Communications in Statistics—Simulation and Computation®, 37: 590–601, 2008

ISSN: 0361-0918 print/1532-4141 online
DOI: 10.1080/03610910701812402

Clinical Trials

Adaptative Designs for Multi-Arm Clinical Trials: The Play-the-Winner Rule Revisited

BRUNO LECOUTRE AND KHADIJA ELQASYR

ERIS, Laboratoire de Mathématiques Raphael Salem,
C.N.R.S. et Université de Rouen, Saint-Etienne-du-Rouvray, France

Adaptative designs for clinical trials that are based on a generalization of the "play-the-winner" rule are considered as an alternative to previously developed models. Theoretical and numerical results show that these designs perform better for the usual criteria. Bayesian methods are proposed for the statistical analysis of these designs.

Keywords Adaptative designs; Bayesian methods; Clinical trials; Generalized Friedman's urn; Play-the-winner rule.

Mathematics Subject Classification Primary 62L05; Secondary 62F15.

1. Introduction

From an ethical point of view, adaptative designs can be desirable for some clinical trials. In such designs, subjects are assumed to arrive sequentially and they are assigned to a treatment with a probability that is updated as a function of the previous events. The intent is to favor the "most effective treatment" given available information. Originally, the *play-the-winner* allocation rule was designed for two treatments with a dichotomous (e.g., success/failure) outcome (Zelen, 1969). It involves an "all-or-none" process: if subject $n-1$ is assigned to treatment t and if the outcome is a success, subject n is assigned to the same treatment; if on the contrary the outcome is a failure, subject n is assigned to the other treatment.

Later, different designs were developed to generalize the rule to the case of three or more treatments and/or to take into account the case of delayed responses (most clinical trials do not result in immediate outcomes and the subject's outcome can be not observable when the next subject arrives); see, e.g., Hoel and Sobel (1972), Wei and Durham (1978), Andersen et al. (1994), Bai et al. (2002b), and Biswas (2003). These designs are generally presented as a *randomized play-the-winner* rule or as a

Received April 1, 2005; Accepted March 9, 2007
Address correspondence to Bruno Lecoutre, ERIS, Laboratoire de Mathématiques Raphael Salem, UMR 6085, C.N.R.S. et Université de Rouen, Avenue de l'Université, BP 12, Saint-Etienne-du-Rouvray 76801, France; E-mail: bruno.lecoutre@univ-rouen.fr

modified version of this rule. We shall see that this name is misleading, because all these designs alter the original all-or-none rule by replacing it with a "linear" adaptive process.

In spite of its apparent determinism, the play-the-winner rule is a stochastic process, since it depends on the probabilities of success on each treatment. However, many people believe that a "less deterministic" rule is better in practice. We shall see that it is not the case.

2. GFU Models and Extensions

The traditional approach is to depict the adaptive rule as a generalized Friedman's urn (also named as generalized Pólya urn) model (GFU model) (Freedman, 1965). A typical GFU model for two treatments can be described as follows. When a new subject n arrives, the urn contains (Y_{n-1}^1, Y_{n-1}^2) balls (or "particles" since the number of balls in the urn can be non integer) that represent the two treatments. A ball is drawn at random and replaced. Then the subject is assigned to the corresponding treatment (say t). When the subject outcome is known, balls are added to the urn. For instance, for a dichotomous outcome $u + v$ balls are added: u type t balls and v balls of the other type in case of success; v type t balls and u balls of the other type in case of failure. Then, if we assume an initial urn composition (Y_0^1, Y_0^2) and immediate outcomes, the urn contains at step n (Y_n^1, Y_n^2) balls, with $Y_n^1 + Y_n^2 = Y_0^1 + Y_0^2 + n(u + v)$. Therefore, the number of balls in the urn at step n is the same, whatever the previous events are.

Bai et al. (2002b) developed a general class of adaptative designs for T treatments and a dichotomous outcome that extend in a straigthforward way the model above. They considered models with $u = 1$ and $v = 0$. Then the models in the class differ only with respect to the repartition of the balls when the response to treatment t is a failure. They proposed, in particular, the three following models.

1. GFU model 1 consists of equally adding 1/(T-1) (fractional) balls of each of the other (T-1) types (see Wei, 1979); of course it is not very satisfactory.
2. GFU model 2 consists of adding balls proportional to the "known" probabilities of success, but this theoretical model is not applicable in practice.
3. In GFU model 3, the unknown probabilities are replaced with the estimated probabilities of success; this looks more satisfactory, but the model is much more complex and is no longer a GFU.

