A Bayesian Sequential Design with Adaptive Randomizations

Qingzhao Yu, Professor in Biostatistics

LSU Health Sciences Center

The Bayesian Causal Inference Workshop, June 3rd, 2019

Outline

Introduction

Adaptive designs in clinical trials

- Group sequential design
  - allow early stopping for efficacy/futility
  - help to allocate resources more efficiently
  - control the overall study-wide Type I error rate
- Adaptive randomization
  - randomization rate - the probabilities of allocating patients to different treatment arms
  - assign patients to a better performing regimen
  - balance prognostic factors among intervention arms
  - increase power over traditional balanced randomization designs and minimize expected treatment failures
- Bayesian adaptive design
  - incorporating prior information
  - reduce the number of required participants
  - adaptive randomization

Bayesian sequential design with adaptive randomization

Bayesian sequential design with adaptive randomization (BSDAR) Yu et al. (2017)

- Use alpha spending function to control the study-wide overall type I error rate
- Randomization rates change adaptively at each interim analysis
- Allow to stop the trial early for efficacy

Settings

Purpose: test the difference between a novel (treatment) and an established (control) treatment.
Assume

$$X_{Ti} \overset{iid}{\sim} N(\mu_T, \sigma_T^2), i = 1, \ldots, n_T, \qquad X_{Ci} \overset{iid}{\sim} N(\mu_C, \sigma_C^2), i = 1, \ldots, n_C,$$

where μ_T, μ_C, σ_T^2 and σ_C^2 are unknown.
The hypotheses to be tested are,

$$H_0 : \mu_T = \mu_C \;\text{ v.s. }\; H_a : \mu_T \neq \mu_C$$

Prior work by Zhu and Yu (2017): a Bayesian sequential design using alpha spending function to control type I error (BSDASF), $H_a : \mu_T > \mu_C$, σ_T^2 and σ_C^2 are known.

Settings

Prior distributions for μ_C, σ_C^2

$$\begin{aligned} \mu_C|\sigma_C^2 &\sim N(\mu_0, \sigma_C^2/\tau), \\ \sigma_C^2 &\sim Inv-\chi^2(\nu_0, \sigma_0^2), \end{aligned}$$

- μ_0, σ_0^2, historical data and knowledge
- τ, control the similarity between μ_C and μ_0
  A small τ indicates large uncertainty of the similarity (Berry et al., 2010).
- σ_0^2, an estimate of the variance σ_C^2
- ν_0, how much we can depend on the prior information

A *non-informative prior* for μ_T and σ_T^2

$$p(\mu_T, \sigma_T^2) \propto (\sigma_T^2)^{-1}$$

Settings

Prior distributions for μ_C, σ_C^2

$$\begin{aligned} \mu_C|\sigma_C^2 &\sim N(\mu_0, \sigma_C^2/\tau), \\ \sigma_C^2 &\sim Inv-\chi^2(\nu_0, \sigma_0^2), \end{aligned}$$

- μ_0, σ_0^2, historical data and knowledge
- τ, control the similarity between μ_C and μ_0
  A small τ indicates large uncertainty of the similarity (Berry et al., 2010).
- σ_0^2, an estimate of the variance σ_C^2
- ν_0, how much we can depend on the prior information

A *non-informative prior* for μ_T and σ_T^2

$$p(\mu_T, \sigma_T^2) \propto (\sigma_T^2)^{-1}$$

LSU Health NEW ORLEANS School of Public Health

Settings

At the jth interim analysis, $n(t_j) = n_T(t_j) + n_C(t_j)$. Given the interim data $\vec{\mathbf{x}}_{Tj}$ and $\vec{\mathbf{x}}_{Cj}$ at t_j, the marginal posterior distributions for σ_T^2 and σ_C^2 are

$$p(\sigma_T^2 \mid \vec{\mathbf{x}}_{Tj}) \sim Inv-\chi^2(n_T(t_j)-1, s_{Tj}^2), \quad (1)$$
$$p(\sigma_C^2 \mid \vec{\mathbf{x}}_{Cj}, \tau, \mu_0, \nu_0, \sigma_0^2) \sim Inv-\chi^2(\nu_{nj}, \sigma_{nj}^2), \quad (2)$$

