Centro de Investigación en Matemáticas, A.C.

Advances in Bayesian Sequential Analysis in Clinical Trials

TESIS

que para obtener el grado de

Doctor en Ciencias

con orientación en

Probabilidad y Estadística

PRESENTA

Luke Akong'o Orawo

DIRECTOR DE TESIS

Dr. José Andrés Christen Gracia

Febrero de 2006, Guanajuato,Gto., México

Agradecimientos

A mi asesor Dr. José Andrés Christen Gracia por su valiosa dirección para culminar esta tesis. Agradezco además su apoyo incondicional durante mis estudios de doctorado.

A los doctores: Peter Müller y Eduardo Gutierrez Peña por sus excelentes comentarios que me ayudaron a mejorar este trabajo y por ser miembros del comité de mis estudios de doctorado.

A mis sinodales, los doctores: Peter Müller, Eduardo Gutierrez Peña, Miguel Nakamura y Belem Trejo.

A la Secretaria de Relaciones Exteriores de México por la beca durante los primeros tres años de mis estudios de doctorado.

Al CONCYTEG por la beca durante un semestre de mis estudios de doctorado.

Al CONACYT por la beca durante el último semestre de mis estudios de doctorado (SEMARNAT 2004-C-01-7).

Al Centro de Investigación en Matemáticas, A.C.(CIMAT) por su apoyo económico por medio de su programa de becas para estudiantes de doctorado. Agradezco además al CIMAT por haberme brindado sus instalaciones para realizar este trabajo.

A mis compañeros del cubículo: Antonio Murillo, Victor Lopez, José Montoya y Oyuki Hermosillo.

Finalmente, a mi esposa Selline por su comprensión y apoyo.

Abstract

Use of the Bayesian decision-theoretic approaches to obtain optimal stopping rules for clinical trial designs requires a method known as Backward Induction. However, implementation of Backward Induction, save for simple trial designs, is generally impossible due to the computational difficulties associated with it. In this thesis our general research topic is the numerical approximation of Backward Induction in three different clinical trial designs. Berry and Chih-Hsiang (1988), Carlin, Kadane and Gelfand (1998), Brockwell and Kadane (2003) and Christen, Müller, Wathen y Wolf (2004) have addressed the same problem under other clinical trial designs. The first two trial designs that we study seem to be efficient clinically but surprisingly have not been considered in the past. First we consider a multiple-arm trial comparing k experimental treatments with standard treatment, where patient response is binary. Our objective here is to propose a novel stopping rule, denoted by τ^p, as an approximation of the optimal stopping rule, using the optimal stopping rule of single-arm clinical trial obtained by Backward Induction. We use a simulation-based algorithm together with τ^p to estimate the expected utility of continuing. An example of a double-arm clinical trial where we compare our estimates with exact values obtained by implementing Backward Induction is presented. Results of this comparison show that our proposed trial design is a good approximation of the optimal stopping rule. Given that Backward Induction cannot be implemented for cases where we have more than two treatment arms, we evaluated our proposed stopping rules by studying its operating characteristics in a three arm trial. Secondly, we consider a trial design involving two double-arm related trials each comparing an experimental treatment with a standard treatment, where patients in trial 1 are in a milder stage of the disease than those in trial 2. Here we aim at eliciting a joint prior distribution for the unknown success probabilities of each treatment such that we can use data from one trial to learn about the other. Similarly, we present novel stopping rules that simultaneously control both trials, commonly borrowing strength from each other to achieve smaller patient accrual and better performance. Our proposed trial design has the advantage of using less

patients as well as improving performance. Lastly, we approximate the optimal stopping rule of one-arm clinical trial where patient response is normal. Our task here is to construct a grid on the space of the sufficient statistic of unknown parameters and approximate expected utilities over the grid using exact predictive probabilities. This is important because our approach provides smooth stopping boundaries with a much less number of grid subdivisions as opposed to other approaches that use simulation to approximate predictive probabilities. Moreover, this may save computing time. These approximate trial designs display attractive properties, through the examples that we take, and hence offer relevant solution to the problem posed by Backward Induction.

Resumen

El uso del enfoque Bayesiano basado en teoria de decisión para obtener reglas de paro óptimas en diseños de ensayos clínicos requiere de un metodo llamado *Backward Induction*. Sin embargo, la implementación de *Backward Induction*, salvo diseños simples, es generalmente imposible, debido a las dificultades computacionales que éste presenta. En esta tesis nuestro tema general de investigación es aproximar *Backward Induction* en diferentes diseños de ensayos clínicos. Berry y Chih-Hsian (1988), Carlin, Kadane y Gelfand (1998), Brockwell y Kadane (2003) y Christen, Müller, Wathen y Wolf (2004) han trabajado en el mismo problema bajo otros diseños de ensayos clínicos. Los primeros dos diseños que consideramos se ven eficientes clinicamente, pero se sorprende que no haya estudios anteriores de éstos en el pasado. Primero consideramos un ensayo clínico de brazos múltiples, comparando k tratamientos experimentales con un tratamiento estándar. Asumimos que las respuestas de los pacientes son binarias. Nuestro objetivo aquí es proponer una regla de paro, denotada por τ^p, usando la regla de paro óptima para un ensayo clínico de un solo brazo obtenido por *Backward Induction*. Usamos un algoritmo basado en simulación junto con τ^p para estimar la utilidad esperada de continuar. Presentamos un ejemplo de un ensayo clínico de dos brazos donde comparamos nuestras estimaciones con valores exactos obtenidos por *Backward Induction*. Los resultados de esta comparación indican que el diseño de ensayo clínico que proponemos es una buena aproximación al diseño óptimo. Dado que Backward Induction no se puede implementar para casos donde tenemos más de dos brazos, evaluamos nuestas reglas de paro propuestas estudiando sus *Operating Characteristics*. Luego consideramos un diseño secuencial que involucra dos ensayos clinicos relacionados, donde cada uno compara un tratamiento experimental con un tratamiento estándar. Se supone que los patientes del ensayo 1 están menos enfermos que los del ensayo 2. El propósito es asignar una distribución conjunta *a priori* a las probabilidades de éxitos desconocidos de cada tratamiento tal que podamos usar datos de un esayo para aprender del otro. Presentamos una nueva regla de paro que controla simultaneamente ambos ensayos y que se prestan fuerza uno al otro

para poder usar un número pequeño de pacientes y lograr un mejor funcionamiento. Por último, aproximamos la regla de paro óptima para un ensayo clínico de un brazo, donde las respuestas siguien una distribución normal. Construiremos un "grid" en el espacio de la estadística suficiente de los parámetros desconocidos y aproximamos la utilidad esperada sobre el "grid" usando las probabilidades predictivas exactas. Este enfoque es importante, porque comparado a otros enfoques, obtenemos reglas de paro con fronteras suaves usando un menor número de subdivisiones mucho menos a las encontradas en la literatura. Además, puede ser que mejore el tiempo de computación. Los diseños de ensayos que proponemos muestran una buena aproximación, lo cual veremos en los ejemplos, y por lo tanto ofrecen una solución alternativa al problema de *Backward Induction*.

Contents

Chapter 1

Introduction

1.1 Problems in Bayesian analysis of clinical trials

We define a clinical trial as a carefully planned study that evaluates the efficacy of a new treatment relative to a standard treatment. In most clinical trials data are reviewed sequentially as they accumulate to enable reaching a decision to stop the trial early if there exist sufficient evidence of treatment difference or harmful side-effects. Depending on the the objective of the study, patient response may be binary or continuous. For example, if a clinical trial is designed to determine the difference between recovery times for pairs of patients treated with two competing treatments, then the patient response will be continuous. Binary responses are preferred in clinical trials where patient recovery takes a short time. A clinical trial design may have a single treatment arm or multiple treatment arms depending on the number of new treatments that are to be compared with the standard treatment.

Use of the Bayesian approach in the design and analysis of clinical trials is appealing, because it enables incorporation of prior expert opinion regarding treatment efficacy into the analysis

and allows for more simplified designs. See Berry (1985, 1987, and 1993) and Spiegelhalter, Freedman and Parmar (1994) for a detailed review of the advantages of the Bayesian approach to clinical trials. Bayesian decision-theoretic approaches to clinical trials seem suitable because they enable quantifying the value of patient outcomes using a numerical measure called utility and hence permits incorporating other factors such as toxicity and cost of a treatment in making decisions. In the past, Ascombe (1963), Berry and Chih-Hsiang (1988), Carlin, Kadane and Gelfand (1998), Stallard, Thall and Whitehead (1999), Brockwell and Kadane (2003), Stallard (2003), Christen, Müller, Wathen and Wolf (2004), just to name a few, have used decision-theoretic approaches to determine which of the two or more treatments is the most efficacious in clinical trials. According to a Bayesian sequential analysis, the decision to terminate or continue at any stage n of a clinical trial is based on the expected utilities evaluated with respect to future outcomes. This requires the algorithm called Backward Induction (see DeGroot, 1970, Chapter 12), which starts at the final stage where the trial must stop and the optimal decisions are obtained for earlier stages working backwards. Unfortunately, Backward Induction only provides theoretical optimal solutions to Bayesian sequential problems and, except for very simple examples, is generally infeasible to implement due to computational and analytical complexities. Carlin *et al.*(1998), Brockwell *et al.* (2003), Christen *et al.* (2003), Christen *et al.* (2004), and Spiegelhalter, Abrams and Myles (2004, p. 220) have pointed out the computational difficulties associated with Backward Induction as an obstacle to successful use of the Bayesian decision-theoretic approaches to sequential analysis. We note that under categorical patient outcome settings the decision tree grows exponentially as we progress through the interim stages towards the horizon. The result is a complex decision tree consisting of a huge number of possible future scenarios that would be extremely difficult to track backwards. Moreover, if we consider that patient response is continuous and that at any stage past data is summarized by a sufficient statistic then we have an infinite space where the statistic takes values. Hence the Backward Induction algorithm becomes impossible to implement. Although numerical methods for approximating optimal sequential decisions have been developed using the idea

of discretization of the space in which a summary statistic assumes values into a grid of points, analyzing sequential clinical trials with multiple arms using the Bayesian approach is complex. Alternative Bayesian approaches, which will be presented in detail in section 1.3, are based on stopping decisions derived from the posterior probability content of the parameter of interest.

1.2 The problems

The three problems that we solve in this thesis are presented follows. First, we consider a multiple-arm clinical trial comparing k experimental treatments $t = E_1, E_2, \ldots, E_k$ with a standard treatment $t = E_0$. We assume that patient response is binary. We suppose that the unknown success probabilities, denoted by $\theta_j, \quad j = 0, 1, \ldots, k$, of the $k+1$ treatments are independent and model their initial uncertainty by assigning them prior distributions $\pi_j(\theta_j), \quad j = 0, 1, \ldots, k$ respectively. We let N be the maximum number of patients enrolled in the trial. At any stage $n < N$ of the trial, we let $d_n = 0, 1$ denote the decisions of continuing and stopping respectively and t_{n+1} be the treatment that we choose for future patients upon stopping or for the next patient if the decision is to continue. Then based on what we have observed, that is the sequence of observations, $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$, obtained from the first n patients upon application of the sequence of treatments $t_n = (t_1, t_2, \ldots, t_n)$, the object is to find the pair in $\{(d_n, t_{n+1}) : d_n = 0, 1, t_{n+1} = E_0, E_1, \ldots, E_k\}$ that maximizes the expected value of of our proposed utility function $u(.)$, where expectation is taken with respect to future observations $x_{n_1}, x_{n+2}, \ldots, x_{N+1}$ and $i = N+1$ represents a future patient outside that trial. We note that for the same reason mentioned in the previous section the exact implementation of Backward Induction is impossible for cases where $k \geq 3$.

We state the second problem that we consider as follows. Suppose we have a double-arm clinical trial designed to evaluate the efficacy of an experimental treatment E with a standard

treatment S, where patients are divided in two cohorts, marked 1 and 2. We assume that the patient response is binary. We suppose that patient in cohort 1 are in a milder stage of the disease than those in cohort 2. Let θ and θ' denote the unknown success probabilities of treatment E in cohorts 1 and 2 respectively. Similarly, we let π and π' be the unknown success probabilities of the standard treatment S in cohorts 1 and 2 respectively. We model the initial uncertainty about each of the pair (θ, θ') and (π, π') by assigning a joint prior distribution, $p(y, w)$, defined over the support $0 < w < y < 1$. We may also regard this problem as one trial divided into two related (sub)trials, one for each cohort. Let N and M be the maximum numbers of patients enrolled in trials 1 and 2 respectively. We use x_i and t_i to denote response and treatment respectively of the i^{th} patient, $i = 1, 2, \ldots, N$ in trial 1. Similarly, we use x'_j and t'_j to denote response and treatment respectively of the j^{th} patient, $j = 1, 2, \ldots, M$ in trial 2. Suppose the two trials are at stages n and m respectively and that $d'_m = 0, 1$ correspond to the decisions that we may take in trial 2. Under this setting the optimal sequential decision is obtained by finding the pairs $(d_n, d'_m) \in \{(1,1), (1,0), (0,1), (0,0)\}$ and $(t_{n+1}, t'_{m+1}) \in \{(E,E), (E,S), (S,E), (S,S)\}$ that maximize the expected value of our proposed utility function, $u(.)$, for the entire trial. We note that the ordering of the success probabilities enables using data from one trial to learning about the other. This means at some stages n and m where the decision is to continue both trias, we may proceed by confining patient entry in one trial and use the accumulating data to learn about the other trial.

