



HAL
open science

Comparing performances of several response-adaptive designs in dose finding studies

Bruno Lecoutre, Derzko Gérard

► **To cite this version:**

Bruno Lecoutre, Derzko Gérard. Comparing performances of several response-adaptive designs in dose finding studies. 2011. hal-00911294

HAL Id: hal-00911294

<https://hal.science/hal-00911294>

Preprint submitted on 29 Nov 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparing performances of several response-adaptive designs in dose finding studies

Bruno Lecoutre¹ and Gérard Derzko²

¹ ERIS, Laboratoire de Mathématique Raphaël Salem, CNRS and Université de Rouen, Avenue de l'Université, BP 12, 76801 Saint-Etienne-du-Rouvray, France. E-mail: bruno.lecoutre@univ-rouen

² Sanofi-Aventis R&D (2011)

May 6, 2011

Abstract

The main objectives of dose-finding studies are (i) to select the best dose(s) for future phase 3 and/or (ii) to evaluate the dose-response relationship. At that phase 2B stage of the drug development, little is known about the drug's efficacy, and only a small- to moderate-size sample can be experimented. Response-Adaptive (RA) designs, where the dose allocation to new patients depends on the information got from previous allocations and responses, do match dose-finding study conditions well. Therefore one can wonder whether they are valuable alternative designs to Balanced Randomization (BR) designs to meet the above-mentioned dose-finding objectives. Here the operating characteristics and the capability of selecting the best dose(s) and of detecting pre-specified dose contrasts are compared in four multiple dose RA designs namely two Play-the-Winner type, Generalized Drop-the-Loser and Doubly-adaptive Biased Coin designs. The restriction of these designs to two treatments have been shown to be the best choices indeed, as well as the Bayesian procedure used here for inference. The critical concept of *target allocation rule* is elaborated, and made explicit to match the specific dose-finding objectives. It is shown that the studied RA designs are efficient competitors to BR designs, but none of them outperforms the others on all purposes. Directions for RA design selection are given to experimenters, depending on the primary objective of their planned dose-finding study.

KEY WORDS: Bayesian inference; dose selection; power; response-adaptive design; treatment failures.

1 INTRODUCTION

Dose-finding clinical studies are phase 2 trials performed in patients with specific medical indication. They are a critical step in drug development, as they determine the design of the next large confirmatory phase 3 studies. Recently the design of dose finding studies have gained a renewed interest, as it was recognized that incorrect dose selection and poorly estimated dose-response relationship obtained from these studies may explain the high failure rate of phase 3 clinical studies. An overview and a comparative study of different designs can be found in Bornkamp et al. (2007) and Dragalin et al. (2010). In particular practical aspects optimal designs were considered in Bretz, Dette and Pinheiro (2010), and specifically studied in Chambaz and Van der Laan (2010), or adaptive allocation rule for grouped sequential trials in Morgan and Coad (2007). The techniques developed here may be classified (see Bornkamp et al., 2007), as another design-focused adaptive approach, based on empirical or optimal dose allocation rules.

Main objectives of dose finding studies are of (i) selecting the most effective dose(s) among a set of candidate doses; (ii) demonstrating dose-effect relationship. Objective (i) generally consists in selecting one or two doses for further investigations in the phase 3 trial. Objective (ii) may enhance in the demonstration of efficacy, if the dose-effect relationship is found increasing, or may detect some non monotonous dose-response profile (Gamma-shaped, bell-shaped, etc). These objectives require that a set of T doses indexed by $t = 1, 2, \dots, T$ be studied, four doses including possibly a placebo being a minimum. This article concentrates on binary responses (success or failure), with success probability φ_t , and number of patients n_t at dose t . Note that, when the term “treatment” is used in this article, it designates some dose of one drug; doses are naturally ordered and expected to induce responses in some order as well. The spacing of the doses is another peculiar feature compared to completely distinct treatments.

At phase 2B development stage, however, very little is known about the drug’s efficacy, and samples can only have small (a few tens patients) to moderate (a few hundreds patients) sizes. Therefore response-adaptive (RA) designs, where the dose allocated to the next patient depends on the already performed allocations and observed responses according to some prespecified rule, seem especially well-fitted. In these designs indeed the treatment allocation rule is devised so that the more effective doses are allocated more often in the trial, and the distribution of the total sample size N to the different dose groups is expected to account as soon as possible for what is learnt from the recorded effectiveness of the doses. How RA designs do perform in dose-finding studies is thus a naturally arising question.

As RA designs have been early advocated for ethical reasons, but study power is always a concern, several authors have developed optimal RA designs, which minimize the expected number of treatment failures $\sum_{t=1}^T n_t(1 - \varphi_t)$ while controlling the study power.

A feature of paramount importance in RA designs is their so-called *target allocation rule*,

$$\rho : t \mapsto \rho_t(\varphi_1, \dots, \varphi_T) \quad [\text{in short } \rho_t] \quad \sum_{t=1}^T \rho_t = 1$$

which depends on the probabilities of success $\{\varphi_t\}$, and whose values are the asymptotic allocation proportions in each dose group t . Since the $\{\varphi_t\}$ are unknown and must be estimated again each time a treatment is to be allocated to a new patient, $\{\rho_t\}$ are only “ideal” (*asymptotic*) allocation proportions.

In some RA designs, the target allocation rule is univocally determined by the allocation rule of this design, and it is said to be *standard*, while in other RA designs, the experimenter can chose the *optimal* target allocation rule, so that it matches the trial objectives. In any case, given any finite sample, the closeness of the expected allocation proportions to their asymptotic values (*accuracy*) and the sample variance of allocation proportions (*precision*) are critical features, which characterize the *convergence* of the RA design.

The first RA designs were proposed for two-group trials, hence they are not suitable for dose-finding studies. These designs have been generalized for more than two treatments, and new designs have been developed. The RA designs can be classified according to the type of their target allocation rule, either standard or optimal.

Standard target

- Play-the-Winner [PW], introduced for 2 treatments by Zelen (1969) and extended to T treatments in Lecoutre and ElQasyr (2008);
- Randomized Play-the-Winner [RPW] (Wei and Durham, 1978; Wei, 1979) and its Friedman’s urn extensions (Bai, Hu and Shen, 2002, Bai, Hu and Rosenberger, 2002);
- Drop-the-Loser [DL] (Ivanova, 2003).

Optimal target

- Doubly-adaptive Biased-Coin [DBC] (Eisele, 1994);
- Generalized Drop-the-Loser [GDL] (Sun, Cheung and Zhang, 2007).

For two treatments, Rosenberger et al. (2001) identified the optimal target allocation rule, which, for a fixed asymptotic variance (reflecting the power of the test), minimizes the expected number of failures. When the rate difference $\varphi_1 - \varphi_2$ is chosen as the effect measure, the solution is

$$\rho_t(\varphi_1, \varphi_2) = \frac{\sqrt{\varphi_t}}{\sqrt{\varphi_1} + \sqrt{\varphi_2}} \quad (1)$$

to be contrasted with the *Neyman allocation*

$$\rho_t^{Neyman}(\varphi_1, \varphi_2) = \frac{\sqrt{\varphi_t(1 - \varphi_t)}}{\sqrt{\varphi_1(1 - \varphi_1) + \varphi_2(1 - \varphi_2)}} \quad (2)$$

which would be the target allocation rule for minimizing the variance, if no condition on treatment failures were imposed.

They also derived the appropriate target allocation rules for other effect measures, which can be used as alternatives to the rate difference; yet they neither study these functions further, nor extend them to more than two treatments. Sun, Cheung and Zhang (2007) considered the “natural extension” of (??) for more than two treatments but did not justified that formula. Tymofyeyev, Rosenberger and Hu (2007) derived two equivalent optimal allocation proportions strategies based on the non-centrality parameter from a multivariate test, and described a targeted RA procedure, which provides some benefit in terms of power over balanced randomization (BR) designs. However, their proposals are difficult to implement in practice. Furthermore, they are not well adapted to the objectives of dose-finding studies. Basak, Biswas and Volkov (2009) also considered a natural (but not justified) extension of the standard target allocation rule for odds-ratio-based urn models to more than two treatments, and compared its performance with DL and RPW with respect to a list of criteria that they defined.

