Play-the-winner rule in clinical trials: Models for adaptative designs and Bayesian methods

Bruno Lecoutre and Khadija Elqasyr

ERIS, Laboratoire de Mathématiques Raphael Salem,
UMR 6085, C.N.R.S. et Université de Rouen,
Mathématiques, Site Colbert, 76821 Mont-Saint-Aignan Cedex, France
(e-mail: bruno.lecoutre@univ-rouen.fr,
khadija.elqasyr@etu.univ-rouen.fr)

Abstract. 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: Clinical trials, Adaptative designs, Play-the-winner rule, Generalized Friedman's urn, Bayesian methods.

1 Introduction

From 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, 1998], [Wei and Durham, 1978], [Andersen and Tamura, 1994], [Bai *et al.*, 2002a], [Biswas, 2003]. These designs are generally presented as a *randomized play-the-winner* rule or as a 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^1_{n-1}, Y^2_{n-1}) balls (or "particles" since the number of balls in the urn can be non integer) that represents 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^1_0, Y^2_0) and immediate outcomes, the urn contains at step n (Y^1_n, Y^2_n) balls, with $Y^1_n + Y^2_n = Y^1_0 + Y^2_0 + n(u + v)$. Therefore the number of balls in the urn at step n is the same, whatever the previous events are.

Bai, Hu and Shen developed a general class of adaptative designs for T treatments and a dichotomous outcome [Bai *et al.*, 2002a] 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. 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. GFU model 2 consists of adding balls proportional to the "known" probabilities of success, but this theoretical model is not applicable in practice. Then in model 3 the unknown probabilities are replaced with the estimated probability 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.*, 2002b]). 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 considerably slowed.

"In order to demonstrate the performance of the new design" the authors gave numerical illustrations. Unfortunately, if we look through their numerical tables ([Bai *et al.*, 2002a], page 17), we can seriously questioned

the real value of their asymptotic results for samples of moderate size, even with immediate outcomes.

For instance, let us consider three treatments with probabilities of success 0.50, 0.80 and 0.90. For the "best design" of the authors, the average allocation proportions in a trial of 100 subjects are respectively 0.165, 0.354 and 0.481, and they are very distant from the asymptotic values 0.089, 0.295 and 0.616. Even in a trial of 10 000 subjects the proportions – 0.099, 0.325 and 0.576 – are not what could be expected. So, we were induced to consider other designs that directly generalize the play-the-winner rule and appear to be preferable.

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 $\mathbf{z}_{n-1} = (z^1_{n-1}, z^2_{n-1})$ where $0 \leq z^i_{n-1} \leq 1$ and $\sum z^i_{n-1} = 1$. 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^i_{n-1}; and (2) the corresponding outcome r_n; $r_n = 1$ (success) with probability φ_1 for t^1 and probability φ_2 for t^2. We assume an initial state $\mathbf{z}_0 = (z^1_0, z^2_0)$.

The probability transition for the GFU model with two treatments described above (named here as Model I) is given in Table 1

t_n	r_n	**Model I**	**Model II**
t^1	1	$z^1_n = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z^1_{n-1} + \frac{u}{n_0+n(u+v)}$	$z^1_n = az^1_{n-1} + (1-a)b$
t^1	0	$z^1_n = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z^1_{n-1} + \frac{v}{n_0+n(u+v)}$	$z^1_n = az^1_{n-1} + (1-a)(1-b)$
t^2	1	$z^1_n = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z^1_{n-1} + \frac{v}{n_0+n(u+v)}$	$z^1_n = az^1_{n-1} + (1-a)(1-b)$
t^2	0	$z^1_n = \frac{n_0+(n-1)(u+v)}{n_0+n(u+v)} z^1_{n-1} + \frac{u}{n_0+n(u+v)}$	$z^1_n = az^1_{n-1} + (1-a)b$

Table 1. Probability transitions for the two classes of models

It must be noted that the initial urn composition (Y^1_0, Y^2_0) is here represented by two parameters with distinct status, on the one hand the initial state $\mathbf{z}_0$ ($z^1_0 = Y^1_0/(Y^1_0 + Y^2_0)$), and on the other hand the parameter n_0 ($= Y^1_0 + Y^2_0$). 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 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 II, 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 1. It must be emphasized that, unlike Model I, Model II 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 $b = 1$ and we shall assume $z_0^1 = z_0^2 = 0.5$.

