Frequentist performance of Bayesian inference with response-adaptive designs

Bruno Lecoutre[1], Gérard Derzko[2] & Khadija ElQasyr[1]

[1]ERIS, UMR 6085, CNRS & Université de Rouen

[2]sanofi-aventis R&D

ISCB 30 - Prague 23-27 August 2009

Outline

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

- Comparing the success rates φ_1 and φ_2 of two treatments with a dichotomous (e.g. success/failure) outcome
- For comparing two treatment success rates, Response-adaptive designs are attractive competitors to 1:1 randomized designs (Pullman & Wang, 2001)
- Play-the-Winner (PW), Randomized Play-the-Winner (RPW) and Drop-the-Loser (DL) rules are considered (for two treatments and immediate outcomes)
- Frequentist inference: the complexity of sampling distributions is a challenge, solved for PW.
- Bayesian inference: how does it perform from a frequentist point-of-view (coverage probabilities of interval estimation procedures)?
- Response-adaptive designs vs 1:1 randomized designs (RD) : what about the loss of power?

Outline

1. Introduction
2. Response-adaptive designs for two treatments
    - Play-the-Winner (Zelen, 1969)
    - Randomized Play-the-Winner (Wei & Durham, 1978)
    - Drop-the-Loser (Ivanova, 2003)
    - Comparisons of the three rules
3. Inferential procedures
    - Frequentist procedures
    - Bayesian procedures
    - Coverage properties
    - Power properties
4. Conclusion
    - Concluding remarks
    - References

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.
- In this presentation the stopping rule is based on the total number of subjects N.

PW dichotomous process

PW rule

If the $(n-1)$-th subject is assigned treatment t:

- If t is a success, the n-th subject is assigned treatment t
- In t is a failure, the n-th subject is assigned 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.
- It is usually believed that "less deterministic" rules should be prefered.

Generalized Friedman's Urn (GFU) model

RPW rule

- Before the *n*-th subject comes in, the urn contains (Y^1_{n-1}, Y^2_{n-1}) balls ((Y^1_0, Y^2_0) initially) that represent either treatments
- A ball is drawn at random and replaced. The subject is assigned the corresponding treatment (say t).
- Balls are added to the urn:
  - If t is a success, u type t balls and v balls of the other type
  - If t is a failure, v type t balls and u balls of the other type

- Whatever the process history, the total number of balls at step n is equal to $Y^1_n + Y^2_n = Y^1_0 + Y^2_0 + n(u+v)$
- RPW process converges more slowly than PW process, sometimes notably so (Lecoutre & ElQasyr, 2008)

Generalized Friedman's Urn (GFU) model

Drop-the-Loser (DL) rule

- Extra "no treatment" balls ("immigration balls") are included.
- Initial urn composition: (Y_0^1, Y_0^2, Z_0)
- if an immigration ball is drawn no subject is treated and the ball is returned to the urn together with one ball of each treatment type
- If a treatment t ball is drawn:
  - If t is a success, the ball is replaced (unchanged composition)
  - If t is a failure, the ball is not replaced (the number of balls is decreased by 1)

DL allocation proportion is less variable than PW's, with power improvement (Ivanova, 2003) but DL process converges more slowly than PW process, sometimes notably so (ElQazyr, 2008).

Expected number of subjects assigned to the better treatment

RPW $Y_0^1 = Y_0^2 = 1 \ u = 1 \ v = 0$ DL $Y_0^1 = Y_0^2 = 3 \ Z_0 = 1$

Probabilities of success: $\varphi_1 = 0.80$ and $\varphi_2 = 0.60$

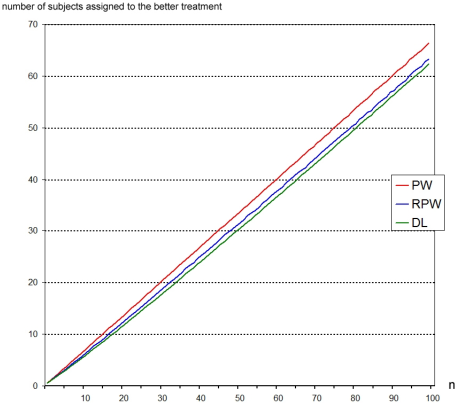

Number of subjects assigned to the better treatment

Expectation and standard deviation (*in italics*).