Frequentist performances, namely (i) coverage probabilities of interval estimation procedures, (ii) power, and (iii) overall failure count, of PW, RPW, DL, GDL and DBC designs were thoroughly compared in the case of two treatments in Lecoutre, Derzko and ElQasyr (2010); they used a Bayesian approach to

inference. They showed or confirmed that (i) the Bayesian approach has excellent frequentist properties, (ii) PW and DBC designs are particularly efficient, and in contrast RPW performs very poorly, (iii) the different target allocation rules have their specific interest, and (iv) the overall failure count is only modestly smaller in RA designs than in BR designs.

In the setting of dose-finding studies, where more than two doses are involved, it appears inadequate to optimize the power of detecting an overall effect, as in Tymofyeyev, Rosenberger and Hu (2007), and it is more in the line of phase 2 dose-finding studies' objectives to optimize: (i) the selection of the best dose, or the two best doses; (ii) the power of detecting a linear contrast of doses' success rates; (iii) the power of detecting a quadratic contrast of doses' success rates. In (i) only standard targets will be studied here. (ii) and (iii) can be obtained on constraint that the overall failure rate is fixed. This should be obtained by selecting the appropriate optimal target allocation rule. Yet it is still interesting and useful to include designs associated with standard targets for comparisons.

In this paper, we project to compare how well several RA designs perform and reach the above-mentioned objectives, and we concentrate on two kind of designs. (i) We will consider two standard target RA designs, which extend the PW design to more than two treatments, one developed in Lecoutre and ElQasyr (2008), and the other a new one; we exclude RPW and DL, as they performed more poorly. (ii) We will consider the GDL and the DBC designs, which can be used both with standard and optimal targets and showed excellent performances when restricted to two treatments (Lecoutre, Derzko and ElQasyr, 2010; Rosenberger and Hu, 2004). As in Lecoutre, Derzko and ElQasyr (2010) we will use non-informative Bayesian inference procedures for comparing those designs, as their (frequentist) operating characteristics were showed to compete advantageously with the frequentist inference procedures themselves in the case of two treatments, not to mention their specific asset in terms of coherence and interpretability.

The paper is organized as follows: in Section 2, the selected RA designs for more than two treatments are either defined or briefly reminded. The expressions of the asymptotic treatment allocation proportions targeted in the studied standard and optimal RA designs are obtained in Section 3; dose schemes for illustrations are also presented in that section. The Bayesian inference technique is reminded and coverage properties obtained in section 4. In section 5, capabilities of RA designs to select the best dose(s) are illustrated. Section 6 focuses on failure counts and power of detecting linear and quadratic contrasts. We conclude that none of the studied RA design is optimal for all the main objectives of a dose-finding study, but their performance depends on the particular study objective considered. A technical appendix is devoted on building the target allocation rules.

2 MULTI-TREATMENTS RESPONSE-ADAPTIVE DESIGNS

2.1 Play-the-Winner type designs

2.1.1 Basic rule for two treatments

In case of success to treatment t the next patient is assigned t ("repeat after success") and in case of failure the next patient is assigned the other treatment ("change after failure"). This ensures two characteristic properties. (i) The allocations are made according to *repeated blocks* of treatments. A typical sequence for two doses 1 and 2 is

block 1 $1^+1^+1^- \mid 2^+2^-$

block 2 $1^- \mid 2^+2^+2^+2^-$

block 3 $1^+1^+1^- \mid 2 \dots$

where $+$ indicates a treatment success and $-$ a failure.

(ii) Sample numbers of failures in the two treatments are equal or differ by only one unit.

2.1.2 Direct extension of Play-the-Winner type designs: Repeated blocks [PWext]

This direct extension, which had never been studied yet, preserves the above characteristic properties. In case of success to treatment t the next patient is assigned t . In case of failure the next patient is assigned one of the other $T - 1$ treatments in such a way that the allocations are made according to *repeated blocks* of the T treatments. Consequently, the sampling numbers of failures for each treatment differ by at most one unit.

The simplest possibility (which will not be considered here) is to always use the same order within blocks. A typical sequence for three doses 1, 2, 3 is

block 1 $1^{+1+1^{-}} | 2^{+2^{-}} | 3^{-}$

block 2 $1^{-} | 2^{+2+2^{-}} | 3^{+3^{-}}$

block 3 $1^{+1+1^{-}} | 2^{-} | 3 \dots$

We will prefer the following alternative: in case of failure the next patient is assigned the one of the other $T - 1$ treatments that has the higher estimated rate of success. It is indeed a better rule with respect to the objective of allocating more often the more effective doses. A typical sequence is

block 1 $1^{+1+1^{-}} | 2^{+2^{-}} | 3^{-}$ [ordered success rates after block 1: $\hat{\varphi}_1 = \frac{2}{3} > \hat{\varphi}_2 = \frac{1}{2} > \hat{\varphi}_3 = \frac{0}{1}$]

block 2 $1^{-} | 2^{+2+2^{-}} | 3^{+3^{-}}$ [ordered success rates after block 2: $\hat{\varphi}_2 = \frac{3}{5} > \hat{\varphi}_1 = \frac{2}{4} > \hat{\varphi}_3 = \frac{1}{3}$]

block 3 $2^{+2^{-}} | 1^{-} | 3 \dots$

2.1.3 Alternative generalization of Play-the-Winner type designs: Randomized after failure [PWraf]

This alternative, studied by Lecoutre and ElQasyr (2008), preserves the basic PW rule, but not its characteristic properties. In case of success to treatment t the next patient is assigned t . In case of failure the next patient is randomly assigned one of the other $T - 1$ treatments, either equiprobably or proportionally to the observed rates of success (other possibilities could be considered), which implies different standard target allocation rules. It must be noted that PWraf with equiprobable assignment has the same standard target as PWext. As demonstrated by Lecoutre and ElQasyr (2008), the convergence speed of PWraf towards the target allocation rule is excellent, both with equiprobable or proportional assignment.

2.1.4 Randomized Play-the-Winner [RPW] and Friedman’s urn extensions

The (completely) randomized Play-the-Winner rule, introduced by Wei and Durham (1978), is generally presented as a Generalized Friedman’s Urn (GFU) model (also called as generalized Pólya urn model (Freedman, 1965)). A typical GFU model can be described as follows. Before the j -th patient comes in, the urn contains balls (or “particles”, since the numbers of balls can be non integers) associated with each treatment t . A ball is drawn at random and replaced. The subject is assigned to the corresponding treatment (say t). When the response is known, balls are added to the urn. In the original RPW rule, one ball is added after success and no ball is added after failure. This does not preserve the basic PW rule, repeat after success and change after failure.

Bai, Hu and Shen (2002) developed GFU extensions for T treatments that generalize in a straightforward way the initial RPW rule. However, by contrast with PWraf, the corresponding GFU extensions of RPW perform poorly. Their convergence speed is always slower and can even be very bad (Lecoutre and ElQasyr, 2008).

2.2 Drop-the-Loser type designs

The Drop-the-Loser [DL] rule, proposed by Ivanova (2003), is also presented as a GFU model. DL has also relatively poor performance compared to PWraf (Lecoutre and ElQasyr, 2008).

Zhang et al. (2007) and Sun, Cheung and Zhang (2007) developed a very general model they called the Generalized Drop-the-Loser [GDL] rule. This rule involves an urn model with additional “immigration” balls.

For practical purposes, these authors recommended the following particular rule (which includes DL). Before the j -th patient comes in, if a treatment ball is drawn, that treatment is assigned, but the ball is never replaced, regardless of the response. If an immigration ball is drawn, no treatment is assigned and the ball is returned into the urn, together with u_t^j balls for each treatment. The total of added balls is fixed C .

For the simple DL rule, $u_t^j = 1$, which implies a standard target. In GDL, the number of added balls u_t^j are chosen for adjusting the desired target allocation rule. A fundamental property is that if $u_t^j = C\rho_t$ the sample allocation proportions converge to the target allocation rule. In practice, for each j $\{\rho_t\}$ are estimated from the estimated probabilities $\{\hat{\varphi}_t\}$, so that u_t^j depends on j . The immigration balls play two roles. They allow to add treatment balls to the urn according to the current estimates of the success rates, and in return they eliminate the possibility that certain type of treatment balls will become extinct. The initial proportion of immigration balls do not yield much difference in terms of the allocation proportions (Sun, Cheung and Zhang, 2007).

2.3 Doubly-adaptive Biased coin Design

Eisele (1994) introduced and studied the asymptotics of the Doubly-adapted Biased Coin Design (DBCD), originally indicated when responses are independent random variables from standard exponential family. This design was considered in the case of two treatments with immediate responses by Rosenberger and Hu (2004) (see also Hu and Zhang, 2004).