where
$\nu_{nj} = \nu_0 + n_C(t_j), \nu_{nj}\sigma_{nj}^2 = \nu_0\sigma_0^2 + (n_C(t_j)-1)s_{Cj}^2 + \frac{\tau n_C(t_j)}{\tau + n_C(t_j)}(\bar{x}_{Cj} - \mu_0)^2$.
The conditional posterior distribution of $\mu_T - \mu_C$ is

$$p(\mu_T - \mu_C \mid \sigma_T^2, \vec{\mathbf{x}}_{Tj}, \sigma_C^2, \vec{\mathbf{x}}_{Cj}, \tau, \mu_0, \nu_0, \sigma_0^2) \sim N(u, \sigma^2), \quad (3)$$

where $u = \bar{x}_{Tj} - \mu_{nj}$ and variance $\sigma^2 = \sigma_T^2/n_T(t_j) + \sigma_C^2/\tau_{nj}$,
$\mu_{nj} = \frac{\tau}{\tau + n_C(t_j)}\mu_0 + \frac{n_C(t_j)}{\tau + n_C(t_j)}\bar{x}_{Cj}, \tau_{nj} = \tau + n_C(t_j)$.

Settings

At the jth interim analysis, $n(t_j) = n_T(t_j) + n_C(t_j)$. Given the interim data $\vec{\mathbf{x}}_{Tj}$ and $\vec{\mathbf{x}}_{Cj}$ at t_j, the marginal posterior distributions for σ_T^2 and σ_C^2 are

$$p(\sigma_T^2 \mid \vec{\mathbf{x}}_{Tj}) \sim Inv - \chi^2(n_T(t_j) - 1, s_{Tj}^2), \quad (1)$$

$$p(\sigma_C^2 \mid \vec{\mathbf{x}}_{Cj}, \tau, \mu_0, \nu_0, \sigma_0^2) \sim Inv - \chi^2(\nu_{nj}, \sigma_{nj}^2), \quad (2)$$

where

$\nu_{nj} = \nu_0 + n_C(t_j), \nu_{nj}\sigma_{nj}^2 = \nu_0\sigma_0^2 + (n_C(t_j) - 1)s_{Cj}^2 + \frac{\tau n_C(t_j)}{\tau + n_C(t_j)}(\bar{x}_{Cj} - \mu_0)^2.$

The conditional posterior distribution of $\mu_T - \mu_C$ is

$$p(\mu_T - \mu_C \mid \sigma_T^2, \vec{\mathbf{x}}_{Tj}, \sigma_C^2, \vec{\mathbf{x}}_{Cj}, \tau, \mu_0, \nu_0, \sigma_0^2) \sim N(u, \sigma^2), \quad (3)$$

where $u = \bar{x}_{Tj} - \mu_{nj}$ and variance $\sigma^2 = \sigma_T^2/n_T(t_j) + \sigma_C^2/\tau_{nj}$,
$\mu_{nj} = \frac{\tau}{\tau + n_C(t_j)}\mu_0 + \frac{n_C(t_j)}{\tau + n_C(t_j)}\bar{x}_{Cj}, \tau_{nj} = \tau + n_C(t_j).$

Control type I error

Type I error under the Bayesian setting - the probability of rejecting the null hypothesis when the null hypothesis is true (Casella and Berger, 2002).
Alpha spending functions (Lan and DeMets, 1983; Kim and DeMets, 1987; Zhu and Yu, 2017; Zhu et al., 2017)

- Information fraction at the jth interim analysis, $t_j^* = n(t_j)/n$, where n is the maximum allowed sample size
- Non-decreasing function $\alpha(t^*)$
- $\alpha(0) = 0, \alpha(1) = \alpha$, where α is the desired significance level