		$N = 50$				$N \to \infty$
φ_1	φ_2	PW	RPW	DL	GDL	
0.30	0.10	28.1 (56.2%) *1.8*	27.9 (55.9%) *3.1*	27.8 (55.7%) *1.8*	27.5 (54.9%) *2.0*	56.2%
0.40	0.20	28.5 (57.0%) *2.3*	28.3 (56.6%) *3.7*	28.2 (56.3%) *2.2*	27.8 (55.5% *2.5*	57.1%
0.70	0.30	34.8 (69.6%) *3.2*	33.6 (67.2%) *4.9*	32.8 (65.6%) *2.8*	32.4 (64.9%) *3.2*	70.0%
0.80	0.60	33.1 (66.1%) *5.0*	30.9 (61.8%) *7.5*	29.8 (59.7%) *3.7*	30.6 (61.3% *4.7*	66.7%
0.90	0.70	36.9 (73.8)% *6.1*	32.1 (64.1%) *9.1*	30.1 (60.2%) *3.8*	32.7 (65.4%) *5.2*	75%
		exact	exact expect.	10^5 replications	10^5 replications	

Note : GDL is Generalized Drop-the-Loser (Sun, Cheung & Zhang, 2007)

Number of subjects assigned to the better treatment

Expectation and standard deviation (*in italics*).

		$N = 50$				$N \to \infty$
φ_1	φ_2	PW	RPW	DL	GDL	
0.30	0.10	28.1 (56.2%) *1.8*	27.9 (55.9%) *3.1*	27.8 (55.7%) *1.8*	27.5 (54.9%) *2.0*	56.2%
0.40	0.20	28.5 (57.0%) *2.3*	28.3 (56.6%) *3.7*	28.2 (56.3%) *2.2*	27.8 (55.5% *2.5*	57.1%
0.70	0.30	34.8 (69.6%) *3.2*	33.6 (67.2%) *4.9*	32.8 (65.6%) *2.8*	32.4 (64.9%) *3.2*	70.0%
0.80	0.60	33.1 (66.1%) *5.0*	30.9 (61.8%) *7.5*	29.8 (59.7%) *3.7*	30.6 (61.3% *4.7*	66.7%
0.90	0.70	36.9 (73.8)% *6.1*	32.1 (64.1%) *9.1*	30.1 (60.2%) *3.8*	32.7 (65.4%) *5.2*	75%
		exact	exact expect.	10^5 replications	10^5 replications	

Note : GDL is Generalized Drop-the-Loser (Sun, Cheung & Zhang, 2007)

Number of subjects assigned to the better treatment

Expectation and standard deviation (*in italics*).

$N = 50$

φ_1	φ_2	PW	RPW	DL	GDL	$N \to \infty$
0.30	0.10	28.1 (56.2%) *1.8*	27.9 (55.9%) *3.1*	27.8 (55.7%) *1.8*	27.5 (54.9%) *2.0*	56.2%
0.40	0.20	28.5 (57.0%) *2.3*	28.3 (56.6%) *3.7*	28.2 (56.3%) *2.2*	27.8 (55.5% *2.5*	57.1%
0.70	0.30	34.8 (69.6%) *3.2*	33.6 (67.2%) *4.9*	32.8 (65.6%) *2.8*	32.4 (64.9%) *3.2*	70.0%
0.80	0.60	33.1 (66.1%) *5.0*	30.9 (61.8%) *7.5*	29.8 (59.7%) *3.7*	30.6 (61.3% *4.7*	66.7%
0.90	0.70	36.9 (73.8)% *6.1*	32.1 (64.1%) *9.1*	30.1 (60.2%) *3.8*	32.7 (65.4%) *5.2*	75%
		exact	exact expect.	10^5 replications	10^5 replications	

Note : GDL is Generalized Drop-the-Loser (Sun, Cheung & Zhang, 2007)

Outline

1. Introduction
2. Response-adaptive designs for two treatments
    - Play-the-Winner (Zelen, 1969)
    - Randomized Play-the-Winner (Wei & Durham, 1978)
    - Drop-the-Loser (Ivanova, 2003)
    - Comparisons of the three rules
3. Inferential procedures
    - Frequentist procedures
    - Bayesian procedures
    - Coverage properties
    - Power properties
4. Conclusion
    - Concluding remarks
    - References

