal allocation is used. We simulated the case of the null distribution with different sample sizes and different values of σ_1 and σ_2, and concluded that the size of the test is well preserved under all the designs.

Simulation results in the case of equal variances are presented in Table 1. Without loss of generality $\mu_2 = 0$, and $\sigma_1 = \sigma_2 = 1$. Since the variances are equal, substantial

Table 1
Comparison of the response-adaptive designs and equal allocation in the case of equal variances with $\mu_2 = 0$, and $\sigma_1 = \sigma_2 = 1$

	Power	E(Y) (SD)	E(X) (SD)	π (SD)	π_0
$\mu_1 = 0.3$, $n = 350$					
Design 1	0.79	172.54 (9.49)	0.17 (0.05)	0.56 (0.03)	0.56
Design 2	0.80	173.01 (9.66)	0.16 (0.06)	0.54 (0.03)	0.54
Design 2E	0.80	173.58 (9.98)	0.16 (0.06)	0.54 (0.03)	0.54
B&B, $M = 1$	0.79	169.88 (9.61)	0.19 (0.06)	0.62 (0.06)	0.62
B&B, $M = 2.99$	0.78	173.35 (9.61)	0.16 (0.06)	0.54 (0.04)	0.54
Equal allocation	0.80	175.00 (9.29)	0.15 (0.05)	0.50 (0.00)	0.50
$\mu_1 = 0.5$, $n = 128$					
Design 1	0.79	61.87 (5.83)	0.29 (0.09)	0.59 (0.03)	0.60
Design 2	0.79	62.43 (5.73)	0.28 (0.09)	0.56 (0.04)	0.57
Design 2E	0.79	62.56 (5.49)	0.28 (0.09)	0.56 (0.04)	0.57
B&B, $M = 1$	0.75	59.33 (6.23)	0.34 (0.10)	0.69 (0.10)	0.69
B&B, $M = 2.83$	0.79	62.22 (5.56)	0.29 (0.09)	0.57 (0.05)	0.57
Equal allocation	0.80	64.00 (5.55)	0.25 (0.09)	0.50 (0.00)	0.50
$\mu_1 = 0.7$, $n = 66$					
Design 1	0.79	31.08 (4.16)	0.42 (0.13)	0.60 (0.05)	0.64
Design 2	0.80	31.73 (4.04)	0.40 (0.13)	0.57 (0.05)	0.60
Design 2E	0.78	31.69 (4.02)	0.40 (0.13)	0.57 (0.05)	0.60
B&B, $M = 1$	0.70	28.74 (4.28)	0.51 (0.15)	0.73 (0.12)	0.73
B&B, $M = 2.76$	0.79	31.36 (3.97)	0.41 (0.13)	0.59 (0.07)	0.60
Equal allocation	0.80	33.00 (3.91)	0.35 (0.12)	0.50 (0.00)	0.50
$\mu_1 = 1.1$, $n = 28$					
Design1	0.77	12.66 (2.55)	0.67 (0.21)	0.60 (0.06)	0.71
Design 2	0.79	13.12 (2.56)	0.63 (0.21)	0.58 (0.06)	0.65
Design 2E	0.78	13.25 (2.54)	0.63 (0.21)	0.57 (0.06)	0.65
B&B, $M = 1$	0.54	10.91 (2.66)	0.84 (0.22)	0.77 (0.11)	0.86
B&B, $M = 2.83$	0.77	12.66 (2.72)	0.68 (0.22)	0.62 (0.09)	0.65
Equal allocation	0.81	14.06 (2.38)	0.55 (0.18)	0.50 (0.00)	0.50

Designs 1 and 2 were used with $k = c = (\mu_1 + \mu_2)/2$, and $T = \sqrt{(\sigma_1^2 + \sigma_2^2)/2}$. Design 2E sequentially estimates parameters c and T from the data. The total number of responses lower than $(\mu_1 + \mu_2)/2$ is denoted by Y. The average response over two treatments is denoted by E(X). The observed allocation proportion to treatment one is denoted by π, and the limiting allocation proportion is denoted by π_0.

deviations from equal allocation resulted in loss of power. For example, the power for the B&B design with $M = 1$ in the last scenario is only 0.54. However, small deviations from equal allocation preserved power pretty well and resulted in observing better outcomes on average. The design where design parameters were sequentially estimated from the data, Design 2E, performed very similarly to the Design 2 with a slight loss of power.