We solve the above two problems by proposing a novel stopping rule and a pair of novel stopping rules, respectively, using the optimal stopping rule of a single-arm clinical trial design obtained by Backward Induction. A detailed description of the construction of each of the stopping rules and their evaluation for better understanding will be presented in chapters 3 and 4 respectively.

Finally, we consider a one-arm clinical trial where an experimental treatment E is compared to a standard treatment S where we assume that under treatments E and S the patient

responses are $x|t = E \sim N(\mu_E, \lambda_E)$ and $x|t = S \sim N(\mu_S, \lambda_S)$ respectively. We will consider two cases: one where we assume that μ_S is unknown and λ known, and the other where both μ and λ_S are unknown. We assign a conjugate normal or normal-gamma prior distribution to the unknown parameter (or parameters) depending on whether λ_E is known or not. At any stage n of the trial if the decision is to continue ($d_n = 0$), then the next patient is assigned to treatment E. However, if the decision is to stop, any of the two treatments may chosen. The object is to find the pair (d_n, t_{n+1}) that maximizes the expected utility, where expectation is taken with future observations. We construct a grid over the space where a sufficient statistic takes values and implementing Backward Induction to obtain the gridding approximation of the expected utility over the grid. We solve this problem by implementing the gridding method using the exact predictive probabilities; unlike in previous implementation of this method (see Brockwell *et al.* 2003).

1.3 Literature review

Sequential methods for use in clinical research stems from an early publication by Armitage (1975). He presents classical sequential methods that we do not use in this thesis. However, a brief review of these methods may be useful for a general presentation of the subject. According to these methods at any interim stage data is reviewed by carrying out a two-sided significance test so that if the absolute value of the test statistic is greater than some fixed value, the trial stops and evidence of treatment difference is declared. Otherwise the trial continues. However, if the trial continues up to the last stage, then one stops without claiming evidence of treatment difference. Berry (1985, 1987, and 1993) and Spiegelhalter, Freedman and Parmar (1994) have pointed out the advantages of the Bayesian approaches over the classical methods with regard to sequential analysis of clinical trials. Ascombe (1963) while reviewing Armitage's book argued for the use of the Bayesian decision-theoretic approaches. Pocock (1977) uses the same approach and proposes a group sequential designs

where patients are divided into equal-sized groups so that an interim analysis is carried out after each group is observed. He discusses these designs for normal and other types of patient response. Such group sequential designs have been used in decision-theoretic approach to reduce the computational burden associated by Backward Induction (See Carlin *et al.* 1998 and Brockwell *et al.* 2003). Other examples of group sequential designs are found in Elfring and Schultz (1973), Freedman, Lowe and Macaskill (1984) and Demets and Ware (1980). See the book by Whitehead (1997) for a more detailed presentation on the design and analysis of sequential clinical trials using the frequentist approaches.

1.3.1 Approaches for implementing Backward Induction in clinical trials

With the discovery of cheap computing power a number of researchers have proposed computational methods for solving different Bayesian sequential problems where Backward Induction is implemented via estimation or approximation of expected utility. Carlin *et al.* (1998) proposed a simulation-based algorithm for finding the optimal sequential decision at any stage of a sequential clinical trial design comparing an experimental treatment with a placebo, where patient outcome is assumed continuous and data are monitored at K predetermined interim stages. The algorithm, as an alternative approach, drastically reduces the computational complexity associated with backward induction, however, it is limited to group-sequential setting where patient outcomes are normal with known variance and allows few interim looks. As few as $K = 2$ interim stages were used to compare their algorithm with backward induction. Brockwell and Kadane, (2003) considers the same design as Carlin *et al.* (1998) and propose a more general algorithm for implementing Backward Induction. The algorithm consists of constructing a grid in the space in which a summary statistic assumes values and evaluating the expected losses corresponding to the decisions stopping and continuing over the grid. Surprisingly, they use a simulation method to calculate the

expected losses for continuing over the grid when it could have been easier to use the common posterior predictive density of future outcomes, which looks more accurate. Another example where gridding method has been used is found in Berry and Chih-Hsiang (1988).

Christen *et al.* (2004) considered a more complex clinical trial design with multiple treatment arms where patient response is categorical and at any stage a set of non-dominated treatments is obtained. They work with a set of utility functions. Their trial design is such that if the decision is to continue, a treatment selected randomly from the non-dominated set of treatments is assigned to the next patient. Otherwise the trial stops and the set of non-dominated treatments is reported. Unlike their trial design, any of the our trial designs work with one utility function throughout, and at any stage of the trial we may stop or continue with one treatment that maximizes the expected utility. They have used a 2-step look-ahead procedure (see Berger 1985, Chapter 7) to implement Backward Induction. We use the same utility function, that they propose, throughout this thesis. One common achievement in all these articles is that the mentioned complexity associated with Backward Induction is eliminated by collapsing down the exponentially growing decision tree to a linearly growing one. Examples where optimal sequential decisions are evaluated by a complete implementation of Backward Induction in clinical trial designs with at most two treatment arms and assuming binary patient responses are provided by Wathen and Christen (2004), Stallard *et al.* (1999), Lewis and Berry (1994) and Petkau (1978), Stallard (2003) proposes decision-theoretic designs where a number of potential new treatments are compared with a standard treatment in a series of single-arm phase II clinical trials. In each successive single-arm trial Backward Induction is used to obtain optimal strategies. This may be expensive and may be avoided by using the multiple-arm trial design similar to the one that we propose in this thesis.

1.3.2 Other Bayesian approaches to clinical trials

Here we briefly describe other Bayesian approaches to clinical trials that we do not consider in this thesis but complements the general framework and may provide documented examples of some of our trial designs. A number of researchers have developed alternative Bayesian approaches to sequential clinical trials where the boundaries of stopping rules are evaluated based on the posterior probability content of the unknown treatment success probability. They avoid the decision-theoretic approaches that we use in this thesis by pointing out the extreme difficulty in implementing Backward Induction and that of choosing meaningful and appropriate utility functions. The probability-only approaches have been used mostly in single-arm trial designs with univariate or multivariate discrete responses. The decision procedure is described briefly as follows. Suppose we consider the case of univariate treatment response. A decision criteria at any stage of the trial would be of the following form. At the beginning of the trial the upper and the lower decision cutoffs, denoted by P_L and P_U respectively, are specified. Suppose that θ_S and θ_E denote the success probabilities of treatments S and E respectively. Let δ be the least targeted improvement over S. Then the trial terminates if the posterior probability $Pr(\theta_S + \delta < \theta_E \mid \text{data})$ is outside the interval (P_L, P_U). Otherwise the trial continues. On the other hand if the treatment response is multivariate, then the monitoring criteria with respect to each compound event of interest is derived from the marginal posterior probability of that event. Berry (1989), Thall and Simon (1994), and Thall,*et al.* (1995) have used these methods to implement Bayesian sequential designs for single-arm clinical trials.

Other probability-only approaches are based either on the relative position of credible interval of the treatment effect relative to some boundary or the Bayesian test of some classical hypotheses. Examples are found in Spiegelhalter *et al.* (1994), Choi and Pepple (1989), Cornfield (1966) and the book by Spiegelhalter, Abrams and Myles (2004). This book not only presents a detailed coverage of the probability-only approaches but also describes the

general Bayesian approaches to clinical trials, outlining prior elicitation, ethical issues and justifying randomization.

1.4 Outline of the thesis

In Chapter 2, we describe the general decision-based Bayesian approach to sequential clinical trials and present a simulation-based algorithm that, for a given stopping rule, permits finding an estimate of the expected utility in clinical trial designs with multiple arms in the last section. In chapter 3, we approximate the optimal stopping rule obtained by Backward Induction by proposing a novel stopping rule, denoted τ^p, whose construction is based on the optimal stopping rule of single-arm clinical trial design. The expected utility estimating algorithm enables us to compare our proposed stopping rule τ^p with the optimal stopping rule in a double-arm clinical trial design. We note that it is common practice to validate stopping rules constructed from a Bayesian view point by evaluating their frequentist properties (see Thall *et al.* (1994,1995), Stallard *et al.* (1999), Christen *el al.* (2004)). For this reason we validate our clinical trial designs with more than two arms, where Backward Induction is impossible to implement, by investigating the operating characteristics of a three-arm clinical trial.

Similarly, we consider two related trials in Chapter 4, each comparing an experimental treatment with a standard treatment, and propose a pair novel proposed stopping rules for approximating the optimal sequential decision at any interim stages. As in Chapter 3, we evaluate our proposed trial design by studying its operating characteristics.

Finally, in Chapter 5, we use a gridding method to approximate the sequential optimal decision at any stage of a one-arm sequential clinical trial where patient response is considered continuous. We use the posterior predictive density to compute the expected utility of

continuing over the grid unlike other approaches where a simulation method has been used. Chapter 6 presents a discussion of our findings.

Chapter 2

Bayesian Sequential Decision

2.1 Introduction

The basic elements of a Bayesian sequential decision problem are a set of possible decisions, a utility function for assessing the consequences of all possible decisions and the prior distributions for the parameters of interest. At any interim stage of a sequential procedure, the optimal sequential decision is obtained by choosing a decision that maximizes the expected value of a utility function, where expectation is taken with respect to future observations. This requires a standard method called Backward Induction (DeGroot, 1970), in which one must start at the last stage and evaluate expected utilities backwards. This method can be extremely difficult to implement if at some stage one has to consider a huge number of future scenarios.

The outline of this chapter is as follows. We describe the Bayesian sequential decision problem with reference to a multiple-arm clinical trial comparing k experimental treatments with a standard treatment in section 2 and briefly outline its solution via the Backward Induction

method in section 3. In sections 4 and 5, we define a stopping rule and the expected utility determined by a stopping rule respectively. In section 6, we present the implementation of Backward Induction for a single-arm trial design. We describe a simulation-based algorithm for estimating expected utility in section 7. This chapter is a preparation for the more original work of Chapters 3, 4 and 5.

2.2 The Utility Function

We consider a clinical trial comparing k experimental treatments $E_1, E_2, \ldots, E_k$ with a standard treatment E_0 where patient response, x, is binary. Suppose a_j and b_j are real numbers in the interval $[0, 1]$ with $a_j < b_j$, $j = 0, 1, \ldots, k$, we define the utility $v(E_j, x)$ of observing a response x under a given treatment E_j, $\quad j = 0, 1, \ldots, k$ as

$$v(E_j, x) = \begin{cases} a_j & \text{if } x = 0; \\ b_j & \text{if } x = 1. \end{cases} \tag{2.1}$$

We note that the two values a_j and b_j are assumed fixed at the beginning of a trial and may be determined based on the level of toxicity, the cost of treatment E_j and other factors that may be considered relevant. Suppose the trial stops at stage n. We have the sequence of observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$ from the first n patients and a sequence of respective treatments $\mathbf{t}_n = (t_1, t_2, \ldots, t_n)$. We define the utility function u for the entire trial, based on the single utility function v in (2.1), as

$$u(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_{N+1}, v) = \\ \alpha v(t_{n+1}, x_{N+1}) + \frac{(1-\alpha)}{N} \left\{ \sum_{i=1}^{n} v(t_i, x_i) + \sum_{i=n+1}^{N} v(t_{n+1}, x_i) \right\} \tag{2.2}$$

where $\alpha = \frac{1}{N+1}$, $x_{n+1}, \ldots, x_{N+1}$ are the future observations under treatment t_{n+1} and $i = N+1$ indicates the inclusion of a future patient outside the trial. That is, $u(.)$ is the weighted average of the utilities $v(.)$ of each of the patients enrolled in the trial and the utility of the

treatment chosen for future patients. At any stage n of the trial we have to choose between the decisions of stopping ($d_n = 1$) and continuing ($d_n = 0$) by comparing their respective expected utilities. If the trial stops, then the best treatment is chosen for al future patients; otherwise the best treatment is allocated to the next patient. The best treatment t^*_{n+1} is obtained by maximizing the expected utility

$$U_n(d_n = 1, \mathbf{t}_{n+1}, \mathbf{x}_n) = E\{u(d_n = 1, \mathbf{t}_{n+1}, \mathbf{x}_{N+1}, v) \mid \mathbf{t}_{n+1}, \mathbf{x}_n\}, \tag{2.3}$$

where the expectation is taken with respect to the future observations. The evaluation of the above expectation requires the joint posterior predictive distribution of future observations $x_{n+1}, \ldots, x_{N+1}$. However, by noting that $u(.)$ is a weighted sum of $v(., x_i)$ and that $x_{n+1}, \ldots, x_{N+1}$ are identically distributed, the expected utility $U_n(.)$ in (2.3) can be evaluated by using the common posterior predictive distribution $p(x \mid \mathbf{t}_{n+1}, \mathbf{x}_n)$ as

$$\begin{aligned} U_n(d_n = 1, \mathbf{t}_{n+1}, \mathbf{x}_n) &= (1-\alpha)\sum_{i=1}^{n} v(t_i, x_i) + \sum_{x=0,1} (\alpha + (1-\alpha)\frac{N-n}{N}) \\ &\quad \times v(t_{n+1}, x) p(x \mid \mathbf{t}_{n+1}, \mathbf{x}_n). \end{aligned} \tag{2.4}$$

