Optimal Response-Adaptive Designs for Continuous Responses in Phase III Trials

Atanu Biswas[*,1], **Rahul Bhattacharya**[2], and **Lanju Zhang**[3]

[1] Applied Statistics Unit, Indian Statistical Institute, 203 B.T. Road, Kolkata 700108, India
[2] Department of Statistics, Asutosh College, 92 S.P. Mukherjee Road, Kolkata 700026, India
[3] Department of Biostatistics, MedImmune Inc., One Medimmune Way, Gaithersburg, MD, 20878, USA

Received 18 October 2006, revised 24 January 2007, accepted 5 April 2007

Summary

Optimal response-adaptive designs in phase III clinical trial set up are gaining more interest. Most of the available designs are not based on any optimal consideration. An optimal design for binary responses is given by Rosenberger et al. (2001) and one for continuous responses is provided by Biswas and Mandal (2004). Recently, Zhang and Rosenberger (2006) proposed another design for normal responses. This paper illustrates that the Zhang and Rosenberger (2006) design is not suitable for normally distributed responses, in general. The approach cannot be extended for other continuous response cases, such as exponential or gamma. In this paper, we first describe when the optimal design of Zhang and Rosenberger (2006) fails. We then suggest the appropriate adjustments for designs in different continuous distributions. A unified framework to find optimal response-adaptive designs for two competing treatments is proposed. The proposed methods are illustrated using some real data.

Key words: Constraints; Ethical allocation; Minimization; Truncated normal distribution; Two parameter exponential family.

1 Introduction

Response-adaptive randomized designs are used in phase III clinical trials to achieve some *ethical* gain by allocating a larger number of patients to the better treatment. The idea is to *skew* the allocation in favor of the better treatment by using the available information, for sequentially entering patients. Several response-adaptive designs are available for this purpose, most of which were introduced from intuitive considerations. The play-the-winner (PW) rule (Zelen, 1969), the randomized play-the-winner (RPW) rule (Wei and Durham, 1978), the generalized Pòlya urn (GPU) design (Wei, 1979), and the drop-the-loser (DL) rule (Ivanova, 2003) are for binary treatment responses. The designs for continuous responses are the linear rank test statistic based design (Rosenberger, 1996), link function based design (Bandyopadhyay and Biswas, 2001), Wilcoxon score based design (Bandyopadhyay and Biswas, 2004), and utility based design (Atkinson and Biswas, 2005), among others. Real life applications of response-adaptive designs are due to Bartlett et al. (1985), Rout et al. (1993), Tamura et al. (1994), Biswas and Dewanji (2004), among others.

Optimal response-adaptive designs drew much attention recently. Rosenberger et al. (2001) extended the approach of Hayre (1979) and introduced the optimal design for binary responses. Biswas and Mandal (2004) provided a response-adaptive design for continuous responses. Specially, they studied normally distributed responses and exponentially distributed responses. The design of Atkinson and

Biswas (2005) was obtained by maximizing some utility function. Recently, Zhang and Rosenberger (2006), as an alternative to the approach of Biswas and Mandal (2004), proposed an optimal design for normally distributed responses. It is important to note that the design of Zhang and Rosenberger (2006) provides an allocation rule which takes square root of the estimates of normal mean, which can be negative. Consequently, this design fails, in general.

In this paper, we propose appropriate modifications to make the design of Zhang and Rosenberger (2006) applicable in a general situation in Section 3. In Section 4, we provide suitable adjustments to apply the Zhang and Rosenberger (2006) rule for positive-valued random variables like the exponential or gamma. Optimal response-adaptive designs for general two-parameter exponential families, in a general situation, are discussed in Section 5. The designs are illustrated using some real data in Section 6. Section 7 concludes.

2 Designs Proposed in the Literature

Zhang and Rosenberger (2006) discussed some of existing response-adaptive designs, namely the *doubly adaptive biased coin design* of Hu and Zhang (2004), the *link function based design* of Bandyopadhyay and Biswas (2001), and the *optimal design* of Biswas and Mandal (2004). In addition, Zhang and Rosenberger (2006) proposed a new optimal design as follows. Let n_A and n_B be the target sample sizes to the two treatments, $n_A + n_B = n$. Zhang and Rosenberger (2006) assumed $X_A \sim N(\mu_A, \sigma_A^2)$ and $X_B \sim N(\mu_B, \sigma_B^2)$, where X_A and X_B are the responses of patients assigned to two treatments A and B, and a *smaller* response is more desirable. Consequently a smaller value of the total expected responses from all the subjects is desirable. Hence they considered the following optimization problem:

$$\min_{n_A/n_B} \{\mu_A n_A + \mu_B n_B\}, \tag{1}$$

subject to

$$\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B} = K, \tag{2}$$

for some constant K. This restriction on $\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B}$ will preserve a specified level of power for the test of treatment equivalence. Solution of (1) yields ρ, the targeted allocation proportion to A, as

$$\rho = \frac{\sigma_A \sqrt{\mu_B}}{\sigma_A \sqrt{\mu_B} + \sigma_B \sqrt{\mu_A}}. \tag{3}$$

