Comparing two success rates with Play-The-Winner designs

Khadija Elqasyr, Bruno Lecoutre

Comparing two success rates with Play-The-Winner designs

Khadija ElQasyr [a] Bruno Lecoutre [b]

[a] *UMR 6085, Laboratoire de Mathématiques Raphaël Salem, C.N.R.S. et Université de Rouen, Mathématiques, Avenue de l'Université, BP 12, 76801 Saint-Etienne-du-Rouvray, France.*
elqasyr@yahoo.fr

[b] *ERIS, UMR 6085, Laboratoire de Mathématiques Raphaël Salem, C.N.R.S. et Université de Rouen, Mathématiques, Avenue de l'Université, BP 12, 76801 Saint-Etienne-du-Rouvray, France.*
bruno.lecoutre@univ-rouen.fr
http://www.univ-rouen.fr/LMRS/Persopage/Lecoutre/Eris.html

October 2009

Abstract

Explicit formulae for the sampling distribution of Play-the-Winner designs are given. These formulae involve familiar and easy-to-compute distributions. Two examples of applications are given. First, conditional test and confidence interval procedures for PW designs are developed. Second, the Bayesian predictive distributions associated with conjugate independent Beta priors are derived. Moreover, the correspondance between Bayesian posterior distributions and conditional tests is made explicit.

Key words: Play-The-Winner rule, Comparison of two proportions, Conditional tests, Confidence intervals, Mid-p approach, Bayesian inference, Predictive distributions,
Play-The-Winner rule; Comparison of two proportions;Conditional tests; Confidence intervals; Mid-p approach; Bayesian inference; Predictive distributions

1 Introduction

In response-adaptive designs newly accrued subjects are assigned a treatment with a probability that is updated as a function of previous outcomes, according to some predefined rule. The intent is to favor the assignment of the "most effective treatment" given available information. Even if these designs can only be reasonably implemented in specific clinical situations, they are attractive competitors to 1:1 randomized designs for comparing the success rates of two treatments (Pullman and Wang, 2001; Rosenberger, 2004).

The *Play-The-Winner* (PW) allocation rule was designed for two treatments with a dichotomous (e.g. success/failure) outcome (Zelen, 1969). It is based on the intuitive and empirical idea of translating the stochastic process of treatment success generation into a stochastic process of treatment allocation. If subject $k-1$ is assigned to treatment t ($t \in \{1, 2\}$) and if the outcome is a success (with probability φ_t), subject k is assigned to the same treatment; if the outcome is a failure, subject k is assigned to the other treatment. 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.

Lecoutre and ElQasyr (2008) demonstrated that the Play-the-Winner rule has optimal properties and always converges more rapidly than the Randomized Play-the-Winner rule, sometimes notably so, with less variability. Moreover, Lecoutre, Derzko and ElQasyr (2009) showed that, in terms of the total number of failures, several response-adaptive designs, including recent ones, generally provide a modest advantage over 1:1 randomized designs. The only notable exceptions were found for the Play-The-Winner design, in the cases where the most effective treatment is clearly superior to the other treatment.

Explicit formulae for the sampling distribution of PW designs are given for two stopping rules. Let n_{11} and n_{21} be the respective numbers of successes to the two treatments, and n_{10} and n_{20} be the numbers of failures. A consequence of the PW rule is that n_{10} and n_{20} are equal or differ by only one unit. Hence, if the experiment stops after a fixed even number of subjects n has been included (stopping rule R1), n_{10} and n_{20} cannot be larger than $n/2$. Thus they cannot exceed the number of failures that can be observed in the 50-50 randomized design. Moreover, the expected total number of failures is less for the PW rule. If treatment t has probability π_t ($0 \leq \pi_t \leq 1$, $\pi_1 + \pi_2 = 1$) of being selected as the first treatment, the expected number of failures to treatment t is for rule R1 (Lecoutre and ElQasyr, 2008)

$$\Big(n\psi_t + (\pi_t - \psi_t)\frac{1-h^n}{1-h}\Big)(1-\varphi_t) \qquad \text{where } h = \varphi_1 + \varphi_2 - 1$$

An alternative stopping rule (R2) is to stop the experiment after a fixed number of failures has been observed for one of the treatment (say treatment 2). This allows to tightly control the number of failures. An equal number of failures can be obtained for the other treatment if the first subject is assigned to this treatment ($\pi_1 = 1$).