2.3 Backward Induction

The Backward Induction algorithm begins at the last stage N where the trial must stop and the optimal decision is chosen by maximizing the expected utility. For any possible sequences of treatments and observations, $(\mathbf{t}_N, \mathbf{x}_N)$, at stage N, the maximum expected utility, denoted by $U^*_N(\mathbf{t}_N, \mathbf{x}_N)$, is obtained as

$$U^*_N(\mathbf{t}_N, \mathbf{x}_N) = \max_{t_{N+1}=E_0,E_1,\ldots,E_k} U_N(d_N = 1, t_{N+1}, \mathbf{t}_N, \mathbf{x}_N). \tag{2.5}$$

By Bayesian sequential analysis, the expected utility of continuing the trial with treatment t_{n+1} at any stage $n < N$ is given by

$$\begin{aligned} U_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n) &= U^*_{n+1}(\mathbf{t}_{n+1}, \mathbf{x}_n, x_{n+1} = 0)(1 - p(x_{n+1} = 1 \mid \mathbf{t}_{n+1}, \mathbf{x}_n)) \\ &\quad + U^*_{n+1}(\mathbf{t}_{n+1}, \mathbf{x}_n, x_{n+1} = 1)p(x_{n+1} = 1 \mid \mathbf{t}_{n+1}, \mathbf{x}_n). \end{aligned} \tag{2.6}$$

Therefore the maximum expected utility, U^*_n, is

$$U^*_n(\mathbf{t}_n, \mathbf{x}_n) = \max_{d_n=0,1} \{ \max_{t_{n+1}=E_0,E_1,\ldots,E_k} U_n(d_n, \mathbf{t}_n, \mathrm{x}_n, t_{n+1}) \}. \tag{2.7}$$

We note that, at any stage n, computing $U_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n)$ requires knowing $U^*_{n+1}(\mathbf{t}_{n+1}, \mathbf{x}_n, x_{n+1})$ for every possible outcome x_{n+1} under treatment t_{n+1} at the next stage $n + 1$ and hence, working backwards, $U^*_n(.)$ can be computed recursively for earlier interim stages $n = N - 1, N - 2, \ldots, 1, 0$. Due to the fact that the decision tree, for our trial design with binary patient outcomes, grows exponentially as the number of interim stages increases and hence resulting into a huge number of future scenarios, backward induction is impossible to implement when we have more than two treatment arms.

2.4 Stopping Rule

The concept of a stopping rule is concerned with choosing a given decision from a set of possible decisions at any stage n of a sequential design based on the observed sequence of observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$. We define a stopping rule relative to our sequential clinical trial design as follows. Suppose that at stage n of the trial we let the sequences of binary observations and treatments $(\mathbf{t}_n, \mathbf{x}_n)$ be summarized by the sufficient statistic $(n_0, n_1, \ldots, n_k, s_0, s_1, \ldots, s_k)$, where n_j and s_j, $j = 0, 1, \ldots, k$ respectively represent the number of observations and the number of observed successes due to treatment E_j. For the purpose of improving our presentation we will use the symbols $n_{(k)}$ and $s_{(k)}$ to denote the vectors $(n_0, n_1, \ldots, n_k)$ and $(s_0, s_1, \ldots, s_k)$ respectively. We note that for every n and j, the

pair $(n_j, s_j) \in (0, 1, \ldots, n)^2$ is the sufficient statistic for the marginal posterior distribution of θ_j and is such that $s_j \leq n_j$ and $\sum_{j=0}^{k} n_i = n$. This implies that the sufficient statistic $(s_{(k)}, n_{(k)})$ belongs to a $2k+2$ dimensional space $\{0, 1, \ldots, n\}^{2k+2}$. We therefore define a stopping rule as a sequence of functions

$$\tau_n : \{0, 1, \ldots, n\}^{2k+2} \rightarrow \{0, 1\} \tag{2.8}$$

such that the trial is stopped when $\tau_n(s_{(k)}, n_{(k)}) = 1$ after observing $(s_{(k)}, n_{(k)})$; otherwise the $(n+1)^{th}$ observation is made.

For example, suppose that at any stage n of a double arm clinical trial ($k = 1$) comparing a new treatment E_1 with a standard treatment E_0 the observed data $(\mathbf{t}_n, \mathbf{x}_n)$ is summarized by the sufficient statistic $(s_{(1)}, n_{(1)})$. As an example we may define a stopping rule τ^e as

$$\tau_n^e(s_{(1)}, n_{(1)}) = \begin{cases} 0 & \text{if } \mid s_1 - s_0 \mid \leq 2; \\ 1 & \text{otherwise.} \end{cases} \tag{2.9}$$

At any stage n of our sequential multiple-arm clinical trial design the optimal stopping rule is obtained by computing the expected utilities of the two decisions of stopping and continuing the trial based on $(s_{(k)}, n_{(k)})$ and choosing the decision that maximizes the expected utility.

2.5 The expected utility determined by a stopping rule

We note that the utility function in (2.2) depends on the observed data $(\mathbf{x}_n, \mathbf{t}_n)$ only through the sufficient statistic $(s_{(k)}, n_{(k)})$ and hence any expected utility can be expressed in terms of $(s_{(k)}, n_{(k)})$. Let the expected utility of stopping in (2.4) be denoted by $U_n(d_n = 1, s_{(k)}, n_{(k)}, t_{n+1})$. If at any stage n of the trial a given stopping rule τ indicates stopping, then the expected utility determined by τ, denoted by $U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1})$, is equal to the

expected utility of stopping. On the other hand, if τ indicates continuing the trial then we define $U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1})$ as follows.

Suppose we summarize the possible sequences of observations after stage n by the sufficient statistic $(r_{(k)}, m_{(k)})$ where $r_{(k)} = (r_0, r_1, \ldots, r_k)$, $m_{(k)} = (m_0, m_1, \ldots, m_k)$ and $\sum_{j=0}^{k} m_j^h = m_j < N - n$. We define m_j and r_j as the number of observations and observed successes respectively, for treatment E_j, $j = 0, 1, \ldots, k$. Suppose that $\tau = 1$ upon observing the values $(r_{(k)}^h, m_{(k)}^h)$, $h = 1, 2, \ldots, L$ of $(r_{(k)}, m_{(k)})$. We note that both the past data $(s_{(k)}, n_{(k)})$ and what we observe after stage n are considered while evaluating the value of τ. If we let $P_h = Pr(r_{(k)} = r_{(k)}^h, m_{(k)} = m_{(k)}^h)$ represent the probability of observing $(r_{(k)}^h, m_{(k)}^h)$, $h = 1, \ldots, L$ and $p_j = Pr(x = 1 \mid s_{(k)}, n_{(k)}, t_{n+1} = E_j)$ denote the posterior success probability of treatment E_j, $j = 0, 1, \ldots, k$ based on past data, then we have $P_h = \prod_{j=0}^{k} p_j^{r_j^h} (1 - p_j)^{m_j^h - r_j^h}$. Using mathematical induction we can show that

$$\sum_{h}^{L} P_h = \sum_{h=1}^{L} \prod_{j=0}^{k} p_j^{r_j^h} (1 - p_j)^{m_j^h - r_j^h} = 1 \tag{2.10}$$

holds for all integers $L \geq 2$.

Without loss of generality we let $\tau^{(L)}$ be any stopping rule that indicates stopping after stage n for L values of $(r_{(k)}, m_{(k)})$. If $L = 2$, then $\tau^{(2)} = 1$ upon observing x_{n+1} under some treatment $t_{n+1} = E_j$, irrespective of the outcome. We observe $(r_{(k)}^h, m_{(k)}^h)$, $h = 1, 2$ with probabilities $P_1 = Pr(r_{(k)} = r_{(k)}^1, m_{(k)} = m_{(k)}^1) = Pr(x_{n+1} = 1 \mid s_{(k)}, n_{(k)}, E_j) = p_j$ and $P_2 = Pr(r_{(k)} = r_{(k)}^2, m_{(k)} = m_{(k)}^2) = Pr(x_{n+1} = 0 \mid s_{(k)}, n_{(k)}, E_j) = 1 - p_j$ respectively. Equation (2.10) holds for the case $L = 2$ since $\sum_{h=1}^{2} P_h = p_j + 1 - p_j = 1$. Next we suppose that equation (2.10) holds for any stopping rule $\tau^{(b)}$ and show that it also holds for $\tau^{(b+1)}$. We note that for every stopping rule $\tau^{(b+1)}$ there exist a sequence $\{(r_{(k)}^h, m_{(k)}^h,), h = 2, 3, \ldots, b\}$ and a stopping rule $\tau^{(b)}$ such that $\tau^{(b)}(r_{(k)}, m_{(k)}) = 1$ for all h and

$$\tau^{(b+1)}(r^h_{(k)}, m^h_{(k)}) = \begin{cases} 0 & \text{if } h = t; \\ 1 & \text{otherwise.} \end{cases} \tag{2.11}$$

This implies that $\tau^{(b+1)} = 1$ when the value $(r^h_{(k)}, m^h_{(k)})$ is combined with the observation x_{m_t+1} from one more patient under some treatment, say E_i. Consequently,

$$\begin{aligned} \sum_{h=1}^{k'+1} P_h &= \sum_{h\neq t} \prod_{j=0}^{k} p_j^{r_j^h} (1-p_j)^{m_j^h - r_j^h} + \prod_{j=0}^{k} p_j^{r_j^t} (1-p_j)^{m_j^t - r_j^t} (1-p_i) + \\ &\quad \prod_{j=0}^{k} p_j^{r_j^t} (1-p_j)^{m_j^t - r_j^t} p_i \\ &= \sum_{h\neq t} \prod_{j=0}^{k} p_j^{r_j^h} (1-p_j)^{m_j^h - r_j^h} + \prod_{j=0}^{k} p_j^{r_j^t} (1-p_j)^{m_j^t - r_j^t} \\ &= \sum_{h=1}^{k'} \prod_{j=0}^{k} p_j^{r_j^h} (1-p_j)^{m_j^h - r_j^h} \\ &= 1. \end{aligned} \tag{2.12}$$

Hence equation (2.10) holds for all integers $L \geq 2$. If we define $u_h = U_{n+m_h}(d_{n+m_h} = 1, s_{(k)}, n_{(k)}, r_{(k)}, m_{(k)}, t^*_{n+m_h+1})$, then by the above result the set of pairs $\{(u_h, P_h), h = 1, 2, \ldots, L\}$ constitutes a probability mass function of a random variable, denoted by U. Hence by equation (2.6) the expected utility $U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1})$ is equal to the expected value of U, and is given as

$$U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1}) = \sum_{h=1}^{L} u_h P_h. \tag{2.13}$$

The development of our expected utility estimating algorithm described in the next chapter is based on this result.

2.6 Implementation of Backward Induction for a single-arm trial

2.6.1 Trial design

Suppose we consider a single-arm trial design where an experimental treatment E is compared with a standard treatment S. As before, we assume that the patient response is binary and that the design is sequential. We also assume that the success probability, denoted by θ_0, of the standard treatment is known and fixed while that of treatment E, denoted by θ_1, is unknown. Suppose that θ_1 has a beta prior distribution with parameters α_1 and β_1. If the trial is at stage n, then we have the sequence of observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$ from the first n patients, all assigned to treatment E. Summarizing past data by the sufficient statistic (n, s), where n and s are the number of patients and observed successes respectively, we obtain the posterior predictive distribution, $p(x_{n+1} \mid n, s)$, of a future observation x_{n+1} under treatment E as a Bernoulli distribution with success probability $\frac{\alpha_1+s}{n+\beta_1-s}$. Therefore, we have that the expected utility in (2.4) of stopping the trial depends on the data, $\mathbf{x}_n$, through the sufficient statistic (n, s) and can be denoted as $U_n(d_n = 1, n, s, t_{n+1})$. We note that the use of (n, s) reduces our exponentially growing decision tree to a two dimensional table, making the implementation of Backward Induction easy and fast. At the last stage N, we compute

$$U_N^*(d_N = 1, N, s, t_{N+1}) = \max_{t_{N+1}=S,E} U_N(d_N, N, s, t_{N+1}) \tag{2.14}$$

for all values of s. Moving one step backwards the expected utility $U_{N-1}(d_{N-1} = 0, N-1, s)$ in (2.6) can also be evaluated for all s and hence the maximum expected utility is obtained as in (2.7). Working backwards that way, $U_n^*(.)$ can be computed recursively for earlier stages $n = N-1, N-2, \ldots, 2, 1, 0$. The result is an $(N+1)\times(N+1)$ table whose $(n+1, s+1)^{th}$ entry in the upper diagonal is equal to the maximum expected utility, $U_n^*(n, s, t_{n+1})$, which may

correspond to the decision of stopping or continuing the trial. Entries below the diagonal represent impossible cases where $s > n$.

