

Contents

1	Introduction	2
1.1	Topics to be covered	2
1.2	Examples of sequential designs	2
2	The sequential probability ratio test (SPRT)	3
2.1	Definitions	3
2.2	Properties of the SPRT	4
2.3	Estimation following the SPRT	8
2.4	Examples of the SPRT	9
3	Group sequential tests	9
3.1	Analysing the data in groups	9
3.2	Designing a group sequential test	10
3.3	Inference following a group sequential test	12
3.4	Examples of group sequential designs	14
4	Adaptive treatment allocation rules	16
4.1	Definitions	16
4.2	Properties of adaptive treatment allocation	17
4.3	Examples of adaptive treatment allocation	18
4.4	Group-sequential response-adaptive tests	20

Topics in the Design of Experiments

Part 2: Sequential Design

Steve Coad*

March 2010

1 Introduction

1.1 Topics to be covered

Part 1 of this course showed how an experiment may be designed optimally when the number of observations is **fixed** in advance. The main focus was the application of this theory to **non-linear models** and the construction of D -optimal designs. In this part, we are concerned with sequential designs, where the number of observations to be taken is not fixed in advance or the design points are chosen sequentially depending on the current data.

We first introduce the sequential probability ratio test for testing two simple hypotheses and study several of its properties. Since this is a fully sequential design, in that a test is performed after every observation, we then introduce group sequential designs and show how these may be carried out in practice. Finally, we study some adaptive treatment allocation rules, where the treatment allocation probabilities are functions of the current data.

1.2 Examples of sequential designs

A sequential design is often **more efficient** than an equivalent fixed-sample one. The examples below demonstrate the wide range of applications of a sequential approach.

Example. Curtailed test.

Suppose that a machine produces items which may be judged good or defective, and that the true proportion of defectives in a large batch is p . Let S_m denote the number of defectives in a random sample of size m . Consider testing $H_0 : p \leq p_0$ against $H_1 : p > p_0$. A reasonable rule is to reject H_0 if $S_m \geq r$ for some constant r .

*School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom

Now let T be the smallest value of k for which $S_k = r$ and put $T' = \min(T, m)$. Consider the procedure which stops sampling at the **random time** T' and decides that $p > p_0$ if and only if $T \leq m$. If one considers these two procedures as tests of H_0 against H_1 , their rejection regions, namely, $\{T \leq m\}$ and $\{S_m \geq r\}$, are the same events, and hence the two tests have the same **power function**. Clearly, the test which stops at the random time T' has a reasonable claim to be regarded as more efficient.

Example. Repeated significance test.

Let X_1, X_2, \dots be independent normal random variables with unknown mean μ and unit variance. Consider testing $H_0 : \mu = 0$ against $H_1 : \mu \neq 0$. The standard fixed-sample 0.05 level significance test rejects H_0 if and only if $|S_n| \geq 1.96\sqrt{n}$, where $S_n = \sum_{k=1}^n x_k$.

Now suppose that, if H_1 is true, a minimum amount of experimentation is desired, but no similar constraint exists under H_0 . Let $b > 0$ and let m be a maximum sample size. Sample sequentially, stopping with rejection of H_0 at the first $n \leq m$, if one exists, such that $|S_n| \geq b\sqrt{n}$. Otherwise, stop sampling at m and accept H_0 . The significance level of this procedure is

$$\alpha = \alpha(b, m) = P_0(|S_n| \geq b\sqrt{n} \text{ for some } n \leq m),$$

where P_0 denotes probability under H_0 . Clearly, b must be somewhat larger than 1.96, depending on m , in order that $\alpha(b, m) = 0.05$.

2 The sequential probability ratio test (SPRT)

2.1 Definitions

Let X_1, X_2, \dots be a sequence of random variables with joint probability density functions

$$P(X_1 \in d\xi_1, \dots, X_n \in d\xi_n) = f_n(\xi_1, \dots, \xi_n) d\xi_1 \dots d\xi_n$$

for $n = 1, 2, \dots$. Consider testing the **simple hypotheses** $H_0 : f_n = f_{0n}$ for all n against $H_1 : f_n = f_{1n}$ for all n . The **likelihood ratio** is

$$\ell_n = \ell_n(x_1, \dots, x_n) = \frac{f_{1n}(x_1, \dots, x_n)}{f_{0n}(x_1, \dots, x_n)}.$$

The **sequential probability ratio test (SPRT)** chooses constants $0 < A < B < \infty$, usually $A < 1 < B$, and samples x_1, x_2, \dots sequentially until the random time

$$\begin{aligned} N &= \text{first } n \geq 1 \text{ such that } \ell_n \notin (A, B) \\ &= \infty \text{ if } \ell_n \in (A, B) \text{ for all } n \geq 1. \end{aligned}$$

The test stops sampling at time N and, if $N < \infty$, rejects H_0 if $\ell_N \geq B$ and accepts H_0 if $\ell_N \leq A$.

Assuming temporarily that $P_i(N < \infty) = 1$ for $i = 0, 1$, where P_i denotes probability under H_i , the above test has **significance level** $\alpha = P_0(\ell_N \geq B)$ and **power** $1 - \beta = P_1(\ell_N \geq B)$.

In the fixed-sample case, the **Neyman-Pearson lemma** tells us that, among all tests with the same significance level, the likelihood ratio test has the highest power.

2.2 Properties of the SPRT

Let B_n denote the subset of n -dimensional space in which $A < \ell_k(\xi_1, \dots, \xi_k) < B$ for $k = 1, 2, \dots, n-1$ and $\ell_n(\xi_1, \dots, \xi_n) \geq B$, so that

$$\{N = n, \ell_n \geq B\} = \{(x_1, \dots, x_n) \in B_n\}.$$

Then

$$\begin{aligned} \alpha = P_0(\ell_N \geq B) &= \sum_{n=1}^{\infty} P_0(N = n, \ell_n \geq B) \\ &= \sum_{n=1}^{\infty} \int_{B_n} f_{0n} d\xi_1 \dots d\xi_n \\ &= \sum_{n=1}^{\infty} \int_{B_n} \frac{f_{0n}}{f_{1n}} f_{1n} d\xi_1 \dots d\xi_n \\ &= \sum_{n=1}^{\infty} E_1(\ell_n^{-1}; N = n, \ell_n \geq B) \\ &= E_1(\ell_N^{-1}; \ell_N \geq B) \\ &\leq B^{-1} P_1(\ell_N \geq B) = B^{-1}(1 - \beta). \end{aligned}$$

Similarly,

$$\beta = P_1(\ell_N \leq A) \leq A P_0(\ell_N \leq A) = A(1 - \alpha).$$

Treating the above inequalities as approximate equalities and solving for α and β leads to the simple approximations

$$\alpha \simeq \frac{1 - A}{B - A} \quad \text{and} \quad \beta \simeq \frac{A(B - 1)}{B - A}.$$

Theorem 1. Wald's equation.