In contrast, Biswas and Mandal (2004) considered the following optimization problem

$$\min_{n_A/n_B} \left\{ n_A \Phi\left(\frac{\mu_A - c}{\sigma_A}\right) + n_B \Phi\left(\frac{\mu_B - c}{\sigma_B}\right) \right\}. \tag{4}$$

subject to (2), where c is some threshold constant. Thus Biswas and Mandal (2004) considered minimization of the total number of responses larger than a threshold c. Since a smaller response is desirable, a sufficiently large response indicates adverse effect of the treatment, which is treated as a *failure*. The threshold c is a boundary between treatment effectiveness and treatment failure. The minimization of (4) can be interpreted as minimization of the total expected failures. The solution of (4) is

$$\rho = \frac{\sigma_A \sqrt{\Phi\left(\frac{\mu_B - c}{\sigma_B}\right)}}{\sigma_A \sqrt{\Phi\left(\frac{\mu_B - c}{\sigma_B}\right)} + \sigma_B \sqrt{\Phi\left(\frac{\mu_A - c}{\sigma_A}\right)}}. \tag{5}$$

In practice, the design is implemented by estimating the parameters using the available data up to the first i patients, and plug them in the expression of ρ to find the allocation probability to treatment A for the $(i+1)$th patient.

It is important to note that for normally distributed responses, μ_A and μ_B can take any value in the real line, positive or negative. If, at any stage, the estimate of μ_A or μ_B becomes negative, the design of Zhang and Rosenberger (2006) fails. For the numerical computations, they considered large positive values of μ_A and μ_B and small values of σ_A and σ_B. Consequently, the design worked. But, in reality, in many cases, μ_A and μ_B, or their estimates can be negative. Consider the trial of fluoxetine hydrochloride reported by Tamura et al. (1994). The responses were the changes in $HAMD_{17}$ (or negative of this change) after the treatment, which are measured on a 53-point scale. The change can be approximated as a continuous variable (see Atkinson and Biswas, 2005). In reality, the changes can be positive or negative. Quite naturally, at some stage, the estimate of μ_A or μ_B, that is the observed mean of changes, can be negative, and the design (3) would fail in such a situation. Thus, the design proposed by Zhang and Rosenberger (2006) is not suitable in reality, in general, specially for normally distributed responses.

We conducted a simulation study using 5,000 simulations for $n = 80$ patients with different values of $(\mu_A, \mu_B, \sigma_A, \sigma_B)$. We obtained the percentage of cases where the ZR rule will crash (due to negative estimate of μ_A and/or μ_B for at least once among the 80 allocations). Figure 1 gives the percentages against the treatment difference, $\mu_B - \mu_A$. It is quite clear that there is considerable probability of such negative estimate unless both μ_A/σ_A and μ_B/σ_B are large positive. For a fixed value of μ_A, the value of μ_B strats from this value and increases gradually to exibit higher treatment difference. Quite natu-

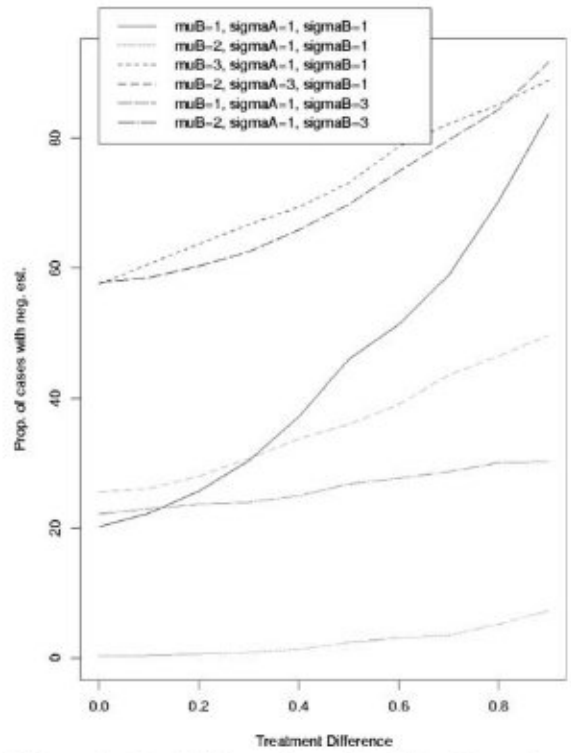

Figure 1 Probability of crashing ZR rule against treatment difference $\mu_B - \mu_A$ for normal responses for different combinations of $(\mu_B, \sigma_A, \sigma_B)$.

rally, the percentage increases when μ_A, μ_B are small. When μ_A and μ_B are close to zero (say $\mu_A - 2\sigma_A$ or $\mu_B - 2\sigma_B$ is negative), this percentage is alarming. This indicates that the ZR rule is not applicable, in general, in reality. Also, the limiting proportion given by (3) does not exist when μ_A and/or μ_B is negative.

The above drawback of the ZR rule can be controlled if a large positive constant d is added to all the responses, as the allocation probability will be the estimate of

$$\rho = \frac{\sigma_A \sqrt{\mu_B + d}}{\sigma_A \sqrt{\mu_B + d} + \sigma_B \sqrt{\mu_A + d}},$$

where the terms under squared-root can be made positive for large value of d. But the allocation probability will largely depend on the choice of d. Thus, although we can get rid of estimated negative mean responses, this will induce further arbitrariness in the procedure. Also, for large value of d, the allocation becomes the Neyman allocation, which is not an ethical allocation.