They investigated the asymptotic properties of this class of models and found them to be "desirable" (see also Bai et al., 2002a). It must be emphasized that the case of delayed outcomes is directly taking into account by the models, the urn being updated when outcomes become available; moreover, this does not affect the limiting distribution, although the adaptation process can be drastically slowed. Recently, considerable attention has been given to asymptotic results for very general urn models (Bai and Hu, 2005; Hu and Zhang, 2004; Hu and Rosenberger, 2006; Zhang et al., 2006). However, the real value of asymptotic results for samples of moderate size, even with immediate outcomes, can be seriously questioned. So, "in order to demonstrate the performance of the new design", Bai et al. (2002b) gave numerical illustrations. Unfortunately, if we look through their numerical tables (p. 17), we find that the average allocation proportions may largely differ from the asymptotic values. Some illustrative results are given in Table 1 for the case of three

treatments with probabilities of success 0.90, 0.80, and 0.50. For the "best design" of the authors (model 3), the average allocation proportions in a trial of 100 subjects are respectively 0.481, 0.354, and 0.165 (10^6 *replications*), and they are very distant from the asymptotic values 0.616, 0.295, and 0.089. Even in a trial of 10,000 subjects, the proportions (0.576, 0.325, and 0.099) are not what could be expected.

One reason for the slowness of the adaptation process is that the number of balls at a given step is constant. An attempt to relax this property was the *drop-the-loser* rule proposed by Ivanova (2003). Its principle is as follows: if a treatment type ball is drawn and if the outcome is a failure, the ball is replaced and hence the urn composition remains unchanged; if, on the contrary, it is a success, the ball is not replaced and hence the number of balls is decreased by 1. Moreover, a constant number of balls of no treatment type ("immigration balls") are included in the urn: if such a ball is drawn no subject is treated and the ball is returned to the urn together with additional balls, one of each treatment type. The drop-the-loser rule was generalized to delayed responses by Sun et al. (2006). This rule has the interesting property to reduce the variability of the allocation proportion. Unfortunately, as for the above class of models, the average allocation proportions may be very distant from the asymptotic values (see Table 1).

So, we were induced to consider alternative designs that directly generalize the play-the-winner rule and appear to be preferable.

Table 1
Examples of average allocation proportions (estimated from 10^6 replications) for the GFU models 1 and 3 and for the drop-the-loser rule in the case of three treatments with probabilities of success 0.90, 0.80, and 0.50

	$N = 100$			$N = 1\,000$			$N \to \infty$		
	t^1	t^2	t^3	t^1	t^2	t^3	t^1	t^2	t^3
GFU model 1*	0.464	0.355	0.181	0.517	0.342	0.141	0.588	0.294	0.118
GFU model 3*	0.481	0.354	0.165	0.544	0.341	0.115	0.616	0.295	0.089
drop-the-loser**	0.465	0.345	0.190	0.563	0.308	0.129	0.588	0.294	0.118

*Initial state: $Y_0^1 = Y_0^2 = 1$.
**Initial state: 3 balls for each treatment and 1 non treatment ball.

3. Alternative Models and Some Basic Results

We shall adopt here an equivalent but slightly different conceptualization. For simplification, we present only the case of two treatments. We represent the *state* of the investigator before subject n arrives by a vector $\mathbf{z}_{n-1} = (z_{n-1}^1, z_{n-1}^2)$ where $0 \leq z_{n-1}^i \leq 1$ and $\sum z_{n-1}^i = 1$. For each subject n, there are two observable events: (1) the treatment t_n to which this subject is assigned; $t_n = t^i$ with probability z_{n-1}^i; and (2) the corresponding outcome r_n; $r_n = 1$ (success) with probability φ_i for t^i. We assume an initial state $\mathbf{z}_0 = (z_0^1, z_0^2)$.

The probability transition for the GFU model with two treatments described above (named here as Model B) is given in Table 2.

It must be noted that the initial urn composition (Y_0^1, Y_0^2) is represented here by two parameters with distinct status; on the one hand, the initial state

Table 2
Probability transition for the two classes of models A and B ($\tau = 2$ treatments)

t_n	r_n	Model B	Model A
t^1	1	$z_n^1 = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z_{n-1}^1 + \frac{u}{n_0+n(u+v)}$	$z_n^1 = az_{n-1}^1 + (1-a)b$
t^1	0	$z_n^1 = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z_{n-1}^1 + \frac{v}{n_0+n(u+v)}$	$z_n^1 = az_{n-1}^1 + (1-a)(1-b)$
t^2	1	$z_n^1 = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z_{n-1}^1 + \frac{v}{n_0+n(u+v)}$	$z_n^1 = az_{n-1}^1 + (1-a)(1-b)$
t^2	0	$z_n^1 = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z_{n-1}^1 + \frac{u}{n_0+n(u+v)}$	$z_n^1 = az_{n-1}^1 + (1-a)b$

$\mathbf{z}_0$ ($z_0^1 = Y_0^1/(Y_0^1 + Y_0^2)$), and on the other hand, the parameter n_0 ($= Y_0^1 + Y_0^2$). Consequently, with the new conceptualization, one can let $n_0 = 0$, so that the initial state only intervenes for the assignment of the first subject, but does not intervene in the probability transition. In that follows, we shall consider only, as is usually done, the particular case $u = 1$ and $v = 0$.