Some frequentist inference procedures

- Permutation test for the RPW design (Wei, 1988)
  - No explicit formula
  - computationally intensive algorithm needed
- confidence intervals for the odds ratio for the RPW design (Wei et al., 1990):
  - Exact conditional CI : "rather conservative and not very useful in practice"
  - An unconditional procedure: "less powerful"
  - Large-sample CIs: "not satisfactory"
- Conditional confidence intervals for the ratio for the PW design (ElQasyr & Lecoutre, 2009)
  - Explicit formulae, valuable solution, easy to compute
  - Mid-p approach, overcome the conservativeness of exact CIs
- However, frequentist methods require ad hoc developments for other criteria or other designs

Bayesian inference

- The likelihood function is proportional to

$$\varphi_1^{n_{11}}(1-\varphi_1)^{n_{10}}\varphi_2^{n_{21}}(1-\varphi_2)^{n_{20}}$$

  i.e. proportional to 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 priors for φ_1 and φ_2

$$\text{Beta}(\nu_{11},\nu_{10}) \text{ and } \text{Beta}(\nu_{21},\nu_{20})$$

- hence the marginal independent Beta posteriors

$$\text{Beta}(\nu_{11}+n_{11},\nu_{10}+n_{10}) \text{ and } \text{Beta}(\nu_{21}+n_{21},\nu_{20}+n_{20})$$

Bayesian approach for PW : Predictive distribution

predictive probability $P(n_{11}, n_{10}, n_{21}, n_{20})$

$$\begin{aligned}&\tfrac{1}{2} p_{\text{B-Bin}}(n_{21}; n_{21}+n_{20}, \nu_{21}, \nu_{20})\, p_{\text{B-NBin}}(n_{11}; n_{10}, \nu_{11}, \nu_{10})\\ &\qquad\qquad \times \mathbb{1}_{\{0,1\}}(n_{10}-n_{20})\\ &+\tfrac{1}{2} p_{\text{B-Bin}}(n_{11}; n_{11}+n_{10}, \nu_{11}, \nu_{10})\, p_{\text{B-NBin}}(n_{21}; n_{20}, \nu_{21}, \nu_{20})\\ &\qquad\qquad \times \mathbb{1}_{\{0,1\}}(n_{20}-n_{10})\end{aligned}$$

(when the first treatment is randomly selected with probability 0.5)

expressed from Beta-Binomial and Beta-Negative-Binomial distributions:

$$p_{\text{B-Bin}}(j; r, a, b) = \binom{r}{j}\frac{B(j+a, r-j+b)}{B(a,b)} = \frac{\binom{j+a-1}{j}\binom{r-j+b-1}{r-j}}{\binom{r+a+b-1}{r}}$$

$$p_{\text{B-NBin}}(j; r, a, b) = \binom{j+r-1}{j}\frac{B(j+a, r+b)}{B(a,b)} = \frac{\binom{j+a-1}{r}\binom{r+b-1}{r}}{\binom{j+r+a+b-1}{j+r}}$$

Note : with the PW rule n_{10} and n_{20} differ by at most one unit

PW : Numerical illustration of the Bayesian inference

Example

- Assume a fixed number of $N = 150$ subjects, and observed success rates :
    - 68 out of 90 attributions for treatment 1
    - 38 out of 60 attributions for treatment 2 (22 failures in each case)
- A simple reasonnable solution for an objective prior is to take two marginal independent priors Beta$(0.5, 0.5)$ (i.e. the Jeffreys prior for the 1:1 randomized design)
- Resulting 90% equal-tailed Bayesian credible intervals :
    - $[.996, 1.457]$ for φ_1/φ_2 (vs conditional mid-p CI [0.989,1.449])
    - $[-0.003, 0.247]$ for $\varphi_1 - \varphi_2$
    - $[.986, 3.255]$ for $(\varphi_1/1 - \varphi_1)/(\varphi_2/1 - \varphi_2)$

φ_1/φ_2 : Bayesian vs frequentist coverage error with PW

Two Beta$(0.5, 0.5)$ independent priors.
φ_1^* and φ_2^*: "true" values of the parameters.
90% credible interval for φ_1/φ_2
Sampling error probability for the CI lower limit :

$N = 50$

			φ_1^*		
	.10	.30	.50	.70	.90
.10	.055 .037	.067 .037	$<.001$ $<.001$	$<.001$ $<.001$	$<.001$ $<.001$
.30	.057 .047	.051 .047	.054 .047	.065 .053	.033 .014
φ_2^* .50	.052 .046	.054 .047	.051 .048	.051 .042	.060 .036
.70	.049 .048	.054 .042	.056 .047	.055 .050	.054 .056
.90	.050 .035	.041 .046	.046 .052	.051 .050	.068 .049