The ZR design may be applicable in some other situations, where X_A and X_B are known to be positive valued random variables, e.g., for $X_A \sim \text{exponential}(\mu_A)$ and $X_B \sim \text{exponential}(\mu_B)$. The objective function (1) is sensible, where $X \sim \text{exponential}(\mu)$ means $E(X) = \mu$. Here the constraint (2) will be replaced by

$$\frac{\mu_A^2}{n_A} + \frac{\mu_B^2}{n_B} = K. \tag{6}$$

The solution reduces to

$$\rho = \frac{\sqrt{\mu_A}}{\sqrt{\mu_A} + \sqrt{\mu_B}}. \tag{7}$$

See Biswas and Mandal (2004, Section 2) for a similar situation for exponential responses. If, for example, X_k follows the gamma distribution with density

$$f(x) = \frac{1}{\mu_k^{\beta_k} \Gamma(\beta_k)} \exp(-x/\mu_k)\, x^{\beta_k - 1}, \quad x > 0,$$

the objective function is

$$\beta_A \mu_A n_A + \beta_B \mu_B n_B,$$

and the constraint will be

$$\frac{\beta_A \mu_A^2}{n_A} + \frac{\beta_B \mu_B^2}{n_B} = K. \tag{8}$$

The solution is given by (7), the same as exponential.

In the following Sections we will develop appropriate designs for general normally distributed responses and a unified approach for distributions in the exponential family.

3 Design Based on Truncated Normal Distribution

For normal responses, the BM rule works well. To apply the Zhang and Rosenberger (2006) approach for normal responses when one or two of μ_A and μ_B are close to 0, we need to ensure that the estimates of μ_k's are always positive. Here we assume that $X_k \sim N(\mu_k, \sigma_k^2)$, $k = A, B$, where $\mu_k > 0$. Thus, $\hat{\mu}_k = \bar{X}_k$ if $\bar{X}_k > 0$, and $\hat{\mu}_k = 0$ if $\bar{X}_k \le 0$.

Alternatively, one may assume that the responses would be always positive. Thus, one needs to consider $X_k \sim TN_{\Re^+}(\mu_k, \sigma_k^2)$, $k = A, B$, which is a truncated normal distribution, where the $N(\mu_k, \sigma_k^2)$

density is truncated in $\Re^+$, the positive part of the real line. Letting ϕ and Φ be the density and distribution function of standard normal, we have

$$\mu_k^* = E_{\Re^+}(X_k) = \mu_k + \frac{\phi(\mu_k/\sigma_k)}{\Phi(\mu_k/\sigma_k)}\,\sigma_k\,,$$

and

$$\sigma_k^{*2} = V_{\Re^+}(X_k) = \sigma_k^2\left[1 - \left(\frac{\phi(\mu_k/\sigma_k)}{\Phi(\mu_k/\sigma_k)}\right)^2 - \frac{(\mu_k/\sigma_k)\,\phi(\mu_k/\sigma_k)}{\Phi(\mu_k/\sigma_k)}\right],$$

both of which are functions of μ_k and σ_k. See Johnson and Kotz (1970, pp. 81–83). The problem (1) and (2) and the solution (3) should be the same except that μ_k and σ_k are replaced by μ_k^* and σ_k^*. In some applications, where the responses will be always positive by nature, this type of truncation may be the method to work. In case of large μ_A and μ_B and small σ_A and σ_B (as in the simulation study of Zhang and Rosenberger (2006)), this truncation will not matter much from the computational point of view, but the background theory will be correct. Depending on the situation, one can think of truncating in some interval $[a, b] \subseteq \Re^+$. See an example in Section 6.

Figure 2 gives the limiting allocation proportion to the better treatment against treatment difference $\Delta = \mu_B - \mu_A$, where $\sigma_A = 1$, $\sigma_B = 1$, for the three rules: (i) ZR: Zhang and Rosenberger rule, (ii) BM: Biswas and Mandal rule, and (iii) TN: Truncated normal rule. For BM rule, $c = (\mu_A + \mu_B)/2$ is taken. The limiting allocation proportion of the BM rule is always much higher. However, this is true for a c which is close to the median of the responses, and not true for any choice of c. If c is in the tail of both the response distributions, which is a very poor choice of c, the BM rule will fail to allocate ethically. The allocation will be close to the Neyman allocation in this case. Our computations show (not given in this paper for the sake of brevity) that for a choice of c near to the inferior treatment mean, the BM rule assigns fewer number of subjects to the better treatment, which is even less than the corresponding number for the ZR rule.

From the Figure 2 we further observe that the ZR rule and the TN rule are almost equivalent up to a treatment difference of 3. It is interesting to note that the allocation proportion to the better treatment decreases in Δ after $\Delta > 4.5$. This is due to the fact that σ_k^{*2} is a complicated function of μ_k, and σ_A^{*2} becomes much larger than σ_B^{*2} for large Δ. This factor dominates the allocation after $\Delta > 4.5$. Thus, we observe that the limiting allocation in the ZR rule is quite different from TN rule for large treatment difference.

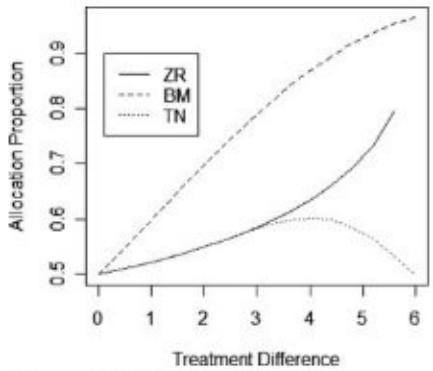

Figure 2 Limiting allocation proportion to treatment A for $\sigma_A = 1$, $\sigma_B = 1$. ZR: Zhang and Rosenberger, BM: Biswas and Mandal, TN: Truncated normal.