The probability z_t^j of assigning treatment t to the j -th patient is a function of both the desired target allocation rule and the proportion P_t^j of patients who have received dose t before the j -th patient. The most frequently proposed allocation function is

$$j, t \mapsto z_t^j = \frac{\rho_t (\rho_t / P_t^j)^\gamma}{\sum_{\ell=1}^T \rho_\ell (\rho_\ell / P_\ell^j)^\gamma} \quad \sum_{t=1}^T z_t^j = 1$$

where γ is a non-negative integer, which is shown to have a limited influence on the performance of the allocation process (Rosenberger and Hu, 2004).

As GDL, under general conditions for the allocation function, the sample allocation proportions converge to the target allocation rule. For each j , $\{\rho_t\}$ are estimated from the estimated probabilities $\{\hat{\varphi}_t\}$.

3 TARGET ALLOCATION RULES

3.1 Standard targets

PWext target allocation rule is

$$\rho_t = \frac{1}{\sum_{s=1}^T \frac{1}{1-\varphi_s}} \quad t = 1, \dots, T \quad (\text{S1})$$

Indeed within each block the rank of the first failure in t has geometric distribution $G(1-\varphi_t)$, with expectation $1/(1-\varphi_t)$. This design can be compared with GDL and DBC, as those two may accommodate any target allocation rule. PWraf (with equiprobable assignment after failure), a version of RPW and DL designs have that same target allocation rule, but, in our experience, RPW and DL

performed more poorly than PWraf, and preliminary results indicate that PWraf was outperformed by PWext, therefore these three designs will not be included in comparisons.

PWraf, with proportional assignment after failure, target allocation rule is

$$\rho_t = \frac{\frac{\varphi_t}{1-\varphi_t}}{\sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}} \quad t = 1, \dots, T \quad (\text{S2})$$

Indeed the expectation $1/(1 - \varphi_t)$ must be weighted by φ_t , because the allocation rule in this design specified proportionality to success rates. This design can be compared with GDL and DBC, as those two may accommodate any target allocation rule. A version of RPW has also the same target allocation rule, but performs more poorly and will not be included in comparisons.

3.2 Optimal targets

Optimal targets aim to control the power of detecting some specified effect, while minimizing the expected number of failures. If controlling the power of the *test of homogeneity of the success rates* is the objective, it can be achieved by minimizing the expected number of failures under the condition of a fixed asymptotic variance for the number of treatment failures, and the target allocation rule can be shown to be

$$\rho_t = \frac{\sqrt{\varphi_t}}{\sum_{s=1}^T \sqrt{\varphi_s}} \quad t = 1, \dots, T \quad (\text{O1})$$

If controlling the power of the *test of a specific contrast of the success rates* is the objective, it can be achieved by minimizing the expected number of failures under the condition of a fixed asymptotic variance for the contrast $\sum_{t=1}^T c_t \varphi_t$ ($\sum_{t=1}^T c_t = 0$), and the target allocation rule can be shown to be

$$\rho_t = \frac{|c_t| \sqrt{\varphi_t}}{\sum_{s=1}^T |c_s| \sqrt{\varphi_s}} \quad t = 1, \dots, T \quad (\text{O2})$$

Proofs are based on Lagrange multiplier technique (see Appendix). Of note, the optimal target also depends on the *effect measure* of interest; optimal target for contrasts of success rates' odds, logarithms or odds logarithms are demonstrated in the appendix. As usual, optimal targets are optimal only if the underlying condition of interest is present; if not, the power may decrease considerably. GDL and DBC designs can be compared when O1 or O2 are the selected target allocation rule. But there is also some interest in assessing PWext and PWraf respectively associated with their standard target allocation rules S1 and S2; indeed this allows appreciating the gain obtained in using optimal target, compared to standard ones.

3.3 Dose schemes for illustration

In the next sections, comparisons between PWext, PWraf, GDL and DBC designs will be illustrated with 4-dose schemes. As the proposed approach to evaluate the dose-response relationship is based on contrasts, all the results will be illustrated with 3 typical scenarios involving linear and quadratic contrasts of equi-spaced doses equal to 1, 2, 3 and 4 times some unit dose u . In this situation, the coefficients $\{c_t^L\}$ of the linear contrast are $(-.3, -.1, +.1, +.3)$, which are such that $\delta^L = \sum_t c_t^L \varphi_t$ is the slope of the regression line. The coefficients $\{c_t^Q\}$ of the quadratic contrast are $(+.25, -.25, -.25, +.25)$, which are such that $\delta^Q = \sum_t c_t^Q \varphi_t$ is the second-degree coefficient of the adjusted quadratic polynomial. Note that these coefficients would be the same for a situation involving a placebo and three doses equal to 1, 2 and 3 times the unit dose. Of course, situations involving unequally spaced doses could easily be dealt with by using the appropriate coefficients.

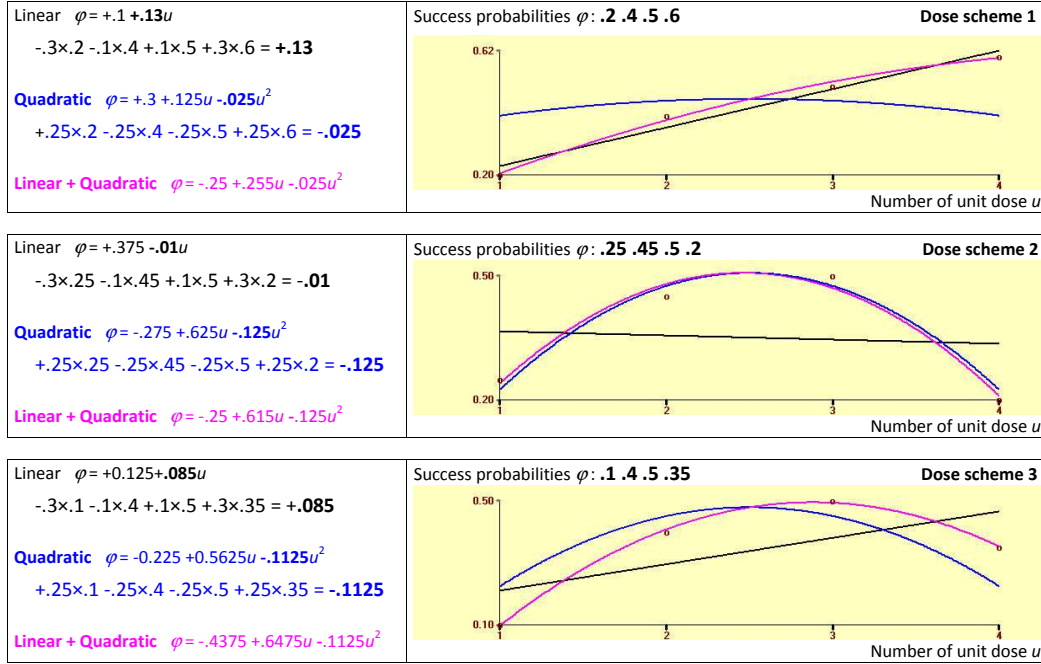


Figure 1: Three doses schemes involving equi-spaced doses equal to 1, 2, 3 and 4 times the unit dose u . The coefficients of the linear and quadratic contrasts are respectively $(-.3, -.1, +.1, +.3)$ and $(+.25, -.25, -.25, +.25)$.

The three scenarios are illustrated in Figure 1. Scheme 1 [success probabilities .2 .4 .5 .6] shows a notable linear but no quadratic effect, scheme 2 [.25 .45 .5 .2] shows a notable quadratic but no linear effect, and scheme 3 [.1 .4 .5 .35] shows both linear and quadratic notable effects. Seven other monotonic dose schemes will be used as illustrations where contrasts are not the concern.