It can be shown that, in order to improve the fastness of the adaptation process, the property of a constant number of balls in the urn at a given step must be relaxed. For this purpose, we can then envisage a new class of models, named as Model A, where z_n^1 is again a linear function of z_{n-1}^1, but with constant coefficients. The corresponding probability transition is given in Table 2. It must be emphasized that, unlike Model B, Model A includes the original play-the-winner rule when $a = 0$ and $b = 1$. In this case, z_n^1 takes only the values 0 and 1 (all-or-none model). In that follows, we shall consider only the particular case $b = 1$.

The two models can be characterized by the recurrence relation

$$\mathrm{E}[z_n^1] = A_n \mathrm{E}[z_{n-1}^1] + B_n, \tag{3.1}$$

where A_n and B_n are constants that are function of the model parameters, and furthermore, of n for Model B. It can be deduced that:

$$\mathrm{E}[z_n^1] = z_0^1 \prod_{i=1}^{n} A_i + \sum_{j=1}^{n} B_j \prod_{i=j+1}^{n} A_i. \tag{3.2}$$

For Model B ($u = 1$ and $v = 0$),

$$A_i = 1 - \frac{2 - \varphi_1 - \varphi_2}{n_0 + i} \quad \text{and} \quad B_j = \frac{1 - \varphi_2}{n_0 + j}. \tag{3.3}$$

For Model A ($b = 0$), A_n and B_n does not depend on n:

$$A_i = a + (1-a)(\varphi_1 + \varphi_2 - 1) \quad \text{and} \quad B_j = (1-a)(1-\varphi_2), \tag{3.4}$$

hence

$$\mathrm{E}[z_n^1] - \psi_1 = (z_0^1 - \psi_1)\left(a + (1-a)\left(1 - \frac{1-\varphi_2}{\psi_1}\right)\right)^n \quad \text{where } \psi_1 = \frac{1-\varphi_2}{1-\varphi_1+1-\varphi_2}. \tag{3.5}$$

For each of the two models, we have asymptoticaly

$$\lim_{n\to\infty} \mathrm{E}[z_n^1] \to \psi_1, \tag{3.6}$$

but the convergence is faster for Model A as shown by the two equalities:

$$\text{Model B: } \mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \psi_1) \prod_{i=1}^{n} \left(1 - \frac{1-\varphi_2}{(n_0+i)\psi_1}\right)^n \tag{3.7}$$

$$\text{Model A: } \mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \psi_1)\left(a + (1-a)\left(1 - \frac{1-\varphi_2}{\psi_1}\right)\right)^n. \tag{3.8}$$

Furthermore, for Model A we have the following properties. The smaller a, the smaller $|\mathrm{E}[z_n^1] - \psi_1|$ is, and when $a = 0$ the minimum is such that

$$\mathrm{E}[z_n^1] - \psi_1 = (z_0^1 - \psi_1)(\varphi_1 + \varphi_2 - 1)^n. \tag{3.8}$$

The closer to one $\varphi_1 + \varphi_2$, the smaller $|\mathrm{E}[z_n^1] - \psi_1|$ is, and for $\varphi_1 + \varphi_2 = 1$, $\mathrm{E}[z_n^1] = \psi_1$ $(\forall n \forall z_0^1)$.

Let S_N^1 be the number of subjects assigned to treatment t^1 in a trial of N subjects. It can be deduced that the expected allocation proportion for treatment t^1 is

$$\mathrm{E}\left[\frac{S_N^1}{N}\right] = \frac{1}{N}\sum_{n=0}^{N-1} \mathrm{E}[z_n^1] = \psi_1 + (z_0^1 - \psi_1)\frac{1-h^N}{N(1-h)}, \tag{3.9}$$

$$\text{where } h = a + (1-a)(\varphi_1 + \varphi_2 - 1); \text{ hence if } a = 0,\ h = \varphi_1 + \varphi_2 - 1. \tag{3.10}$$

Note that for the case $a = 0$ formulas for higher factorial moments can be derived from recurrence relations (Elqasyr, 2008).