In contrast, the design (5) will work well for any kind of responses, exponential or normal. In fact, the objective function in (4) has the interpretation that it is the expected number of responses greater than the threshold c.

4 Design for Positive Valued Random Variables

The optimal design of ZR provides the allocation proportion (7) for exponential or gamma distributions. Although apparently this is alright (no possibility of negative estimates in the process), it has serious drawback. Since we are consistently assuming a smaller value indicates a better response, $\mu_A < \mu_B$ should provide $\rho > 0.5$. Also, ρ should be an increasing function of $\Delta = \mu_B/\mu_A$. But, the expression provides a completely opposite pattern of allocation, a larger μ_B provides a smaller allocation proportion to A. The reason is that for exp*onential* (μ_k), the variance is μ_k^2, while the expectation is μ_k. Hence, $\sqrt{\mu_A}\,\sigma_B$ is dominated by $\sqrt{\mu_B}$, which is the square-root of the expectation of the responses by treatment B. Thus, the variances dominate the allocation procedure, and the nature of the distribution make the allocation completely *unethical*. A similar scenario arises for gamma distributions. Thus, the optimization problem (1) subject to (6) is not even suitable for exponential or gamma responses. We need to set this optimal rule in some other way.

To achieve this, to put more weight on the mean (than the variance), we set the optimal rule as to miniize

$$\mu_A^{\alpha} n_A + \mu_B^{\alpha} n_B\,, \tag{9}$$

subject to (6), where α is a design parameter which is to be determined by satisfying some (ethical) requirement. The solution of (9) subject to (6) reduces to

$$\rho = \frac{\mu_B^{\alpha/2-1}}{\mu_A^{\alpha/2-1} + \mu_B^{\alpha/2-1}}\,. \tag{10}$$

A choice of $\alpha = 1$ in (10) gives the ZR rule. For ethical allocation, we need $\alpha > 2$. One may fix the value of α, for example, by fixing the allocation proportion to $\rho_0 > 0.5$ at a treatment difference $\delta = \mu_B/\mu_A = \delta_0 > 1$. In that case,

$$\alpha = 2\left[\frac{\log\left(\dfrac{\rho_0}{1-\rho_0}\right)}{\log \delta_0} + 1\right].$$

For gamma responses, the design minimizes

$$(\beta_A \mu_A)^{\alpha} n_A + (\beta_B \mu_B)^{\gamma} n_B \tag{11}$$

subject to (8). The solution for $\alpha = \gamma$ in (11) is

$$\rho = \frac{\beta_B^{(\alpha-1)/2}\mu_B^{(\alpha-2)/2}}{\beta_A^{(\alpha-1)/2}\mu_A^{(\alpha-2)/2} + \beta_B^{(\alpha-1)/2}\mu_B^{(\alpha-2)/2}}\,. \tag{12}$$

Here α can be determined by setting an allocation proportion ρ_0 at $\beta_B\mu_B/(\beta_A\mu_A) = \delta_1$ and $\beta_B\mu_B^2/(\beta_A\mu_A^2) = \delta_2$, that is

$$\alpha = \frac{2\left[\dfrac{1}{2}\log \delta_2 + \log\left(\dfrac{\rho_0}{1-\rho_0}\right)\right]}{\log \delta_1}\,.$$

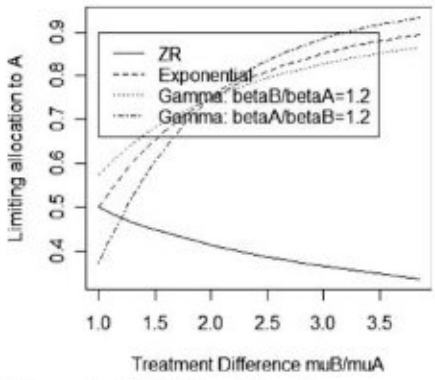

Figure 3 Limiting allocation proportion to treatment A for exponential and gamma distributions using different rules.

Figure 3 provides the limiting allocation proportions for the exponential and gamma distributions ($\beta_A = \beta_B = 1$ in the gamma case provides the exponential) against $\delta = \mu_B/\mu_A$. In particular, we consider (i) ZR: Zhang-Rosenberger-type rule of (7), (ii) Exponential: optimal design of the exponential model provided by (10), here $\alpha = 2\left(\frac{\log 3}{\log 2} + 1\right)$ is chosen by setting $\rho_0 = 0.75$ at $\mu_B/\mu_A = \delta_0 = 2$, (iii) Gamma: optimal design of the gamma model provided by (12), here $\alpha = 2(\log 2 + 0.5 \log 1.2 + \log 3)/(\log 2 + \log 1.2)$ is chosen by setting $\rho_0 = 0.75$ at $\mu_B/\mu_A = 2$ and $\beta_B/\beta_A = 1.2$, and (iv) same as (iii) for $\beta_A/\beta_B = 1.2$. Here $\alpha = 2(\log 2 - 0.5 \log 1.2 + \log 3)/(\log 2 - \log 1.2)$. From Figure 3, we observe that the allocation proportion to A for the ZR rule decreases as δ increases, for both the exponential and gamma models (allocation proportions are same for ZR rule), and it is unethical. The optimal rule for the exponential distribution has limiting allocation 0.5 at $\delta = 1$, and it increases steadily with δ. For gamma distribution the allocation proportion at $\delta = 1$ may be greater than or less than 0.5 depending on the value of β_B/β_A. If $\beta_B/\beta_A > 1$, the allocation proportion at $\delta = 1$ is greater than 0.5, and it is less than 0.5 at $\delta = 1$ when $\beta_B/\beta_A < 1$. But, in any case, unlike the ZR-type rule, the allocation proportion increases with the increase of δ. Less than 50% allocation to A for a δ slighly greater than 1, is due to the fact that the allocation at this stage is more driven by the variances.