Bayes
Mid-*p*

Probability of error for the upper limit = probability for the lower limit associated with the reverse combination of parameters

$\varphi_1 - \varphi_2$: Bayesian vs frequentist coverage error with PW

Two Beta$(0.5, 0.5)$ independent priors.
φ_1^* and φ_2^*: “true” values of the parameters.
90% credible interval for $\varphi_1 - \varphi_2$.
Sampling error probability for the CI lower limit :

$N = 50$

φ_2^* \ φ_1^*	.10	.30	.50	.70	.90	
.10	.055 .073	.039 .054	.034 .076	.003 .129	0 .473	
.30	.061 .060	.051 .058	.050 .065	.053 .079	.016 .214	
.50	.063 .048	.053 .053	.052 .059	.050 .068	.048 .123	Bayes Wei et al
.70	.068 .036	.060 .043	.058 .054	.056 .047	.059 .062	
.90	.064 .020	.055 .040	.051 .054	.053 .062	.067 .027	

Probability of error for the upper limit = probability for the lower limit associated with the reverse combination of parameters

$\varphi_1 - \varphi_2$: coverage error of Bayesian CI with PW and DL

Two $\text{Beta}(0.5, 0.5)$ independent priors.

φ_1^* and φ_2^* : "true" values of the parameters.

90% credible interval for $\varphi_1 - \varphi_2$.

Sampling error probability for the CI lower limit :

$N = 50$

		φ_1^*					
		.10	.30	.50	.70	.90	
	.10	.055 .058	.039 .038	.034 .037	.003 .022	0 .002	
	.30	.061 .060	.051 .050	.050 .048	.053 .047	.016 .045	
φ_2^*	.50	.063 .067	.053 .055	.052 .054	.050 .052	.048 .052	PW DL
	.70	.068 .070	.060 .058	.058 .057	.056 .056	.059 .053	
	.90	.064 .073	.055 .054	.051 .054	.053 .058	.067 .062	

Probability of error for the upper limit = probability for the lower limit associated with the reverse combination of parameters

sampling probability of $P(\varphi_1 > \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20}) > 1-\alpha$

- Suppose that we aim to demonstrate that $\varphi_1 > \varphi_2$
- It can be concluded that $\varphi_1 > \varphi_2$ when the posterior probability $P(\varphi_1 > \varphi_2 \,|\, n_{11}, n_{10}, n_{21}, n_{20})$ is larger than a given $1-\alpha$
- The sampling probability of this event evaluates the power of the procedure
- Furthermore when $\varphi_2^* > \varphi_1^*$, the sampling probability evaluates the risk of concluding in the wrong direction

PW ($N = 50$) and 1:1 randomized design ($N = 25 + 25$)

φ_2^*	Bayes					mid-p					
			φ_1^*					φ_1^*			
	.10	.30	.50	.70	.90	.10	.30	.50	.70	.90	
.10	.055	.552	.947	.999	.997	.037	.555	.940	.999	.997	PW
	.056	.572	.950	.999	$>$.999	.026	.512	.942	.999	$>$.999	1:1 RD
.30		.051	.413	.886	.994		.047	.389	.873	.992	PW
		.050	.446	.915	.999		.045	.426	.891	.999	1:1 RD
.50		.001	.051	.417	.915		.001	.048	.388	.889	PW
		.001	.059	.446	.950		.001	.047	.426	.942	1:1 RD
.70			.001	.055	.532			.001	.050	.467	PW
			.001	.050	.572			.001	.045	.512	1:1 RD
.90					.068					.049	PW
					.056					.026	1:1 RD

Probabilities $< .0005$ omitted

Power : PW vs RD, Bayesian vs frequentist

- When $\varphi_1^* > \varphi_2^*$, the power is generally larger with RD than with PW, as expected
- However, the difference in power is always small
- Yao and Wei (1996) obtained similar conclusions with RD versus RPW
- When $\varphi_1^* < \varphi_2^*$, the PW design does not inflate the probability of a conclusion in the wrong direction.
- The Bayesian procedure appears to be more powerful than the conditional test. In other words, a similar power can be obtained by slightly increasing the sample size of the PW design.