Alpha spending function

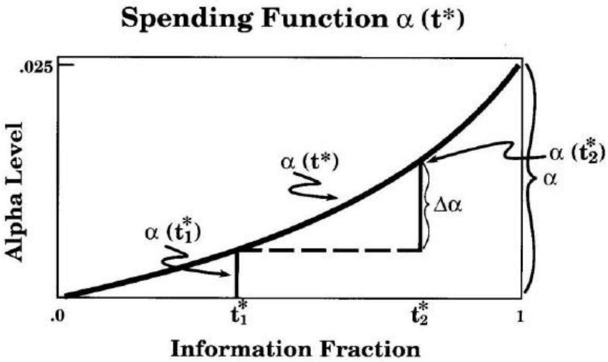

Figure 1: Alpha spending function indicating additional type I error rate $\triangle\alpha$, allocated between interim analyses (DeMets and Lan, 1995).

Control type I error

Four types of alpha spending functions:

- O'Brien–Fleming alpha spending function (OF)
  $\alpha_1(t^*) = 2 - 2\Phi(z_{\alpha/2}/\sqrt{t^*})$,
  where Φ is the cumulative distribution function of the standard normal distribution.
- Pocock alpha spending function
  $\alpha_2(t^*) = \alpha \log\{1 + (e - 1)t^*\}$
- Uniform alpha spending function
  $\alpha_3(t^*) = t^*\alpha$
- Equal alpha spending function
  the traditional method that sets equal critical values for all t^*, predetermined through simulations.

Control type I error

Four types of alpha spending functions:

- O'Brien–Fleming alpha spending function (OF)
  $\alpha_1(t^*) = 2 - 2\Phi(z_{\alpha/2}/\sqrt{t^*})$,
  where Φ is the cumulative distribution function of the standard normal distribution.
- Pocock alpha spending function
  $\alpha_2(t^*) = \alpha \log\{1 + (e-1)t^*\}$
- Uniform alpha spending function
  $\alpha_3(t^*) = t^*\alpha$
- Equal alpha spending function
  the traditional method that sets equal critical values for all t^*, predetermined through simulations.

Control type I error

Four types of alpha spending functions:

- O'Brien–Fleming alpha spending function (OF)
  $\alpha_1(t^*) = 2 - 2\Phi(z_{\alpha/2}/\sqrt{t^*})$,
  where Φ is the cumulative distribution function of the standard normal distribution.
- Pocock alpha spending function
  $\alpha_2(t^*) = \alpha \log\{1 + (e-1)t^*\}$
- Uniform alpha spending function
  $\alpha_3(t^*) = t^*\alpha$
- Equal alpha spending function
  the traditional method that sets equal critical values for all t^*, predetermined through simulations.

Control type I error

Four types of alpha spending functions:

- O'Brien–Fleming alpha spending function (OF)
  $\alpha_1(t^*) = 2 - 2\Phi(z_{\alpha/2}/\sqrt{t^*})$,
  where Φ is the cumulative distribution function of the standard normal distribution.
- Pocock alpha spending function
  $\alpha_2(t^*) = \alpha \log\{1 + (e - 1)t^*\}$
- Uniform alpha spending function
  $\alpha_3(t^*) = t^*\alpha$
- Equal alpha spending function
  the traditional method that sets equal critical values for all t^*, predetermined through simulations.

Alpha spending function

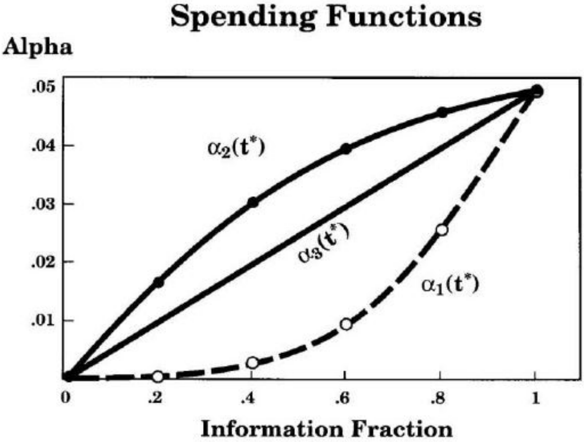

Figure 2: Comparison of alpha spending functions (DeMets and Lan, 1995).