2 Sampling distributions

The sampling probability to observe (n_{11}, n_{10}, n_{21}, n_{20}) can be expressed from Binomial and Negative Binomial probabilities.

Stopping rule R1

The experiment stops after a fixed number of subjects $n = n_{11}+n_{10}+n_{21}+n_{20}$ has been included. Recall that $n_{10} = n_{20}$ or $n_{10} = n_{20} + 1$ or $n_{20} = n_{10} + 1$.

$$\begin{aligned}
&P(n_{11}, n_{10}, n_{21}, n_{20} | \, \pi_1, \varphi_1, \varphi_2) \\
&\quad = \pi_1 p_{\text{Bin}}(n_{21}; n_{21} + n_{20}, \varphi_2) p_{\text{NBin}}(n_{11}; n_{10}, \varphi_1) \;\text{ if } n_{10} = n_{20} + 1 \\
&\quad = \pi_2 p_{\text{Bin}}(n_{11}; n_{11} + n_{10}, \varphi_1) p_{\text{NBin}}(n_{21}; n_{20}, \varphi_2) \;\text{ if } n_{20} = n_{10} + 1 \\
&\quad = \pi_1 p_{\text{Bin}}(n_{11}; n_{11} + n_{10}, \varphi_1) p_{\text{NBin}}(n_{21}; n_{20}, \varphi_2) \\
&\quad + \pi_2 p_{\text{Bin}}(n_{21}; n_{21} + n_{20}, \varphi_2) p_{\text{NBin}}(n_{11}; n_{10}, \varphi_1) + + \;\text{ if } n_{10} = n_{20},
\end{aligned}$$

where, for $t = 1, 2$,

$$\begin{aligned}
p_{\text{Bin}}(n_{t1}; n_{t1} + n_{t0}, \varphi_t) &= \tbinom{n_{t1}+n_{t0}}{n_{t1}} \varphi_t^{n_{t1}} (1 - \varphi_t)^{n_{t0}} \\
p_{\text{NBin}}(n_{t1}; n_{t0}, \varphi_t) &= \tbinom{n_{t1}+n_{t0}-1}{n_{t1}} \varphi_t^{n_{t1}} (1 - \varphi_t)^{n_{t0}} \\
&= \frac{n_{t0}}{n_{t1} + n_{t0}} \, p_{\text{Bin}}(n_{t1}; n_{t1} + n_{t0}, \varphi_t), \qquad (1)
\end{aligned}$$

with the conventions

$0! = 1, \tbinom{N}{0} = 1$ if $N \geq 0, \tbinom{N}{k} = 0$ if $k > N, \; \frac{0}{0} = 1$.

Stopping rule R2

The experiment stops after a fixed number of failures n_{20} has been observed for treatment 2. Recall that $n_{10} = n_{20}$ or $n_{20} = n_{10} + 1$.

$$P(n_{11}, n_{10}, n_{21}, n_{20} | \, \pi_1, \varphi_1, \varphi_2) = \pi_t p_{\text{NBin}}(n_{11}; n_{10}, \varphi_1) p_{\text{NBin}}(n_{21}; n_{20}, \varphi_2) \quad (2)$$

where $t = 1$ if $n_{10} = n_{20}$ and $t = 2$ if $n_{20} = n_{10} + 1$.

Proof

The sampling distribution involves the standard combinatory problem of distributing N nondistinguishable objects to k numbered drawers, hence the number of possible filling patterns $\tbinom{N+k-1}{N}$.

Stopping rule R1. For each treatment t ($t = 1, 2$), for given numbers n_{t1} and n_{t0}, there are n_{t0} ordered drawers of type t that contain exactly one failure and from zero to n_{t1} successes. In addition, there is a final drawer, possibly empty, that contains no failure and contains also from zero to n_{t1} successes. Hence, for each treatment t, there are $N = n_{t1}$ objects to distribute, either to $k = n_{t0} + 1$ drawers if the final drawer is of type t, or to $k = n_{t0}$ drawers if the final drawer is of the other type. The number of corresponding samples is the binomial coefficient $\binom{n_{t1}+n_{t0}}{n_{t1}}$ in the first case, and the negative binomial coefficient $\binom{n_{t1}+n_{t0}-1}{n_{t1}}$ in the second case. If the first drawer is of type 1 (with probability π_1), this implies that either $n_{10} = n_{20}+1$ (the last, with no success, drawer being of type 2) or $n_{10} = n_{20}$ (the last drawer being of type 1). If the first drawer is of type 2 (with probability π_2), this implies that either $n_{20} = n_{10} + 1$ (the last drawer being of type 1) or $n_{10} = n_{20}$ (the last drawer being of type 2).