2.6.2 Example

We consider an example of the above trial design by setting $N = 12$, $\theta_0 = 0.65$, assuming that $v(.)$ in (2.1) is such that $v(t, 1) = 1$ and $v(t, 0) = 0$ for any $t = S, E$, and that $\theta_1 \sim beta(0.75, 0.25)$. Computing the maximum expected utility, $U_n^*(n, s, t_{n+1})$, at any stage n of the trial for each possible value of (n, s), as outlined in the above subsection, we obtain the optimal stopping rule given in the form of the table below.

	n												
s	0	1	2	3	4	5	6	7	8	9	10	11	12
0	C	S	S	S	S	S	S	S	S	S	S	S	S
1	I	C	C	S	S	S	S	S	S	S	S	S	S
2	I	I	C	C	C	S	S	S	S	S	S	S	S
3	I	I	I	C	C	C	S	S	S	S	S	S	S
4	I	I	I	I	C	C	C	S	S	S	S	S	S
5	I	I	I	I	I	C	C	C	C	S	S	S	S
6	I	I	I	I	I	I	C	C	C	C	S	S	S
7	I	I	I	I	I	I	I	E	C	C	C	C	S
8	I	I	I	I	I	I	I	I	E	E	E	E	E
9	I	I	I	I	I	I	I	I	I	E	E	E	E
10	I	I	I	I	I	I	I	I	I	I	E	E	E
11	I	I	I	I	I	I	I	I	I	I	I	E	E
12	I	I	I	I	I	I	I	I	I	I	I	I	E

Table 2.1: In the table, the symbols S and E indicate the decision to stop with the standard and the experimental treatments respectively, and C indicates the decision to continue. I represent impossible cases.

2.7 Estimation of expected utility given a stopping rule

2.7.1 The algorithm

Consider a clinical trial aimed at comparing k experimental treatments with a standard treatment. Suppose the trial is at stage n and that based on past data $(\mathbf{x}_n, \mathbf{t}_n)$, summarized by the sufficient statistic $(s_{(k)}, n_{(k)})$, a given stopping rule τ indicates the decision to continue. The expected utility determined by τ, denoted by $U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1})$, is estimated as follows. For each $q = 1, 2, \ldots, G$, simulate a set of $k+1$ values, denoted by $\theta_j^{(q)}$, $j = 0, 1, \ldots, k$, from probability distributions $\pi_j(\theta_j)$, $j = 0, 1, \ldots, k$ respectively. Simulate the observations $x_r^{(q)}$ from a $ber(\gamma_r^{(q)})$ distribution, $r = 1, 2, \ldots, n_q$ for each q until τ prescribes stopping at stage $n_q \leq N$. Suppose that, based on the $l-1$ previous simulated observations $x_1^{(q)}, \ldots, x_{l-1}^{(q)}$, τ indicates continuation; then the observation $x_l^{(q)}$ is simulated from a $ber(\gamma_l^{(q)})$ distribution where the success probability, $\gamma_l^{(q)}$, is equal to a value in $\{\theta_j^{(q)}, j = 0, 1, \ldots, k\}$ corresponding to the treatment that maximizes the expected utility. The above step yields G sequences $\{\mathbf{x}_{n_q}^{(q)} = (x_1^{(q)}, x_2^{(q)}, \ldots, x_{n_q}^{(q)}), \quad q = 1, \ldots, G\}$ of simulated observations. Then summarize each sequence $\mathbf{x}_{n_q}^{(q)}$ by the sufficient statistic $(s_{(k)}^{(q)}, n_{(k)}^{(q)})$ and compute the value, denoted by $u^{(q)}$, of the expected utility $U_{n_q}(d_{n_q} = 1, s_{(k)}^{(q)}, n_{(k)}^{(q)}, t_{n_q+1})$ of stopping. Based on the G simulations $u^{(q)}$, $q = 1, 2, \ldots, G$, of expected utilities, a point estimate of the expected utility determined by τ is given by

$$\widehat{U}_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1}) = \frac{1}{G} \sum_{q=1}^{G} u^{(q)}. \tag{2.15}$$

Note that, for a given stopping rule τ, our algorithm permits simulating the values $u^{(q)}$, $q = 1, 2, \ldots, G$, from the discrete probability distribution defined in section 2.4 whose mean is equal to the expected utility $U_n^{(\tau)}(s_{(k)}, n_{(k)}, t_{n+1})$, and thus $\widehat{U}_n^{(\tau)} \to U_n^{(\tau)}$ as $G \to \infty$.

Suppose s_G denotes the sample standard deviation of the above G simulations of the expected utilities. We have by the central limit theorem that as $G \to \infty$ the random variable

$\frac{\sqrt{G}(\widehat{U}_n^{(\tau)}-U_n^{(\tau)})}{s_G}$ is approximately distributed as a standard normal random variable. Using this normal approximation, a suitable choice of G that would yield a good point estimate in (2.15) is chosen big enough such that the expected utility of stopping lies outside the interval $\widehat{U}_n^{(\tau)} \pm z_{\frac{\phi}{2}} \frac{s_G}{\sqrt{G}}$ where $z_{\frac{\phi}{2}}$ is obtained as the solution to the equation

$$Pr(\mid \frac{\sqrt{G}(\widehat{U}_n^{(\tau)} - U_n^{(\tau)})}{s_G} \mid \leq z_{\frac{\phi}{2}}) \approx 1 - \phi. \tag{2.16}$$

Once this interval has been attained, we can use our estimate to tell with probability $1 - \phi$ that the expected utility determined by τ, $U_n^{(\tau)}$, is less or greater than the expected utility of stopping. We note that, although this simulation-based algorithm is a known procedure for estimating expected utility it has not been consistently documented.

We will use our algorithm together with the proposed stopping rule, to be described in the next section, to estimate the expected utility of continuing at a given stage in a double-arm clinical trial design and compare our estimates with the exact values to see how well the proposed stopping rule, τ^p, performs as an approximation of the optimal stopping rule τ^{BI} obtained by backward induction. Note that for a double-arm trial design we obtain the exact values of the expected utility of continuing at any stage by implementing backward induction using a software coded in C++ and developed by Wathen Christen and Christen (2004).

2.7.2 Example

For the purposes of illustration we estimate the expected utility of continuing at stage $n = 0$ of a double-arm clinical trial using the simulation-based algorithm, outlined the previous section, together with the stopping rule

$$\tau_n^e(s_{(1)}, n_{(1)}) = \begin{cases} 0 & \text{if } \mid s_1 - s_0 \mid \leq 2; \\ 1 & \text{otherwise,} \end{cases}$$

where $n_{(1)} = (n_0, n_1)$ and $s_{(k)} = (s_0, s_1)$. Note that $\tau^e = 0$ at stage $n = 0$. We fix the maximum number of patients at $N = 12$, and suppose that probabilities of success θ_0 and θ_1 have beta distributions $beta(0.75, 0.25)$ and $beta(0.65, 0.35)$ respectively. Also, we assume that for any treatment t we have the utility function v defined as $v(0, t) = 0$ and $v(1, t) = 1$. The expected utility of stopping at stage $n = 0$, computed as outlined in the previous chapter, is equal to 0.75. Using our algorithm we generated $G = 10,000$ simulations of the the expected utilities and obtained a point estimate of the expected utility of continuing as 0.7811. Hence, based on this estimate, the decision would be continue the trial to the next stage.

Chapter 3

Multiple-arm Clinical Trial with Binary Response

3.1 Introduction

Conducting a clinical trial with multiple treatment arms is appropriate in situations where many new treatments are to be compared with a standard treatment. For instance, in dose-finding studies where several dose levels are compared with a control level (see Christen *et al.*, 2004) or in cases where there are many potential new treatments of a severe disease without a standard treatment. It has the advantage of using less patients and saving both time and resources in comparison to conducting a series of single-arm trials.

In this chapter, we consider multiple-arm clinical trials with binary responses where data are monitored regularly using a Bayesian approach. We note that, for the case of clinical trials with two treatment arms the optimal sequential decision can be found by implementing Backward Induction using a software written in C++ and developed by Wathen and Christen

(2004). However, for trial designs with more than two treatment arms Backward Induction is impossible to implement. We approximate the optimal stopping rules for such trial designs by proposing a novel stopping rule using the optimal stopping rule of a single-arm trial obtained by Backward Induction. We also estimate the expected utility using the simulation-based algorithm outlined in last section of chapter 2 together with our proposed stopping rule. Obtaining an approximate stopping rule in this manner is original and has never been used or published elsewhere.

We describe the construction of our proposed stopping in section 2. In section 3, we take an example of a double-arm clinical trial where we compare our approximations of the expected utility of continuing, at the initial stage, with the exact values. We present the results of the operating characteristics of our proposed trial design in section 4. The chapter ends with a discussion of our proposed trial design in section 5.

3.2 The Proposed Stopping Rule τ^p

The construction of the proposed stopping rule τ^p for multiple-arm clinical trials is based on the optimal stopping rule for a single-arm clinical trial design for comparing an experimental treatment with a standard treatment whose success probability is assumed fixed and known. Suppose a single-arm trial is at stage n and that the observed data $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ from the first n patients are summarized by the sufficient statistic (n, s) where n and s denote the number of observations and the number of observed successes respectively. Using the sufficient statistic (n, s) transforms the decision tree to a simple two dimensional table and hence we can implement backward induction algorithm as outlined in section 2.5 without any difficulty.

Suppose we consider the multiple-arm clinical trial where k experimental treatments are

compared with a standard treatment. At any stage n of the trial the sequential decision according to the proposed stopping rule τ^p based on past data, $(s_{(k)}, n_{(k)})$, is evaluated as follows. The expected utility, denoted by $U_n(d_n = 1, s_{(k)}, n_{(k)}, t_{n+1} = E_j)$, of stopping the trial at stage n with treatment E_j, is computed for all $j = 0, 1, ..., k$. Suppose that from stage n onwards the multiple-arm clinical trial continues with one of the $(k+1)$ treatments, say E_l, until stopping. It follows by the above supposition that the multiple-arm clinical trial is converted to a single-arm clinical trial comparing treatment E_l with the treatment that corresponds to $\max_{j \neq l} U_n(d_n = 1, s_{(k)}, n_{(k)}, t_{n+1} = E_j)$. We then apply the method of backward induction in the resulting single-arm trial to compute the expected utility of continuing denoted by $u^*_{n(l)}$. Repeating this process for $l = 0, 1, \ldots, k$ gives the values $u^*_{n(0)}, u^*_{n(1)}, \ldots, u^*_{n(k)}$. Suppose for the purpose of illustration we consider a clinical trial with multiple arms that is two stages away from the horizon N. The decision tree at stage n with the branch indicating continuation with treatment E_l extended up to the last stage is given in figure 3.1. The value $u^*_{n(l)}$ is obtained by beginning at the last stage with the shaded square nodes and working backwards along the branches marked with doted lines.

If we let $u^*_n = \max\{u^*_{n(0)}, u^*_{n(1)}, \ldots, u^*_{n(k)}\}$, then the value u^*_n becomes our estimate of the expected utility of continuing the trial $U_n(d_n = 0, s_{(k)}, n_{(k)}, t^*_{n+1})$. Therefore, based the observed data $(s_{(k)}, n_{(k)})$, the sequential decision according to τ^p is to continue the trial if the value u^*_n is greater than the maximum expected utility $U_n(d_n = 1, s_{(k)}, n_{(k)}, t^*_{n+1})$ of stopping the multiple-arm trial. Otherwise the decision is to stop the trial.