Table 3 illustrates the superiority of the all-or-none model A ($a = 0$). For this model, not only the convergence is faster, but also the standard deviation of the

Table 3
Comparison of Models A and B for two treatments with probability of success φ_1 and φ_2: average allocation proportions (exact) for treatment t^1 (standard deviations, exact for Model A with $a = 0$ and estimated from 10^5 replications for the three other models)

		$N = 50$ subjects				
φ_1	φ_2	Model B $n_0 = 1$	Model B $n_0 = 0$	Model A $a = 0.15$	Model A $a = 0$	$N \to \infty$
0.30	0.10	0.559 (0.062)	0.562 (0.063)	0.562 (0.044)	0.562 (0.036)	0.562
0.40	0.20	0.566 (0.074)	0.571 (0.077)	0.570 (0.054)	0.570 (0.046)	0.571
0.50	0.40	0.540 (0.097)	0.545 (0.105)	0.544 (0.072)	0.545 (0.064)	0.545
0.70	0.30	0.671 (0.098)	0.696 (0.108)	0.695 (0.073)	0.696 (0.065)	0.700
0.60	0.50	0.546 (0.115)	0.554 (0.132)	0.554 (0.086)	0.554 (0.078)	0.556
0.80	0.60	0.618 (0.149)	0.649 (0.186)	0.660 (0.111)	0.664 (0.101)	0.667
0.90	0.70	0.642 (0.181)	0.692 (0.246)	0.735 (0.132)	0.738 (0.122)	0.750

allocation proportion is notably decreased comparatively to Model B. Model A also gives satisfactory results for a moderate value of a, with, however, an increase of the standard deviation. This is illustrated for $a = 0.15$ in Table 3. For Model B, the possibility of setting $n_0 = 0$ improves the average allocation proportion, but may dramatically increase the variability. The standard deviation of the number of subjects assigned to each treatment is obviously an important criterion for comparing adaptative designs (Hu and Rosenberger, 2003). As pointed out by Ivanova (2003), decrease in variability translates into a gain in statistical power.

4. Generalizations

The two classes of models can be easily generalized to the case of $\mathrm{T} > 2$ treatments, the state of the investigator at step n being a vector $\mathbf{z}_{n-1} = (z_n^1, \ldots z_n^{\mathrm{T}})$ where $0 \leq z_n^i \leq 1$ and $\sum z_n^i = 1$. We can translate as a Model B each of the particular models (1, 2, and 3) considered by Bai et al. (2002a) (and other related models proposed). We can also associate a Model A to each of these models; these models differ with respect to the probability transition in case of failure, while for $a = 0$ they comply with the original play-the-winner rule which is to repeat the treatment in case of success. As for Model B, delayed outcomes are directly taking into account.

The probability transition for the two models is given in Table 4. For simplicity, we consider only without loss of generality the particular cases $u = 1, v = 0$ for Model B, and $b = 1$ for Model A. $q_t^{t^*}$ ($\sum_{t^* \in T \setminus \{t\}} q_t^{t^*} = 1$) is a function of t and t^* which is dependent of the particular model. For instance, for Models B1 and A1, $q_t^{t^*} = 1/(\mathrm{T} - 1)$, and for Models B2 and A2, $q_t^{t^*} = \varphi_t/(\phi - \varphi_{t^*})$, where $\phi = \sum_{t' \in T} \varphi_{t'}$ (see Sec. 2). We assume here that $q_t^{t^*}$ does not depend of the past events.

It can be demonstrated that each particular Model A has the same asymptotic properties as the corresponding Model B (see Appendix). But it always perform better, the adaptation process being the fastest when $a = 0$. This is illustrated in Table 5 for the probability of success $\varphi_1 = 0.90$, $\varphi_2 = 0.80$, and $\varphi_3 = 0.50$.

We have also computed the proportions of the different orders of treatment allocations in each replication. Table 6 illustrates again the manifest superiority of Model A. For the same probability of success, when $N = 300$, for a given trial there is for instance about a 98% chance with Model A3 that a majority of subjects is assigned to the most effective treatment against only about a 74% chance with Model B3.