The two models can be characterized by the recurrence relation

$$\mathrm{E}(z_n^1) = A_n \mathrm{E}(z_{n-1}^1) + B_n$$

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

$$\mathrm{E}(z_n^1) = z_0^1 \prod_{i=1}^{n} A_i + \sum_{j=1}^{n} B_j \prod_{i=j+1}^{n} A_i$$

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

$$A_i = 1 - \frac{2 - \varphi_1 - \varphi_2}{n_0 + i} \quad \text{and} \quad B_i = \frac{1 - \varphi_2}{n_0 + i}$$

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

$$A_i = a + (1 - a)(\varphi_1 + \varphi_2 - 1) \quad \text{and} \quad B_i = (1 - a)(1 - \varphi_2)$$

hence

$$\mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \varphi_1)\Big(a + (1 - a)\big(1 - \frac{1 - \varphi_2}{\psi_1}\big)\Big)^n$$

For each of the two models, we have asymptoticaly

$$\text{when } n \to \infty,\ \mathrm{E}(z_n^1) \to \psi_1 = \frac{1 - \varphi_2}{1 - \varphi_1 + 1 - \varphi_2}$$

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

$$\text{Model I:}\ \ \mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \varphi_1) \prod_{i=1}^{n} \Big(1 - \frac{1 - \varphi_2}{(n_0 + i)\psi_1}\Big)^n$$

$$\text{Model II:}\ \ \mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \varphi_1)\Big(a + (1 - a)\big(1 - \frac{1 - \varphi_2}{\psi_1}\big)\Big)^n$$

Furthermore, for Model II 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

$$\mathrm{E}(z_n^1) - \psi_1 = (z_0^1 - \varphi_1)(\varphi_1 + \varphi_2 - 1)^n$$

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 T_N^1 be the number of subjects assigned to treatment t^1 in a trial of N subjects. It can be deduced that

$$\mathrm{E}(T_N^1) = \frac{1}{N}\sum_{n=0}^{N-1}\mathrm{E}(z_n^1) = \psi_1 + \frac{1}{N}(z_0^1 - \varphi_1)\frac{1-h^N}{1-h}$$

$$\text{where } h = a + (1-a)(\varphi_1 + \varphi_2 - 1)$$

$$= \varphi_1 + \varphi_2 - 1 \quad \text{if } a = 0$$

Table 2 illustrates the superiority of the all-or-none model for the probability of success $\varphi_1 = 0.60$ and $\varphi_2 = 0.80$. The possibility of setting $n_0 = 0$ in Model I improves the average allocation proportion, but notably increases the standard deviation.

$\varphi_1 = 0.60 \quad \varphi_2 = 0.80$

$N = 50$ subjects

Model I $n_0 = 1$	Model I $n_0 = 0$	Model II $a = 0$	$N \to \infty$
0.618 (0.149)	0.649 (0.186)	0.661 (0.101)	0.667

Table 2. Comparison of Models I and II with two treatments: average allocation proportions (exact) for treatment t^2 (standard deviations estimated from 10^6 replications)

4 Generalizations

The two class of models can be easily generalized to the case of $\mathrm{T} > 2$ treatments. We can translate as a Model I each of the particular models (1, 2 and 3) considered by Bai *et al.* (and other related models proposed). We can also associate a Model II 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 I, delayed outcomes are directly taking into account. Moreover, it can be demonstrated that each particular Model II has the same asymptotic properties as the corresponding Model I. But it always perform better, the adaptation process being the fastest when $a = 0$.

$\varphi_1 = 0.50 \quad \varphi_2 = 0.80 \quad \varphi_3 = 0.90$

	Model I-1		Model II-1 ($a = 0$)		
	$N = 100$	$N = 300$	$N = 100$	$N = 300$	$N \to \infty$
t^1	0.181 (0.088)	0.135 (0.063)	0.122 (0.053)	0.089 (0.036)	0.118
t^2	0.355 (0.152)	0.349 (0.127)	0.299 (0.119)	0.296 (0.073)	0.294
t_3	0.464 (0.165)	0.516 (0.137)	0.579 (0.134)	0.615 (0.079)	0.588
	Model I-3		Model II-3 ($a = 0$)		
	$N = 100$	$N = 300$	$N = 100$	$N = 300$	$N \to \infty$
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

Table 3. Comparison of Models I and II with three treatments: average allocation proportions and standard deviations between parentheses (estimated from 10^6 replications)

This is illustrated in table 3 for the probability of success $\varphi_1 = 0.50$, $\varphi_2 = 0.80$ and $\varphi_3 = 0.90$.

We have also computed the proportions of the different orders of treatment allocations in each replication. Table 4 illustrates again the manifest superiority of Model II. For the same probability of success, when $N = 300$, for a given trials there is for instance about a 98% chance with Model II-3 that a majority of subjects is assigned to the most effective treatment against only about a 74% chance with Model I-3.

$\varphi_1 = 0.50 \quad \varphi_2 = 0.80 \quad \varphi_3 = 0.90$

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

Table 4. Comparison of Models I and II with three treatments: Proportions of the different orders of treatment allocations (estimated from 10^6 replications)

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. So it is desirable in clinical research to assume noninformative priors for objective report in publication, the 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., [Spiegelhalter *et al.*, 1986, Lecoutre *et al.*, 1995, Lecoutre *et al.*, 2002]). 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 II (with $a = 0$). We also assume immediate outcomes. The sequel of treatment allocations (t_1, $t_2 \dots t_n$, $t_{n+1} \dots t_{N+1}$) 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

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

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. 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: respectively $\beta(\nu_{11}, \nu_{10})$ and $\beta(\nu_{21}, \nu_{20})$. The marginal posterior distribution are again two independent beta distributions: $\beta(\nu_{11} + n_{11}, \nu_{10} + n_{10})$ and $\beta(\nu_{21} + n_{21}, \nu_{20} + n_{20})$. The predictive distributions for future observations are two independent beta-binomial distributions ([Lecoutre *et al.*, 1995]).