Here we discuss the possible choice of α from another consideration. In any clinical trial, there is an upper bound of the allowable sample size due to funding and time constraint. We examine the possible size of the trial for different values of α. The approximate sample size to detect a departure Δ from the null with power $1 - \beta$ for a size λ test is

$$n = \frac{(\sigma_0 \tau_\lambda + \sigma_1 \tau_\beta)^2}{\Delta^2},$$

where σ_i^2/n is the asymptotic variance of the estimated treatment difference under H_i, $i = 0, 1$, H_0 and H_1 be the null and alternative hypotheses. For numerical illustration, we determine the approximate sample size required in each procedure to maintain 80% power to detect a given shift $\Delta = 0.4, 0.6, 0.8$ when $\mu_A = 1.0$. We always consider a test with size 0.05. Figure 5 gives the plot for proposed procedure (P), ZR rule (which is same as that for the equal allocation (E) rule). We observe that the size of the trial is lower in the proposed procedure whenever $\alpha < 2$, is equal for $\alpha = 2$, and is higher otherwise. However, this amount of additional samples required for $\alpha > 2$ is almost negligible (an excess of 1–2% than in the equal allocation case). But for such a sample size the ethical gain for the proposed procedure is maximum. Thus any given α exceeding 2 can be a sensible choice.

5 General Optimal Allocation Rule

As a general approach, one may choose to minimize

$$n_A \Psi(\mu_A) + n_B \Psi(\mu_B), \tag{13}$$

subject to (2), where $\Psi(\cdot)$ is a function such that $\Psi(x)$ is increasing in x, and $\Psi : \Re \to S(\Re^+)$, $S(\Re^+)$ being a subset of $\Re^+$.

Note that, the minimization problem (13) subject to (2) is quite similar to the formulation of Jennison and Turnbull (2000, p. 328), where the formulation was for continuous treatment responses, and Ψ_A and Ψ_B were functions of treatment differences. In particular, Jennison an Turnbull (2000) considered minimizing the loss function of the form

$$L(\eta) = u(\eta)\, n_A + v(\eta)\, n_B\,,$$

where $\eta = \mu_A - \mu_B$, and $u(\eta)$ and $v(\eta)$ are strictly positive with $v(\eta)$ increasing in η for $\eta > 0$ and $u(\eta)$ increasing as η decreases for $\eta < 0$, when a high response is desirable. But, we present this as a function of individual treatment parameters, which is easy to interpret.

Now, based on this Ψ, the optimal allocation is

$$\rho = \frac{\sigma_A \sqrt{\Psi(\mu_B)}}{\sigma_A \sqrt{\Psi(\mu_B)} + \sigma_B \sqrt{\Psi(\mu_A)}}\,. \tag{14}$$

One may wish to choose an appropriate Ψ. This can be done in several ways. One may wish to achieve a specified allocation proportion $\rho_0 > 0.5$ for a treatment difference $\Delta = \mu_B - \mu_A = \delta > 0$. If we consider $\Psi(x) = \exp(dx)$, then from (14) we immediately get

$$d = \frac{2}{\delta} \log\left(\frac{\rho_0 \sigma_B}{(1-\rho_0)\,\sigma_A}\right).$$

For other forms of Ψ, e.g., $\Psi(x) = \Phi(dx)$ or $\Psi(x) = \Psi(x - d)$, one may choose d by fixing $\rho = \rho_0 > 0.5$ at $(\Delta = \delta, \mu_A = \mu_A^0)$. As an example, in Figure 4, we plot the limiting allocation propor-

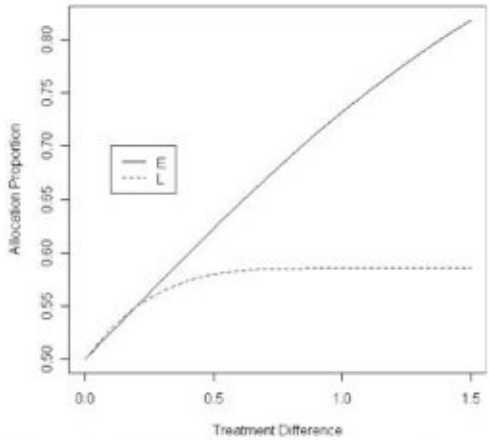

Figure 4 Limiting allocation proportion to treatment A for $\sigma_A = 1$, $\sigma_B = 1$. E: $\Psi(x) = \exp(dx)$, L: $\Psi(x) = \Phi(dx)$. Here $\rho_0 = 0.55$ at $\delta = 0.2$.