Randomization method

Randomization rate $r_T(t_j)$ - the distribution rate of the newly recruited $n(t_{j+1}) - n(t_j)$ patients to be assigned to the treatment group after the jth interim analysis.

$$r_T(t_j) = \min\left\{\max\left(\frac{\hat{\sigma}_{Tj} n_C(t_j) + \hat{\sigma}_{Tj}\tau + \hat{\sigma}_{Tj}(n(t_{j+1}) - n(t_j)) - \hat{\sigma}_{Cj} n_T(t_j)}{(\hat{\sigma}_{Tj} + \hat{\sigma}_{Cj})(n(t_{j+1}) - n(t_j))}, 0\right), 1\right\}, \tag{4}$$

where $\hat{\sigma}_{Tj}$ and $\hat{\sigma}_{Cj}$ are the estimates of σ_T and σ_C from the jth interim analysis (see Equations (1)–(2)).

Lemma

Under the settings described in Slide 5–7, and given the information obtained up till the jth interim analysis, assigning patients to the treatment group at the randomization rate defined by Equation (4) after the jth interim analysis can achieve the minimum variance estimation for the testing statistic, $\hat{\mu}_T - \hat{\mu}_C = \bar{x}_{Tj} - \mu_{nj}$ (the posterior mean by Equation (3)).

Randomization method

Randomization rate $r_T(t_j)$ - the distribution rate of the newly recruited $n(t_{j+1}) - n(t_j)$ patients to be assigned to the treatment group after the jth interim analysis.

$$r_T(t_j) = \min\left\{\max\left(\frac{\hat{\sigma}_{Tj} n_C(t_j) + \hat{\sigma}_{Tj}\tau + \hat{\sigma}_{Tj}(n(t_{j+1}) - n(t_j)) - \hat{\sigma}_{Cj} n_T(t_j)}{(\hat{\sigma}_{Tj} + \hat{\sigma}_{Cj})(n(t_{j+1}) - n(t_j))}, 0\right), 1\right\}, \tag{4}$$

where $\hat{\sigma}_{Tj}$ and $\hat{\sigma}_{Cj}$ are the estimates of σ_T and σ_C from the jth interim analysis (see Equations (1)–(2)).

Lemma

Under the settings described in Slide 5–7, and given the information obtained up till the jth interim analysis, assigning patients to the treatment group at the randomization rate defined by Equation (4) after the jth interim analysis can achieve the minimum variance estimation for the testing statistic, $\hat{\mu}_T - \hat{\mu}_C = \bar{x}_{Tj} - \mu_{nj}$ (the posterior mean by Equation (3)).

Sensitivity analysis

The impact of the parameters μ_0, τ, ν_0, σ_0^2 in the prior distributions of μ_C and σ_C^2 on the randomization rates and decision bounds

- $\mu_C = 0$, $\sigma_C = 1$, $\mu_T = 0$, $\sigma_T = 5$,
- $\mu_0 = 0$, $\tau = 0.1$, $\nu_0 = 6$, $\sigma_0^2 = 1$,
- $r_T(t_0) = 0.5$, $n = 100$
- Five equal-interval interim analyses planned at $t_1^* = 0.2$, $t_2^* = 0.4$, $t_3^* = 0.6$, $t_4^* = 0.8$, and $t_5^* = 1$