Achieving 90% power with the Bayesian procedure

Number of subjects needed/number of subjects assigned to the less effective treatment

	$\tau^* = \varphi_1^*/\varphi_2^*$						
φ_2^*	1.5	2.0	2.5	3.0	3.5	4.0	
.10	1518/737.3	428/201.4	213/96.8	130/56.9	89/37.4	67/26.9	PW
	1502/728.6	427/201.0	213/97.0	130/57.1	90/38.0	67/27.1	RPW
	1506/731.5	426/200.6	211/96.1	130/57.2	88/36.9	66/26.9	DL
	1504/752	426/213	212/106	130/65	88/44	64/32	1:1 RD
.20	648/302.4	177/75.9	84/32.4	49/16.5	33/9.2	25/5.3	PW
	641/299.3	177/76.2	85/33.2	50/17.5	34/10.5	25/6.7	RPW
	639/298.4	176/75.8	83/32.7	49/17.6	32/10.7	20/6.6	DL
	642/321	176/88	82/41	48/24	32/16	20/10	1:1 RD
.30	359/158.0	93/33.9	42/11.3	28/4.0			PW
	355/156.6	93/34.7	43/13.0	25/5.9			RPW
	354/156.2	91/34.3	44/13.3	20/6.5			DL
	352/176	90/45	42/21	20/10			1:1 RD
.40	216/86.5	52/13.3					PW
	216/87.4	55/16.3					RPW
	207/83.9	49/16.1					DL
	206/103	48/24					1:1 RD
.50	129/43.2						PW
	133/47.0						RPW
	128/43.7						DL
	122/61						1:1 RD
.60	74/15.4						PW
	79/22.7						RPW
	66/22.0						DL
	64/32						1:1 RD

Outline

1. Introduction
2. Response-adaptive designs for two treatments
    - Play-the-Winner (Zelen, 1969)
    - Randomized Play-the-Winner (Wei & Durham, 1978)
    - Drop-the-Loser (Ivanova, 2003)
    - Comparisons of the three rules
3. Inferential procedures
    - Frequentist procedures
    - Bayesian procedures
    - Coverage properties
    - Power properties
4. Conclusion
    - Concluding remarks
    - References

Bayesian inference for PW in summary

- The Bayesian inference is efficient for PW designs
- The moderate loss of power renders these designs an attractive alternative to the conventional 1:1 randomized design, when it is desirable to minimize the number of subjects assigned to the less effective treatment
- Moreover, the predictive approach enables stopping the trial early, or on the contrary extending it to an adequate size, in a sequential perspective that fits with the methodology of adaptive designs (Lecoutre, Derzko & Grouin, 1995)
- Conditional mid-p CIs may be a valuable and easy-to-compute alternative, yet restricted to the ratio of proportions

Bayesian inference for adaptive designs & Practicals

- In contrast, the Bayesian approach allows obtaining the distributions of any derived parameter of interest from the joint posterior distribution using numerical methods : for instance, efficient inferences can be made similarly for proportion difference or odds-ratio (ElQasyr, 2008)
- Bayesian procedures appeared equally appropriate for other response-adaptive designs, such as the RPW and DL
- The posterior distribution can be approximated by simulating a large sample from two independent Beta distributions, or alternatively by numerical integration (Novick & Jackson, 1974), as implemented in a statistical computer program "LesProportions", freely available at address

http://www.univ-rouen.fr/LMRS/Persopage/Lecoutre/Eris.html

Computer program LesProportions

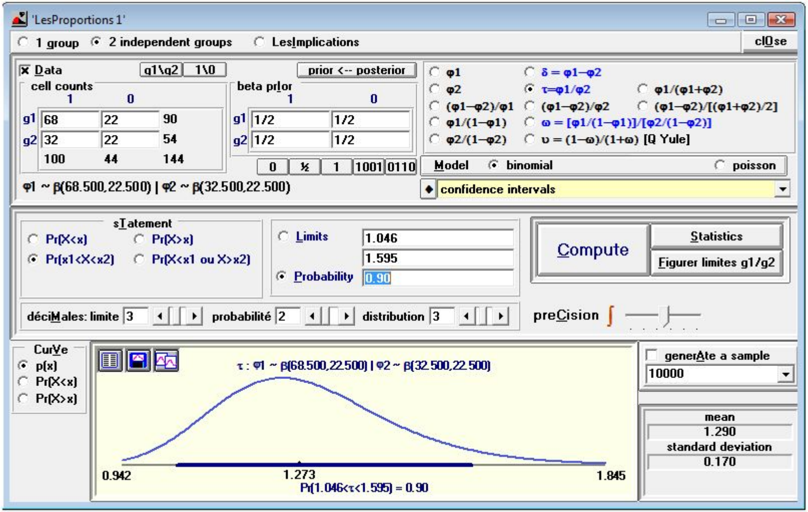
'LesProportions 1'