Note that, for this stopping rule, various marginal and conditional probabilities, as well as factorial moments derived from recurrence relations, were given in ElQasyr (2008).

Stopping rule R2. In this case, there is no additional drawer, the final drawer being of type 2 by definition. Hence, for each treatment t, there are $N = n_{t1}$ objects to distribute to $k = n_{t0}$ drawers. The number of corresponding samples is the negative binomial coefficient $\binom{n_{t1}+n_{t0}-1}{n_{t1}}$ The fact that the final drawer is of type 2 implies that either $n_{10} = n_{20}$ if the first drawer is of type 1 (with probability π_1) or $n_{20} = n_{10}+1$ if the first drawer is of type 2 (with probability π_2).

3 Conditional sampling distributions

Conditional procedures are obtained from the conditional sampling distribution of n_{11}. Owing to the strong dependence between the numbers of failures for the two treatments, it appears appropriate to condition on n_{10} and n_{20}, rather than on fixed margins ($n_{.1} = n_{11} + n_{21}$ is not fixed).

Stopping rule R1

The conditional probability is derived from the marginal probability (1), replacing n_{21} with $n - n_{11} - n_{10} - n_{20}$. It involves only the ratio φ_1/φ_2. We get after simplification (the notation $[n]$ recall that n is fixed by the stopping rule):

$$P(n_{11}|\,[n], n_{.0}, \pi_1, \varphi_1, \varphi_2) = P(n_{11}|\,[n], n_{10}, n_{20}, \rho = \varphi_1/\varphi_2)$$
$$= \frac{P(n_{11}, n_{10}, n - n_{11} - n_{10} - n_{20}, n_{20})}{\sum_{j=0}^{n-n_{10}-n_{20}} P(j, n_{10}, n - j - n_{10} - n_{20}, n_{20})}$$

$$= \frac{Q_1(n_{11}, n_{10}, n_{20}, n)\rho^{n_{11}}}{\sum_{j=0}^{n-n_{10}-n_{20}} Q_1(j, n_{10}, n_{20}, n)\rho^j} \qquad (0 \le n_{11} \le n - n_{10} - n_{20}),$$

where $Q_1(j, n_{10}, n_{20}, n)$

$$= \tbinom{j+n_{10}-1}{j}\tbinom{n-j-n_{10}}{n-j-n_{10}-n_{20}} \text{ if } n_{10} = n_{20}+1$$
$$= \tbinom{j+n_{10}}{j}\tbinom{n-j-n_{10}-1}{n-j-n_{10}-n_{20}-1} \text{ if } n_{20} = n_{10}+1$$
$$= \pi_1\tbinom{j+n_{10}}{j}\tbinom{n-j-n_{10}-1}{n-j-n_{10}-n_{20}-1} + \pi_2\tbinom{j+n_{10}-1}{j}\tbinom{n-j-n_{10}}{n-j-n_{10}-n_{20}} \text{ if } n_{10} = n_{20} \quad (3)$$

In the particular case $\rho = 1$ ($\varphi_1 = \varphi_2$), the conditional distribution of n_{11} is a Beta-Binomial distribution when $n_{10} \neq n_{20}$, and is a mixture of Beta-Binomial distributions when $n_{10} = n_{20}$:

$$P(n_{11}|\,[n], n_{10}, n_{20}, \varphi_1 = \varphi_2)$$
$$= p_{\text{B-Bin}}(n_{11}; n - n_{10} - n_{20}, n_{10}, n_{20}+1) \text{ if } n_{10} = n_{20}+1$$
$$= p_{\text{B-Bin}}(n_{11}; n - n_{10} - n_{20}, n_{10}+1, n_{20}) \text{ if } n_{20} = n_{10}+1$$
$$= \pi_1 p_{\text{B-Bin}}(n_{11}; n - n_{10} - n_{20}, n_{10}+1, n_{20})$$
$$+ \pi_2 p_{\text{B-Bin}}(n_{11}; n - n_{10} - n_{20}, n_{10}, n_{20}+1) \text{ if } n_{10} = n_{20},$$

where $p_{\text{B-Bin}}(x; N, a, b)$ is the probability function of the Beta-Binomial distribution with parameters N, a, b:

$$p_{\text{B-Bin}}(x; N, a, b) = \tbinom{N}{x}\frac{B(x+a, N-x+b)}{B(a,b)} = \frac{\tbinom{x+a-1}{x}\tbinom{N-j+b-1}{N-j}}{\tbinom{N+a+b-1}{N}}.$$