Let X_1, X_2, \dots be independent and identically distributed random variables with finite mean μ . Let M be any integer-valued random variable such that $\{M = n\}$ is an event determined only by X_1, \dots, X_n for all $n = 1, 2, \dots$, and assume that $E(M) < \infty$. Then

$$E\left(\sum_{k=1}^M X_k\right) = \mu E(M).$$

Proof. Suppose initially that $X \geq 0$. Write

$$\sum_{k=1}^M X_k = \sum_{k=1}^{\infty} 1_{\{M \geq k\}} X_k$$

and note that

$$\{M \geq k\} = \left[\bigcup_{j=1}^{k-1} \{M = j\} \right]^c$$

is independent of X_k, X_{k+1}, \dots . Hence, by the **monotone convergence theorem**, we have that

$$E\left(\sum_{k=1}^M X_k\right) = \sum_{k=1}^{\infty} E(X_k; M \geq k) = \mu \sum_{k=1}^{\infty} P(M \geq k) = \mu E(M).$$

For the general case, write

$$\sum_{k=1}^M X_k = \sum_{k=1}^M X_k^+ - \sum_{k=1}^M X_k^-,$$

where $a^+ = \max(a, 0)$ and $a^- = -\min(a, 0)$, and apply the above case to these two terms separately. \square

Suppose that X_1, X_2, \dots are independent and identically distributed random variables, so that

$$\ell_n = \prod_{k=1}^n \frac{f_1(x_k)}{f_0(x_k)},$$

where f_i is the probability density function of X under H_i for $i = 0, 1$. Then

$$\log \ell_n = \sum_{k=1}^n \log \left\{ \frac{f_1(x_k)}{f_0(x_k)} \right\}$$

is a sum of independent and identically distributed random variables. Further, the **stopping rule** for the SPRT may be written as

$$\begin{aligned} N &= \text{first } n \geq 1 \text{ such that } \log \ell_n \notin (a, b) \\ &= \infty \text{ if } \log \ell_n \in (a, b) \text{ for all } n, \end{aligned}$$

where $a = \log A$ and $b = \log B$. Now, by Theorem 1,

$$E_i(\log \ell_N) = \mu_i E_i(N),$$

where $\mu_i = E_i[\log\{f_1(X)/f_0(X)\}]$ for $i = 0, 1$. Also, we may write

$$E_i(\log \ell_N) \simeq aP_i(\ell_N \leq A) + bP_i(\ell_N \geq B).$$

Combining these two equations yields the approximations

$$E_0(N) \simeq \frac{1}{\mu_0} \left\{ a \frac{(B-1)}{B-A} + b \frac{(1-A)}{B-A} \right\}$$

and

$$E_1(N) \simeq \frac{1}{\mu_1} \left\{ a \frac{A(B-1)}{B-A} + b \frac{B(1-A)}{B-A} \right\}.$$

Alternatively, we may write

$$E_0(N) \simeq \frac{1}{\mu_0} \left\{ (1 - \alpha) \log \left(\frac{\beta}{1 - \alpha} \right) + \alpha \log \left(\frac{1 - \beta}{\alpha} \right) \right\}$$

and

$$E_1(N) \simeq \frac{1}{\mu_1} \left\{ \beta \log \left(\frac{\beta}{1 - \alpha} \right) + (1 - \beta) \log \left(\frac{1 - \beta}{\alpha} \right) \right\}.$$

The following result, which is stated without proof, shows that the SPRT terminates with probability one.

Theorem 2. Stein's lemma.

Let Y_1, Y_2, \dots be independent and identically distributed random variables with the property $P(Y = 0) < 1$. Let $-\infty < a < b < \infty$ and $S_n = \sum_{k=1}^n Y_k$, and define

$$\begin{aligned} M &= \text{first } n \geq 1 \text{ such that } S_n \notin (a, b) \\ &= \infty \text{ if } S_n \in (a, b) \text{ for all } n. \end{aligned}$$

Then there exist constants $C > 0$ and $0 < \rho < 1$ such that $P(M > n) \leq C\rho^n$ for $n = 1, 2, \dots$. In particular, $E(M^k) < \infty$ for all $k = 1, 2, \dots$ and $E(e^{\lambda M}) < \infty$ for $\lambda < \log(\rho^{-1})$.

Theorem 3. Wald's likelihood ratio identity.

Let X_1, X_2, \dots be an arbitrary sequence of random variables and suppose that there exists a likelihood ratio ℓ_n for x_1, \dots, x_n under P_1 relative to P_0 such that

$$E_1(Y_n) = E_0(Y_n \ell_n),$$

where $Y_n = g(X_1, \dots, X_n)$ for some function g . Then, for any stopping time N and non-negative random variable $Y = g(X_1, \dots, X_N)$, say,

$$E_1(Y; N < \infty) = E_0(Y \ell_N; N < \infty).$$

In particular, if $Y = 1_A$, then

$$P_1[A \cap \{N < \infty\}] = E_0[\ell_N; A \cap \{N < \infty\}].$$

Proof. We have that

$$\begin{aligned} E_1(Y; N < \infty) &= \sum_{n=1}^{\infty} E_1(Y; N = n) \\ &= \sum_{n=1}^{\infty} E_0(Y \ell_n; N = n) \\ &= E_0(Y \ell_N; N < \infty), \end{aligned}$$

as required. □

Note that we have already used Theorem 3 to obtain the approximations for the error probabilities α and β . The following result is a special case of Theorem 3 and is stated without proof.

Corollary. Wald's fundamental identity.

Let X_1, X_2, \dots be a sequence of independent and identically distributed random variables and suppose that $\phi(t) = E(e^{tX}) < \infty$ for some $t \neq 0$. Then, if N is a stopping time such that $P(N < \infty) = 1$,

$$E \left[\{\phi(t)\}^{-N} e^{tS_N} \right] = 1,$$

where $S_N = \sum_{k=1}^N X_k$.

If $\phi(t) < \infty$ for $|t| < \delta$, where $\delta > 0$, then the above identity may be differentiated with respect to t at $t = 0$ to reproduce Theorem 1, provided that differentiation under the expectation can be justified: see Exercise 2.

When testing a simple hypothesis against a simple alternative with independent and identically distributed observations, the **Wald-Wolfowitz theorem** states that the SPRT minimises $E_i(N)$ for $i = 0, 1$, among all tests having no larger error probabilities. For cases where the results for expected sample size are exact, Theorem 5 contains a complete proof.

Theorem 5. Let T be the stopping time of any test of $H_0 : f = f_0$ against $H_1 : f = f_1$ with error probabilities α and β , $0 < \alpha, \beta < 1$. Assume that $E_i(T) < \infty$ for $i = 0, 1$. Then

$$E_1(T) \geq \frac{1}{\mu_1} \left\{ (1 - \beta) \log \left(\frac{1 - \beta}{\alpha} \right) + \beta \log \left(\frac{\beta}{1 - \alpha} \right) \right\}$$

and

$$E_0(T) \geq \frac{1}{\mu_0} \left\{ \alpha \log \left(\frac{1 - \beta}{\alpha} \right) + (1 - \alpha) \log \left(\frac{\beta}{1 - \alpha} \right) \right\},$$

where

$$\mu_i = E_i \left[\log \left\{ \frac{f_1(X)}{f_0(X)} \right\} \right]$$

for $i = 0, 1$.