Table 2 displays simulation results for the case of unequal variances. As expected, adaptive designs resulted in some gain of power in the first three scenarios where the larger variance corresponds to the better treatment. All adaptive designs had power less than 0.80 in the

Table 2
Comparison of the response-adaptive designs and equal allocation in the case of unequal variances with $\mu_2 = 0$

	Power	E(Y) (SD)	E(X) (SD)	π(SD)	π_0
$\mu_1 = 0.5, \sigma_1 = 2, \sigma_2 = 1, n = 316$					
Design 1	0.81	162.47 (9.12)	0.28 (0.10)	0.57 (0.03)	0.57
Design 2	0.82	163.32 (8.89)	0.28 (0.10)	0.55 (0.03)	0.55
Design 2E	0.81	163.80 (9.10)	0.28 (0.10)	0.55 (0.03)	0.55
B&B, $M = 1$	0.81	158.02 (11.04)	0.33 (0.13)	0.67 (0.12)	0.69
B&B, $M = 4.38$	0.81	163.26 (8.91)	0.27 (0.10)	0.54 (0.03)	0.55
Equal allocation	0.80	150.26 (8.78)	0.25 (0.09)	0.50 (0.00)	0.50
$\mu_1 = 1, \sigma_1 = 2, \sigma_2 = 1, n = 79$					
Design 1	0.83	40.75 (4.56)	0.61 (0.21)	0.61 (0.04)	0.63
Design 2	0.81	41.23 (4.47)	0.58 (0.20)	0.58 (0.05)	0.59
Design 2E	0.81	41.58 (4.55)	0.57 (0.20)	0.57 (0.05)	0.59
B&B, $M = 1$	0.77	36.70 (5.74)	0.78 (0.28)	0.78 (0.15)	0.84
B&B, $M = 4.44$	0.79	41.28 (4.60)	0.58 (0.21)	0.58 (0.07)	0.59
Equal allocation	0.80	36.29 (4.26)	0.50 (0.18)	0.50 (0.00)	0.50
$\mu_1 = 1, \sigma_1 = 3, \sigma_2 = 1, n = 158$					
Design 1	0.82	84.78 (6.56)	0.61 (0.21)	0.61 (0.03)	0.61
Design 2	0.82	85.95 (6.46)	0.57 (0.20)	0.57 (0.04)	0.57
Design 2E	0.81	86.60 (6.55)	0.56 (0.21)	0.56 (0.04)	0.57
B&B, $M = 1$	0.81	78.80 (11.21)	0.75 (0.34)	0.75 (0.21)	0.84
B&B, $M = 6.09$	0.81	86.41 (6.81)	0.56 (0.21)	0.56 (0.05)	0.56
Equal allocation	0.80	69.10 (6.02)	0.50 (0.18)	0.50 (0.00)	0.50
$\mu_1 = 1, \sigma_1 = 1, \sigma_2 = 3, n = 158$					
Design 1	0.69	64.21 (6.50)	0.62 (0.15)	0.63 (0.04)	0.65
Design 2	0.77	66.03 (5.95)	0.57 (0.16)	0.57 (0.04)	0.57
Design 2E	0.77	66.42 (5.97)	0.56 (0.15)	0.57 (0.04)	0.57
B&B, $M = 1$	0.30	55.18 (7.34)	0.84 (0.13)	0.84 (0.14)	0.84
B&B, $M = 5.76$	0.77	66.15 (5.90)	0.57 (0.15)	0.57 (0.05)	0.57
Equal allocation	0.80	84.17 (6.22)	0.25 (0.18)	0.50 (0.00)	0.50