Stopping rule R2

In this case, the distribution is also conditioned on the observed n. The stopping rule implies that either $n_{10} = n_{20}$ ($t^1 = 1$) or $n_{20} = n_{10}+1$ ($t^1 = 2$). From (2), we get in the same way ($[n_{20}]$ recall that n_{20} is fixed by the stopping rule):

$$P(n_{11}|\,n, n_{.0}, [n_{20}], \pi_1, \varphi_1, \varphi_2) = P(n_{11}|\,n, n_{10}, [n_{20}], \rho)$$
$$= \frac{Q_2(n_{11}, n_{10}, n_{20}, n)\rho^{n_{11}}}{\sum_{j=0}^{n-n_{10}-n_{20}} Q_2(j, n_{10}, n_{20}, n)\rho^j} \qquad (0 \le n_{11} \le n - n_{10} - n_{20}), \qquad (4)$$

where $Q_2(j, n_{10}, n_{20}, n) = \tbinom{j+n_{10}-1}{j}\tbinom{n-j-n_{10}-1}{n-j-n_{10}-n_{20}}$.

In the particular case $\rho = 1$, it is a Beta-Binomial distribution with parameters $n - n_{10} - n_{20}$, n_{10}, and n_{20}:

$$P(n_{11}|\, n, n_{10}, [n_{20}], \varphi_1 = \varphi_2) = p_{\text{B-Bin}}(n_{11}; n - n_{10} - n_{20}, n_{10}, n_{20})$$

4 Application: Conditional tests and intervals

Using the appropriate conditional distribution, the null hypothesis $\rho = \rho_0$ can be tested against the alternative $\rho > \rho_0$ by declaring the result significant at level α if

$$\overline{p}_{inc}^{\rho_0} = \sum_{j=n_{11}}^{n-n_{10}-n_{20}} P(n_{11} = j \,|\, t^1, n, n_{10}, n_{20}, \rho_0) \leq \alpha.$$

In the same way, $\rho = \rho_0$ can be tested against the alternative $\rho < \rho_0$, by declaring the result significant at level α if the observed level

$$\underline{p}_{inc}^{\rho_0} = \sum_{j=0}^{n_{11}} P(n_{11} = j \,|\, t^1, n, n_{10}, n_{20}, \rho_0) \leq \alpha.$$

The summation is over all values of j, consistent with the fixed numbers of failures, which are "more extreme" than the observed frequency. The observed frequency n_{11} is included in the summation, hence the subscript *inc* for *inclusive* test. Since this inclusive test can be highly conservative, an alternative solution would be to exclude n_{11}, hence the subscript *exc* for *exclusive* test. Let $\overline{p}_{exc}^{\rho_0} = 1 - \underline{p}_{inc}^{\rho_0}$ and $\underline{p}_{exc}^{\rho_0} = 1 - \overline{p}_{inc}^{\rho_0}$ be the corresponding probabilities. This exclusive test is anti-conservative. Consequently, a solution which gives probabilities typically closer to the nominal level than "inclusive" or "exclusive" approaches consists in considering a mid-p-value, defined as the mean of the inclusive and exclusive p-values (Berry and Armitage, 1995; Lecoutre, 2008; Routledge, 1994)

$$\overline{p}_{mid}^{\rho_0} = (\overline{p}_{inc}^{\rho_0} + \overline{p}_{exc}^{\rho_0})/2 \qquad \text{and} \qquad \underline{p}_{mid}^{\rho_0} = (\underline{p}_{inc}^{\rho_0} + \underline{p}_{exc}^{\rho_0})/2.$$

In the particular case $\rho_0 = 1$ ($\varphi_1 = \varphi_2$), the conditional test is the Fisher's randomization test. Its p-values are given by the corresponding Beta-Binomial distribution. More generally, for $\varphi_1 \neq \varphi_2$, they are given by (3) or (4), which are generalized Beta-Binomial distributions.