Proof. Let $R = \{\text{Reject } H_0\}$ and $R^c = \{\text{Accept } H_0\}$. Then, by Theorem 3,

$$\begin{aligned} \alpha = P_0(R) &= E_1 \left(\ell_T^{-1}; R \right) \\ &= E_1 \left(e^{-\log \ell_T} | R \right) P_1(R) \\ &\geq \exp \{ -E_1(\log \ell_T | R) \} (1 - \beta) \\ &= \exp \{ -E_1(\log \ell_T; R) / (1 - \beta) \} (1 - \beta), \end{aligned}$$

where the penultimate line is due to **Jensen's inequality**. Thus,

$$(1 - \beta) \log \left(\frac{\alpha}{1 - \beta} \right) \geq -E_1(\log \ell_T; R).$$

Similarly,

$$\beta \log \left(\frac{1 - \alpha}{\beta} \right) \geq -E_1(\log \ell_T; R^c).$$

Hence, by Theorem 1,

$$\begin{aligned} (1 - \beta) \log \left(\frac{\alpha}{1 - \beta} \right) + \beta \log \left(\frac{1 - \alpha}{\beta} \right) &\geq -E_1(\log \ell_T) \\ &= -\mu_1 E_1(T). \end{aligned}$$

Since $\mu_1 > 0$, this completes the proof of the first assertion. The second assertion is proved similarly. \square

2.3 Estimation following the SPRT

The estimation of a parameter when the data have been obtained from a SPRT is a difficult problem. Even sequentially stopped versions of ordinarily unbiased estimators are biased, and their sampling distributions are often quite complicated. The following result, which is stated without proof, shows that randomly stopped averages are asymptotically normally distributed under quite general conditions.

Theorem 6. Anscombe-Doebelin theorem.

Let X_1, X_2, \dots be independent and identically distributed random variables with finite mean μ and finite positive variance σ^2 . Let $S_n = \sum_{k=1}^n X_k$ and suppose that $M_c, c \geq 0$, are positive integer-valued random variables such that, for some constants $m_c \rightarrow \infty$, $M_c/m_c \rightarrow 1$ in probability as $c \rightarrow \infty$. Then

$$P \left(\frac{S_{M_c} - \mu M_c}{\sigma M_c^{\frac{1}{2}}} \leq x \right) \rightarrow \Phi(x)$$

as $c \rightarrow \infty$, where Φ denotes the standard normal distribution function.

It follows from Theorem 6 that $N^{-\frac{1}{2}}(S_N - N\mu)/\sigma$ is approximately standard normal. Thus, an approximate 95% confidence interval for μ for large a and b is given by

$$\frac{S_N}{N} \pm 1.96 \frac{\sigma}{N^{\frac{1}{2}}}.$$

However, this approximation is very poor for moderate values of a and b .

2.4 Examples of the SPRT

Let N be the stopping time for the SPRT. Then, from Theorem 2, $P_i(N < \infty) = 1$ for $i = 0, 1$.

Example. Let X_1, X_2, \dots be independent normal random variables with mean μ and unit variance. Consider testing $H_0 : \mu = \mu_0$ against $H_1 : \mu = \mu_1$, where $\mu_0 < \mu_1$. The likelihood ratio is

$$\begin{aligned} \ell_n &= \prod_{k=1}^n \frac{\phi(x_k - \mu_1)}{\phi(x_k - \mu_0)} \\ &= \exp \left\{ (\mu_1 - \mu_0) S_n - \frac{n}{2} (\mu_1^2 - \mu_0^2) \right\}, \end{aligned}$$

where $\phi(x) = (2\pi)^{-\frac{1}{2}} \exp(-x^2/2)$ and $S_n = \sum_{k=1}^n x_k$. Hence, the stopping rule for the SPRT is given by

$$N = \text{first } n \geq 1 \text{ such that } S_n - \frac{n}{2}(\mu_1 + \mu_0) \notin (a, b),$$

where $a = \log A/(\mu_1 - \mu_0)$ and $b = \log B/(\mu_1 - \mu_0)$. Note that, in the **symmetric case**, $\mu_1 = -\mu_0$ and $b = -a$, and we have that

$$N = \text{first } n \geq 1 \text{ such that } |S_n| \geq b.$$

Example. Let X_1, X_2, \dots be independent random variables with $P_p(X = 1) = p$ and $P_p(X = -1) = q$, where $p + q = 1$. Consider testing $H_0 : p = p_0$ against $H_1 : p = p_1$, where $p_0 < p_1$. The likelihood ratio is

$$\begin{aligned} \ell_n &= \left(\frac{p_1}{p_0} \right)^{\frac{n+S_n}{2}} \left(\frac{q_1}{q_0} \right)^{\frac{n-S_n}{2}} \\ &= \left(\frac{p_1 q_0}{p_0 q_1} \right)^{\frac{S_n}{2}} \left(\frac{p_1 q_1}{p_0 q_0} \right)^{\frac{n}{2}}, \end{aligned}$$

where $S_n = \sum_{k=1}^n x_k$. Note that, in the symmetric case, $p_0 = q_1$ and $B = A^{-1}$, the stopping rule for the SPRT is given by

$$N = \text{first } n \geq 1 \text{ such that } |S_n| \geq b,$$

where $b = \log B / \log(q_0/p_0)$.

3 Group sequential tests

3.1 Analysing the data in groups

The SPRT is an example of a fully sequential test, since a test is performed after every observation. In practice, especially in the context of clinical trials, it is more convenient

to analyse the data after groups of observations. In fact, a group sequential approach can often achieve most of the efficiency gains of an analogous fully sequential one. In this course, group sequential tests are described in the context of two-treatment clinical trials.

Let X_{A1}, X_{A2}, \dots and X_{B1}, X_{B2}, \dots denote the responses of subjects assigned to two treatments, A and B . Interest lies in testing the null hypothesis of no treatment difference $H_0 : \theta = 0$ against the two-sided alternative $H_1 : \theta \neq 0$ that there is a treatment difference with type I error probability α and power $1 - \beta$ when $\theta = \pm\delta$. Suppose that there are a maximum of K groups and that m denotes the **group size**.

For $k = 1, 2, \dots, K$, a **standardised test statistic** Z_k is calculated from the first k groups of observations and H_0 is rejected if $Z_k \notin (a_k, b_k)$, where a_k and b_k denote the critical values for the k th analysis. If the test continues to the K th analysis and $Z_K \in (a_K, b_K)$, it terminates and H_0 is accepted. The critical values are chosen to achieve the required type I error probability and the power condition determines the group size.

3.2 Designing a group sequential test

Suppose that a group sequential test with a maximum of K analyses yields the sequence of test statistics $\{Z_1, \dots, Z_K\}$. These statistics are said to have the **canonical joint distribution** with information levels $\{\mathcal{I}_1, \dots, \mathcal{I}_K\}$ for the parameter θ if (i) (Z_1, \dots, Z_K) is multivariate normal, (ii) $E(Z_k) = \theta\sqrt{\mathcal{I}_k}$, $k = 1, 2, \dots, K$, and (iii) $\text{cov}(Z_{k_1}, Z_{k_2}) = \sqrt{\mathcal{I}_{k_1}/\mathcal{I}_{k_2}}$, $1 \leq k_1 \leq k_2 \leq K$. The fact that this implies that $\{Z_1, \dots, Z_K\}$ is a **Markov sequence** simplifies the calculations. Let $\Delta_k = \mathcal{I}_k - \mathcal{I}_{k-1}$ for $k = 2, 3, \dots, K$. Then $Z_1 \sim N(\theta\sqrt{\mathcal{I}_1}, 1)$, and, for each $k = 2, 3, \dots, K$,

$$Z_k\sqrt{\mathcal{I}_k} - Z_{k-1}\sqrt{\mathcal{I}_{k-1}} \sim N(\theta\Delta_k, \Delta_k)$$

independently of Z_1, \dots, Z_{k-1} .