Designs 1 and 2 were used with $k = c = (\mu_1 + \mu_2)/2$, and $T = \sqrt{(\sigma_1^2 + \sigma_2^2)/2}$. Design 2E sequentially estimates parameters c and T from the data. The total number of responses lower than $(\mu_1 + \mu_2)/2$ is denoted by Y. The average response over two treatments is denoted by E(X). The observed allocation proportion to treatment one is denoted by π, and the limiting allocation proportion is denoted by π_0.

last scenario where the better treatment has smaller variance. The loss of power is not substantial when Design 2 is used because its allocation proportion is not as extreme as for other designs. For this reason we recommend Design 2.

Comparison of Design 2 and the B&B design that yields the same limiting allocation proportion shows that the allocation proportion for the B&B design is more variable (Tables 1 and 2). This excessive variability resulted in some loss of power for the B&B design. The distribution of the allocation proportion for the B&B design even for large sample sizes is skewed (scenario 4 in Table 2). At the same time the distribution of the allocation proportion for Designs 1 and 2 is close to normal for large sample sizes (Ivanova, 2003).

The B&B procedure, and Designs 1 and 2 assign more subjects to the better treatment on average. However, occasionally less than 50% of the subjects can be assigned to the treatment with the better true mean even for large sample sizes. Consider scenario 1 from Table 1. In 10% of the runs the B&B design with $M = 2.99$ resulted in assigning fewer subjects to the treatment with the better true mean. In several runs as little as 151 subjects (43%) were assigned to the better treatment. In comparison, Design 2 resulted in assigning fewer subjects to the treatment with the better true mean in 5% of the runs, with the minimum number of subjects assigned to the better treatment equalled to 160 (46%). The reasons are the variability of response to treatment and the variability of the design itself. For two designs with the same limiting allocation proportion, the larger the variability of the allocation proportion the more likely the design will assign fewer subjects to the better treatment. Ivanova (2003) pointed that out when comparing the drop-the-loser design with more variable the play-the-winner rule.

5. Conclusions

In this paper we introduced two drop-the-loser-type designs for continuous responses. According to simulation results these designs yield adaptive allocation for continuous responses with smaller variability than the design of Bandyopadhyay and Biswas (2001). Theoretical comparison is not possible because the variability of allocation proportion of the design of Bandyopadhyay and Biswas (2001) is not known. Simulations show that the proposed designs can be advantageous compared to the equal allocation since they yield better treatment responses on average without substantial loss of power (or with some gain in power). Present work assumes the structure where there are no delayed responses and no covariate information in the data. Incorporation of all of the above in the design is a topic of future research.

References

Bandyopadhyay, U., Biswas, A., 2001. Adaptive designs for normal responses with prognostic factors. Biometrika 88, 409–419.

Eisele, J.R., 1994. The doubly adaptive biased coin design for sequential clinical trials. J. Statist. Plann. Inference 38, 249–261.

Hu, F., Ivanova, A., 2004. Adaptive design. In: Chow, S.-C. (Ed.), Encyclopedia of Biopharmaceutical Statistics, second ed. Marcel Dekker, New York.

Ivanova, A., 2003. A play-the-winner-type urn design with reduced variability. Metrika 58, 1–13.

Melfi, V.F., Page, C., Geraldes, M., 2001. An adaptive randomized design with applications to estimation. Canad. J. Statist. 29, 107–116.

Rosner, B., 1995. Fundamentals of Biostatistics. Duxbury, Boston.

Wei, L.J., 1979. The generalized Polya's urn design for sequential medical trials. Ann. Statist. 7, 291–296.

Wei, L.J., Durham, S.D., 1978. The randomized play-the-winner rule in medical trials. J. Amer. Statist. Assoc. 73, 840–843.