We note that the estimate u^*_n , as figure 3.1 indicates, is computed by using the branches of the decision tree at stage n that corresponds to the possibility of continuing the trial with one of the $k + 1$ treatments until stopping and is therefore less than the expected utility of continuing, $U_n(d_n = 0, s_{(k)}, n_{(k)}, t^*_{n+1})$, at stage n obtained by backward induction. This implies that if at a given stage of a clinical trial the proposed stopping rule, τ^p, indicates the decision to continue the trial, then the optimal stopping rule, τ^{BI}, will also indicate continuation. That is, our decision to continue a trial according to τ will be optimal and an

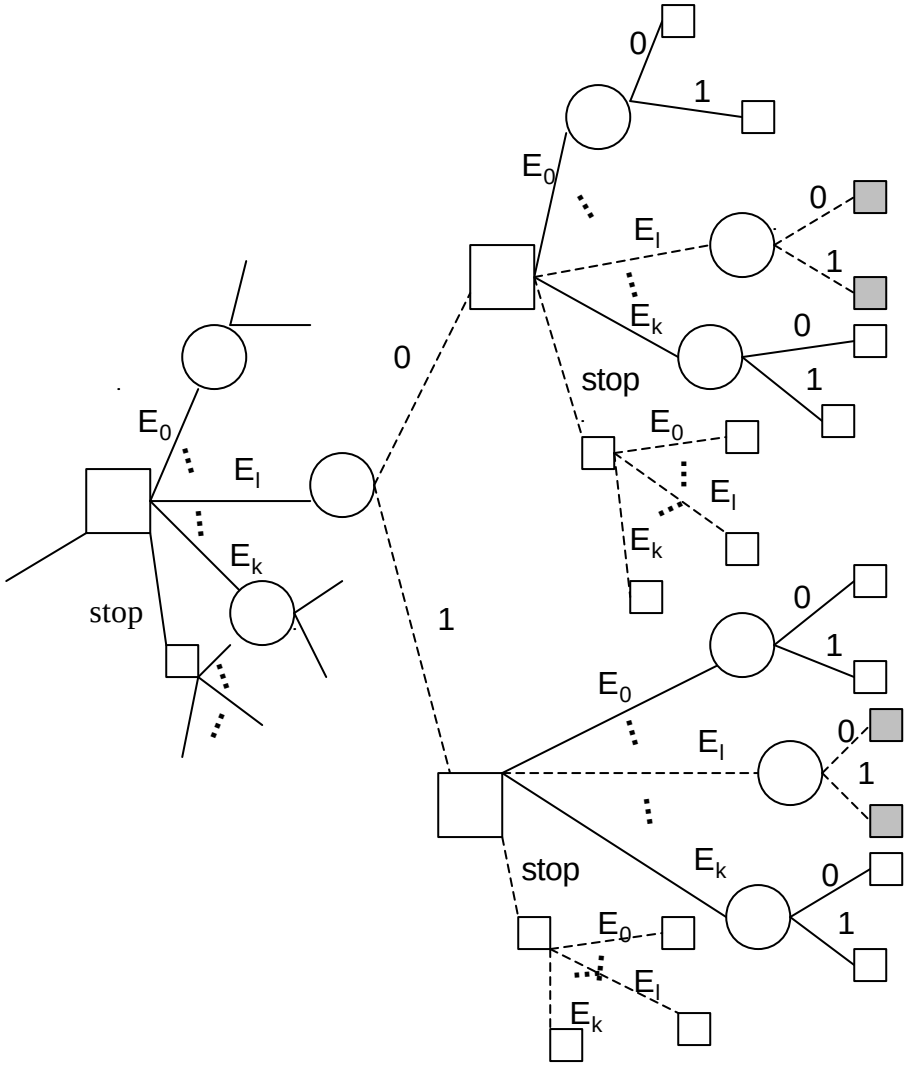

Figure 3.1: Decision tree at stage n. The square nodes represent treatment allocation decisions and the circles stand for binary patient outcomes.

error may only be committed when τ^p indicates stopping.

3.3 Example

We consider the estimation of the expected utility of continuing by our algorithm using the proposed stopping rule τ^p in a double-arm clinical trial design comparing an experimental treatment E with a standard treatment S. We then compare our estimates with the exact values, obtained by a software developed by Wathen and Christen (2004), to evaluate the accuracy of our approximation of the optimal stopping rule by the proposed stopping rule τ^p. We assume that the success probabilities θ_0 and θ_1 for treatments S and E are both unknown and have $beta(\alpha_0, \beta_0)$ and $beta(\alpha_1, \beta_1)$ distributions respectively. We consider three cases determined by values chosen for the parameters $(\alpha_0, \beta_0, \alpha_1, \beta_1)$ for all of which the trial continues at stage $n = 0$. In the first and the second cases we estimate the optimal sequential decision when the differences between expected utilities of stopping and continuing are very big and very small, respectively. We note that when the said difference of expected utilities is very small it becomes difficult to detect it through estimation and hence the second case is considered to evaluate the accuracy of our approximation by τ^p. The third case presents estimation of expected utility with noninformative priors for θ_0 and θ_1 where the likelihood dominates. Suppose the trial is at stage n and that the observed data $(\mathbf{x}_n, \mathbf{t}_n)$ is summarized by the sufficient statistic $(s_{(1)}, n_{(1)})$. The posterior predictive distribution of a future observation x_{n+1}, $\quad p(x_{n+1} \mid \mathbf{t}_{n+1}, \mathbf{x}_n)$, is obtained as a Bernoulli distribution (as explained in section 1.3) with success probability

$$\theta = \begin{cases} \frac{\alpha_0 + s_0}{n_0 + \alpha_0 + \beta_0} & \text{if } t_{n+1} = S; \\ \frac{\alpha_1 + s_1}{n_1 + \alpha_1 + \beta_1} & \text{if } t_{n+1} = E. \end{cases} \tag{3.1}$$

$(\alpha_0, \beta_0, \alpha_1, \beta_1)$	$U_n^{(1)}$	$U_n^{(0)}$	u_n^*	$\widehat{U}_n^{(\tau^p)}$
(0.10,0.90,0.75,0.25)	0.75	0.7523	0.7504	0.7507
(0.5,0.5,0.5,0.5)	0.50	0.6505	0.6234	0.6487
(0.75,0.25,0.65,0.35)	0.75	0.8426	0.8233	0.8412

Table 3.1: Estimation of the expected utility of continuing: In this table the first column indicates the fixed parameter values for the respective prior distributions $beta(\alpha_0, \beta_0)$ and $beta(\alpha_1, \beta_1)$ for θ_0 and θ_1 and the columns $U_n^{(1)}$ and $U_n^{(0)}$ report the exact expected utilities of stopping and continuing respectively. The estimates obtained by evaluating the value of τ^p and by our algorithm at stage n are denoted u_n^* and $\widehat{U}_n^{(\tau^p)}$ respectively.

Suppose in equation (2.1) we let $a_j = 0$ and $b_j = 1$ for all j. Then the expected utility of stopping given in equation (2.4) becomes

$$U_n(d_n = 1, t_{n+1}, s_{(1)}, n_{(1)}) = (\alpha + (1-\alpha)\frac{N-n}{N})\theta + \frac{1-\alpha}{N}(s_0 + s_1). \tag{3.2}$$

We fix the horizon to $N = 12$ and report our estimates in the table 3.1. The values have been selected purely for the purpose of illustration. For every set of parameter values the proposed stopping rule, τ^p, is constructed as described in section 3.2 and the estimate u_0^* of the expected utility of continuing is indicated. Also, we indicate in the table the corresponding values of the expected utility of stopping, the expected utility of continuing obtained by backward induction and the estimate of the expected utility obtained by our simulation-based algorithm using τ^p, denoted by the symbols $U_n^{(1)}$, $U_n^{(0)}$ and $\widehat{U}_n^{(\tau^p)}$ respectively. Note that in each case estimation by our algorithm was based on G=10,000 simulations.

We observe from table 3.1 that the estimates u_n^* and $\widehat{U}_n^{(\tau^p)}$ are very close to the exact value for all choices of $(\alpha_0, \beta_0, \alpha_1, \beta_1)$, and by the fact that both estimates lead to optimal continuation of the double-arm trial at stage $n = 0$ even when the difference between the expected utilities of stopping and continuing is very small, we conclude that τ^p is a good approximation of the optimal stopping rule at this stage.

3.4 Operating characteristics for a three-arm clinical trial

In this section we study the behavior of our proposed clinical trial design when the number of treatment arms is more than two. Here we do not have the benefit of an exact computation by Backward Induction. We consider a clinical trial with three treatment arms where two experimental treatments E_1 and E_2 are compared with a standard treatment E_0 and evaluate the operating characteristics of our trial design. We consider five different scenarios determined by fixing the values of the success probabilities, denoted by θ_0, θ_1 and θ_2. Under each scenario we simulate 1,000 possible observations of the entire trial and give a tabular report indicating the average number of patients allocated to each treatment, the corresponding standard deviation, the probability of stopping with the best treatment evaluated under repeated simulations and the fixed success probabilities. We also report in the tables the average number of patients observed before the trial stops.

We assume in all scenarios that each of the prior parameters for the success probabilities of the three treatments has a $beta(0.5, 0.5)$ distribution, that the utility function v is defined as $v(0, t) = 0$ and $v(1, t) = 1$ for any treatment t and that a maximum of $N = 50$ patients are enrolled. Scenario 1 depicts a case where all the three treatments are equally efficacious, i.e. no treatment provides treatment advance over the others. Here we set $\theta_0 = \theta_1 = \theta_2 = 0.3$ and obtained an average of 41.56 observed patients and the trial stops with each of the three treatments at approximately equal probability. Scenario 2 depicts a case where the standard treatment is superior to both experimental treatments. We fixed the probabilities θ_0, θ_1 and θ_2 to $0.5, 0.3$ and 0.3 respectively. The trial stopped with average of 36.44 patients and did so with the standard treatment 62.7% of the time. Scenarios 3 and 4 depicts cases where an experimental treatment is superior. The probabilities are fixed according to the sets $(0.3, 0.6, 0.4)$ and (0.1,0.2,0.6) respectively. The trial stopped with 34.66 and 28.98 patients

in scenarios 3 and 4 respectively and did so in both cases with the superior treatment with the highest probability (70.1%, 92.0% respectively). In scenario 5 depicts a situation where both the experimental treatments are equally effective and both are superior to the standard treatment. We fixed probabilities at 0.3, 0.6 and 0.6 respectively. An average of 34.96 patients were observed and with each of the treatments E_1 and E_2 resulting to be the best 46.6% and 49.6% of the times.

We note that at any given stage of the trial design the maximum expected utility may correspond to at least two treatments and in this case if continuation is optimal then we use randomization with equal probability to assign a treatment to the next patient.

In all the five scenarios, based on the clinical trial with three arms, our design demonstrates attractive properties in the sense that the results obtained for any fixed set of values of the treatment success probabilities agree with what might realistically be anticipated. In other words, if the treatments are equally efficacious as in scenario 1 columns $\tilde{n}_t$ and P_t indicate that the three treatments were equally preferred and in each of the other scenarios the same columns indicate that the three treatments were preferred according to the assumed order of superiority.

Scenario 1

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	13.82	16.60	0.3	0.333
E_1	13.89	16.46	0.3	0.341
E_2	13.85	16.67	0.3	0.327
$\tilde{n}$	41.56	9.94		

Scenario 2

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	18.08	15.81	0.5	0.627
E_1	9.12	14.40	0.3	0.184
E_2	9.24	14.39	0.3	0.189
$\tilde{n}$	36.44	11.89		

Scenario 3

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	5.49	10.86	0.3	0.084
E_1	19.64	16.90	0.6	0.701
E_2	9.53	14.78	0.4	0.215
$\tilde{n}$	34.66	13.27		

Scenario 4

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	2.09	3.91	0.1	0.009
E_1	5.03	10.34	0.2	0.071
E_2	21.86	14.58	0.6	0.920
$\tilde{n}$	28.98	13.31		

Scenario 5

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	3.66	8.66	0.3	0.038
E_1	15.05	17.51	0.6	0.466
E_2	16.24	17.89	0.6	0.496
$\tilde{n}$	34.96	13.45		

Table 3.2: Operating characteristics. In all tables, E_0, E_1 and E_2 denote treatments, $\tilde{n}_t$ indicates the average number of patients assigned to each treatment, St.dev is the corresponding standard deviation, θ_t denotes the assumed true value of the treatment success probability and the column P_t indicates the probability of choosing each treatment as the best on stopping the trial. See Appendix B for more details on the interpretation of operating characteristics.

3.5 Discussion

Here we have considered the approximation of the optimal stopping rules for clinical trials with more than two treatment arms using the optimal stopping rule by our proposed stopping rule, τ^p, whose construction is based on the optimal stopping rule of a single-arm clinical trial. The comparison study displayed in table 3.1 and the results of the operating characteristics in the previous section, show that our proposed stopping rule, τ^p, is a good approximation of the optimal stopping rule obtained by Backward Induction. Our approximation will enable using the Bayesian decision theoretic approach to clinical trials with more than two treatment arms. The rationale behind our approximation can be understood by noting that at any stage n of a multiple-arm trial, the use of the single-arm conversion provides a systematic way of choosing some of the numerous future scenarios; hence starting at the last stage and working backwards, the approximate expected utility of continuing is computed. Our proposed trial design is easy to implement as computing the value of τ^p at any stage involves simple programming and hence allows high number of interim looks.