For each conditional test, a lower and an upper confidence limits for ρ can be found by solving respectively $\overline{p}^{\rho_0} = \alpha$ and $\underline{p}^{\rho_0} = \alpha$.

5 Application: Bayesian predictive probabilities

The likelihood function is proportional to $\varphi_1^{n_{11}}(1-\varphi_1)^{n_{10}}\varphi_2^{n_{21}}(1-\varphi_2)^{n_{20}}$, hence identical (up to a multiplicative constant) with the likelihood function associated with the comparison of two independent binomial (or negative binomial) proportions. A simple and usual Bayesian solution assumes two independent Beta prior distributions for φ_1 and φ_2: respectively Beta(ν_{11}, ν_{10}) and Beta(ν_{21}, ν_{20}) (Lecoutre, Derzko and Grouin, 1995).

5.1 Posterior and predictive distributions

This is a conjugate prior and the marginal posterior distributions 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 probability of observing $(n_{11}, n_{10}, n_{21}, n_{20})$ is a mixture of the sampling probabilities. Consequently, it can be expressed from Beta-Binomial and Beta-Negative-Binomial distributions.

Stopping rule R1

$$\begin{aligned}
&P(n_{11}, n_{10}, n_{21}, n_{20})\\
&= \pi_1 p_{\text{B-Bin}}(n_{21}; n_{21}+n_{20}, \nu_{21}, \nu_{20})\, p_{\text{B-NBin}}(n_{11}; n_{10}, \nu_{11}, \nu_{10}) \quad \text{if } n_{10} = n_{20}+1\\
&= \pi_2 p_{\text{B-Bin}}(n_{11}; n_{11}+n_{10}, \nu_{11}, \nu_{10})\, p_{\text{B-NBin}}(n_{21}; n_{20}, \nu_{21}, \nu_{20}) \quad \text{if } n_{20} = n_{10}+1\\
&= \pi_1 p_{\text{B-Bin}}(n_{11}; n_{11}+n_{10}, \nu_{11}, \nu_{10})\, p_{\text{B-NBin}}(n_{21}; n_{20}, \nu_{21}, \nu_{20})\\
&+ \pi_2 p_{\text{B-Bin}}(n_{21}; n_{21}+n_{20}, \nu_{21}, \nu_{20})\, p_{\text{B-NBin}}(n_{11}; n_{10}, \nu_{11}, \nu_{10}) \quad \text{if } n_{10} = n_{20},
\end{aligned}$$

where $p_{\text{B-NBin}}(j; r, a, b)$ is the probability function of the Beta-Negative-Binomial distribution of parameters r, a and b

$$\begin{aligned}
p_{\text{B-NBin}}(j; r, a, b) &= \tbinom{j+r-1}{j}\frac{B(j+a, r+b)}{B(a,b)} = \frac{\binom{j+a-1}{r}\binom{r+b-1}{r}}{\binom{j+r+a+b-1}{j+r}}\\
&= \frac{r}{j+r}\, p_{\text{B-Bin}}(j; j+r, a, b).
\end{aligned}$$

Stopping rule R2

$$\begin{aligned}
&P(n_{11}, n_{10}, n_{21}, n_{20})\\
&\quad = \pi_t p_{\text{B-NBin}}(n_{11}; n_{10}, \nu_{11}, \nu_{10}) p_{\text{B-NBin}}(n_{21}; n_{20}, \nu_{21}, \nu_{20})
\end{aligned}$$

where $t = 1$ if $n_{10} = n_{20}$ and $t = 2$ if $n_{20} = n_{10}+1$.

The predictive probability of observing $(n'_{11},\ n'_{10},\ n'_{21},\ n'_{20})$ in a future independent sample of size n' can be obtained in the same way, replacing the prior Beta distributions with the posterior Beta distributions.

5.2 *Correspondence with the p-values of the conditional tests*

There is a correspondence between the above exclusive and inclusive p-values of the conditional tests for $\rho_0 = 0$ and the posterior Bayesian probabilities that $\varphi_1 < \varphi_2$ associated with particular choices of the prior. We have the following equalities.