4 BAYESIAN INFERENCE

Let n_{t1} and n_{t0} be the respective numbers of successes and failures to treatment t . For all considered adaptive designs, the likelihood function is proportional to

$$\prod_{t=1}^{t=T} \varphi_t^{n_{t1}} (1 - \varphi_t)^{n_{t0}} \quad (3)$$

i.e. proportional to the likelihood function associated with the comparison of T independent Binomial (or Negative Binomial) proportions. A simple and usual Bayesian solution assumes T independent Beta priors for φ_t , $\text{Beta}(\nu_{t1}, \nu_{t0})$. This is a conjugate prior and the marginal posterior distributions are again independent Beta distributions

$$\text{Beta}(\nu_{t1} + n_{t1}, \nu_{t0} + n_{t0}) \quad (4)$$

The Bayesian approach allows obtaining the distributions of any derived parameter of interest from the joint posterior distribution (Lecoutre, Derzko and Grouin, 1995): in particular, in the particular case of two treatments, efficient inferences can be made for the ratio, the difference or the odds-ratio (Lecoutre and ElQasyr, 2008). The posterior distribution for any parameter of interest can be

Table 1: Simultaneous 95 percent Bayesian credibility region for φ_t : Probability of error for $N = 100$ (simulated from 10^5 replications).

Target				S1 $\frac{1}{1-\varphi_t} / \sum_{s=1}^T \frac{1}{1-\varphi_s}$			S2 $\frac{\varphi_t}{1-\varphi_t} / \sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}$			BR
φ_1	φ_2	φ_3	φ_4	PWext	GDL	DBC	PWraf	GDL	DBC	
.2	.4	.5	.6	.045	.052	.048	.048	.052	.048	.049
.25	.45	.5	.2	.054	.057	.053	.047	.042	.053	.054
.1	.4	.5	.35	.045	.045	.047	.051	.053	.047	.056
.1	.15	.2	.3	.043	.047	.046	.030	.036	.035	.058
.1	.2	.3	.4	.051	.050	.051	.043	.044	.041	.046
.2	.3	.4	.5	.057	.062	.055	.045	.049	.046	.051
.5	.6	.7	.8	.052	.051	.052	.061	.060	.066	.050
.4	.6	.8	.9	.049	.054	.050	.046	.054	.045	.045
.6	.7	.8	.9	.060	.054	.056	.057	.070	.060	.046
.7	.8	.85	.9	.061	.055	.058	.067	.065	.074	.056

PW: play-the-winner; GDL: generalized drop-the-loser
 DBC: doubly-adaptive biased coin; BR: balanced randomization

approximated by simulating a large sample from two independent Beta distributions, i.e. repeatedly drawing random φ_t from their marginal posterior Beta distributions and computing the associated value for the parameter of interest. The $\alpha/2$ and $1 - \alpha/2$ quantiles of the generated values give the (approximate) desired confidence limits.

In all rigor, the Jeffreys rule gives different priors for the different designs, since it is based on the Fisher information. So, it could be argued that an objective prior should depend on the design (e.g., de Cristofaro, 2004; Bunouf and Lecoutre, 2006; Sun and Berger, 2008). Most possible refinements of the inference would however need more complicated numerical tools for a small expected improvement in terms of efficiency. This is in agreement with the solution proposed by Berger (2004, pages 5-6) for a problem of medical diagnosis involving three proportions. While expressing his theoretical preference for the *reference prior* approach (Berger and Bernardo, 1992), involving a different prior for each particular parameter of interest, he uses the simple Jeffreys prior in practice.

4.1 Coverage properties

95 percent simultaneous Bayesian credible region for $\phi = (\varphi_1, \dots, \varphi_T)$ were simulated from 10^5 replications for the sample size $N = 100$. The probability of error – the region does not contain ϕ – is reported in Table ?? for the following designs: PWext and PWraf, two versions of DBC with the corresponding standard targets (S1 and S2), BR.

The Bayesian credibility regions have always reasonable coverage probabilities, in most cases close to the nominal level. Similar results have been obtained for all designs considered in this paper, including the case of optimal targets (with GDL and DBC designs). Interval estimates for specific contrasts between the $\{\varphi_t\}$ were also investigated and have also good coverage properties. This confirms the conclusions previously obtained for the Jeffreys prior in other situations (see Lecoutre, Derzko and ElQasyr, 2010; Berger, 2004; Lecoutre and Charron, 2000; Agresti and Min, 2005; Lecoutre, 2008).

5 COMPARISON OF DESIGNS AND ILLUSTRATIONS

5.1 Simulation study

The initial conditions and parameter values recommended by the authors were used (we found no clear differences with other settings). In all models, the first treatment was randomly selected. For PWext and DBC, the first four patients were randomly assigned to each of the four treatments. GDL was used with no initial treatment balls and one immigration ball. The parameter values were $C = 2$ for GDL and $\gamma = 2$ for DBC (PW-type designs have no parameter).

When needed, the $\{\varphi_t\}$ were estimated by the Bayesian posterior mean associated with the Jeffreys prior for a Binomial proportion.

5.2 Allocation proportions for standard targets

Table ?? gives, for the 3+7 doses schemes, the expectation and standard deviation of the proportion of subjects assigned to each treatment, as well as the asymptotics (target values) for the standard targets S1 and S2. Numbers in Table ?? were obtained from simulations with 10^5 replications. The sample size was fixed to $N = 100$.

With target S1 all compared designs have an excellent accuracy, i.e. expected sample allocation proportions nicely close to their target allocation rule values, although PWext is always better performer than GDL and DBC. The convergence speed of PWext, and in a less degree of DBC, is especially good.

For illustration, simulations for sample size $N = 25$ and success probabilities .5, .6, .7, .8 give the respective expected sample allocations:

[.161 .199 .262 .378] (PWext) [.211 .230 .258 .301] (GDL) [.178 .210 .261 .351] (DBC)

which are to be compared to the asymptotic values [.156 0.195 0.260 0.390]. The superiority of PWext also appears in terms of precision (see the standard deviations of sample allocation proportions).

Target S2 is more significant of an already noticed trade-off between precision and accuracy: GDL has better precision (standard deviation) than PWraf and DBC designs, but worsened accuracy (expected allocation proportion). Indeed, as commented in Lecoutre, Derzko and ElQasyr (2010) “the gain in variance reduction obtained with GDL is highly correlated with the loss in convergence rate”. In fact, this gain appears to be an artefact, due to the poor performance of the GDL rule in terms of allocation proportion. For illustration, simulations obtained by increasing the sample size to $N = 500$ for the probabilities of success .5, .6, .7, .8 give the respective expected sample allocations:

[.123 .178 .266 .434] (PWraf) [.132 .183 .265 .420] (GDL) [.124 .180 .266 .430] (DBC)

with standard-deviations

[.036 .042 .050 .058] (PWraf) [.034 .042 .051 .061] (GDL) [.032 .037 .044 .052] (DBC)

So, when the GDL rule allocation proportions become closer to the other rules (it is nevertheless always less efficient), it is not less variable, if not more.

5.3 Selecting doses

The performance of RA designs for dose selection is evaluated by the chance of finding a correct ordering of the allocation and success sample proportions and by the chance of observing more success for the best dose and for the two best doses.

Success rates φ_t are not necessarily increasing with t , but it can be easily verified that, for standard targets S1 and S2, the application $\varphi_t \mapsto \rho_t$ is monotonic, so that dose ordering can be consistently considered through either $\hat{\varphi}$ or $\hat{\rho}$ orderings. Of course, the sample estimates $\hat{\varphi}_t$ do not always have the

Table 2: Allocation proportions of treatments for standard target RA designs ($N = 100$). For each dose scheme, target allocation rule values and expected sample allocation proportions obtained with selected RA designs (simulated from 10^5 replications).