Table 4
Probability transition for the two classes of models A and B ($\mathrm{T} > 2$ treatments)

t_n	r_n	Model B	Model A
t	1	$z_n^t = \frac{n_0+n-1}{n_0+n} z_{n-1}^t + \frac{1}{n_0+n}$	$z_n^t = a z_{n-1}^t + 1 - a$
t	0	$z_n^t = \frac{n_0+n-1}{n_0+n} z_{n-1}^t$	$z_n^t = a z_{n-1}^t$
$t^* \neq t$	1	$z_n^t = \frac{n_0+n-1}{n_0+n} z_{n-1}^t$	$z_n^t = a z_{n-1}^t$
$t^* \neq t$	0	$z_n^t = \frac{n_0+n-1}{n_0+n} z_{n-1}^t + \frac{1}{n_0+n} q_t^{t^*}$	$z_n^t = a z_{n-1}^t + (1-a) q_t^{t^*}$

Table 5
Comparison of Models A and B with three treatments and probability of success $\varphi_1 = 0.90$, $\varphi_2 = 0.80$, $\varphi_3 = 0.50$: average allocation proportions and standard deviations between parentheses (estimated from 10^6 replications)

	$N = 100$	$N = 300$	$N = 100$	$N = 300$	$N \to \infty$
	Model B1		Model A1 ($a = 0$)		
t^1	0.181 (0.088)	0.135 (0.063)	0.122 (0.053)	0.119 (0.030)	0.118
t^2	0.355 (0.152)	0.349 (0.127)	0.299 (0.119)	0.296 (0.070)	0.294
t_3	0.464 (0.165)	0.516 (0.137)	0.579 (0.134)	0.585 (0.078)	0.588
	Model B3		Model A3 ($a = 0$)		
t^1	0.165 (0.092)	0.135 (0.063)	0.097 (0.056)	0.089 (0.036)	0.089
t^2	0.354 (0.157)	0.349 (0.127)	0.296 (0.127)	0.296 (0.073)	0.295
t_3	0.481 (0.167)	0.516 (0.137)	0.607 (0.136)	0.615 (0.079)	0.616

5. Bayesian Methods

In conclusion, our results are very incentive: the simplest model is the best! This greatly facilitates both a thoughtful planning into the design phase and the use of efficient inference procedures.

For this purpose, the Bayesian statistical methodology can be used for designing the study (how many subjects?) and for comparing the treatments. A clinical trial is generally expected to bring evidence by itself. Therefore, it is desirable in clinical research to assume non informative priors for objective report in publication, the

Table 6
Comparison of Models A and B with three treatments and probability of success $\varphi_1 = 0.90$, $\varphi_2 = 0.80$, $\varphi_3 = 0.50$: Proportions of the different orders of treatment allocations (estimated from 10^6 replications)

		$N = 100$	$N = 300$	$N = 100$	$N = 300$
		Model B1		Model A1	
$t^1 \geq t^2 \geq t^3$	$+++$	0.489	0.676	0.801	0.975
$t^1 \geq t^3 \geq t^2$	$+--$	0.136	0.060	0.064	0.007
$t^3 \geq t^2 \geq t^1$	$-+-$	0.015	0.001	0.000	0.000
$t^2 \geq t^1 \geq t^3$	$--+$	0.287	0.253	0.129	0.018
$t^2 \geq t^3 \geq t^1$ or $t^3 \geq t^1 \geq t^2$	$---$	0.073	0.010	0.003	0.000
		Model B3		Model A3	
$t^1 \geq t^2 \geq t^3$	$+++$	0.511	0.675	0.814	0.975
$t^1 \geq t^3 \geq t^2$	$+--$	0.136	0.061	0.071	0.006
$t^3 \geq t^2 \geq t^1$	$-+-$	0.012	0.001	0.000	0.000
$t^2 \geq t^1 \geq t^3$	$--+$	0.281	0.254	0.113	0.018
$t^2 \geq t^3 \geq t^1$ or $t^3 \geq t^1 \geq t^2$	$---$	0.061	0.010	0.003	0.000

posterior distribution being based solely on the data. But alternative choices of priors may be used to refining inference.

Moreover, the Bayesian predictive approach is a very appealing method for monitoring the study and, in particular, for stopping it early if necessary (e.g., Lecoutre et al., 1995, 2002; Spiegelhalter et al., 1986). It simulates the probability of achieving the trial target, conditionally on available data and simple conjectures about the future observations. The simulations can be explicitly based on either the hypotheses used to design the study, expressed in terms of the prior distribution, or on available data, or on both.

We shall briefly illustrate Bayesian methods for the basic situation of an adaptative design with two treatments using Model A (with $a = 0$). We also assume immediate outcomes. The sequel of treatment allocations $(t_1, t_2 \ldots t_n, t_{n+1} \ldots t_{N+1})$ contains all the information in the data. Indeed, $t_n = t_{n+1}$ implies that a success to t_n has been observed and $t_n \neq t_{n+1}$ implies that a failure to t_n has been observed. Moreover, the likelihood function is simply

$$l(\varphi_1, \varphi_2)|(t_1, \ldots t_{N+1}) = \frac{1}{2}\varphi_1^{n_{11}}(1-\varphi_1)^{n_{10}}\varphi_2^{n_{21}}(1-\varphi_2)^{n_{20}}, \tag{5.1}$$

where n_{ij} is the number of pairs (t_n, t_{n+1}) equal to (t^i, t^j), so that n_{11} and n_{21} are the respective numbers of success to treatments t^1 and t^2, and n_{10} and n_{20} are the numbers of failure ($1/2$ is the probability of t_1).