Stopping rule R1

$$
\begin{aligned}
\overline{p}^{1}_{inc} &= 1 - \underline{p}^{1}_{exc} \\
&= P_{0,0,1,1}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \ \text{ if } n_{10} = n_{20} + 1 \\
&= P_{0,1,1,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}). \ \text{ if } n_{20} = n_{10} + 1 \\
&= \pi_1\, P_{0,1,1,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \\
&+ \pi_2\, P_{0,0,1,1}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \ \text{ if } n_{10} = n_{20}.
\end{aligned}
$$

$$
\begin{aligned}
\overline{p}^{1}_{exc} &= 1 - \underline{p}^{1}_{inc} \\
&= P_{1,0,0,1}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \ \text{ if } n_{10} = n_{20} + 1 \\
&= P_{1,1,0,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \ \text{ if } n_{20} = n_{10} + 1 \\
&= \pi_1\, P_{1,1,0,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \\
&+ \pi_2\, P_{1,0,0,1}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \ \text{ if } n_{10} = n_{20}.
\end{aligned}
$$

where $P_{\nu_{11},\nu_{10},\nu_{21},\nu_{20}}(\varphi_1 < \varphi_2) \,|\, n_{11}, n_{10}, n_{21}, n_{20})$ is the posterior probability associated with the prior defined by $(\nu_{11},\ \nu_{10},\ \nu_{21},\ \nu_{20})$. This extends the previous correspondence obtained for the Binomial and Negative Binomial sampling (Lecoutre, 2008) and for the multinomial sampling (Altham, 1969; **?**).

Stopping rule R2

$$
\begin{aligned}
\overline{p}^{1}_{inc} &= 1 - \underline{p}^{1}_{exc} = P_{0,0,1,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \\
\overline{p}^{1}_{exc} &= 1 - \underline{p}^{1}_{inc} = P_{1,0,0,0}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20})
\end{aligned}
$$

Proof

Given φ_2, the conditional posterior probability $P(\varphi_1 < \varphi_2 \,|\, \varphi_2)$ (omitting the references to the prior weights and to the data counts) is the incomplete Beta function $I_{\varphi_2}(n_{11} + \nu_{11}, n_{10} + \nu_{10})$. Consequently, if ν_{11}, ν_{10}, ν_{21} and ν_{20} are non-null integers, it is equal to the probability that a Binomial distribution with parameters $n_{11}+n_{10}+\nu_{11}+\nu_{10}-1$ and φ_2 is greater or equal to $n_{11}+\nu_{11}$ (Johnson, Kemp and Kotz, 1993, p. 117). From the marginal Beta distribution of φ_2, it results that the posterior probability $P(\varphi_1 < \varphi_2)$ is the probability that a Beta-Binomial distribution with parameters $n_{11} + n_{10} + \nu_{11} + \nu_{10} - 1$, $n_{21} + \nu_{21}$ and $n_{20} + \nu_{20}$ is greater or equal to $n_{11} + \nu_{11}$. Hence,

$$\begin{aligned} &P_{\nu_{11},\nu_{10},\nu_{21},\nu_{20}}(\varphi_1 < \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) \\ &= 1 - F_{\text{B-Bin}}(n_{11} + \nu_{11} - 1; n_{11} + n_{10} + \nu_{11} + \nu_{10} - 1, n_{21} + \nu_{21}, n_{20} + \nu_{20}), \end{aligned}$$

where $F_{\text{B-Bin}}(x; N, a, b)$ is the distribution function of the Beta-Binomial distribution with parameters N, a and b.

$$F_{\text{B-Bin}}(x; N, a, b) = \sum_{j=0}^{x} p_{\text{B-Bin}}(j; N, a, b).$$

The correspondence with the p-values of the conditional tests can be deduced from the fact that

$$F_{\text{B-Bin}}(x; x + c, a, b) = F_{\text{B-Bin}}(x; x + a, c, b),$$

This equality results from the link with the distribution function of the Hypergeometric distribution (Johnson, Kemp and Kotz, 1993, p. 254)

$$\begin{aligned} F_{\text{B-Bin}}(x; x + c, a, b) &= F_{\text{HG}}(x; x + a + b + c - 1, x + a, x + c) \\ F_{\text{B-Bin}}(x; x + a, c, b) &= F_{\text{HG}}(x; x + a + b + c - 1, x + c, x + a), \end{aligned}$$

using the fact that

$$F_{\text{HG}}(x; N, K, n) = F_{\text{HG}}(x; N, n, K).$$

It follows that

$$\begin{aligned} &F_{\text{B-Bin}}(n_{11} + \nu_{11}; n_{11} + n_{10} + \nu_{11} + \nu_{10} - 1, n_{21} + \nu_{21}, n_{20} + \nu_{20}) \\ &\quad = F_{\text{B-Bin}}(n_{11} + \nu_{11}; n - n_{10} - n_{20} + \nu_{11} + \nu_{21} - 1, n_{10} + \nu_{10}, n_{20} + \nu_{20}), \end{aligned}$$

from which we can easily deduce the appropriate priors.