	Target	S1 $\frac{1}{1-\varphi_t} / \sum_{s=1}^T \frac{1}{1-\varphi_s}$			Target	S2 $\frac{\varphi_t}{1-\varphi_t} / \sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}$		
		PWext	GDL	DBC		PWraf	GDL	DBC
$\varphi_1 = .2$.169	.169(.024)	.185(.026)	.174(.029)	.092	.112(.053)	.134(.044)	.106(.051)
$\varphi_2 = .4$.225	.225(.037)	.231(.038)	.226(.042)	.212	.209(.082)	.218(.067)	.209(.080)
$\varphi_3 = .5$.267	.270(.045)	.267(.046)	.269(.046)	.293	.286(.092)	.281(.077)	.289(.088)
$\varphi_4 = .6$.337	.336(.053)	.318(.054)	.331(.058)	.403	.394(.098)	.367(.084)	.397(.092)
$\varphi_1 = .25$.208	.209(.027)	.217(.029)	.211(.033)	.162	.170(.070)	.185(.056)	.165(.071)
$\varphi_2 = .45$.284	.284(.039)	.277(.041)	.282(.045)	.329	.318(.089)	.306(.075)	.323(.086)
$\varphi_3 = .5$.312	.312(.043)	.300(.045)	.309(.049)	.381	.368(.091)	.346(.078)	.375(.086)
$\varphi_4 = .2$.195	.196(.024)	.206(.026)	.199(.030)	.127	.144(.062)	.163(.050)	.137(.062)
$\varphi_1 = .1$.176	.176(.019)	.190(.020)	.180(.025)	.064	.104(.044)	.126(.036)	.094(.042)
$\varphi_2 = .4$.264	.264(.036)	.262(.038)	.263(.042)	.293	.280(.087)	.277(.071)	.284(.086)
$\varphi_3 = .5$.317	.316(.043)	.303(.045)	.313(.048)	.393	.376(.091)	.354(.077)	.382(.086)
$\varphi_4 = .35$.244	.244(.033)	.245(.035)	.244(.039)	.249	.240(.083)	.243(.067)	.240(.083)
$\varphi_1 = .1$.224	.224(.017)	.229(.018)	.225(.025)	.142	.175(.064)	.190(.048)	.165(.067)
$\varphi_2 = .15$.237	.237(.020)	.239(.021)	.238(.027)	.208	.215(.074)	.222(.056)	.210(.079)
$\varphi_3 = .2$.252	.252(.022)	.251(.024)	.252(.029)	.270	.258(.080)	.257(.062)	.260(.085)
$\varphi_4 = .3$.288	.287(.027)	.280(.030)	.285(.034)	.379	.353(.082)	.331(.066)	.365(.083)
$\varphi_1 = .1$.204	.204(.018)	.212(.019)	.206(.025)	.100	.139(.055)	.159(.043)	.128(.055)
$\varphi_2 = .2$.229	.229(.023)	.233(.025)	.230(.030)	.200	.202(.074)	.213(.058)	.198(.077)
$\varphi_3 = .3$.262	.262(.029)	.260(.031)	.261(.035)	.300	.283(.083)	.277(.067)	.288(.085)
$\varphi_4 = .4$.305	.305(.034)	.295(.037)	.302(.041)	.400	.377(.084)	.352(.071)	.386(.082)
$\varphi_1 = .2$.197	.197(.024)	.208(.026)	.200(.030)	.128	.146(.063)	.164(.051)	.139(.063)
$\varphi_2 = .3$.225	.225(.030)	.230(.032)	.227(.035)	.202	.203(.078)	.213(.063)	.200(.078)
$\varphi_3 = .4$.263	.262(.036)	.260(.038)	.262(.041)	.285	.277(.087)	.273(.072)	.281(.086)
$\varphi_4 = .5$.315	.315(.043)	.302(.045)	.311(.048)	.385	.374(.092)	.350(.078)	.380(.087)
$\varphi_1 = .5$.156	.157(.042)	.179(.042)	.162(.045)	.125	.129(.071)	.153(.060)	.129(.066)
$\varphi_2 = .6$.195	.196(.053)	.211(.052)	.200(.056)	.179	.179(.089)	.197(.073)	.181(.082)
$\varphi_3 = .7$.260	.260(.068)	.260(.065)	.260(.072)	.265	.264(.112)	.264(.090)	.265(.100)
$\varphi_4 = .8$.390	.387(.086)	.351(.081)	.378(.090)	.430	.428(.131)	.386(.106)	.425(.116)
$\varphi_1 = .4$.087	.089(.031)	.126(.032)	.099(.034)	.054	.063(.043)	.094(.039)	.064(.039)
$\varphi_2 = .6$.130	.133(.050)	.167(.049)	.141(.053)	.111	.113(.073)	.147(.062)	.117(.067)
$\varphi_3 = .8$.261	.264(.094)	.272(.084)	.264(.097)	.267	.265(.136)	.279(.105)	.269(.119)
$\varphi_4 = .9$.522	.514(.117)	.434(.102)	.496(.121)	.569	.559(.158)	.480(.122)	.550(.137)
$\varphi_1 = .6$.120	.123(.048)	.158(.046)	.131(.050)	.098	.102(.069)	.136(.099)	.105(.062)
$\varphi_2 = .7$.160	.163(.064)	.191(.060)	.169(.066)	.146	.149(.094)	.178(.075)	.152(.083)
$\varphi_3 = .8$.240	.242(.091)	.252(.080)	.245(.094)	.240	.241(.133)	.253(.099)	.243(.116)
$\varphi_4 = .9$.480	.472(.121)	.399(.102)	.455(.124)	.515	.509(.167)	.433(.123)	.500(.143)
$\varphi_1 = .7$.133	.136(.059)	.170(.054)	.144(.061)	.118	.121(.085)	.153(.068)	.124(.074)
$\varphi_2 = .8$.200	.202(.086)	.221(.074)	.206(.087)	.194	.196(.124)	.215(.091)	.198(.106)
$\varphi_3 = .85$.267	.267(.107)	.266(.087)	.268(.108)	.269	.269(.152)	.269(.107)	.269(.129)
$\varphi_4 = .9$.400	.394(.129)	.343(.102)	.383(.130)	.419	.413(.180)	.363(.123)	.408(.152)

PW: play-the-winner; GDL: generalized drop-the-loser
 DBC: doubly-adaptive biased coin; BR: balanced randomization

true ordering, and the application $\hat{\varphi}_t \mapsto \hat{\rho}_t$ is not always monotonic. Hence looked-for performance measures are possibly the capability of RA design to generate as large as possible proportions of correctly ordered sample values $\hat{\varphi}_t$ and $\hat{\rho}_t$. As ties may occur in sample $\hat{\varphi}_t$ and/or $\hat{\rho}_t$, a correct order may be considered either in the strict sense (counted as correct orderings are only those with no tied doses), or in the extended sense (all correct orderings up to possible tied doses are considered as correct). Numbers are always larger with the latter approach. Similar issue occurs for best dose(s) selection: for example if the estimated success rates of dose 2 and 3 are maximal and identical, we can consider this selection of the best dose as incorrect, because of the tie (strict sense), or correct (extended sense). Since all considered dose schemes include only distinct doses, we will only report results for correct order in the strict sense. In consequence, special care should be taken when comparing RA designs versus BR, as the balance in the latter increases the probability of generating tied estimates.

For the dose schemes, target allocation rules and RA designs considered in Table ??, Table ?? gives the following statistics (simulated from 10^5 replications): the expected proportion of treatment failures (standard error); the proportions of sample with correct (strict) ordering of the allocation proportions and of success proportions; the proportion of samples with more success for the best dose and for the two best doses.

The total numbers of failures are quite similar, and minimal, in the three designs compared under target S2; in contrast with what has been observed with only two treatments (see Lecoutre, Derzko and ElQasyr, 2010; Rosenberger and Hu, 2004), the gain in the failure counts obtained with RA designs is notable compared with the one obtained with an equally sized balanced randomization design (BR). A similar observation is also valid, yet to a lesser degree, for the three designs compared under target S1. Of note in these two cases, GDL design has slightly poorer performance than PW-types or DBC designs. Accuracy in failure counts is quite similar in all six RA designs. The picture is quite different in terms of ranking correctly the allocated proportions or the success proportions: PWext outperforms all the other designs, including BR. Finally, the PWext shows a very significant advantage over all others for selecting the best dose, and a significant advantage over all the other RA designs for selecting the two best doses.

In order to refine the comparisons of correct rankings of allocations and successes, the sample sizes N required to get 80 per cent chance of selecting the best dose (i.e. such that 80 per cent of simulated samples have a strictly higher estimated success rate for the best dose) were estimated for each compared design.

For this sample size, Table ?? gives the following statistics (simulated from 10^5 replications): the expected number of failures; the expected number of sample with correct ordering of allocation proportions and of success proportions; the proportion of samples with more success for the two best doses. Results are only reported for the dose schemes involving the linear and quadratic contrasts.

The sample size required is generally minimal for DBC design (target S2), which also outperforms others in terms of failure counts (very notably so compared BR). Second performer for sample size and failure count is the PWext design, which outperforms other RA designs for obtaining the correct ordering of allocated doses and successes. BR remains the best performer for selecting the two best doses, and the six RA designs perform as well as each other under these conditions.

Table 3: Performance of RA designs for selecting doses ($N = 100$). For each dose scheme, line 1: expected proportion of treatment failures (standard error), line 2: proportion of samples with correct ordering of the allocation proportions / of the success proportions, line 3: proportion of samples with more success for the best dose / for the two best doses (simulated from 10^5 replications).