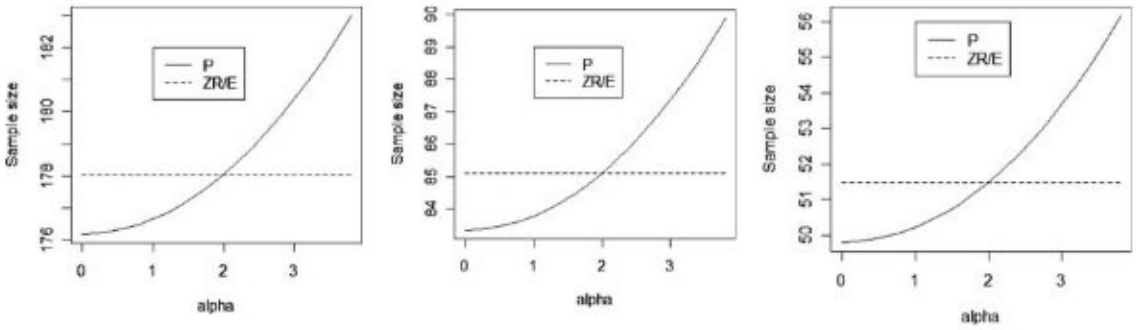

Figure 5 Size of the trial against α for the propose (P), ZR (or equal (E)) allocation rules. The three figures are for $\Delta = 0.4, 0.6, 0.8$ respectively.

tion for the two cases: (i) $\Psi(x) = \exp(dx)$, and (ii) $\Psi(x) = \Phi(dx)$, where $\rho_0 = 0.55$ at $\delta = 0.2$ and $\mu_A^0 = 0$. The performance of $\Psi(x) = \exp(dx)$ is much better than the other.

We consider a general two-parameter exponential family where X_k has the density

$$f(x_k) = \exp\left\{\frac{x_k\mu_k - b(\mu_k)}{a(\phi_k)} + c(x_k, \phi_k)\right\} \tag{15}$$

for some specific functions $a(\cdot)$, $b(\cdot)$ and $c(\cdot)$. For details about such a family of distributions, see, for example, Lindsey (1997) and McCullagh and Nelder (1989). The variance of the usual estimator of the risk difference $\Delta = b'(\mu_A) - b'(\mu_B)$ can be expressed as $\text{Var}(\hat{\Delta}) = \dfrac{b''(\mu_A)}{n_A}\, a(\phi_A) + \dfrac{b''(\mu_B)}{n_B}\, a(\phi_B)$. Therefore, finding the optimal allocation proportion is equivalent to minimize (13) subject to the restriction $\text{Var}(\hat{\Delta}) = K$. Solution to this problem yields the optimal ratio

$$\rho = \frac{n_A}{n} = \frac{\tau_A\sqrt{\Psi_B}}{\tau_A\sqrt{\Psi_B} + \tau_B\sqrt{\Psi_A}}, \tag{16}$$

where $\tau_k^2 = a(\phi_k)\, b(\mu_k)$. A convex linear combination of overall number of failures and total expected number of allocations to the inferior treatment can be used as the objective function. The weights in the convex combination should be determined keeping some criterion (say cost) in mind.

For asymptotic variance, we consider the most general form (16). Now suppose, after possible reparameterization, $\rho = \rho(\theta)$, where θ is a vector of the first two or all of $\{\mu_A, \mu_B, \sigma_A, \sigma_B\}$, dependent on which specific distribution in the exponential family. Let $\hat{\theta}$ is a consistent estimator of θ with covariance matrix Σ. We use the allocation function $g(x, y)$, defined on $[0,1] \times [0,1]$, recommended in the Hu and Zhang (2004) in the DBCD procedure

$$g(0,y) = 1,$$
$$g(1,y) = 0,$$
$$g(x,y) = \frac{y\left(\frac{y}{x}\right)^{\gamma}}{y\left(\frac{y}{x}\right)^{\gamma} + (1-y)\left(\frac{1-y}{1-x}\right)^{\gamma}}, \quad 0 < x < 1,$$

where $\nabla\rho(\theta)$ is the gradient of ρ. Here x is the current value of allocation proportion, and y is the target allocation proportion. See Hu and Zhang (2004) for details. Then we can obtain the asymptotic variance of the procedure as

$$\frac{\rho(1-\rho) + 2(1+\gamma)(\nabla\rho(\theta)'\,\Sigma\nabla\rho(\theta))}{1+2\gamma}.$$

6 Illustration with Real Trial

To illustrate the need of the adaptive procedures, we consider the real clinical trial conducted by Dworkin et al. (2003). This data set was also used by Zhang and Rosenberger (2006) for illustration. It was a randomized, placebo-controlled trial with an objective to evaluate the efficacy and safety of *pregabalin* in the treatment of *postherpetic neuralgia (PHN)*. There were $n = 173$ patients of which 84 received the standard therapy placebo and 89 were randomized to *pregabalin*. The primary efficacy measure was the mean of the last 7 daily pain ratings, as maintained by patients in a daily diary using the 11 point numerical pain rating scale (0 = no pain, 10 = worst possible pain) and, therefore, a lower score (response) indicates a favorable situation. After the 8 week duration of the trial, it was observed that *pregabalin*-treated patients experienced a higher decrease in pain score than patients treated with placebo. We use the final mean scores, i.e. 3.60 (with SD = 2.25) for *pregabalin* and 5.29 (with SD = 2.20) for *placebo* as the true ones for our purpose with an appropriate assumption regarding the distribution for pain scores.