Our approximation of the optimal stopping rule is not limited to a specific utility function and in principle may be implemented with other utility functions, leading to similar optimal stopping rules for a single-arm trial and is suited to a problem at hand. Similarly, aside from the beta prior distributions that we have used to model the success probabilities, other distributions that may be considered adequate can be used. Detailed methods of prior elicitation are presented by Kadane and Wolfson (1998), Chaloner K. (1996)and Geisser,S. (1984).

We carried out a sensitivity analysis of our proposed trial design by repeating scenario 3 considered in the previous section with a slightly changed beta prior distribution, $beta(0.6, 0.4)$, and obtained the following results in table 3.3. Compared to the results of scenario 3 in table 3.2, the difference between the columns marked $\tilde{n_t}$, St.dev and P_t are minimal. Similarly,

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	5.99	11.25	0.3	0.079
E_1	19.20	14.95	0.6	0.717
E_2	9.54	14.53	0.4	0.204
$\tilde{n}$	34.73	12.29		

Table 3.3: Sensitivity analysis

.

with reference to the informative prior distributions stated in (3.3), we observed minimal differences in the results upon making a number of slight changes of the priors and repeated scenario 3. Scenario 3 in section 3.4 was considered using noninformative prior distributions for the unknown success probabilities. However, if repeated using informative beta prior distributions,

$$\theta_0 \sim beta(1.5, 3.5),\ \theta_1 \sim beta(3, 2) \quad \text{and} \quad \theta_2 \sim beta(2, 3), \tag{3.3}$$

indicating evidence that treatment E_1 is superior to both E_0 and E_2 and that E_2 is the least effective (see figure 3.2), a much better performance is obtained, as seen in table 3.4. The trial recommends the best treatment upon stopping 96.7% of the time and stops earlier, compared to scenario 3 in table $3, 2$, by a big margin of eight patients. We note that, use of such informative prior distributions, as is the case in most clinical trials, will produce a remarkable increase in the efficiency of our proposed trial design to assign the best treatment to a patient and recommend the same upon stopping. See Appendix B for more analysis of scenario 3.

We remark that based on the general construction of our proposed stopping rule τ^p, any multiple-arm trial design with more than three treatment arms will produce impressive results of the study of operating characteristics just like the ones displayed here and in section 3.4. Consequently, based the numerical evaluation of our trial design presented in this section and in sections 3.3 and 3.4, we conclude that our proposed stopping rule τ^p

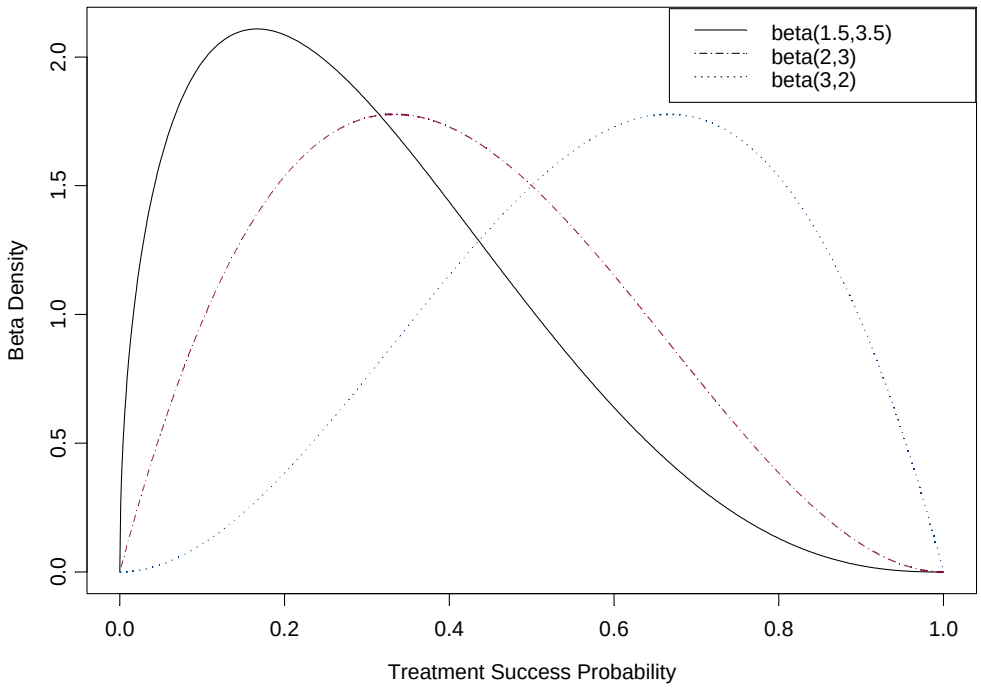

Figure 3.2: Informative beta prior distributions for the unknown treatment success probabilities θ_0, θ_1 and θ_2.

is a good approximation of the optimal stopping rule for multiple-arm trial designs with more than two arms where backward induction is infeasible to implement. Simulations for representative scenarios 1, 3 and 5 and a software written in R for generating them are found in $http://www.cimat.mx/\sim jac/material/Orawothesis.zip$.

Trt	$\tilde{n}_t$	St.dev.	θ_t	P_t
E_0	0.08	1.44	0.3	0.002
E_1	24.91	8.28	0.6	0.967
E_2	1.59	6.38	0.4	0.031
$\tilde{n}$	26.58	8.17		

Table 3.4: Scenario 3 with informative prior distributions assigned to the unknown success probabilities.

Chapter 4

Related Trials

4.1 Introduction

Suppose we have a sequential clinical trial where a new treatment E is to be compared with a standard treatment S such that a successful treatment of a patient depends on some covariates, for instance, weight, age of a patient, the stages of development of the disease etc. Assuming that the patient response is binary, then any of the two treatments has different unknown success probabilities corresponding to the various levels of the covariate under consideration. For instance, the unknown success probabilities may be ordered, from the highest to the lowest, according to the stages of the disease. Using the initial information about how the unknown success probabilities depend on the levels of the covariate, it may be appropriate to divide patients into cohorts in such a way that we can *borrow strength* or learn about certain cohorts from the data in the other cohorts. Hence gaining performance in comparison to conducting a series of independent trials, in the sense that fewer patients may be used and both time and resources may be saved. The problem may be regarded as one trial divided in a series of related (sub)trials, one for each cohort. We note that solving

this sequential problem using Bayesian decision theory is complex even when the patients are divided into only two cohorts (see the decision tree in figure 4.1). Hence to illustrate the importance of conducting related trials we consider a case of two related trials.

In this chapter, we approximate the optimal sequential decision at any paired stages of two related trials by proposing two novel stopping rules, one for each trial. We outline the trial design in the section 2. Then in section 3, we define the prior distributions and outline how the posterior distributions are obtained. In section 4, we present the version of the simulation-based algorithm for estimating the expected utility, in the last section of Chapter 2, suited for the trial design that we consider here. Section 5 outlines the construction of our proposed stopping rules. Section 6 presents an example illustrating how we approximate the optimal sequential decision at the initial stages of two related trials. In section 7, we present the results of the operating characteristics of our trial design. We discuss our findings in section 8.

4.2 The trial design

Suppose the effectiveness of an experimental treatment E is to be evaluated relative to a standard treatment in a clinical trial where patients are divided into two cohorts. Patients in cohort 1 are in a milder stage of the disease than those in cohort 2. Suppose we denote by θ, π and θ', π' the unknown probabilities of success for treatments E, S in cohorts 1 and 2 respectively. Then, based on the above mentioned relation between the two cohorts, we have that $\theta' < \theta$ and $\pi' < \pi$.

We assume that $p(\theta, \theta', \pi, \pi') = p(\theta, \theta')p(\pi, \pi')$. However, we will restrict ourselves to a class of joint prior distributions for which the two trials can learn from each other as they progress. We use x to represent a patient response under treatment t in trial (cohort) 1 and similarly

let x' to denote patient response under treatment t' in trial (cohort) 2. Let the values $v(t,x)$ and $v'(t',x')$, as in chapter 2, indicate for the corresponding trials the value of the utility for having made a given observation under a given treatment. Also, we let N and M denote the fixed maximum numbers of patients in trials 1 and 2 respectively. At any stages $n < N$ and $m < M$ of trials 1 and 2, respectively, we have to decide whether to stop both trials or continue with at least one trial for one more period. In other words, if we let 0 and 1 denote the decisions of stopping and continuing a trial respectively, then at any stages of the two trials we are faced with a set of four paired decisions $\{(1,1),(0,1),(1,0),(0,0)\}$ to choose from. For any pair of decisions out of the four that turns out to be optimal there are four possible treatment pairs, namely $(E,E),(E,S),(S,E),(S,S)$, that may be indicated by the design as being the best pair of treatments for the two trials respectively. We present part of the decision tree in figure 4.1 to illustrate the decision process described above.

The notation we use here will be the same as in the previous chapters except that we use a "prime" superscript on symbols denoting treatments, observations and decisions for trial 2. We use $U_{n,m}(.)$ to denote the expected utility at stages n and m of trials 1 and 2 respectively. We let $U^*_{n,m}(.)$ to represent the corresponding maximum expected utility. Also, unlike in the previous chapters, we use $\mathbf{x}_{(j,k)} = (x_j, x_{j+1}, \ldots, x_{j+k-1})$ and $\mathbf{t}_{(j,k)} = (t_j, t_{j+1}, \ldots, t_{j+k-1})$ to denote sequences of k observations and the applied treatments beginning from the j^{th} stage in trial 1. The corresponding sequences of observations and treatments in trial 2 will be denoted in a similar manner. Suppose the two trials are at stages n and m, then we have the sequences of binary observations $\mathbf{x}_{(1,n)} = (x_1, x_2, \ldots, x_n)$ and $\mathbf{x}'_{(1,m)} = (x'_1, x'_2, \ldots, x'_m)$ from the first n and m patients in trials 1 and 2 respectively resulting from applying the sequences of treatments $\mathbf{t}_{(1,n)} = (t_1, t_2, \ldots, t_n)$ and $\mathbf{t}'_{(1,m)} = (t'_1, t'_2, \ldots, t'_m)$ where each t_i and each t'_h, $i = 1,2,\ldots,n$, $h = 1,2,\ldots,m$ may be either treatment E or S. We use $\mathbf{D}_{n,m}$ to represent the observed data $(\mathbf{t}_{(1,n)}, \mathbf{x}_{(1,n)}, \mathbf{t}'_{(1,m)}, \mathbf{x}'_{(1,m)})$ in both trials. If both trials stop, then based on the utilities $v(t_i, x_i)$ and $v'(t'_j, t'_j)$ the utility function, denoted by $u(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{x}_{(n+1,N-n+1)}, \mathbf{x}'_{(m+1,N-m+1)}, \mathbf{D}_{(n,m)})$, for the entire trial 1 and trial 2 is analogous

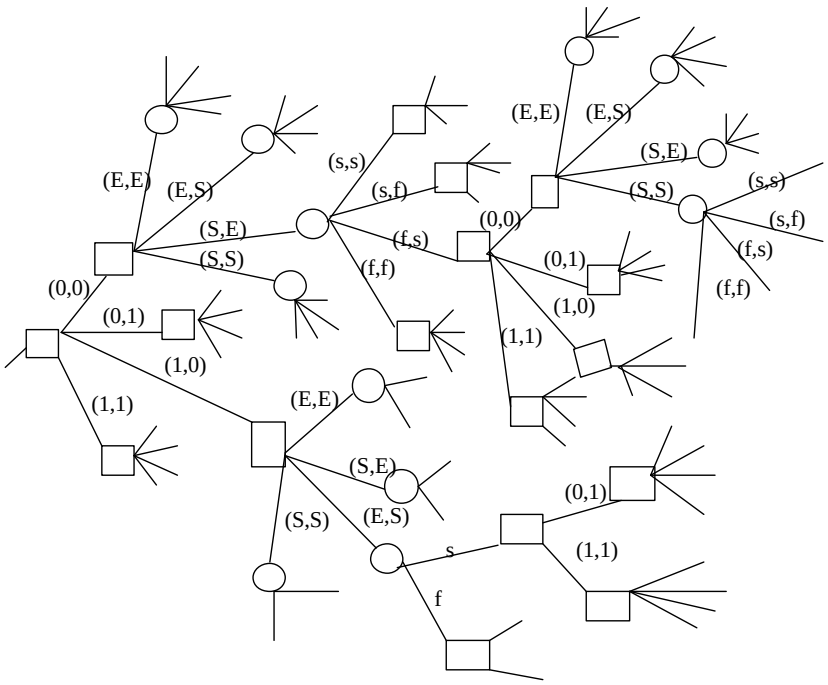

Figure 4.1: Decision tree at stages n and m. The square nodes stand for possible paired decisions and treatments and the circles stand for patient outcomes classified as success(s) and failure(f).