Target		S1 $\frac{1}{1-\varphi_t} / \sum_{s=1}^T \frac{1}{1-\varphi_s}$			S2 $\frac{\varphi_t}{1-\varphi_t} / \sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}$			
		PWext	GDL	DBC	PWraf	GDL	DBC	BR
.2 .4 .5 .6	failures	.539(.053)	.547(.052)	.541(.053)	.515(.055)	.525(.054)	.513(.054)	.575(.047)
	ordering	.427/.434	.302/.408	.336/.418	.326/.372	.319/.387	.367/.389	/.407
	+ success	.738/.695	.707/.675	.723/.679	.734/.660	.718/.666	.745/.681	.681/.677
.25 .45 .5 .2	failures	.625(.051)	.630(.050)	.626(.051)	.602(.054)	.610(.052)	.598(.054)	.650(.046)
	ordering	.290/.295	.214/.287	.230/.290	.251/.270	.245/.279	.274/.277	/.282
	+ success	.617/.859	.609/.845	.611/.842	.624/.812	.618/.823	.630/.828	.576/.858
.1 .4 .5 .35	failures	.633(.051)	.639(.050)	.635(.051)	.606(.053)	.615(.052)	.602(.053)	.663(.045)
	ordering	.372/.378	.271/.358	.283/.361	.305/.340	.305/.354	.330/.345	/.344
	+ success	.695/.560	.671/.550	.680/.549	.700/.548	.687/.551	.711/.558	.646/.537
.1 .15 .2 .3	failures	.805(.041)	.807(.040)	.806(.040)	.793(.044)	.797(.042)	.791(.044)	.813(.038)
	ordering	.192/.195	.125/.190	.133/.189	.196/.188	.186/.191	.209/.188	/.185
	+ success	.715/.508	.705/.505	.706/.505	.728/.504	.722/.507	.742/.514	.685/.504
.1 .2 .3 .4	failures	.733(.046)	.737(.046)	.734(.046)	.710(.050)	.718(.048)	.707(.050)	.750(.042)
	ordering	.353/.361	.234/.353	.236/.352	.296/.322	.292/.340	.324/.324	/.352
	+ success	.735/.704	.723/.697	.725/.697	.746/.682	.739/.689	.758/.691	.698/.700
.2 .3 .4 .5	failures	.630(.050)	.634(.050)	.632(.051)	.612(.053)	.619(.052)	.610(.053)	.650(.047)
	ordering	.317/.324	.217/.307	.233/.311	.257/.288	.249/.298	.288/.299	/.304
	+ success	.728/.671	.706/.658	.713/.658	.728/.638	.717/.645	.736/.655	.681/.662
.5 .6 .7 .8	failures	.312(.050)	.322(.049)	.315(.050)	.301(.052)	.311(.050)	.301(.051)	.350(.046)
	ordering	.333/.338	.237/.297	.284/.315	.252/.278	.242/.283	.289/.306	/.306
	+ success	.789/.674	.737/.634	.769/.651	.762/.614	.738/.618	.780/.644	.710/.660
.4 .6 .8 .9	failures	.211(.051)	.241(.048)	.218(.051)	.192(.051)	.219(.047)	.194(.048)	.325(.043)
	ordering	.578/.588	.470/.565	.513/.564	.417/.448	.454/.521	.491/.510	.638
	+ success	.885/.885	.828/.874	.863/.871	.851/.812	.830/.854	.874/.852	.789/.918
.6 .7 .8 .9	failures	.194(.046)	.211(.044)	.198(.046)	.184(.048)	.201(.044)	.186(.046)	.250(.042)
	ordering	.362/.368	.275/.326	.321/.346	.267/.285	.276/.314	.317/.331	/.349
	+ success	.861/.690	.795/.658	.839/.668	.819/.614	.796/.644	.848/.657	.766/.695
.7 .8 .85 .9	failures	.161(.040)	.169(.038)	.163(.040)	.157(.041)	.166(.039)	.158(.040)	.187(.038)
	ordering	.238/.242	.179/.207	.210/.222	.172/.182	.175/.196	.206/.215	/.183
	+ success	.659/.543	.598/.498	.640/.518	.620/.474	.598/.486	.642/.514	.540/.503

PW: play-the-winner; GDL: generalized drop-the-loser; DBC: doubly-adaptive biased coin; BR: balanced randomization

Table 4: Comparisons of the ordering of allocations and successes. N is selected in order to get 80 per cent chance of selecting the best dose. For each dose scheme, line 1: expected number of treatment failures, line 2: proportion of samples with correct ordering of the allocation proportions / of the success proportions, line 3: proportion of samples with more success for the two best doses (simulated from 10^5 replications).

Target	S1 $\frac{1}{1-\varphi_t} / \sum_{s=1}^T \frac{1}{1-\varphi_s}$			S2 $\frac{\varphi_t}{1-\varphi_t} / \sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}$			BR
	PWext	DLG	DBC	PWraf	DLG	DBC	
.2 .4 .5 .6	$N = 147$	$N = 175$	$N = 162$	$N = 151$	$N = 170$	$N = 142$	$N = 192$
Failures	79.3	95.2	87.6	77.3	88.4	72.5	110.4
Correct ordering	.550/.557	.455/.581	.475/.568	.436/.498	.452/.541	.467/.501	/.550
More success pairs	.765	.777	.767	.740	.757	.747	.801
.25 .45 .5 .2	$N = 505$	$N = 598$	$N = 549$	$N = 483$	$N = 551$	$N = 459$	$N = 656$
Failures	315.6	374.5	343.2	285.4	329.9	272.3	426.4
Correct ordering	.628/.629	.548/.643	.553/.632	.530/.574	.533/.596	.558/.578	/.673
More success pairs	.999	1	.999	.997	.998	.998	1
.1 .4 .5 .35	$N = 177$	$N = 203$	$N = 195$	$N = 177$	$N = 191$	$N = 165$	$N = 220$
Failures	113.3	129.3	123.7	106.3	116.1	98.6	145.8
Correct ordering	.505/.509	.425/.516	.440/.514	.431/.481	.430/.500	.448/.473	/.520
More success pairs	.641	.645	.644	.632	.634	.632	.653

PW: play-the-winner; GDL: generalized drop-the-loser
 DBC: doubly-adaptive biased coin; BR: balanced randomization

6 NUMBER OF TREATMENT FAILURES AND POWER FOR CONTRASTS

For each target, we consider the power for the two contrasts of interest, the linear contrast δ^L with coefficients $(-.3, -.1, +.1, +.3)$ and the quadratic contrast δ^Q with coefficients $(+.25, -.25, -.25, +.25)$. Suppose we aim to demonstrate that $\delta^L > 0$. With the Bayesian approach, it can be concluded that $\delta^L > 0$ when the posterior probability $\Pr(\delta^L > 0 \mid \text{data})$ is larger than a given $1 - \alpha$ (the lower limit of the $100(1 - 2\alpha)$ per cent credible interval for δ^L is larger than zero). The sampling probability of this event, assuming true values φ_t^* , evaluates the power of the procedure. The same computation applies to the quadratic contrast, the aim being in this case to demonstrate that $\delta^Q < 0$. Power arguments should be balanced against the performance in terms of allocation proportions and consequently in terms of expected total numbers of treatment failures.

Both the standard and optimal targets are considered here. Since the coefficients c_t^Q are equal in absolute value, the optimal target for the quadratic contrast $O2^Q$ happens in this particular case to be identical to the optimal target $O1$, which controls the power of the test of homogeneity of the success rates. For each target, results for the best designs are given in Table ???. The expected sample allocation proportions and the expected power (probability of the event above, averaged over samples) are reported. The expected standard deviation of the posterior distribution of each contrast is also reported; it evaluates the precision of inference.

For target $O2^L$, optimal for the linear contrast, the allocation proportions of the smallest and the largest doses are the largest, while those of intermediate doses are smaller: this is a reminder that $O2^L$ is optimal if the linear contrast is truly notable; but if this condition happens to be wrong, the target may be far from optimal. In contrast allocation proportions for standard target $S1$ or $S2$, and incidentally here for $O2^Q$ (quadratic contrast) are monotonous in φ_t , so that the information collected

Table 5: Power study for linear and quadratic contrasts ($N = 100$, 25 subjects for each dose in BR). For each dose scheme, lines 1-4: expected sample allocation proportions, line 5: expected number of treatment failures (standard error), lines 6 and 7: estimated power (standard error of the *posterior* distribution of δ), respectively for the linear and quadratic contrasts (simulated from 10^4 replications).