Sensitivity analysis

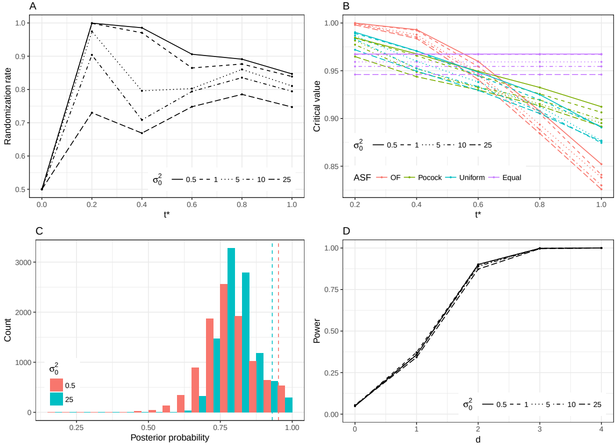

Figure 3: Comparison of (A) randomization rates, (B) critical values, $N_{rep} = 2000$, critical values are averaged over 20 replicates; (C) histograms of posterior probabilities at the 1st interim analysis in 10000 simulated trials (OF); (D) powers of BSDAR, $N_{rep} = 10000$.

Simulation studies

Compare the powers and required sample sizes on testing the hypotheses between BSDAR and a Bayesian sequential design without adaptive randomization

- $\mu_C = 0$, $\sigma_C = 1$, $\sigma_T = 5$, $d = \mu_T - \mu_C$
- $\mu_0 = 0$, $\tau = 0.1$, $\nu_0 = 6$, $\sigma_0^2 = 1$
- $r_T(t_0) = 0.5$

Compare the powers, required sample sizes, and randomization rates of BSDAR when different alpha spending functions are used.

Simulation studies

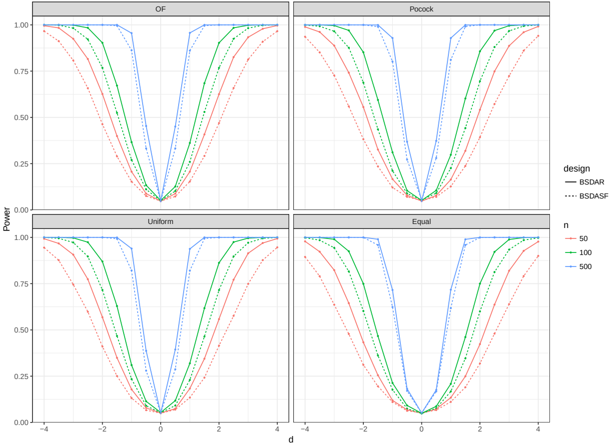

Figure 4: Compare the powers between BSDAR (solid) and a Bayesian sequential design without adaptive randomization (dashed) with different alpha spending functions at $n = 50$ (red), 100 (green), and 500 (blue) when $\delta = 0.64$, $J = 5$, obtained using $N_{rep} = 10000$.

Real data application

A clinical trial for diabetic patients

- Primary endpoint: the change from baseline in HbA1c (glycosylated hemoglobin) after 24 weeks of treatment
- Objective: to test if the treatment is different from the control in reducing HbA1c
- 508 patients were enrolled
- 168 patients in the control group with a mean reduction in HbA1c of 0.0042 mmol (variance $= 0.6394$)
- 340 patients in the treatment group with a mean reduction of 0.5218 mmol (variance $= 1.5672$)
- Conclusion: compared with the control group, the HbA1c was significantly reduced in the treatment group (p-value < 0.0001) by an ANOVA analysis.

Real data application

To apply BSDAR to this trial,

- 10 interim analyses are evenly planned over the clinical trial, i.e. $t_1^* = 50/508, t_2^* = 100/508, \ldots, t_{10}^* = 1$
- Prior parameters, $\mu_0 = 0$, $\tau = 0.1$, $\nu_0 = 6$, $\sigma_0^2 = 1$
- $r_T(t_0) = 0.5$

Table 1: Compare the required sample sizes of BSDAR and that of a Bayesian sequential design without adaptive randomization using different alpha spending functions.