The results in the following were obtained by simulations with 10,000 repetitions of a response adaptive trial for $n = 173$ patients with $N(3.60, 2.25^2)$ distribution for *pregabalin* and $N(5.29, 2.20^2)$ distribution for *placebo*. Allocation probabilities are updated according to the rule considered.

The BM procedure has an expected allocation probability to A of 0.5985 (SD = 0.0901). If we carry out the ZR procedure, the probability of getting negative response at at least once is 0.0154. On the other hand, if we carry out the TN rule (where we assume that the responses are truncated only in the domain [0, 10], the expected allocation is 0.550 (SD = 0.0882). Thus, the BM procedure is best as far as the allocation proportion is concerned. There is a small, but sizable probability of difficulty in implementation of the ZR rule in this data set at least once. One remedy may be to use $\rho = 1/2$ in case of a negative estimate of μ_A and μ_B.

It is clear that the BM rule has a slightly larger variance. This is also clear from our detailed simulation study, which is not reported for the sake of brevity. This variability should also be taken into account for a fair and legitimate comparison. This slightly higher variability is true in real trials, but the difference in variability with other designs is not so alarming, specially when almost all the adaptive trials are subject to high variability.

For the balanced random allocation, both the allocation proportion and failure proportion are 0.50 (with standard error 0.04 for both the cases). As one referee has suggested, we carried out a comparison of the BM procedure with the modified BM (MBM) procedure where the initial 15% patients are randomly assigned to either treatment followed by a BM procedure for the remaining 85% patients. We carried out a numerical study with $n = 80$, and the initial 15% (12 in number) are randomly allocated for the MBM procedure. A simulation study of 5000 repetitions yields the comparisons in Tables 1–3. It is clear that the BM procedure works better (worse) than the MBM procedure in terms of allocation proportion, failure proportion and power when σ_A is larger (smaller) than σ_B. The performances of these two designs are almost similar when $\sigma_A = \sigma_B$.

7 Conclusions

The main contribution of the present paper is to provide a general optimal design through the general objective function (13), which embraces almost all the existing designs and also provides a greater scope of generating a variety of new designs. This is further generalized to a broad family of distributions (15).

The development of this paper was motivated by the limitation of the ZR rule in case of negative responses. Optimal response-adaptive designs depend on the underlying distribution of responses. If there is reason to believe that the responses are normal or any other distribution which can take negative values, the ZR (2006) cannot be used in a straightforward way. Either one needs to consider some kind of modification in the objective function like the BM rule, or one needs to assume a truncated distribution, which is truncated in the appropriate positive domain.

Table 1 Allocation proportions to treatment A and standard errors (in parentheses) when $\mu_A = 1$.

	μ_B					
	1.0	1.2	1.4	1.6	1.8	2.0
Design	$(\sigma_A = 1, \sigma_B = 1)$					
MBM	0.500 (0.141)	0.526 (0.137)	0.554 (0.133)	0.581 (0.131)	0.605 (0.126)	0.627 (0.123)
BM	0.500 (0.134)	0.527 (0.131)	0.552 (0.136)	0.576 (0.134)	0.604 (0.134)	0.623 (0.128)
	$(\sigma_A = 2, \sigma_B = 1)$					
MBM	0.644 (0.072)	0.656 (0.071)	0.669 (0.069)	0.680 (0.071)	0.689 (0.066)	0.699 (0.067)
BM	0.671 (0.124)	0.685 (0.118)	0.700 (0.117)	0.714 (0.118)	0.729 (0.116)	0.742 (0.111)
	$(\sigma_A = 1, \sigma_B = 2)$					
MBM	0.357 (0.072)	0.369 (0.072)	0.383 (0.073)	0.398 (0.074)	0.413 (0.075)	0.432 (0.067)
BM	0.331 (0.125)	0.347 (0.127)	0.371 (0.125)	0.390 (0.124)	0.412 (0.124)	0.435 (0.125)

Table 2 Overall failure proportions to treatment A and standard errors (in parentheses) when $\mu_A = 1$.

	μ_B					
	1.0	1.2	1.4	1.6	1.8	2.0
Design	$(\sigma_A = 1, \sigma_B = 1)$					
MBM	0.500 (0.055)	0.499 (0.056)	0.492 (0.058)	0.479 (0.063)	0.466 (0.065)	0.452 (0.069)
BM	0.500 (0.056)	0.499 (0.056)	0.491 (0.059)	0.482 (0.064)	0.467 (0.067)	0.452 (0.072)
	$(\sigma_A = 2, \sigma_B = 1)$					
MBM	0.501 (0.055)	0.502 (0.055)	0.498 (0.055)	0.498 (0.056)	0.493 (0.057)	0.487 (0.057)
BM	0.500 (0.055)	0.498 (0.055)	0.494 (0.057)	0.491 (0.058)	0.485 (0.061)	0.475 (0.062)
	$(\sigma_A = 1, \sigma_B = 2)$					
MBM	0.501 (0.056)	0.498 (0.056)	0.491 (0.057)	0.487 (0.058)	0.481 (0.057)	0.474 (0.057)
BM	0.497 (0.054)	0.498 (0.057)	0.491 (0.056)	0.489 (0.059)	0.482 (0.061)	0.472 (0.064)

Table 3 Power of a test for treatment equivalence.