to the utility function in (2.2) and is given by

$$u(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{x}_{(n+1,N-n+1)}, \mathbf{x}'_{(m+1,M-m+1)}, \mathbf{D}_{n,m}) =$$
$$\alpha\{v(t_{n+1}, x_{N+1}) + v'(t'_{m+1}, x'_{M+1})\} + (1-\alpha)\frac{1}{N+M}\left\{\sum_{i=1}^{n} v(t_i, x_i) + \right.$$
$$\left.\sum_{j=1}^{m} v'(t'_j, x'_j) + \sum_{i=n+1}^{N} v(t_{n+1}, x_i) + \sum_{j=m+1}^{M} v'(t'_{m+1}, x'_j)\right\}, \tag{4.1}$$

where $\alpha = \frac{1}{N+M+2}$ and $\mathbf{x}_{(n+1,N-n+1)} = (x_{n+1}, x_{n+2}, \ldots, x_{N+1})$ and $\mathbf{x}'_{(m+1,M-m+1)} = (x'_{m+1}, x'_{m+2}, \ldots, x'_{M+1})$ are the future observations under treatments t_{n+1} and t'_{m+1} respectively. Note that we are dealing with one single trial to test the same new treatment, and therefore we should use one single utility function. The different characteristic here is that we have two cohorts (groups 1 and 2) dividing the study into two related trials. Given the current history $\mathbf{D}_{n,m}$, the optimal pair of treatments (t^*_{n+1}, t'^*_{m+1}) is obtained by maximizing the expected utility

$$U_{n,m}(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) =$$
$$E\{u(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{x}_{(n+1,N-n+1)}, \mathbf{x}'_{(m+1,M-m+1)}, \mathbf{D}_{n,m})\} \tag{4.2}$$

over $(t_{n+1}, t_{m+1}) \in \{(E,E), (E,S), (S,E), (S,S)\}$, where the expectation is with respect to future observations. The utility function in (4.1) is a weighted sum of the utilities $v(t_i, x_i)$ and $v'(t'_i, x'_i)$ of the patients in both trials and the utilities of future patients assigned to treatments t_{n+1} and t'_{m+1} respectively. We note that the future observations under a given treatment in a given trial are identically distributed and therefore, instead of using the respective joint posterior predictive distributions $P(\mathbf{x}_{(n+1,N-n+1)} \mid t_{n+1}, \mathbf{D}_{n,m})$ and $P'(\mathbf{x}'_{(m+1,M-m+1)} \mid t'_{m+1}, \mathbf{D}_{n,m})$ of sequences of future observations $\mathbf{x}_{(n+1,N-n+1)}$ and $\mathbf{x}'_{(m+1,M-m+1)}$ to evaluate the above expectation, we use their respective common predictive

distributions $P(x \mid t_{n+1}, \mathbf{D}_{n,m})$ and $P'(x' \mid t'_{m+1}, \mathbf{D}_{n,m})$. We obtain

$$\begin{aligned} U_{n,m}(d_n = 1, d'_m &= 1, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) = \frac{1-\alpha}{N+M}\left\{\sum_{i=1}^{n} v(t_i, x_i) + \sum_{j=1}^{m} v'(t'_j, x'_j)\right\} + \\ &\left\{\alpha + (1-\alpha)\frac{N-n}{N+M}\right\} \sum_{x=0,1} v(t_{n+1}, x) P(x \mid t_{n+1}, \mathbf{D}_{n,m}) + \\ &\left\{\alpha + (1-\alpha)\frac{M-m}{N+M}\right\} \sum_{x'=0,1} v'(t'_{m+1}, x') P'(x' \mid t_{m+1}, \mathbf{D}_{n,m}). \end{aligned} \tag{4.3}$$

If at least one trial continues or, equivalently, if $(d_n, d'_m) \in \{(0,0), (0,1), (1,0)\}$, then by Backward induction the expected utility $U_{n,m}(.)$ is defined as

$$\begin{aligned} U_{n,m}(d_n, \quad & d_m, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) = \\ & c_1 \sum_{x_{n+1}=0,1} \sum_{x'_{m+1}=0,1} U^*_{n+1,m+1}(\mathbf{D}_{n+1,m+1}) P(x_{n+1}, x'_{m+1} \mid t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) \\ + \quad & c_2 \sum_{x_{n+1}=0,1} U^*_{n+1,m}(\mathbf{D}_{n+1,m}) P(x_{n+1} \mid t_{n+1}, \mathbf{D}_{n,m}) \\ + \quad & c_3 \sum_{x'_{m+1}=0,1} U^*_{n,m+1}(\mathbf{D}_{n,m+1}) P'(x'_{m+1} \mid t'_{m+1}, \mathbf{D}_{n,m}), \end{aligned} \tag{4.4}$$

where $c_1 = (1-d_n)(1-d'_m), c_2 = d_n$,$c_3 = d'_m$, $\mathbf{D}_{n+1,m} = \mathbf{D}_{n,m} \bigcup \{x_{n+1}, t_{n+1}\}$ and $\mathbf{D}_{n,m+1} = \mathbf{D}_{n,m} \bigcup \{x_{m+1}, t_{m+1}\}$. We note that the quantities $U^*_{n+1,m}(.)$ and $U^*_{n,m+1}(.)$ in the above equation are random variables. Beginning at the last stage N where both trials must terminate and the optimal pair of treatments (t^*_{N+1}, t'^*_{M+1}) obtained by maximizing the expected utility of stopping both trials, the expected utilities defined in (4.4) can be evaluated working backward to the initial stages. Consequently, at any stages n and m of trials 1 and 2 respectively the maximum expected utility $U^*_{n,m}(\mathbf{D}_{n,m})$ is obtained as

$$U^*_{n,m}(\mathbf{D}_{n,m}) = \max_{d_n, d'_m} \{ \max_{t_{n+1}, t'_{m+1}} \{U_{n,m}(d_n, d'_m, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m})\}\}. \tag{4.5}$$

The optimal solution described above is infeasible to implement because of the exponential explosion of the decision tree, which looks complex even for $N = M = 2$. We therefore construct a stopping rule as in section 3.2 using the optimal stopping rule of a single-arm clinical trial.

4.3 Prior and posterior distributions

We now describe the general Bayesian framework relevant to our trial design outlined above. If we assume that the random vectors (θ, θ') and (π, π') are independent, then the joint prior density of θ, θ', π and π', denoted by $p(\theta, \theta', \pi, \pi')$, can be expressed as

$$p(\theta, \theta', \pi, \pi') = p(\theta, \theta')p(\pi, \pi') \tag{4.6}$$

where $p(\theta, \theta')$ and $p(\pi, \pi')$ are the marginal joint densities of the random vectors (θ, θ') and (π, π') respectively. If the two trials are at stages n and m and the observed data is denoted by $\mathbf{D}_{n,m}$ as in the previous section, the posterior distribution $p(\theta, \theta', \pi, \pi \mid \mathbf{D}_{n,m})$ is given by

$$p(\theta, \theta', \pi, \pi \mid \mathbf{D}_{n,m}) \propto f(\mathbf{D}_{n,m} \mid \theta, \theta', \pi, \pi)p(\theta, \theta')p(\pi, \pi) \tag{4.7}$$

where $f(\mathbf{D}_{n,m} \mid \theta, \theta', \pi, \pi')$ is the likelihood function. Suppose we let $\mathbf{D}^E_{n,m}$ and $\mathbf{D}^S_{n,m}$ be the data from both trials corresponding to treatments E and S respectively. We note that by assuming the result in (4.6), and assuming the conditional independence of data given the parameters, the posterior can be expressed as a product of the marginal likelihoods, denoted by $f(\mathbf{D}^E_{n,m} \mid \theta, \theta')$ and $f(\mathbf{D}^S_{n,m} \mid \pi, \pi')$, of (θ, θ') and (π, π') respectively. Hence the marginal posterior distributions can be computed as

$$p(\theta, \theta' \mid \mathbf{D}^E_{n,m}) \propto f(\mathbf{D}^E_{n,m} \mid \theta, \theta')p(\theta, \theta') \tag{4.8}$$

and

$$p(\pi, \pi' \mid \mathbf{D}^S_{n,m}) \propto f(\mathbf{D}^S_{n,m} \mid \pi, \pi)p(\pi, \pi). \tag{4.9}$$

Using these two marginal posterior distributions, the joint posterior predictive distribution, $P(x_{n+1}, x_{m+1} \mid t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m})$, of future observations x_{n+1} and x'_{m+1} under treatments t_{n+1} and t'_{m+1} respectively at stages $n+1$ and $m+1$ can be obtained as

$$P(x_{n+1}, x'_{m+1} \mid t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) = \int\int f(x_{n+1}, x'_{m+1} \mid \mu, \lambda)p(\mu, \lambda \mid \mathbf{D}_{n,m})d\mu d\lambda \tag{4.10}$$

where

$$\mu = \begin{cases} \theta & \text{if } t_{n+1} = E; \\ \pi & \text{if } t_{n+1} = S, \end{cases} \tag{4.11}$$

$$\lambda = \begin{cases} \theta' & \text{if } t'_{m+1} = E; \\ \pi' & \text{if } t'_{m+1} = S, \end{cases} \tag{4.12}$$

and $f(x_{n+1}, x'_{m+1} \mid \mu, \lambda)$ is the joint probability mass function of x_{n+1} and x_{m+1}. If we consider continuation of one the two trials, say trial 1, then under treatment t_{n+1} the posterior predictive distribution of x_{n+1} is given by

$$P(x_{n+1} \mid t_{n+1}, \mathbf{D}_{n,m}) = \int f(x_{n+1} \mid \mu) p(\mu \mid \mathbf{D}_{n,m}) d\mu. \tag{4.13}$$

The posterior predictive distribution of x'_{m+1} under treatment t'_{m+1} can be obtained in a similar manner.

The general form of our marginal joint prior distributions for (θ, θ') and (π, π') is described as follows. We suppose that θ' is known to have a $beta(\alpha, \beta)$ distribution and for a given value of θ', define

$$\theta \mid \theta' = \theta' + (1 - \theta')z \tag{4.14}$$

where z is a random variable with a $beta(\alpha', \beta')$ distribution. The conditional density $p(\theta \mid \theta')$ is obtained as

$$p(\theta \mid \theta') = \frac{1}{(1-\theta')^{\alpha'+\beta'-1} B(\alpha', \beta')} (\theta - \theta')^{\alpha'-1} (1-\theta)^{\beta'-1}, \; \theta' < \theta < 1. \tag{4.15}$$

Therefore the joint prior density $p(\theta, \theta')$ is given by

$$\begin{aligned} p(\theta, \theta') &= p(\theta \mid \theta') p(\theta') \\ &= \frac{1}{(1-\theta')^{\alpha'+\beta'-1} B(\alpha', \beta')} (\theta - \theta')^{\alpha'-1} (1-\theta)^{\beta'-1} \frac{1}{B(\alpha, \beta)} \theta'^{\alpha-1} (1-\theta')^{\beta-1}, \\ &\qquad\qquad 0 < \theta' < \theta < 1. \end{aligned} \tag{4.16}$$

Note that the above process of defining $p(\theta, \theta')$ can be reversed by supposing a known distribution for θ and that $\theta' \mid \theta = \theta + (1-\theta)z$. Similarly, we can derive the joint probability density in (4.16) for (π, π'). Using joint prior distributions of this general form for the treatment success probabilities, the dependence between the two trials is such that one trial can learn from the other and vice versa. Observe this dependence in the posterior predictive distributions for the future outcomes x_{n+1} and x'_{m+1} under the treatments t_{n+1} and t'_{m+1}, respectively, in section 4.5 where we propose a common prior distribution from this family of distributions for both (θ, θ') and (π, π'), in an example illustrating estimation of the optimal sequential decision at the initial stage.