A key quantity to calculate for a group sequential test is the probability of crossing a specific stopping boundary at a particular analysis. For each $k = 1, 2, \dots, K$, let

$$\psi_k(a_1, b_1, \dots, a_k, b_k; \theta) = P_\theta(a_1 < Z_1 < b_1, \dots, a_{k-1} < Z_{k-1} < b_{k-1}, Z_k \geq b_k)$$

and

$$\xi_k(a_1, b_1, \dots, a_k, b_k; \theta) = P_\theta(a_1 < Z_1 < b_1, \dots, a_{k-1} < Z_{k-1} < b_{k-1}, Z_k \leq a_k).$$

Now, Z_1 has density

$$f_1(z_1; \theta) = \phi(z_1 - \theta\sqrt{\mathcal{I}_1}).$$

Further, the conditional density of Z_k given $Z_1 = z_1, \dots, Z_{k-1} = z_{k-1}$ depends only on z_{k-1} and is

$$f_k(z_{k-1}, z_k; \theta) = \frac{\sqrt{\mathcal{I}_k}}{\sqrt{\Delta_k}} \phi\left(\frac{z_k\sqrt{\mathcal{I}_k} - z_{k-1}\sqrt{\mathcal{I}_{k-1}} - \theta\Delta_k}{\sqrt{\Delta_k}}\right).$$

Hence, for each $k = 2, 3, \dots, K$,

$$\begin{aligned}\psi_k(a_1, b_1, \dots, a_k, b_k; \theta) &= \int_{a_1}^{b_1} \dots \int_{a_{k-1}}^{b_{k-1}} \int_{b_k}^{\infty} f_1(z_1; \theta) f_2(z_1, z_2; \theta) \dots f_k(z_{k-1}, z_k; \theta) dz_k \dots dz_1 \\ &= \int_{a_1}^{b_1} \dots \int_{a_{k-1}}^{b_{k-1}} f_1(z_1; \theta) f_2(z_1, z_2; \theta) \dots f_{k-1}(z_{k-2}, z_{k-1}; \theta) \\ &\quad \times e_{k-1}(z_{k-1}, b_k; \theta) dz_{k-1} \dots dz_1,\end{aligned}$$

where

$$e_{k-1}(z_{k-1}, b_k; \theta) = \Phi \left(\frac{z_{k-1} \sqrt{\mathcal{I}_{k-1}} + \theta \Delta_k - b_k \sqrt{\mathcal{I}_k}}{\sqrt{\Delta_k}} \right).$$

A similar expression can be obtained for $\xi_k(a_1, b_1, \dots, a_k, b_k; \theta)$.

Although the above multiple integral appears difficult, it can be rewritten to simplify the computation. To see this, let $g_k(z_k; \theta)$, $k = 1, 2, \dots, K$, denote the **sub-densities** of Z_1, \dots, Z_K , their integrals being less than unity for $k > 1$ due to early stopping at stages $1, 2, \dots, k-1$. In other words,

$$g_1(z_1; \theta) = f_1(z_1; \theta)$$

and

$$g_k(z_k; \theta) = \int_{a_{k-1}}^{b_{k-1}} g_{k-1}(z_{k-1}; \theta) f_k(z_{k-1}, z_k; \theta) dz_{k-1}$$

for $k = 2, 3, \dots, K$. It then follows that we can write

$$\begin{aligned}\psi_k(a_1, b_1, \dots, a_k, b_k; \theta) &= \int_{b_k}^{\infty} g_k(z_k; \theta) dz_k \\ &= \int_{a_{k-1}}^{b_{k-1}} \int_{b_k}^{\infty} g_{k-1}(z_{k-1}; \theta) f_k(z_{k-1}, z_k; \theta) dz_k dz_{k-1} \\ &= \int_{a_{k-1}}^{b_{k-1}} g_{k-1}(z_{k-1}; \theta) e_{k-1}(z_{k-1}, b_k; \theta) dz_{k-1}.\end{aligned}$$

Thus, the computation only requires a succession of univariate integrations.

The values of $\psi_k(a_1, b_1, \dots, a_k, b_k; \theta)$ and $\xi_k(a_1, b_1, \dots, a_k, b_k; \theta)$ for $k = 1, 2, \dots, K$ determine the distribution of the stopping time and associated decision for a group sequential test. From these, we can obtain the test's error probabilities for any θ . For example, the test's type I error probability is

$$P_{\theta=0}(\text{Reject } H_0) = \sum_{k=1}^K \{\psi_k(a_1, b_1, \dots, a_k, b_k; 0) + \xi_k(a_1, b_1, \dots, a_k, b_k; 0)\}.$$

Similarly, the test's power when $\theta = \delta$ is

$$\begin{aligned}P_{\theta=\delta}(\text{Reject } H_0) &= \sum_{k=1}^K \{\psi_k(a_1, b_1, \dots, a_k, b_k; \delta) + \xi_k(a_1, b_1, \dots, a_k, b_k; \delta)\} \\ &\simeq \sum_{k=1}^K \psi_k(a_1, b_1, \dots, a_k, b_k; \delta)\end{aligned}$$

if $\delta > 0$ is large. The approximate power when $\theta = -\delta$ has the same form, but with ψ_k replaced with ξ_k and δ with $-\delta$. For specified values of K and the type I error probability α , a **numerical search** can be used to find the a_k and b_k .

Recall that a fixed-sample test of $H_0 : \theta = 0$ against $H_1 : \theta \neq 0$ with type I error probability α and power $1 - \beta$ at $\theta = \pm\delta$ has information

$$\mathcal{I}_{f,2} = \frac{\{\Phi^{-1}(1 - \alpha/2) + \Phi^{-1}(1 - \beta)\}^2}{\delta^2}.$$

Then a group sequential test requires a larger maximum sample size and we set a **maximum information level**, $\mathcal{I}_{max} = R\mathcal{I}_{f,2}$, where $R > 1$ and depends on K , α , β and the type of group sequential boundary being used. With equally-spaced information levels, we have

$$\mathcal{I}_k = \frac{k}{K}\mathcal{I}_{max}, \quad k = 1, 2, \dots, K.$$

By finding the value \mathcal{I}_{max} such that the test's power is $1 - \beta$ for this sequence of information levels, we can obtain R .

The above calculations can be used to design group sequential tests with specific properties. For example, the **Wang and Tsatis family** of two-sided tests are indexed by a parameter Δ , which gives boundaries of different shapes. Members of this family include the Pocock test with constant critical values and the O'Brien and Fleming test with converging critical values. The test with parameter Δ has boundaries of the form

$$a_k = -c \left(\frac{k}{K}\right)^{\Delta - \frac{1}{2}} \quad \text{and} \quad b_k = c \left(\frac{k}{K}\right)^{\Delta - \frac{1}{2}}.$$

Taking $\Delta = 1/2$ gives Pocock's test and $\Delta = 0$ gives the O'Brien and Fleming test.