	μ_B					
	1.0	1.2	1.4	1.6	1.8	2.0
Design	$(\sigma_A = 1, \sigma_B = 1)$					
MBM	0.050	0.208	0.501	0.805	0.946	0.991
BM	0.050	0.194	0.526	0.813	0.956	0.990
	$(\sigma_A = 2, \sigma_B = 1)$					
MBM	0.050	0.136	0.289	0.526	0.753	0.897
BM	0.050	0.107	0.235	0.438	0.645	0.842
	$(\sigma_A = 1, \sigma_B = 2)$					
MBM	0.050	0.121	0.268	0.502	0.709	0.881
BM	0.050	0.127	0.258	0.486	0.720	0.884

For the Tamuara et al. (1994) data set, the responses are the changes in some score, which can be either positive or negative. In such a situation, the truncated distribution is not a good assumption. One should simply use the BM-type rule, where the Ψ's in the general type objective function (13) are always positive.

A careful investigation on which objective function is the most appropriate for specific distributions in the exponential family may be desitable. Also, in principle, our present methodology can be extended to more than two treatments at hand. But the choice of the variance constraint is not quite immediate in such a case. A careful choice of such variance constrant (one or more) is the first step for such problem. An iterative solution of the allocation probabilities can then be obtained. We skip the details for some future study.

Acknowledgement The authors wish to thank the Associate Editor and two anonymous referees for their valuable comments which led to some improvement over an earlier version of the paper.

References

Atkinson, A. C. and Biswas, A. (2005). Bayesian adaptive biased coin design for sequential clinical trials. *Biometrics* **61**, 118–125.

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

Bandyopadhyay, U. and Biswas, A. (2004). An adaptive allocation for continuous response using Wilcoxon-Mann-Whitney score. *Journal of Statistical Planning and Inference* **123**, 207–224.

Bartlett, R. H., Roloff, D. W., Cornell, R. G., Andrews, A. F., Dillon, P. W., and Zwischenberger, J. B. (1985). Extracorporeal circulation in neonatal respiratory failure: a prospective randomized trial. *Pediatrics* **76**, 479–487.

Biswas, A. and Dewanji, A. (2004). Inference for a RPW-type clinical trial with repeated monitoring for the treatment of rheumatoid arthritis. *Biometrical Journal* **46**, 769–779.

Biswas, A. and Mandal, S. (2004). Optimal adaptive designs in phase III clinical trials for continuous responses with covariates. In A. Di. Bucchianico, H. Läuter and H. P. Wynn (Eds.), *mODa 7 – Advances in Model-Oriented Design and Analysis*, pp. 51–58. Heidelberg: Physica-Verlag.

Dworkin, R. H., Corbin, A. E., Young, J. P., Sharma, U., LaMoreaux, L., Bockbrader, H., Garofalo, E. A., and Poole, R. M. (2003). Pregabalin for the treatment of posther-petic neuralgia. A randomized, placebo-controlled trial. *Neurology* **60**, 1274–1283.

Hayre, L. S. (1979). Two-population sequential tests with three hypotheses. *Biometrika* **66**, 465–474.

Hu, F. and Zhang, L. X. (2004). Asymptotic properties of doubly adaptive biased coin design for multi-treatment clinical trials. *Annals of Statistics* **32**, 268–301.

Ivanova, A. (2003). A play-the-winner type urn model with reduced variability. *Metrika* **58**, 1–13.

Jennison, C. and Turnbull, B. W. (2000). *Group Sequential Methods with Applications to Clinical Trials*. Boca Raton, Florida: Chapman and Hall/CRC.

Johnson, N. L. and Kotz, S. (1970). *Continuous Univariate Distributions – 1*. John Wiley & Sons, New York.

Lindsey, J. K. (1997). *Applying Generalized Linear Models*. Springer, New York.

McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*. 2nd ed., Chapman and Hall, London.

Rosenberger, W. F. (1993). Asymptotic inference with response-adaptive treatment allocation designs. *Annals of Statistics* **21**, 2098–2107.

Rosenberger, W. F., Stallard, N., Ivanova, A., Harper, C., and Ricks, M. (2001). Optimal adaptive designs for binary response trials. *Biometrics* **57**, 173–177.

Rout, C. C., Rocke, D. A., Levin, J., Gouw's, E., and Reddy, D. (1993). A reevaluation of the role of crystalloid preload in the prevention of hypotension associated with spinal anesthesia for elective cesarean section. *Anesthesiology* **79**, 262–269.

Tamura, R. N., Faries, D. E., Andersen, J. S., and Heiligenstein, J. H. (1994). A case study of an adaptive clinical trial in the treatment of out-patients with depressive disorder. *Journal of the American Statistical Association* **89**, 768–776.

Wei, L. J. (1979). The generalized Pòlya's urn design for sequential medical trials. *Annals of Statistics* **7**, 291–296.

Wei, L. J. and Durham, S. (1978). The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association* **73**, 840–843.

Zelen, M. (1969). Play the winner rule and the controlled clinical trial. *Journal of the American Statistical Association* **64**, 131–146.

Zhang, L. and Rosenberger, W. F. (2006). Response-adaptive randomization for clinical trials with continuous outcomes. *Biometrics* **62**, 562–569.