Target	S1	S2	O1 $\sqrt{\varphi_t}/\sum_{s=1}^T \sqrt{\varphi_s}$		O2 ^L $ c_t^L /\sqrt{\varphi_t}/\sum c_s^L /\sqrt{\varphi_s}$		BR
	PWext	DBC	GDL	DBC	GDL	DBC	
$\varphi_1^* = .2$.169	.106	.198	.174	.287	.267	.25
$\varphi_2^* = .4$.225	.209	.245	.246	.130	.128	.25
$\varphi_3^* = .5$.270	.289	.268	.277	.139	.142	.25
$\varphi_4^* = .6$.336	.397	.288	.304	.444	.462	.25
Failure	.539(.053)	.513(.054)	.555(.048)	.547(.048)	.555(.047)	.547(.048)	.575(.047)
Linear	.903(.039)	.885(.043)	.921(.039)	.921(.039)	.964(.035)	.959(.035)	.930(.039)
Quadratic	.124(.045)	.110(.048)	.141(.045)	.148(.045)	.125(.050)	.122(.050)	.139(.045)
$\varphi_1^* = .25$.209	.165	.224	.214	.356	.354	.25
$\varphi_2^* = .45$.284	.323	.277	.288	.155	.160	.25
$\varphi_3^* = .5$.312	.375	.290	.305	.160	.169	.25
$\varphi_4^* = .2$.196	.137	.209	.192	.329	.318	.25
Failure	.625(.051)	.598(.054)	.631(.047)	.626(.047)	.696(.044)	.691(.044)	.650(.046)
Linear	.032(.039)	.038(.043)	.035(.038)	.043(.038)	.033(.033)	.028(.034)	.029(.036)
Quadratic	.837(.044)	.839(.045)	.851(.043)	.856(.043)	.796(.047)	.820(.047)	.855(.044)
$\varphi_1^* = .1$.176	.094	.181	.152	.275	.234	.25
$\varphi_2^* = .4$.264	.284	.269	.278	.151	.157	.25
$\varphi_3^* = .5$.316	.382	.294	.311	.163	.175	.25
$\varphi_4^* = .35$.244	.240	.255	.259	.412	.435	.25
Failure	.633(.051)	.602(.053)	.638(.047)	.627(.048)	.686(.046)	.675(.045)	.662(.045)
Linear	.687(.037)	.544(.041)	.707(.036)	.693(.037)	.806(.032)	.784(.032)	.745(.035)
Quadratic	.785(.042)	.745(.044)	.803(.042)	.802(.042)	.738(.046)	.761(.046)	.803(.043)

PW: play-the-winner; GDL: generalized drop-the-loser; DBC: doubly-adaptive biased coin; BR: balanced randomization

S1: standard target $\frac{1}{1-\varphi_t}/\sum_{s=1}^T \frac{1}{1-\varphi_s}$; S2: standard target $\frac{\varphi_t}{1-\varphi_t}/\sum_{s=1}^T \frac{\varphi_s}{1-\varphi_s}$

O1: optimal target for the test of homogeneity

O2^L: optimal target for the linear contrast; O2^Q: optimal target for the quadratic contrast

is all the more important, since the dose is truly more effective; in consequence S1 and S2 can be expected to be robust against discrepancy to any particular dose-response relationship assumption. In addition the failure counts are notably smaller with the standard targets S1 (PWext) and even more so with S2 (DBC), versus the two optimal targets, or versus BR. As expected $O2^L$ gives the best power when the linear contrast is truly notable, and so does $O2^Q$ (or O1) when the quadratic contrast is truly notable; the power may even be larger than BR's one sometimes.

7 CONCLUDING REMARKS

In this article, we focused on RA design's operating characteristics, which match the particular objectives of phase 2B dose-finding studies well, namely dose selection and dose-response relationship. Restricting attention to success-failure responses, four promising multi-treatment RA designs were selected on the basis of the excellent operating characteristics of their restriction to two treatments; noticeably, based on extended experience, RPW and DL have been discarded as less performing. Then appropriate standard or optimal target allocation rules have been elicited; finally the general (coverage, convergence), and specific to dose-finding (ability for dose selection, power of detecting contrasts and associated failure count) properties of the selected RA designs have been compared. All the considered designs, PWext, PWraf, GDL and DBC designs had close-to-nominal coverage probabilities, using the well-tried Bayesian approach to confidence. It was found that in PW-type and DBC (with the same target allocation rule) designs, the sample allocation proportions converge excellently to their targets, especially the PWext design.

It was demonstrated that none of the studied RA design is optimal for all the main objectives of a dose-finding study, but their performance depends on the particular study objective considered: if the main objective is to select the best dose(s) for further phase 3 trial, the PWext outperforms the competitors, yet optimal designs minimize the numbers of failures in these conditions. If the objective is to infer linear or quadratic contrasts, optimal designs perform the best as expected. Yet, the required allocation proportions may not reflect the true success rates; this exposes to the risk of decreased power if the dose-effect relationship had been misspecified. In addition they perform worse than standard designs in terms of expected number of failures, and the relevant target allocation rule depends on the effect measure used for inference. Therefore a more robust option may consist in using the slightly less performing standard designs, whose allocations proportions always reflect the success rates correctly. Of note the conclusion that RA designs generally provide a modest advantage to BR in the case of two treatments (see Lecoutre, Derzko and ElQasyr, 2010; Rosenberger and Hu, 2004) is not confirmed with more than two treatments: sizeable savings of failures can even be obtained in some settings.

Comparison with the most usual BR should not be skipped: the main argument in favour of RA designs is that doses are to be allocated in conformity with some specific phase 2B objective, but one generally expects BR to have better power. In this article the discrepancy from balanced samples obtained with the diverse RA designs, and any possible worsening of power in comparison to BR can be evaluated. It was found that the best RA designs are at least equally powerful as BRs, while failure counts are smaller. When the main objective is to select the best dose(s), RA designs seem to be interesting competitors to BR group sequential designs, where involved techniques may be used to drop the worst dose(s) early.

The two extensions of PW rule to more than two treatments, PWext and PWraf, are of course subjected to potentially biasing the dose allocation process: if the trial takes place in a single experimental center indeed, any success with one dose is followed by the same allocated dose. In multi-center with centralized dose allocation, or when the evaluations are independent of the inclusion process, any selection bias becomes very unlikely. But those two designs are useful mainly as references for design comparisons. It was shown that DBC designs perform excellently, but cannot be suspected to selection bias: they can be used when complete protection of selection bias is required. The implementation

of RA designs is more involved than that of BR, as in particular a centralized interactive system is required for dose allocation and response recording, and the drug supply management is more difficult. It is not the place in this article to develop these important practical aspects, but a detailed analysis of them can be found in Quinlan et al. (2010).

GDL and DBC designs cope with delayed responses. PW designs can be adapted for this circumstance. Either some default random allocation, or preferably an allocation based on available responses can be used when the adaptive rule cannot be applied.

APPENDIX: DERIVATION OF THE TARGET ALLOCATION RULE

Let $\{n_t\}_{t=1}^T$ the sizes of T treatment groups and $N = \sum_{t=1}^T n_t$ the total sample size. Assume the overall failure count $\sum_{t=1}^T n_t(1 - \varphi_t)$ is to be minimized, while keeping constant ($= K$) the variance $\sum_{t=1}^T c_t^2 \frac{g(\varphi_t)}{n_t}$ of a contrast $\{c_t\}$, with $\sum_{t=1}^T c_t = 0$. Function $g(\cdot)$ depends only on the effect measure $f(\cdot)$ and on φ_t . The table below shows a few examples.