3.3 Inference following a group sequential test

The group sequential test stops at stage

$$T = \min\{\text{first } k \geq 1 \text{ such that } Z_k \notin (a_k, b_k), K\}.$$

Now, the sequence of test statistics $\{Z_1, \dots, Z_K\}$ has the same joint distribution as the sequence $\{(Y_1 + \dots + Y_k)/\sqrt{\mathcal{I}_k}; k = 1, 2, \dots, K\}$, where the Y_k are independent such that $Y_k \sim N(\Delta_k\theta, \Delta_k)$. The structure of the joint density of (Z_1, \dots, Z_K) shows that (T, Z_T) is a pair of **sufficient statistics** for θ and that the maximum likelihood estimator of θ is given by $\hat{\theta} = Z_T/\sqrt{\mathcal{I}_T}$: see Exercise 3. Although the form of the maximum likelihood estimator is the same as for a fixed-sample test, its sampling distribution is more complicated.

The sampling density of $\hat{\theta}$ at $\hat{\theta} = y$ is given by

$$\sum_{k=1}^K g_k(y\sqrt{\mathcal{I}_k}; \theta) \sqrt{\mathcal{I}_k} 1_{y\sqrt{\mathcal{I}_k} \notin (a_k, b_k)},$$

where the contributions from the K sub-densities yield a multi-modal density with a peak for each value of $T \in \{1, 2, \dots, K\}$. This means that the density is completely different to a normal with mean θ and variance \mathcal{I}_T^{-1} , which it would be for a fixed-sample test. As a result, $\hat{\theta}$ is now a biased estimator of θ . More specifically, we can write

$$E_{\theta}(\hat{\theta}) = \sum_{k=1}^K \left\{ \int_{-\infty}^{a_k} g_k(z_k; \theta) \frac{z_k}{\sqrt{\mathcal{I}_k}} dz_k + \int_{b_k}^{\infty} g_k(z_k; \theta) \frac{z_k}{\sqrt{\mathcal{I}_k}} dz_k \right\} dz_k.$$

An expression can also be obtained for the variance of $\hat{\theta}$.

As before, the above integrals can be rewritten to simplify the computation. For example, the first integral can be written as

$$\begin{aligned} \int_{-\infty}^{a_k} g_k(z_k; \theta) \frac{z_k}{\sqrt{\mathcal{I}_k}} dz_k &= \int_{a_{k-1}}^{b_{k-1}} \int_{-\infty}^{a_k} g_{k-1}(z_{k-1}; \theta) f_k(z_{k-1}, z_k; \theta) \frac{z_k}{\sqrt{\mathcal{I}_k}} dz_k dz_{k-1} \\ &= \int_{a_{k-1}}^{b_{k-1}} g_{k-1}(z_{k-1}; \theta) r_{k-1}(z_{k-1}, a_k; \theta) dz_{k-1}, \end{aligned}$$

where

$$\begin{aligned} r_{k-1}(z_{k-1}, a_k; \theta) &= -\frac{\sqrt{\Delta_k}}{\mathcal{I}_k} \phi \left(\frac{a_k \sqrt{\mathcal{I}_k} - z_{k-1} \sqrt{\mathcal{I}_{k-1}} - \theta \Delta_k}{\sqrt{\Delta_k}} \right) \\ &\quad + \frac{(z_{k-1} \sqrt{\mathcal{I}_{k-1}} + \theta \Delta_k)}{\mathcal{I}_k} \Phi \left(\frac{a_k \sqrt{\mathcal{I}_k} - z_{k-1} \sqrt{\mathcal{I}_{k-1}} - \theta \Delta_k}{\sqrt{\Delta_k}} \right). \end{aligned}$$

The second integral can be computed in a similar way.

Upon termination of the group sequential test, rather than just concluding that we accept or reject H_0 , we can report the p -value of the observed data for testing H_0 . Now, the **sample space** Ω defined by the group sequential design is the set of all pairs (k, z) , where $z \notin (a_k, b_k)$, so that the test can terminate with $(T, Z_T) = (k, z)$. Let the observed value of (T, Z_T) be denoted by (k^*, z^*) . Then the p -value is

$$P_{\theta=0}\{\text{Obtain } (k, z) \text{ as extreme or more extreme than } (k^*, z^*)\}.$$

In order to calculate this, we need to specify the ordering of Ω . We write $(k', z') \succ (k, z)$ to denote that (k', z') is above (k, z) in a given ordering.

There are a number of orderings available. In **stage-wise ordering**, $(k', z') \succ (k, z)$ if any of the following conditions hold: (i) $k' = k$ and $z' \geq z$; (ii) $k' < k$ and $z' \geq b_{k'}$; (iii) $k' > k$ and $z \leq a_k$. As an example, suppose that the test terminates after crossing the upper boundary. Then the one-sided upper p -value is

$$P_{\theta=0}\{(T, Z_T) \succeq (k^*, z^*)\} = \sum_{j=1}^{k^*-1} \psi_j(a_1, b_1, \dots, a_j, b_j; 0) + \psi_{k^*}(a_1, b_1, \dots, a_{k^*-1}, b_{k^*-1}, a_{k^*}, z^*; 0).$$

One-sided lower p -values are found in the same manner and the two-sided p -value is twice the smaller of these two quantities.

Equal-tailed $100(1 - \alpha)\%$ confidence intervals for θ can be obtained by inverting a family of hypothesis tests with two-sided type I error probability α . For any given value θ_0 , we can find pairs $(k_u(\theta_0), z_u(\theta_0))$ and $(k_\ell(\theta_0), z_\ell(\theta_0))$ such that

$$P_{\theta=\theta_0}\{(T, Z_T) \succeq (k_u(\theta_0), z_u(\theta_0))\} = \frac{\alpha}{2}$$

and

$$P_{\theta=\theta_0}\{(T, Z_T) \preceq (k_\ell(\theta_0), z_\ell(\theta_0))\} = \frac{\alpha}{2}.$$

It follows that the acceptance region

$$A(\theta_0) = \{(k, z) : (k_\ell(\theta_0), z_\ell(\theta_0)) \prec (k, z) \prec (k_u(\theta_0), z_u(\theta_0))\}$$

defines a two-sided hypothesis test of $\theta = \theta_0$ with type I error probability α . This implies that the set $\{\theta : (T, Z_T) \in A(\theta)\}$ obtained by inverting this family of tests is a $100(1 - \alpha)\%$ equal-tailed **confidence set** for θ . If $P_\theta\{(T, Z_T) \succeq (k, z)\}$ is an increasing function of θ for each $(k, z) \in \Omega$, then this set is an interval.

3.4 Examples of group sequential designs

Two examples are now given to show how group sequential designs may be constructed in practice.

Example. Let X_{A1}, X_{A2}, \dots be independent normal random variables with mean μ_A and unit variance and let X_{B1}, X_{B2}, \dots be independent normal random variables with mean μ_B and unit variance. For $k = 1, 2, \dots, K$, let n_{Ak} and n_{Bk} denote the cumulative number of observations on treatments A and B , respectively, at the time of the k th analysis. Then the parameter of interest is $\theta = \mu_A - \mu_B$ and its natural estimator is

$$\bar{X}_A^{(k)} - \bar{X}_B^{(k)} = \frac{1}{n_{Ak}} \sum_{i=1}^{n_{Ak}} X_{Ai} - \frac{1}{n_{Bk}} \sum_{i=1}^{n_{Bk}} X_{Bi} \sim N(\theta, \mathcal{I}_k^{-1}),$$

where

$$\mathcal{I}_k = \left(\frac{1}{n_{Ak}} + \frac{1}{n_{Bk}} \right)^{-1}$$

is the information for θ . So the standardised statistic at analysis k for testing $H_0 : \theta = 0$ is

$$Z_k = \left\{ \bar{X}_A^{(k)} - \bar{X}_B^{(k)} \right\} \sqrt{\mathcal{I}_k}$$

for $k = 1, 2, \dots, K$.