○ 1 group ● 2 independent groups ○ LesImplications

☒ Data

cell counts

	1	0	
g1	68	22	90
g2	32	22	54
	100	44	144

beta prior

	1	0
g1	1/2	1/2
g2	1/2	1/2

$\varphi1 \sim \beta(68.500,22.500) \mid \varphi2 \sim \beta(32.500,22.500)$

○ $\varphi1$ ○ $\delta=\varphi1-\varphi2$
○ $\varphi2$ ● $\tau=\varphi1/\varphi2$ ○ $\varphi1/(\varphi1+\varphi2)$
○ $(\varphi1-\varphi2)/\varphi1$ ○ $(\varphi1-\varphi2)/\varphi2$ ○ $(\varphi1-\varphi2)/[(\varphi1+\varphi2)/2]$
○ $\varphi1/(1-\varphi1)$ ○ $\omega=[\varphi1/(1-\varphi1)]/[\varphi2/(1-\varphi2)]$
○ $\varphi2/(1-\varphi2)$ ○ $\upsilon=(1-\omega)/(1+\omega)$ [Q Yule]

Model ● binomial ○ poisson

confidence intervals

sTatement
○ Pr[X<x] ○ Pr[X>x]
● Pr[x1<X<x2] ○ Pr[X<x1 ou X>x2]

○ Limits 1.046
1.595
● Probability 0.90

déciMales: limite 3 probabilité 2 distribution 3

CurVe
● p(x)
○ Pr[X<x]
○ Pr[X>x]

$\tau : \varphi1 \sim \beta(68.500,22.500) \mid \varphi2 \sim \beta(32.500,22.500)$

0.942 1.273 1.845

$\Pr[1.046<\tau<1.595] = 0.90$

☐ generAte a sample
10000

mean
1.290
standard deviation
0.170

Delayed responses

- The allocation rule for subject k must be modified when the outcome is not available
- A randomized allocation scheme, either a 1:1 allocation or an allocation based on the estimates of the success rates can then be used
- Frequentist procedures must be modified to take into account the effect of these responses on the sampling distribution
- In contrast, the same Bayesian procedures can be used. Moreover, it can be expected that the resulting gain in power due to randomized allocations will partly compensate the loss in efficiency of the adaptive process

More on objective priors and stopping rules

- The independence of objective priors and the experimental design has been recently challenged (e.g., de Cristofaro, 2004; Bunouf & Lecoutre, 2006; Sun & Berger, 2008)
- A consequence is that the Jeffreys rule should give different priors for the different designs
- Most possible refinements of the inference would however need more complicated numerical tools for a small expected improvement in terms of efficiency
- Actually, for all designs with fixed number of subjects we have mentioned, independent priors Beta(0.5,0.5) for each of the two success rates were shown to work very well for the different parameters (ElQasyr, 2008)

References I

Bunouf, P., Lecoutre, B. (2006) - Bayesian priors in sequential binomial design. *Comptes Rendus de l'Académie des Sciences, Série I*, **343**, 339–344.

de Cristofaro, R. (2004) - On the foundations of likelihood principle. *Journal of Statistical Planning and Inference*, **126**, 401–411.

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

ElQasyr, K., Lecoutre, B. (2009) - Comparing two success rates with Play-The-Winner designs. Submitted for publication

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

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.

Novick, M.R., Jackson, P.H. (1974) - *Statistical Methods for Educational and Psychological Research*. New York: McGraw-Hill.

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

References II

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.

Wei, L.J. (1988) - Exact two-sample permutation tests based on the randomized play-the-winner rule. *Biometrika*, **75**, 603–606.

Wei, L.J, Smythe R.T, Lin D.Y., Park T.S. (1990) - Statistical inference with data-dependent treatment allocation rules. *Journal of the American Statistical Association*, **85**, 156–162.

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

Yao, Q., Wei, L.J. (1996) - Play the winner for phase II/III clinical trials. *Statistics in Medicine*, **15**, 2413–2423.

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