	BSDAR				w/o adaptive randomization		
	n	n_T	n_C	*ratio*	n	n_T	n_C
O'Brien–Fleming	150	91	59	0.78	200	100	100
Pocock	100	62	38	0.66	150	75	75
Uniform	100	60	40	0.66	150	75	75
Equal	100	61	39	0.66	150	75	75

Real data application

To study how the allocation of interim analyses can influence the required sample sizes, we adopt the O'Brien–Fleming alpha spending function to control the overall type I error rate and assume the following six scenarios of interim analysis:

A. 10 evenly spaced interim analyses over 200 total sample size, i.e. $t_1^* = 20/200, t_2^* = 40/200, \ldots, t_{10}^* = 1$;
B. 4 evenly spaced interim analyses over 200 total sample size, i.e. $t_1^* = 50/200, t_2^* = 100/200, t_3^* = 150/200, t_4^* = 1$;
C. 4 unevenly spaced interim analyses over 200 total sample size, with $t_1^* = 80/200, t_2^* = 100/200, t_3^* = 140/200, t_4^* = 1$;
D. 3 evenly spaced interim analyses over 150 total sample size, i.e. $t_1^* = 50/150, t_2^* = 100/150, t_3^* = 1$;
E. 6 evenly spaced interim analyses over 150 total sample size, i.e. $t_1^* = 30/150, t_2^* = 60/150, \ldots, t_6^* = 1$;
F. 4 unevenly spaced interim analyses over 150 total sample size, with $t_1^* = 60/150, t_2^* = 100/150, t_3^* = 130/150, t_4^* = 1$.

Real data application

Table 2: Compare the required sample sizes used by BSDAR and that by a Bayesian sequential design without adaptive randomization for different scenarios when O'Brien–Fleming alpha spending function is used.

	BSDAR				w/o adaptive randomization		
Scenario	n	n_T	n_C	*ratio*	n	n_T	n_C
A	100	63	37	0.64	120	60	60
B	100	58	42	0.66	150	75	75
C	100	60	40	1.00	140	70	70
D	100	53	47	0.66	150	75	75
E	90	53	37	0.78	120	60	60
F	100	58	42	0.75	130	65	65

Discussions

Advantages

- attribute newly recruited patients to different treatment arms more efficiently
- reduce required sample size
- improve the power of tests at a given sample size

Discussions

- BSDAR with O'Brien–Fleming alpha spending function has the largest power but is also related with the largest required sample size
- choose a τ smaller than 1 when not enough information of μ_C is available
- change randomization rate in favor of the treatment arm that is currently empirically superior

Discussions

Advantages

- attribute newly recruited patients to different treatment arms more efficiently
- reduce required sample size
- improve the power of tests at a given sample size

Discussions

- BSDAR with O'Brien–Fleming alpha spending function has the largest power but is also related with the largest required sample size
- choose a τ smaller than 1 when not enough information of μ_C is available
- change randomization rate in favor of the treatment arm that is currently empirically superior

A Bayesian sequential design for time-to-event outcomes

Bayesian sequential design for time-to-event outcomes (BSD4TEO)

- Use alpha spending function to control the study-wide overall type I error rate
- Put a prior on log hazard ratio, without imposing further assumptions on the distribution of times to event
- Bayes factor is adapted for decision-making at interim analyses
- Allow to stop the trial early for efficacy

Settings

Purpose: to test whether the novel treatment (treatment group) has a better treatment effect compared to the established treatment (control group) w.r.t. time to event

Denote (X_{li}, δ_{li}) for the ith patient, $i = 1, 2, \dots n$,
$l = 1$ for treatment group, $l = 0$ for control group,

$$X_{li} = \min(T_{li}, C_{li}),$$
$$\delta_{li} = 1 \,\text{if}\, X_{li} = T_{li},$$
$$\delta_{li} = 0 \,\text{if}\, X_{li} = C_{li}.$$

Suppose events occur at D ordered times $t_1 \leq t_2 \leq \dots \leq t_D$
At time t_k, $k = 1, 2, \dots, D$

$$d_{lk} = 0 \,\text{or}\, 1, \quad \text{number of events observed}$$
$$Y_{lk}, \quad \text{number of subjects at risk}$$