It is easily verified that the above statistics have the canonical joint distribution with information levels $\{\mathcal{I}_1, \dots, \mathcal{I}_K\}$ for θ . Firstly, (Z_1, \dots, Z_K) is multivariate normal, since each

Z_k is a linear combination of the independent normal random variables X_{Ai} and X_{Bi} for $i = 1, 2, \dots$. Secondly, we know that $Z_k \sim N(\theta\sqrt{\mathcal{I}_k}, 1)$. Lastly, for $k_1 \leq k_2$,

$$\begin{aligned} \text{cov}(Z_{k_1}, Z_{k_2}) &= \text{cov} \left[\left\{ \bar{X}_A^{(k_1)} - \bar{X}_B^{(k_1)} \right\} \sqrt{\mathcal{I}_{k_1}}, \left\{ \bar{X}_A^{(k_2)} - \bar{X}_B^{(k_2)} \right\} \sqrt{\mathcal{I}_{k_2}} \right] \\ &= \left(\frac{1}{n_{Ak_1}} \frac{1}{n_{Bk_2}} n_{Ak_1} + \frac{1}{n_{Bk_1}} \frac{1}{n_{Ak_2}} n_{Bk_1} \right) \sqrt{\mathcal{I}_{k_1}} \sqrt{\mathcal{I}_{k_2}} \\ &= \mathcal{I}_{k_2}^{-1} \sqrt{\mathcal{I}_{k_1}} \sqrt{\mathcal{I}_{k_2}} = \sqrt{\mathcal{I}_{k_1}/\mathcal{I}_{k_2}}, \end{aligned}$$

as required.

Now suppose that we wish to test $H_0 : \theta = 0$ against $H_1 : \theta \neq 0$ with type I error probability $\alpha = 0.05$ and power $1 - \beta = 0.90$ when $\theta = \pm 0.5$. We will use a Pocock test with a maximum of $K = 5$ analyses. The information for θ required by a fixed-sample test with these error probabilities is

$$\mathcal{I}_{f,2} = \frac{\{\Phi^{-1}(0.975) + \Phi^{-1}(0.9)\}^2}{0.5^2} = 42.032$$

and the maximum information level for the group sequential test can be shown to be

$$R_P(5, 0.05, 0.1) \times \mathcal{I}_{f,2} = 1.207 \times 42.032 = 50.7.$$

Assuming that $n_{Ak} = n_{Bk}$ for each $k = 1, 2, \dots, 5$, we see that $\mathcal{I}_5 = n_5/2$, where n_5 denotes the common value of n_{A5} and n_{B5} . Thus, solving $\mathcal{I}_5 = 50.7$ yields $n_5 = 101.4$, which we round to 110 to obtain a multiple of 10. This means that five groups of 11 observations per treatment should be planned. It may also be shown that H_0 is rejected at analysis k if $|Z_k| \geq 2.413$, $k = 1, 2, \dots, 5$.

Example. Let X_{A1}, X_{A2}, \dots be independent Bernoulli random variables with parameter p_A and let X_{B1}, X_{B2}, \dots be independent Bernoulli random variables with parameter p_B . Then the parameter of interest is $\theta = p_A - p_B$ and its natural estimator is $\hat{p}_A^{(k)} - \hat{p}_B^{(k)} = \bar{X}_A^{(k)} - \bar{X}_B^{(k)}$. Let $\bar{p} = (p_A + p_B)/2$. Then, under H_0 , $p_A = p_B = \bar{p}$, and the information for θ is

$$\mathcal{I}_k = \left\{ \bar{p}(1 - \bar{p}) \left(\frac{1}{n_{Ak}} + \frac{1}{n_{Bk}} \right) \right\}^{-1}.$$

Estimating \bar{p} by

$$\tilde{p}_k = \frac{\sum_{i=1}^{n_{Ak}} X_{Ai} + \sum_{i=1}^{n_{Bk}} X_{Bi}}{n_{Ak} + n_{Bk}},$$

we obtain the estimated information level at analysis k given by

$$\hat{\mathcal{I}}_k = \left\{ \tilde{p}_k(1 - \tilde{p}_k) \left(\frac{1}{n_{Ak}} + \frac{1}{n_{Bk}} \right) \right\}^{-1}$$

for $k = 1, 2, \dots, K$. So the standardised statistics for testing H_0 are

$$Z_k = \left\{ \hat{p}_A^{(k)} - \hat{p}_B^{(k)} \right\} \sqrt{\hat{\mathcal{I}}_k}$$

for $k = 1, 2, \dots, K$. If θ is small, these statistics can be shown to follow approximately the canonical joint distribution with information levels $\{\hat{\mathcal{I}}_1, \dots, \hat{\mathcal{I}}_K\}$ for θ .

Now suppose that we wish to test $H_0 : \theta = 0$ against $H_1 : \theta \neq 0$ with type I error probability $\alpha = 0.05$ and power $1 - \beta = 0.8$ when $\theta = \pm 0.2$. We will use an O'Brien and Fleming test with a maximum of $K = 10$ analyses. The information for θ required by a fixed-sample test with these error probabilities is $\mathcal{I}_{f,2} = 196.224$ and the maximum information level for the group sequential test can be shown to be

$$R_B(10, 0.05, 0.2) \times \mathcal{I}_{f,2} = 1.040 \times 196.224 = 204.1.$$

Taking $\bar{p} = 0.5$ and assuming that $n_{Ak} = n_{Bk}$ for $k = 1, 2, \dots, 10$, we see that $\mathcal{I}_{10} = 2n_{10}$. Thus, solving $\mathcal{I}_{10} = 204.1$ yields $n_{10} = 102.1$, which we round to 120 to obtain a multiple of 20. This means that 10 groups of six observations per treatment should be planned. It may also be shown that H_0 is rejected at analysis k if $|Z_k| \geq 2.087\sqrt{10/k} = 6.600/\sqrt{k}$, $k = 1, 2, \dots, 10$.

4 Adaptive treatment allocation rules

4.1 Definitions

So far, we have been concerned with how to construct sequential tests of some null hypothesis H_0 against an alternative H_1 which have certain error probabilities. We now turn our attention to the problem of how to assign patients to treatments in the context of a fixed-sample clinical trial. The incorporation of stopping rules will be addressed later. Suppose initially that there are $t \geq 2$ treatments.

If **complete randomisation** is used, the next patient is equally likely to be assigned to any of the t treatments, so that the treatment allocation probabilities are all $1/t$. So this randomisation rule does not take into account the previous treatment assignments and responses, or any other information. Consequently, complete randomisation is a non-adaptive treatment allocation rule.

Since complete randomisation can lead to treatment group imbalances, a **restricted randomisation** rule can be used to ensure that each treatment group has roughly the same number of patients. For one of the simplest such rules, the treatment which most reduces the imbalance is assigned with probability p , $0.5 < p < 1$, and the other $t - 1$ treatments are assigned with probability $(1 - p)/(t - 1)$.

As the trial progresses, some treatments may look more promising than others and it would be desirable to allocate a higher proportion of patients to these treatments. In such cases, a **response-adaptive randomisation** rule is used. The simplest such rules may be represented as urn models, in which balls of different types are added to or removed from the urn according to the previous assignments and responses.