Let us consider for illustration the results of a trial with $N = 100$ subjects. The observed rates of success are respectively 17 out of 31 attributions for treatment t^1 and 56 out of 69 attributions for treatment t^2.

A joint probability statement is, in a way, the best summary of the posterior distribution. For instance, if we conventionally adopt the Jeffreys prior ($\nu_{11} = \nu_{10} = \nu_{21} = \nu_{20} = .5$), the joint posterior probability that $\varphi_1 < 0.712$ and $\varphi_2 > 0.708$ 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 $\varphi_2 > \varphi_1$. 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 $\varphi_2 - \varphi_1$, $[1.10, 2.18]$ for φ_2/φ_1 and $[1.41, 9.07]$ for $(\varphi_2/(1-\varphi_2))/\varphi_1/(1-\varphi_1))$.

For the Jeffreys prior, Bayesian methods have fairly good frequentist coverage properties for interval estimates, even in the the cases of moderate sample sizes and small parameter values (see e.g., [Lecoutre and Charron, 2000]). As an illustration, 10^5 samples of size $N = 50$ were generated for different set of parameter values. We considered the inference about the difference $\varphi_2 - \varphi_1$. The proportion of samples for which respectively the 95% lower and 95% upper limits were respectively greater and smaller than the true difference are reported in Table 5.

	Lower limit	Upper limit
$\varphi 2 = 0.80$ $\varphi 1 = 0.80$	0.059	0.057
$\varphi 2 = 0.60$ $\varphi 1 = 0.60$	0.053	0.053
$\varphi 2 = 0.50$ $\varphi 1 = 0.50$	0.052	0.051
$\varphi 2 = 0.80$ $\varphi 1 = 0.70$	0.058	0.052
$\varphi 2 = 0.70$ $\varphi 1 = 0.60$	0.056	0.049
$\varphi 2 = 0.60$ $\varphi 1 = 0.50$	0.053	0.051
$\varphi 2 = 0.80$ $\varphi 1 = 0.60$	0.055	0.053
$\varphi 2 = 0.70$ $\varphi 1 = 0.50$	0.058	0.050
$\varphi 2 = 0.60$ $\varphi 1 = 0.40$	0.058	0.047

Table 5. Coverage properties of Bayesian credible intervals for the comparison of two treatments: proportions of errors for the 95% lower and 95% upper limits (10^5 replications)

These 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.

References

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

[Bai *et al.*, 2002a]Z. D. Bai, F. Hu, and W. F. Rosenberger. Asymptotic properties of adaptive designs for clinical trials with delayed response. *The Annals of Statistics*, pages 122–139, 2002.

[Bai *et al.*, 2002b]Z. D. Bai, F. Hu, and L. Shen. An adaptive design for multi-arm clinical trials. *Journal of Multivariate Analysis*, pages 1–18, 2002.

[Biswas, 2003]A. Biswas. Generalized delayed response in randomized play-the-winner rule. *Communications in Statistics, Simulation and Computation*, pages 259–274, 2003.

[Freedman, 1965]D. Freedman. Bernard friedman's urn. *The Annals of Mathematical Statistics*, pages 956–970, 1965.

[Hoel and Sobel, 1998]D. G. Hoel and M. Sobel. Comparison of sequential procedures for selecting the best binomial population. In *Proceedings of the Sixth Berkeley Symposium on Probability and Statistics*, pages 53–69, 1998.

[Lecoutre and Charron, 2000]B. Lecoutre and C. Charron. Bayesian procedures for prediction analysis of implication hypotheses in 2×2 contingency tables. *Journal of Educational and Behavioral Statistics*, pages 185–201, 2000.

[Lecoutre *et al.*, 1995]B. Lecoutre, G. Derzko, and J.-M. Grouin. Bayesian predictive approach for inference about proportions. *Statistics in Medicine*, pages 1057–1063, 1995.

[Lecoutre *et al.*, 2002]B. Lecoutre, B. Mabika, and G. Derzko. Assessment and monitoring in clinical trials when survival curves have distinct shapes in two groups: a bayesian approach with weibull modeling. *Statistics in Medicine*, pages 663–674, 2002.

[Spiegelhalter *et al.*, 1986]D.J. Spiegelhalter, L.S. Freedman, and P.R. Blackburn. Monitoring clinical trials: Conditional or predictive power? *Controlled Clinical Trials*, pages 8–17, 1986.

[Wei and Durham, 1978]L. J. Wei and S. Durham. The randomized play-the-winner rule in medical trial. *Journal of the American Statistical Association*, pages 840–843, 1978.

[Wei, 1979]L. J. Wei. The generalized pólya urn design for sequential medical trials. *Annals of Statistics*, pages 291–296, 1979.

[Zelen, 1969]M. Zelen. Play the winner rule and the controlled clinical trial. *Journal of the American Statistical Association*, pages 131–146, 1969.