Settings

Assumption: independent censoring and proportional hazard

The hypotheses to be tested are,

$$H_0 : \theta = 0 \quad \text{v.s.} \quad H_a : \theta < 0$$

θ - log hazard ratio of the treatment relative to the control

Assume

$$d_{1k} \sim \text{Ber}(p_{1k}), \quad k = 1, 2, \ldots, D,$$

where p_{1k} is the probability of observing the event in the treatment group at t_k, estimated by

$$\hat{p}_{1k} = \frac{\exp(\theta) y_{1k}}{y_{0k} + \exp(\theta) y_{1k}}.$$

Settings

Prior distributions for θ

$$\theta \quad \sim \quad N(\theta_0, \sigma^2/\tau),$$

- σ^2, assumed to be known
- θ_0, an estimate of the log hazard ratio based on historical data
- τ, how much we trust the prior information

Plan the jth interim analysis at $t_j^* = n_e(t_j)/n_e$, n_e the total targeted number of events

Posterior distribution is log-concave in θ - adaptive rejection sampling

Research on BSD4TEO

Zhu et al. (2019)

- Algorithms to calculate critical values and power
- Sensitivity analysis, the prior mean θ_0 and the precision parameter τ
- Simulations to compare BSD4TEO with the frequentist group sequential design
- Apply on the Chronic Granulotomous Disease (CGD) Dataset (Gallin et al., 1991)

Future research

Future research

- Extension of the Bayesian sequential design with adaptive randomization to outcome with a distribution from the exponential family
- Extend the Bayesian sequential design with adaptive randomization to multi-arm multi-stage clinical trials
- A R package for Bayesian sequential designs with alpha spending function to control type I error rate

References I

Scott M Berry, Bradley P Carlin, J Jack Lee, and Peter Muller. *Bayesian adaptive methods for clinical trials*. CRC press, Boca Raton, FL, 2010.

George Casella and Roger L Berger. *Statistical inference*. Duxbury, Pacific Grove, CA, second edition, 2002.

David L. DeMets and Gordon Lan. *The alpha spending function approach to interim data analyses*, pages 1–27. Springer, Boston, MA, 1995. doi: 10.1007/978-1-4615-2009-2_1. URL https://doi.org/10.1007/978-1-4615-2009-2_1.

JI Gallin, HL Malech, WJT CURNUTTE, PG QUIE, HS JAFFE, and RAB EZKOWITZ. A controlled trial of interferon gamma to prevent infection in chronic granulomatous disease. *The New England journal of medicine*, 324(8): 509–516, 1991.

Kyungmann Kim and David L. DeMets. Design and analysis of group sequential tests based on the type i error spending rate function. *Biometrika*, 74(1): 149–154, mar 1987. doi: 10.2307/2336029. URL https://doi.org/10.2307/2336029.

K. K. Gordon Lan and David L. DeMets. Discrete sequential boundaries for clinical trials. *Biometrika*, 70(3):659–663, dec 1983. doi: 10.2307/2336502. URL https://doi.org/10.2307/2336502.

References II

Qingzhao Yu, Lin Zhu, and Han Zhu. A bayesian sequential design with adaptive randomization for 2-sided hypothesis test. *Pharmaceutical Statistics*, 16(6): 451–465, 2017. ISSN 1539-1612. doi: 10.1002/pst.1830. URL http://dx.doi.org/10.1002/pst.1830.

Han Zhu and Qingzhao Yu. A bayesian sequential design using alpha spending function to control type i error. *Statistical Methods in Medical Research*, 26 (5):2184–2196, 2017. doi: 10.1177/0962280215595058. URL https://doi.org/10.1177/0962280215595058.

Han Zhu, Qingzhao Yu, and Donald E. Mercante. A bayesian sequential design with binary outcome. *Pharmaceutical Statistics*, 16(3):192–200, 2017. ISSN 1539-1612. doi: 10.1002/pst.1805. URL http://dx.doi.org/10.1002/pst.1805.

Lin Zhu, Qingzhao Yu, and Donald Mercante. A bayesian sequential design for clinical trials with time-to-event outcomes. *Statistics in Pharmaceutical Research*, Under Review, 2019.