4.2 Properties of adaptive treatment allocation

Let N_{jn} denote the number of patients on treatment j after n assignments for $j = 1, 2$. First suppose that complete randomisation is used. Then it is easy to see that $N_{1n}/n \rightarrow 1/2$ almost surely and

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{1}{2} \right) \rightarrow \text{N} \left(0, \frac{1}{4} \right)$$

in distribution as $n \rightarrow \infty$. If restricted randomisation is used instead, then we still have $N_{1n}/n \rightarrow 1/2$, but, since the aim is now to balance the treatment groups, N_{1n}/n will be **less variable**. This means that any tests will have greater power.

As an example of a restricted randomisation rule, suppose that the responses on the two treatments are independent with variance σ^2 . Then, after n assignments, the covariance matrix of the estimated means is $\sigma^2 \text{diag}(1/n_{1n}, 1/n_{2n})$. Since the parameter of interest is the difference between the means, the D_A -optimal design minimises $\sigma^2(1/n_{1n} + 1/n_{2n})$. This design assigns the next patient to treatment 1 if $n_{2n}^2 > n_{1n}^2$ and to treatment 2 if $n_{1n}^2 > n_{2n}^2$. From this deterministic design, we can construct a **biased coin design** which assigns the n th patient to treatment 1 with probability

$$\phi_n = \frac{N_{2,n-1}^2}{N_{1,n-1}^2 + N_{2,n-1}^2}.$$

It can be shown that

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{1}{2} \right) \rightarrow \text{N} \left(0, \frac{1}{20} \right)$$

in distribution as $n \rightarrow \infty$. This means that N_{1n}/n is now asymptotically 80% less variable.

The above biased coin design is a special case of a **generalised biased coin design** which assigns the n th patient to treatment 1 with probability

$$\phi_n = \frac{N_{2,n-1}^\gamma}{N_{1,n-1}^\gamma + N_{2,n-1}^\gamma},$$

where $\gamma \geq 0$. It can be shown that

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{1}{2} \right) \rightarrow \text{N} \left\{ 0, \frac{1}{4(1+2\gamma)} \right\}$$

in distribution as $n \rightarrow \infty$. When $\gamma = 0$, we have complete randomisation. The recommended design is $\gamma = 5$, for which the asymptotic variance is $1/44$.

Now suppose that response-adaptive randomisation is used. In order to assess how good an allocation rule is, we need to study the behaviour of N_{1n}/n for large n . Although response-adaptive randomisation will assign a higher proportion of patients to the better treatment, it induces correlation among treatment assignments, so that N_{1n}/n may be **more variable**. This means that any tests may have lower power. If the distribution of the responses on

treatment j depends on the parameter θ_j for $j = 1, 2$ and $\rho_1(\theta_1, \theta_2)$ denotes the **target allocation** for treatment 1, then we can use the variance of N_{1n}/n to compare rules with the same target allocation.

For a given target allocation, a Cramér-Rao lower bound on the asymptotic variance of the allocation proportions can be obtained. Suppose that

$$\frac{N_{n1}}{n} \rightarrow \rho_1(\theta_1, \theta_2)$$

almost surely and

$$\sqrt{n} \left\{ \frac{N_{1n}}{n} - \rho_1(\theta_1, \theta_2) \right\} \rightarrow N\{0, V_1(\theta_1, \theta_2)\}$$

in distribution as $n \rightarrow \infty$. Then it can be shown that

$$V_1(\theta_1, \theta_2) \geq B(\theta_1, \theta_2) = \frac{\left\{ \frac{\partial \rho_1(\theta_1, \theta_2)}{\partial \theta_1} \right\}^2}{\rho_1(\theta_1, \theta_2) I_1(\theta_1)} + \frac{\left\{ \frac{\partial \rho_2(\theta_1, \theta_2)}{\partial \theta_2} \right\}^2}{\rho_2(\theta_1, \theta_2) I_2(\theta_2)},$$

where $I_j(\theta_j)$ denotes the Fisher information for a single observation on treatment j for $j = 1, 2$. Any rule that attains this lower bound is called **asymptotically best**.

There are essentially two approaches to response-adaptive randomisation, one based on a class of urn models and the other on a class of adaptive biased coin designs. As an example of the latter, suppose that we wish to minimise a weighted average of the numbers of patients on the two treatments subject to attaining a fixed power, where the weights are functions of the θ_j for $j = 1, 2$. In the binary case, the weights would be the failure probabilities. Then the **optimal** treatment allocation probabilities can be derived and the θ_j replaced by their current maximum likelihood estimates. Thus, we obtain a sequential maximum likelihood estimation rule.

4.3 Examples of adaptive treatment allocation

There are a wide variety of adaptive treatment allocation rules available. Four of the more popular ones are described below.

Example. Efron's biased coin design.

Let the **treatment imbalance** after n assignments be $D_n = N_{1n} - N_{2n}$ and let $0.5 < p < 1$ be a constant. Then the probability that the n th patient is assigned to treatment 1 is $1/2$ if $D_{n-1} = 0$, p if $D_{n-1} < 0$ and $1 - p$ if $D_{n-1} > 0$. It can be shown that

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{1}{2} \right) \rightarrow 0$$

in probability as $n \rightarrow \infty$. This means that $\text{var}(N_{1n}/n) = o(1/n)$, which shows why this design is so effective in terms of balancing the numbers of patients on the two treatments. In fact, Efron's biased coin design gives a uniformly more powerful Z or t test than complete randomisation.

Example. Adjustable biased coin design.

Let $F(\cdot)$ be a function $F : \mathbf{Z} \rightarrow [0, 1]$, where \mathbf{Z} is the set of integers, such that F is non-increasing and $F(-x) = 1 - F(x)$. Then the probability that the n th patient is assigned to treatment 1 is $F(D_{n-1})$. Thus, with this design, the tendency towards balance becomes **stronger** the more we move away from it. It can be shown that

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{1}{2} \right) \rightarrow 0$$

almost surely as $n \rightarrow \infty$. When $F(x) = p$ for $x < 0$, we have Efron's biased coin design. The adjustable biased coin design yields a uniformly more powerful Z or t test than Efron's biased coin design.

Example. Drop-the-loser rule.

Consider an **urn model** in which there are initially a balls for both treatment types and b immigration balls. The immigration balls are present in order to ensure that the urn does not empty. When a treatment ball is drawn, it is only replaced if the response is a success. If an immigration ball is drawn, it is replaced along with one ball of each treatment type. Assume that the probability of success for treatment j is p_j for $j = 1, 2$. Then it can be shown that

$$\frac{N_{1n}}{n} \rightarrow \frac{q_2}{q_1 + q_2}$$

in probability and

$$\sqrt{n} \left(\frac{N_{1n}}{n} - \frac{q_2}{q_1 + q_2} \right) \rightarrow N \left\{ 0, \frac{q_1 q_2 (p_1 + p_2)}{(q_1 + q_2)^3} \right\}$$

in distribution as $n \rightarrow \infty$, where $q_j = 1 - p_j$. Thus, the target allocation for treatment 1 is $\rho_1(p_1, p_2) = q_2 / (q_1 + q_2)$. Since $I_j(p_j) = 1 / (p_j q_j)$ for $j = 1, 2$, it is easily verified that

$$B(p_1, p_2) = \frac{q_1 q_2 (p_1 + p_2)}{(q_1 + q_2)^3}.$$

Consequently, the drop-the-loser rule is an asymptotically best procedure for the above target allocation.