Let us consider for illustration the results of an experiment with a fixed number of $n = 150$ subjects. The first treatment is randomly selected ($\pi_1 = 1/2$). The observed rates of success are respectively 68 out of 90 attributions for treatment 1 and 38 out of 60 attributions for treatment 2 (22 failures in each case). The p-values of the conditional test of the null hypothesis $\varphi_1 = \varphi_2$ against the alternative $\varphi_1 > \varphi_2$ are

$$\begin{aligned}
\overline{p}^1_{inc} &= \frac{1}{2}\Big(1 - F_{\text{B-Bin}}(67; 106, 23, 22)\Big) + \frac{1}{2}\Big(1 - F_{\text{B-Bin}}(67; 106, 22, 23)\Big) \\
&= \frac{1}{2}P_{0,1,1,0}(\varphi_1 < \varphi_2 \,|\, 68, 22, 38, 22) + \frac{1}{2}P_{0,0,1,1}(\varphi_1 < \varphi_2 \,|\, 68, 22, 38, 22) \\
&= (.0772 + .0468)/2 = 0.0620 \\
\overline{p}^1_{exc} &= \frac{1}{2}\Big(1 - F_{\text{B-Bin}}(68; 106, 23, 22)\Big) + \frac{1}{2}\Big(1 - F_{\text{B-Bin}}(68; 106, 22, 23)\Big) \\
&= \frac{1}{2}P_{1,1,0,0}(\varphi_1 < \varphi_2 \,|\, 68, 22, 38, 22) + \frac{1}{2}P_{1,0,0,1}(\varphi_1 < \varphi_2 \,|\, 68, 22, 38, 22) \\
&= (.0625 + .0370)/2 = 0.0497 \\
\overline{p}^1_{mid} &= (.0620 + .0497)/2 = .0559.
\end{aligned}$$

A simple reasonnable solution for an objective prior is to take the average of the prior weights of the four involved different Beta priors, hence the two marginal independent priors Beta$(1/2, 1/2)$. Note that it is the Jeffreys prior for the 1:1 randomized design. We get the intermediate probability $P_{1/2,1/2,1/2,1/2}(\varphi_1 < \varphi_2 \,|\, 68, 22, 38, 22) = .0542$, close to the mid-$p$-value. The 90% equal-tailed Bayesian credible interval for $\rho = \varphi_1/\varphi_2$ is $[.996, 1.457]$, close to the conditional mid-p confidence interval [0.989,1.449].

References

Altham, P.M.E., 1969. Exact Bayesian analysis of a 2×2 contingency table and Fisher's "exact"significance test. Journal of the Royal Statistical Society B, 31, 261-269.

Berry, G., Armitage, P., 1995 Mid-P confidence intervals: A brief review. The Statistician, 44, 417-423.

ElQasyr, K., 2008. Modélisation et Analyse Statistique des Plans d'Expérience Séquentiels. Unpublished doctoral thesis, Université de Rouen.

Johnson, N.L., Kotz, S. and Kemp, A.W., 1993. Univariate Discrete Distributions (2nd Edition). New York: John Wiley.

Lecoutre, B., 2008. Bayesian methods for experimental data analysis. In C.R.

Rao, J. Miller & D.C. Rao (Eds.), Handbook of statistics: Epidemiology and Medical Statistics (Vol 27), Amsterdam: Elsevier, 775-812.

Lecoutre, B., Derzko G., ElQasyr, K., 2009. Frequentist performance of Bayesian inference with response-adaptive designs. Submitted for publication.

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. Adaptative designs for multi-arm clinical trials: The play-the-winner rule revisited. Communications in Statistics - Simulation and Computation, 37, 590-601.

Pullman, D., Wang, X., 2001. Adaptive designs, informed consent, and the ethics of research. Controlled Clinical Trials, 22, 203-210.

Rosenberger, W.F., Hu, F., 2004. Maximizing power and minimizing treatment failures. Clinical Trials, 1, 141-147.

Routledge, R.D., 1994. Practicing safe statistics with the mid-p*. The Canadian Journal of Statistics, 22, 103-110.

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