Bayesian methods only involve the likelihood function and are immediately available. This results from the fact that the likelihood function is identical (up to a multiplicative constant) with the likelihood function associated with the comparison of two independent binomial proportions. However, here only $n_{11} + n_{10} + n_{n21} + n_{20} = N$ is fixed. Therefore we can apply the same Bayesian procedures. A simple and usual solution assumes two independent beta prior distributions for φ_1 and φ_2: $\beta(v_{11}, v_{10})$ and $\beta(v_{21}, v_{20})$, respectively. The marginal posterior distribution are again two independent beta distributions: $\beta(v_{11} + n_{11}, v_{10} + n_{10})$ and $\beta(v_{21} + n_{21}, v_{20} + n_{20})$.

Let us consider, for illustration, the results of a trial with $N = 100$ subjects. The observed rates of success are, respectively, 56 out of 69 attributions for treatment t^1 and 17 out of 31 attributions for treatment t^2. A joint probability statement is, in a way, the best summary of the posterior distribution. For instance, if we adopt the same conventional Jeffreys prior as for two binomial proportions ($v_{11} = v_{10} = v_{21} = v_{20} = .5$), the joint posterior probability that $\varphi_1 > 0.708$ and $\varphi_2 < 0.712$ is 0.95.

However, a statement that deals with the comparison of the two treatments directly would be preferable. So we have a probability 0.996 that $\varphi_1 > \varphi_2$. Moreover, the main classical criteria for comparing two proportions can be dealt with. This is easily solved in the Bayesian approach, since the distribution of any derived parameter of interest can be easily obtained from the joint posterior distribution using numerical methods. For instance, we find the 95% credible intervals $[+0.068, +0.453]$ for $\varphi_1 - \varphi_2$, $[1.10, 2.18]$ for φ_1/φ_2, and $[1.41, 9.07]$ for $(\varphi_1/(1-\varphi_1))/\varphi_2/(1-\varphi_2))$.

These Bayesian intervals have fairly good frequentist coverage properties, even in the cases of moderate sample sizes and small parameter values (see, e.g., Lecoutre and Charron, 2000). As an illustration, we considered the inference about the difference, ratio, and odds-ratio for a sample of size $N = 50$. We generated the set of all possible samples for different values of φ_1 and φ_2 and computed the sampling

probability that the 95% lower and 95% upper limits were, respectively, greater and smaller than the true parameter value. These frequentist probabilities of an error ("non coverage") are reported in Table 7.

The above methods can be easily generalized with virtually no more conceptual difficulties to the case of several treatments and/or delayed outcomes. The Bayesian approach is appropriate as well for a definitely decisional trial (e.g., for selecting the best treatment) as for estimation (e.g., for assessing the difference in efficacy between two treatments). Moreover, the predictive approach enables the trial to be stopped early, or on the contrary, to be extended to an adequate size, in a sequential perspective that fits with the methodological principle of adaptative designs.

Table 7
Coverage properties of Bayesian credible intervals for the difference, the ratio and the odds-ratio: frequentist probabilities of an error for the 95% lower and 95% upper limits ($N = 50$)

		$\varphi_1 - \varphi_2$		φ_1/φ_2		$\frac{\varphi_1(1-\varphi_2)}{\varphi_2(1-\varphi_1)}$	
φ_1	φ_2	Lower	Upper	Lower	Upper	Lower	Upper
0.80	0.80	0.0576	0.0576	0.0576	0.0576	0.0576	0.0576
0.60	0.60	0.0536	0.0536	0.0536	0.0536	0.0536	0.0536
0.50	0.50	0.0517	0.0517	0.0517	0.0517	0.0517	0.0517
0.80	0.70	0.0582	0.0517	0.0589	0.0543	0.0521	0.0550
0.70	0.60	0.0566	0.0488	0.0537	0.0492	0.0517	0.0574
0.60	0.50	0.0542	0.0508	0.0539	0.0509	0.0501	0.0558
0.80	0.60	0.0562	0.0544	0.0544	0.0558	0.0496	0.0619
0.70	0.50	0.0579	0.0500	0.0567	0.0510	0.0515	0.0554
0.60	0.40	0.0575	0.0480	0.0537	0.0512	0.0549	0.0553