		$f(\varphi_t)$	$g(\varphi_t)$
Arithmetic	Rate	φ_t	$\varphi_t(1 - \varphi_t)$
	Odds	$\frac{\varphi_t}{1 - \varphi_t}$	$\frac{\varphi_t}{(1 - \varphi_t)^3}$
Geometric	Rate	$\log(\varphi_t)$	$\frac{1 - \varphi_t}{\varphi_t}$
	Odds	$\log\left(\frac{\varphi_t}{1 - \varphi_t}\right)$	$\frac{1}{\varphi_t(1 - \varphi_t)}$

Let λ be a Lagrange multiplier, the function

$$L(n_t, \lambda) = \sum_{t=1}^T n_t(1 - \varphi_t) + \lambda \left(\sum_{t=1}^T c_t^2 \frac{g(\varphi_t)}{n_t} - K \right)$$

is to be minimized in n_t 's. One gets

$$\begin{aligned} \frac{\partial L(n_t, \lambda)}{\partial n_t} &= (1 - \varphi_t) \left(1 - \frac{\lambda g(\varphi_t) c_t^2}{(1 - \varphi_t) n_t^2} \right) = 0 \\ \frac{\partial L(n_t, \lambda)}{\partial \lambda} &= \sum_{t=1}^T c_t^2 \frac{g(\varphi_t)}{n_t} - K = 0 \end{aligned}$$

and solutions

$$n_t = |c_t| \sqrt{\frac{\lambda g(\varphi_t)}{1 - \varphi_t}}$$

The asymptotic treatment allocation rates are

$$\begin{aligned} N &= \sum_{t=1}^T n_t = \sum_{t=1}^T |c_t| \sqrt{\frac{\lambda g(\varphi_t)}{1 - \varphi_t}} \\ \frac{n_t}{N} &= \frac{|c_t| \sqrt{\frac{g(\varphi_t)}{1 - \varphi_t}}}{\sum_{s=1}^T |c_s| \sqrt{\frac{g(\varphi_s)}{1 - \varphi_s}}} \end{aligned}$$

In particular

- Contrasts of rates

$$\frac{n_t}{N} = \frac{|c_t| \sqrt{\varphi_t}}{\sum_{s=1}^T |c_s| \sqrt{\varphi_s}}$$

- Contrasts of odds

$$\frac{n_t}{N} = \frac{|c_t| \frac{\sqrt{\varphi_t}}{(1-\varphi_t)^2}}{\sum_{s=1}^T |c_s| \frac{\sqrt{\varphi_s}}{(1-\varphi_s)^2}}$$

- Contrasts of rate logarithms

$$\frac{n_t}{N} = \frac{\frac{|c_t|}{\sqrt{\varphi_t}}}{\sum_{s=1}^T \frac{|c_s|}{\sqrt{\varphi_s}}}$$

- Contrasts of odds logarithms (log(odds))

$$\frac{n_t}{N} = \frac{\frac{|c_t|}{(1-\varphi_t)\sqrt{\varphi_t}}}{\sum_{s=1}^T \frac{|c_s|}{(1-\varphi_s)\sqrt{\varphi_s}}}$$

REFERENCES

- Agresti, A., Min, Y. (2005). Frequentist performance of Bayesian confidence intervals for comparing proportions in 2×2 contingency tables. *Biometrics* 61:515–523.
- Bai, Z. D., Hu, F., Rosenberger, W. F. (2002). Asymptotic properties of adaptive designs for clinical trials with delayed response. *The Annals of Statistics* 30:122–139.
- Bai, Z. D., Hu, F., Shen, L. (2002). An adaptive design for multi-arm clinical trials. *Journal of Multivariate Analysis* 81:1–18.
- Basak, G. P., Biswas, A., Volkov, F. (2009). An urn model for odds-ratio-based response-adaptive phase III clinical trials for two or more treatments *Journal of Biopharmaceutical Statistics* 19:838–856.
- Berger, J. (2004). The case for objective Bayesian analysis. *Bayesian Analysis* 1:1–17.
- Berger, J. O., Bernardo, J. M. (1992). On the development of the reference prior method (with discussion). In *Bayesian Statistics 4*, Bernardo JM, Berger JO, Dawid AP, Smith AFM (eds.). Oxford: Oxford University Press, 35–60.
- Bornkamp, B., Bretz, F., Dmitrienko, A., Enas, G., Gaydos, B., Hsu, C.-H., Koumnlign, F., Krams, M., Liu, Q., Neuenschwander, B., Parke, T., Pinheiro, J., Roy, A., Sax, R., Shen, F. (2007). Innovative approaches for designing and analyzing adaptive dose-ranging trials. *Journal of Biopharmaceutical Statistics* 17: 965–995.
- Bretz, F., Dette, H., Pinheiro, J. (2010). Practical considerations for optimal designs in clinical dose finding studies. *Statistics in Medicine* 29:731–742.
- Bunouf, P., Lecoutre, B. (2006). Bayesian priors in sequential binomial design. *Comptes Rendus de l'Académie des Sciences Série I* 343:339–344.
- Chambaz, A., Van der Laan, M. J. (2010). Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate. Retrieved October 10, 2010, from <http://hal.archives-ouvertes.fr/hal-00457749/en/>.
- de Cristofaro, R. (2004). On the foundations of likelihood principle. *Journal of Statistical Planning and Inference* 126:401–411.

- Dragalin, V., Bornkamp, B., Bretz, F., Miller, F., Padmanabhan, S. K., Patel, N., Perevozskaya, I., Pinheiro J, Smith, J. R. (2010). A simulation study to compare new adaptive dose-ranging designs. *Statistics in Biopharmaceutical Research* 2: 487–512.
- Eisele, J. R. (1994). The doubly adaptive biased coin design for sequential clinical trials. *Journal of Statistical Planning and Inference* 38:249–61.
- Freedman, D. (1965). Bernard Friedman’s urn. *The Annals of Mathematical Statistics* 36:956–970.
- Hu, F., Zhang, L.-X. (2004). Asymptotic properties of Doubly-adaptive Biased Coin Designs for multitreatment clinical trials. *The Annals of Statistics* 32:268–301.
- Ivanova, A. (2003). A play-the-winner-type urn design with reduced variability. *Metrika* 58:1–13.
- Lecoutre, B. (2008). Bayesian methods for experimental data analysis. In *Handbook of Statistics: Epidemiology and Medical Statistics*, vol. 27, Rao CR, Miller J, Rao DC (eds), Amsterdam: Elsevier, 775–812.
- Lecoutre, B., Charron, C. (2000). Bayesian procedures for prediction analysis of implication hypotheses in 2×2 contingency tables. *Journal of Educational and Behavioral Statistics* 25:185–201.
- Lecoutre, B., Derzko, G., ElQasyr, K. (2010). Frequentist performance of Bayesian inference with response-adaptive designs. *Statistics in Medicine* 29:3219–3231.
- Lecoutre, B., Derzko, G., Grouin, J.-M. (1995). Bayesian predictive approach for inference about proportions. *Statistics in Medicine* 14:1057–1063.
- Lecoutre, B., ElQasyr, K. (2008). Adaptive designs for multi-arm clinical trials: The play-the-winner rule revisited. *Communications in Statistics - Simulation and Computation* 37:590–601.
- Morgan, C. C., Coad, D. S. (2007). A comparison of adaptive allocation rules for group-sequential binary response clinical trials. *Statistics in Medicine* 26:1937–1954.
- Quinlan, J., Gaydos, B., Maca, J., Krams, M. (2010). Barriers and opportunities for implementation of adaptive designs in pharmaceutical product development. *Clinical Trials* 7:167–173.
- Rosenberger, W. F., Hu, F. (2004). Maximizing power and minimizing treatment failures. *Clinical Trials* 1:141–147.
- Rosenberger, W. F., Stallard, N., Ivanova, A., Harper, C. N., Ricks, M. L. (2001). Optimal adaptive designs for binary response trials. *Biometrics* 57:909–913.
- Sun, D., Berger, J. O. (2008). Objective Bayesian analysis under sequential experimentation. *IMS Collections, Pushing The Limits of Contemporary Statistics: Contributions in Honour of Jayanta K Ghosh* 3:19–32.
- Sun, R., Cheung, S. H., Zhang, L.-X. (2007). A generalized drop-the-loser rule for multi-treatment clinical trials. *Journal of Statistical Planning and Inference* 137:2011–2023.
- Tymofyeyev, Y., Rosenberger, W. F., Hu, F. (2007). Implementing Optimal Allocation in Sequential Binary Response Experiments. *Journal of the American Statistical Association* 102:224–234.
- Wei, L. J. (1979). The generalized Pólya urn design for sequential medical trials. *Annals of Statistics* 7:291–296.
- Wei, L. J., Durham, S. (1978). The randomized play-the-winner rule in medical trial. *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. X., Chan, W. S., Cheung, S. H., Hu, F. (2007). Generalized drop-the-loser urn clinical trials with delayed response. *Statistica Sinica* 17:387–409.