Example. Sequential maximum likelihood estimation rule.

Suppose that responses are binary and interest lies in minimising the number of treatment failures for a fixed power. Then, if we use the usual large-sample Z test, this means that we need to find the allocation $\rho = \rho_1(p_1, p_2)$ that minimises $q_1 n_1 + q_2 n_2$ subject to

$$\frac{p_1 q_1}{n_1} + \frac{p_2 q_2}{n_2} = C,$$

where C is a constant. Letting $n_1 = \rho n$ and $n_2 = (1 - \rho)n$, we see that

$$n = \frac{p_1 q_1}{\rho C} + \frac{p_2 q_2}{(1 - \rho)C}.$$

Substituting for n in the formula for the number of treatment failures and differentiating with respect to ρ , we obtain the equation

$$\frac{q_1 p_2 q_2}{(1 - \rho)^2} - \frac{p_1 q_1 q_2}{\rho^2} = 0.$$

Solving yields

$$\rho = \frac{\sqrt{p_1}}{\sqrt{p_1} + \sqrt{p_2}}.$$

Since this is a function of the unknown success probabilities, after $n - 1$ assignments, we replace p_1 and p_2 by their maximum likelihood estimates based on the first $n - 1$ responses. It can be shown that, for this sequential maximum likelihood estimation rule,

$$\frac{N_{1n}}{n} \rightarrow \frac{\sqrt{p_1}}{\sqrt{p_1} + \sqrt{p_2}}$$

almost surely as $n \rightarrow \infty$ and

$$V_1(p_1, p_2) = \frac{p_1^{\frac{3}{2}}(p_2 + q_2/2) + p_2^{\frac{3}{2}}(p_1 + q_1/2)}{(\sqrt{p_1} + \sqrt{p_2})^3 \sqrt{p_1 p_2}}.$$

It is also easily verified that

$$B(p_1, p_2) = \frac{1}{4(\sqrt{p_1} + \sqrt{p_2})^3} \left(\frac{p_2 q_1}{\sqrt{p_1}} + \frac{p_1 q_2}{\sqrt{p_2}} \right).$$

Consequently, this is not an asymptotically best procedure for the above target allocation.

4.4 Group-sequential response-adaptive tests

Up to now, we have considered adaptive treatment allocation in the context of a fixed trial size. Since it is often more efficient to conduct a trial group sequentially, it is natural to investigate the consequences of incorporating adaptive treatment allocation. The formulation of such a group sequential test requires the determination of the joint distribution of sequentially computed test statistics. Because of the **dependencies** induced by adaptive treatment allocation, this is difficult in general.

Let X_{A1}, X_{A2}, \dots be independent normal random variables with mean μ_A and known variance σ_A^2 and let X_{B1}, X_{B2}, \dots be independent normal random variables with mean μ_B and known variance σ_B^2 . Then response-adaptive randomisation can be incorporated into a general family of group sequential tests **without affecting** the error probabilities if the group sizes do not depend on the estimated mean responses at the previous stage in any other way but through their difference.

Let $\theta = \mu_A - \mu_B$ and suppose that we wish to test $H_0 : \theta = 0$. Then the standardised statistic at analysis k for testing H_0 is

$$Z_k = \left\{ \bar{X}_A^{(k)} - \bar{X}_B^{(k)} \right\} \sqrt{\mathcal{I}_k},$$

where

$$\mathcal{I}_k = \left(\frac{\sigma_A^2}{n_{Ak}} + \frac{\sigma_B^2}{n_{Bk}} \right)^{-1}$$

is the information level. Let m_{Ak} and m_{Bk} denote the group sizes on treatments A and B , respectively, at stage k . Then these are allowed to depend on the accumulated data through the current estimate of θ given by $\hat{\theta}^{(k-1)} = \bar{X}_A^{(k-1)} - \bar{X}_B^{(k-1)}$ and chosen to achieve the specified value of \mathcal{I}_k . Under such an adaptive sampling scheme, the above statistics have the canonical joint distribution with information levels $\{\mathcal{I}_1, \dots, \mathcal{I}_K\}$ for θ .

To see how such a **group-sequential response-adaptive test** is constructed in practice, suppose that we wish to minimise $u(\theta)n_{A\tau} + v(\theta)n_{B\tau}$, where $u(\theta)$ and $v(\theta)$ are specified weights, and τ denotes the stage at which the test terminates. Then the allocation ratio which minimises this weighted average is

$$\frac{n_{A\tau}}{n_{B\tau}} = \frac{\sigma_A}{\sigma_B} w(\theta),$$

where

$$w(\theta) = \sqrt{\frac{v(\theta)}{u(\theta)}}.$$

Since this is a function of the unknown θ , after $k - 1$ stages, we replace θ with its maximum likelihood estimate $\hat{\theta}^{(k-1)}$. This means that we choose m_{Ak} and m_{Bk} so that

$$\frac{n_{Ak}}{n_{Bk}} = \frac{\sigma_A}{\sigma_B} w(\hat{\theta}^{(k-1)})$$

for $k = 1, 2, \dots, K$. If $w(\theta) = 1$, then sampling is non-adaptive and only gives equal group sizes if $\sigma_A = \sigma_B$.

At the first stage, we take

$$m_{A1} = \sigma_A \{\sigma_A + \sigma_B w(\hat{\theta}^{(0)})\} \mathcal{I}_1$$

and

$$m_{B1} = \sigma_B \{\sigma_A + \sigma_B w(\hat{\theta}^{(0)})\} \mathcal{I}_1 / w(\hat{\theta}^{(0)}),$$

where $\hat{\theta}^{(0)}$ is a preliminary estimate of θ and the group sizes are rounded to integers. If no such estimate is available, we can use $w(\hat{\theta}^{(0)}) = 1$. At stage $2 \leq k \leq K$, assuming that the test has not yet terminated, we take

$$m_{Ak} = \sigma_A \{\sigma_A + \sigma_B w(\hat{\theta}^{(k-1)})\} \mathcal{I}_k - n_{A,k-1}$$

and

$$m_{Bk} = \sigma_B \{\sigma_A + \sigma_B w(\hat{\theta}^{(k-1)})\} \mathcal{I}_k / w(\hat{\theta}^{(k-1)}) - n_{B,k-1}.$$

Again, these group sizes are rounded to integers. If either of their values is negative, that group size is set to zero and sufficient observations are taken on the other treatment to achieve the specified information level. The calculation of the information levels and boundaries for a group-sequential response-adaptive test is the **same** as for its non-adaptive analogue.

References

- Atkinson, A.C. (1982). Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika* **69**, 61-67. [4]
- Baldi Antognini, A. (2008). A theoretical analysis of the power of biased coin designs. *Journal of Statistical Planning and Inference* **138**, 1792-1798. [4]
- Ghosh, M., Mukhopadhyay, N. and Sen, P.K. (1997). *Sequential Estimation*. Wiley. [2]
- Hu, F. and Rosenberger, W.F. (2006). *The Theory of Response-Adaptive Randomization in Clinical Trials*. Wiley. [4]
- Jennison, C. and Turnbull, B.W. (2000). *Group Sequential Methods with Applications to Clinical Trials*. Chapman and Hall. [3,4]
- Siegmund, D. (1985). *Sequential Analysis: Tests and Confidence Intervals*. Springer. [1,2]