Appendix

We sketch the demonstration that each particular Model A with a set of T treatments $T = \{t_1, t_2, \ldots, t_T\}$ has the same asymptotic properties as the corresponding Model B. We have the recurrence relation

$$\mathrm{E}[z_n^t] = (a + (1-a)\varphi_t)\mathrm{E}[z_{n-1}^t] + (1-a)\left(\sum_{t^* \in T\setminus\{t\}} q_t^{t^*}(1-\varphi_{t^*})\mathrm{E}[z_{n-1}^{t^*}]\right). \quad \text{(A.1)}$$

For each $t \in T\setminus\{t_T\}$, using the identity $z_{n-1}^T = 1 - z_{n-1}^t - \sum_{t^* \in T\setminus\{t,t_T\}} z_{n-1}^{t^*}$, we can write:

$$\begin{aligned}\mathrm{E}[z_n^t] = {} & \big(a + (1-a)\big(\varphi_t - q_t^K(1-\varphi_K)\big)\big)\mathrm{E}[z_{n-1}^t] \\ & + (1-a)\left(\sum_{t^* \in T\setminus\{t,t_T\}} \big(q_t^{t^*}(1-\varphi_{t^*}) - q_t^T(1-\varphi_T)\big)\mathrm{E}[z_{n-1}^{t^*}]\right) \\ & + (1-a)q_t^T(1-\varphi_T). \end{aligned} \quad \text{(A.2)}$$

Letting $\overline{\mathbf{Z}}_n = \mathrm{E}(z_n^1, z_n^2, \ldots, z_n^{\mathrm{T}-1})$, the overall recurrence relation can be written

$$\overline{\mathbf{Z}}_n = \mathbf{A}\overline{\mathbf{Z}}_{n-1} + \mathbf{B}$$

$$\text{where } \mathbf{A} = \begin{pmatrix} a + (1-a)\big(\varphi_1 - q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}})\big) & \cdots & \cdots \\ \cdots & a + (1-a)\big(\varphi_2 - q_2^{\mathrm{T}}(1-\varphi_{\mathrm{T}})\big) & \cdots \\ \cdots & \cdots & \cdots \end{pmatrix}$$

$$\text{and } \mathbf{B} = \begin{pmatrix} (1-a)\, q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) \\ \cdots \\ (1-a)\, q_{\mathrm{T}-1}^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) \end{pmatrix}. \tag{A.3}$$

It can be deduced that ($\mathbf{I}$ is the identity matrix):

$$\overline{\mathbf{Z}}_n = \mathbf{A}^n\overline{\mathbf{Z}}_0 + \mathbf{B}(\mathbf{I}-\mathbf{A})^{-1}(\mathbf{I}-\mathbf{A}^n). \tag{A.4}$$

The matrix $\mathbf{A}$ can be written as $\mathbf{A} = a\mathbf{I} - (1-a)\mathbf{H}$, where

$$\mathbf{H} = \begin{pmatrix} \varphi_1 - q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) & \cdots & \cdots & q_{\mathrm{T}-1}^1(1-\varphi_1) - q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) \\ q_1^2(1-\varphi_2) - q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) & \varphi_2 - q_2^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ q_1^{\mathrm{T}-1}(1-\varphi_{\mathrm{T}-1}) - q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) & \cdots & \cdots & \varphi_{\mathrm{T}-1} - q_{\mathrm{T}-1}^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) \end{pmatrix}. \tag{A.5}$$

Since $\|\mathbf{H}\| = \sup_{i,j} |H(i,j)| < 1$, $(\mathbf{I}-\mathbf{H})$ is inversible. It follows that $\mathbf{I}-\mathbf{A} = (1-a)(\mathbf{I}-\mathbf{H})$ is also inversible and that $\mathbf{I}-\mathbf{A} = \sum_{n=0}^{+\infty}\mathbf{A}^n$. Then the series $\sum_{n=0}^{+\infty}\mathbf{A}^n$ converges and $\|\mathbf{A}\| < 1$. This implies that $\mathbf{A}^n \to \mathbf{0}$ when $n \to \infty$. It follows that

$$\lim_{n\to\infty} \overline{\mathbf{Z}}_n = \mathbf{B}(\mathbf{I}-A)^{-1}. \tag{A.6}$$

For instance, we find for Model A2

$$\lim_{n\to\infty} z_n^t = \frac{\varphi_t(1-\varphi_t)^{-1}(\phi-\varphi_t)}{\sum_{t'}\varphi_{t'}(1-\varphi_{t'})^{-1}(\phi-\varphi_{t'})} \qquad \left(\phi = \sum_{t'\in T}\varphi_{t'}\right). \tag{A.7}$$

Let S_N^t be the number of subjects assigned to treatment t in a trial of N subjects and $\overline{\mathbf{S}}_N = \mathrm{E}(S_N^1, S_N^2, \ldots, S_N^{\mathrm{T}-1})$. From $\mathrm{E}[S_N^t] = \sum_{n=1}^N \mathrm{E}[z_n^t]$, we deduce

$$\overline{\mathbf{S}}_N = \sum_{n=1}^{N}\mathbf{M}^n\overline{\mathbf{Z}}_0 = (\mathbf{I}-\mathbf{M})^{-1}(\mathbf{I}-\mathbf{M}^n)\overline{\mathbf{Z}}_0 \tag{A.8}$$

where

$$\mathbf{M} = \begin{pmatrix} a + (1-a)\varphi_1 & \cdots & \cdots & (1-a)\, q_1^{\mathrm{T}}(1-\varphi_{\mathrm{T}}) \\ (1-a)q_2^1(1-\varphi_1) & a + (1-a)\varphi_2 & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ (1-a)q_{\mathrm{T}}^1(1-\varphi_1) & \cdots & \cdots & a + (1-a)\varphi_{\mathrm{T}} \end{pmatrix},$$

and

$$\lim_{n\to\infty} \overline{\mathbf{S}}_N = (\mathbf{I} - \mathbf{M})^{-1}\mathbf{Z}_0. \tag{A.9}$$

We get:

$$\text{for Model A1,} \quad \lim_{n\to\infty} S_N^t = \frac{(1-\varphi_t)^{-1}}{\sum_{t'\in \mathbf{T}}(1-\varphi_{t'})^{-1}} \tag{A.10}$$

$$\text{for Model A2,} \quad \lim_{n\to\infty} S_N^t = \frac{\varphi_t(1-\varphi_t)^{-1}}{\phi - \varphi_t} \tag{A.11}$$

By adapting the demonstrations given by Bai et al. (2002a,b), the results can be extended to the models where $q_t^{t^*}$ depends of the past events. In particular, the limits for Model A3 are the same as for Model A2. In the same way, the results can be extended to the case of delayed outcomes.

References

Andersen, D., Faries, J., Tamura, R. (1994). A randomized play-the-winner design for multiarm clinical trials. *Communications in Statistics, Theory and Methods* 23:309–323.

Bai, Z. D., Hu, F. (2005). Asymptotics in randomized urn models. *The Annals of Applied Probability* 15:914–940.

Bai, Z. D., Hu, F., Rosenberger, W. F. (2002a). 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. (2002b). An adaptive design for multi-arm clinical trials. *Journal of Multivariate Analysis* 81:1–18.

Biswas, A. (2003). Generalized delayed response in randomized play-the-winner rule. *Communications in Statistics, Simulation and Computation* 32:259–274.

Elqasyr, K. (2008). Plans Séquentiels pour les Essais Cliniques et Procédures Bayésiennes. Unpublished doctoral thesis, Université de Rouen, France.

Freedman, D. (1965). Bernard Friedman's urn. *The Annals of Mathematical Statistics* 36:956–970.

Hoel, D. G., Sobel, M. (1972). Comparison of sequential procedures for selecting the best binomial population. *Proceedings of the Sixth Berkeley Symposium on Probability and Statistics* 4:53–69.

Hu, F., Rosenberger, W. F. (2003). Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association* 98:671–678.

Hu, F., Zhang, L. X. (2004). The asymptotic normality of urn models for clinical trials with delayed response. *Bernoulli* 10:447–463.

Hu, F., Rosenberger, W. F. (2006). *The Theory of Response-Adaptative Randomization in Clinical Trials*. New York: Wiley.

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

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., Grouin, J.-M. (1995). Bayesian predictive approach for inference about proportions. *Statistics in Medicine* 14:1057–1063.

Lecoutre, B., Mabika, B., Derzko, G. (2002) Assessment and monitoring in clinical trials when survival curves have distinct shapes in two groups: a Bayesian approach with Weibull modeling. *Statistics in Medicine* 21:663–674.

Sun, R., Cheung, S. H., Zhang, L.-X. (2006). A generalized drop-the-loser rule for multi-treatment clinical trials. *Journal of Statistical Planing and Inference* 137:2011–2023.

Spiegelhalter, D. J., Freedman, L. S., Blackburn, P. R. (1986). Monitoring clinical trials: conditional or predictive power? *Controlled Clinical Trials* 7:8–17.

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

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

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

Zhang, L. X., Hu, F., Cheung, S. H. (2006). Asymptotic theorems of sequential estimation-adjusted urn models. *The Annals of Applied Probability* 16:340–369.