4.4 Estimation of Expected Utility

We modify our expected utility estimating algorithm described in section 3.1 so that, given two stopping rules τ and τ', we can estimate the expected utility of continuing both trials at any stages n and m. The two stopping rules that we consider are such that, at any corresponding stages of the two trials, each is evaluated based on the observed data $\mathbf{D}_{n,m}$ from both trials. Suppose that based on $\mathbf{D}_{n,m}$, both τ and τ' indicate the decision to continue each trial at the respective stages n and m. The expected utility determined by both τ and τ', and denoted by $U_{n,m}(t_{n+1}, t'_{m+1}, \mathbf{D}_{(n,m)}, \tau, \tau')$, is estimated as follows. Simulate $(\theta^{(q)}, \theta'^{(q)})$ and $(\pi^{(q)}, \pi'^{(q)})$ from the probability densities $P(\theta, \theta')$ and $P(\pi, \pi')$ respectively, for $q = 1, 2, \ldots, G$. We note that any of the joint probability densities $P(\theta, \theta')$ and $P(\pi, \pi')$ could be a prior or a posterior probability density depending on whether data has been accumulated or not, with respect to the corresponding treatment. For each q we simulate x_i from $ber(\gamma_i^{(q)})$, $i = 1, 2, \ldots, n_q$ and x'_j from $ber(\gamma_j^{(q)})$, $j = 1, 2, \ldots, m_q$ using the stopping rules τ and τ' respectively. Suppose that, after we have simulated the sequences of observations $x_1, x_2, \ldots, x_{l-1}$ and $x'_1, x'_2, \ldots, x'_{r-1}$ the stopping rule, say τ, indicates continuation. We proceed to obtain the treatment combination $(t_l^*, t_r'^*)$ that maximizes the expected utility

$U_{l-1,r-1}(d_{l-1}=1, d'_{r-1}=1, t_l, t_r, \mathbf{x}_{(1,l-1)}, \mathbf{t}_{(1,l-1)}, \mathbf{x}'_{(1,r-1)}, \mathbf{t}'_{(1,r-1)})$ of stopping, simulating both trials and simulating x_l from $ber(\gamma_l^{(q)})$ where

$$\gamma^{(q)} = \begin{cases} \theta^{(q)} & \text{if } t_l^* = E; \\ \pi^{(q)} & \text{if } t_l^* = S. \end{cases} \tag{4.17}$$

Similarly, x_r can be simulated if τ' indicates continuation based on the simulated observations $\mathbf{x}_{(1,l-1)}, \mathbf{t}_{(1,l-1)}, \mathbf{x}'_{(1,r-1)}, \mathbf{t}'_{(1,r-1)}$. The above simulations stop at stages n_q and m_q when stopping is prescribed by τ^p and τ'^p, respectively. The expected utility of stopping $U_{n_q,m_q}(d_{n_q}=1, d'_{m_q}=1, t_{n_q+1}, t'_{m_q+1}, \mathbf{x}_{(1,n_q)}, \mathbf{t}_{(1,n_q)}, \mathbf{x}'_{(1,m_q)}, \mathbf{t}'_{(1,m_q)})$ is then computed and denoted by u_q. The point estimate of the expected utility are therefore obtained as

$$\widehat{U}_{n,m}(\mathbf{x}_{(1,n)}, \mathbf{t}_{(1,n)}, \mathbf{x}'_{(1,m)}, \mathbf{t}'_{(1,m)}, t_{n+1}, t'_{m+1}, \tau, \tau') = \frac{1}{G}\sum_{t=1}^{G} u_q. \tag{4.18}$$

We note that, if at the corresponding stages n and m of the two trials one trial stops and the other continues, then for a given stopping rule the expected utility of continuing is computed by implementing our expected utility estimating algorithm as in section 3.1.

4.5 The Proposed Stopping Rules

We now describe the proposed stopping rules, denoted by τ^p and τ'^p respectively, for the two related clinical trials, each with two treatment arms, as an approximation for the optimal stopping rules obtained by backward induction. The construction of the two proposed stopping rules is based on the alternating conversion of the two trials into single-arm clinical trials whose optimal stopping rules can be evaluated by using the backward induction algorithm. Suppose trials 1 and 2 are at stages n and m respectively. We wish to establish the sequential decisions according to the proposed stopping rules based on the observed data ($\mathrm{D}_{n,m}$ from both trials). Let $\mathrm{D}_{n,m}$ be summarized by the sufficient statistics (n_E, s_E, n_S, s_S) and (m_E, s'_E, m_S, s'_S) where, for instance, the pairs (n_E, s_E) and (m_E, s'_E) denote the number

of observations and the number of observed successes due to treatment E in trials 1 and 2 respectively. The other symbols corresponding to treatment S are defined similarly. We proceed as follows. We assume that one trial stops and the other continues with only one of the two treatments assigned to patients until stopping. This implies that for the remaining stages, beginning from the present stage, the continuing trial is equivalent to a single-arm clinical trial comparing the two treatments. Note that the known probability of success of any treatment not assigned to patients in the resulting single-arm trial is assumed equal to the mean of posterior predictive probability distribution of a future observation under this treatment conditioned on the available observed data from both trials.

For example, if we consider that trial 1 continues and trial 2 stops at stages n and m, and that treatment E is allocated to patients in trial 1 until stopping. The resulting single-arm trial compares treatment E with S where the probability of success of treatment S, given by $E(x_{n+1} \mid \mathbf{D}_{n,m}, t_{n+1} = S)$, is assumed fixed and based on the utility function defined in (4.1) the expected utility of stopping at j^{th} stage of the single-arm trial is obtained by maximizing the expected utility of stopping both trials $U_{n+j,m}(d_{n+j} = 1, d'_m = 1, t_{n+j+1}, t'_{m+1}, \mathbf{D}_{n+j,m})$ at stages $n+j$ and m as in (4.3). We then calculate, using backward induction, a value denoted by u_E corresponding to the expected utility of continuing the resulting single-arm trial at the present stage. Similarly, assuming that the trial with treatment S until stopping in trial 1 we compute the corresponding expected utility u_S. Alternating to trial 2 we use a similar process to calculate the values u'_E and u'_S.

Suppose we define $u^* = \max\{u_E, u_S\}$ and $u'^* = \max\{u'_E, u'_S\}$, then u^* and u'^* are our estimates for the expected utilities corresponding to the paired decisions $(0, 1)$ and $(1, 0)$ respectively. Our estimate for the sequential optimal decision at stages n and m is stated as follows.

•If $\min(u^*, u'^*) > U_{n,m}(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m})$, continue both trials.
•Else if $u'^* \leq U_{n,m}(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) < u^*$, continue trial 1 and stop trial 2.

•Else if $u^* \leq U_{n,m}(d_n = 1, d'_m = 1, t_{n+1}, t'_{m+1}, \mathbf{D}_{n,m}) < u'^*$, continue trial 2 and stop trial 1.
•Else stop both trials.

We note that when the two trials are alternately converted to a single-arm trial, beginning from stages n and m, their assumed progress is restricted to the branches of the complete decision tree indicating the paired decisions $(0, 1)$ and $(1, 0)$ (See fig. 4.1). This leads to excluding completely the branches indicating $(0, 0)$. Thus, our estimates u^* and u'^*, obtained by implementing backward induction using the resulting incomplete decision tree, are respectively less than the optimal expected utilities of continuation under the corresponding paired decisions $(0, 1)$ and $(1, 0)$. This again implies that if, for instance, u^* is greater than the expected utility of stopping at some stages n and m, then the expected utility of continuing under the decision $(0, 1)$ obtained by backward induction is also greater the expected utility of stopping. Similar implication can be made about the value u'^*. The formulation of the above decision rule is based on this result.

4.6 Example

We consider an example of the sequential trial design for two related trials and use the proposed stopping rules τ^p and τ'^p to estimate the optimal sequential decisions at stages n and m of the trials 1 and 2 respectively, based on the observed data (n_E, n_S, s_E, s_S) and (m_E, m_S, s'_E, s'_S). For the purpose of illustration we set the maximum number, N, of enrolled patients in the trial equal to 12 and specify the respective joint prior distributions of (θ, θ') and (π, π'). Also, we let $v(0, t) = v'(0, t) = 0$ and $v(1, t) = v'(1, t) = 1$ for any treatment $t = E, S$. Suppose we assume that θ' and π' have the common $beta(1, 2)$ distribution and that z in (4.14) has the $U(0, 1)$ distribution. Therefore each of the joint prior probability densities $p(\theta, \theta')$ and $p(\pi, \pi')$, whose general form is given in (4.16), simplifies to a probability

density of the form

$$f(w,y) = \begin{cases} 2 & \text{if } 0 < y < 1, y < w < 1; \\ 0 & \text{otherwise.} \end{cases} \tag{4.19}$$

Based on this common prior probability density and conditioning on the data, the joint posterior probability densities $p(\theta, \theta' \mid n_E, m_E, s_E, s'_E)$ and $p(\pi, \pi' \mid n_S, m_S, s_S, s'_S)$ of (θ, θ') and (π, π') are obtained as

$$p(\theta, \theta' \mid n_E, m_E, s_E, s_E) \propto \theta^{s_E}(1-\theta)^{n_E - s_E}\theta'^{s'_E}(1-\theta')^{m_E - s'_E},\ 0 < \theta' < \theta < 1, \tag{4.20}$$

and

$$p(\pi, \pi' \mid n_S, m_S, s_S, s'_S) \propto \pi^{s_S}(1-\pi)^{n_S - s_S}\pi'^{s'_S}(1-\pi')^{m_S - s'_S},\ 0 < \pi' < \pi < 1. \tag{4.21}$$

Suppose we let $\theta(n_E, m_E, s_E, s'_E)$, $\theta'(n_E, m_E, s_E, s'_E)$, $\pi(n_S, m_S, s_S, s'_S)$ and $\pi'(n_S, m_S, s_S, s'_S)$ denote the posterior expected values of $\theta_1, \theta_2, \pi_1$ and π_2 respectively. Writing $\theta^{s_E}(1-\theta)^{n_E - s_E} = \sum_{i=0}^{n_E - s_E}(-1)^i\binom{n_E - s_E}{i}\theta^{i+s_E}$ and $\pi^{s_S}(1-\pi)^{n_S - s_S} = \sum_{j=0}^{n_S - s_S}(-1)^j\binom{n_S - s_S}{j}\pi^{j+s_S}$ and integrating with respect to θ and π respectively, we obtain the marginal distributions $p(\theta' \mid n_E, m_E, s_E, s'_E)$ and $p(\pi' \mid n_S, m_S, s_S, s'_S)$ as

$$p(\theta' \mid n_E, s_E, m_E, s'_E) \propto \theta'^{s'_E}(1-\theta')^{m_E - s'_E}\sum_{i=0}^{n_E - s_E}(-1)^i\binom{n_E - s_E}{i}\left\{\frac{1}{s_E + i + 1} - \frac{\theta'^{s_E + i + 1}}{s_E + i + 1}\right\} \tag{4.22}$$

and

$$p(\pi' \mid n_S, s_S, m_S, s'_S) \propto \pi'^{s'_S}(1-\pi')^{m_S - s'_S}\sum_{j=0}^{n_S - s_S}(-1)^j\binom{n_S - s_S}{j}\left\{\frac{1}{s_S + j + 1} - \frac{\pi'^{s_S + j + 1}}{s_S + j + 1}\right\}. \tag{4.23}$$

Consequently, we get

$$\begin{aligned} \theta'(n_E, m_E, s_E, s'_E) &= \int_0^1 \theta' p(\theta' \mid n_E, m_E, s_E, s'_E)\, d\theta' \\ &= \frac{1}{C_\theta}\sum_{i=0}^{n_E - s_E}(-1)^i\binom{n_E - s_E}{i}\frac{1}{s_E + i + 1}\left\{\frac{\Gamma(s'_E + 2)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + 3)}\right. \\ &\quad \left. -\frac{\Gamma(s_E + s'_E + i + 3)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + s_E + i + 4))}\right\} \end{aligned} \tag{4.24}$$

and

$$\begin{aligned}\pi'(n_S, m_S, s_S, s'_S) &= \int_0^1 \pi' p(\pi' \mid n_S, m_S, s_S, s'_S)\, d\pi' \\ &= \frac{1}{C_\pi} \sum_{j=0}^{n_S - s_S} (-1)^j \binom{n_S - s_S}{j} \frac{1}{s_S + j + 1} \left\{ \frac{\Gamma(s'_S + 2)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + 3)} \right. \\ &\quad \left. - \frac{\Gamma(s_S + s'_S + j + 3)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + s_S + j + 4))} \right\}, \end{aligned} \tag{4.25}$$

where the normalizing constants C_θ and C_π are obtained as

$$\begin{aligned} C_\theta &= \sum_{i=0}^{n_E - s_E} (-1)^i \binom{n_E - s_E}{i} \frac{1}{s_E + i + 1} \left\{ \frac{\Gamma(s'_E + 1)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + 2)} \right. \\ &\quad \left. - \frac{\Gamma(s_E + s'_E + i + 2)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + s_E + i + 3))} \right\} \end{aligned} \tag{4.26}$$

and

$$\begin{aligned} C_\pi &= \sum_{i=0}^{n_S - s_S} (-1)^i \binom{n_S - s_S}{i} \frac{1}{s_S + i + 1} \left\{ \frac{\Gamma(s'_S + 1)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + 2)} \right. \\ &\quad \left. - \frac{\Gamma(s_S + s'_S + i + 2)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + s_S + i + 3))} \right\}. \end{aligned} \tag{4.27}$$

Next, by taking expectation of the conditional expectations $E(\theta \mid \theta', n_E, s_E, m_E, s'_E)$ and $E(\pi \mid \pi', n_S, s_S, m_S, s'_S)$ with respect to θ' and π' respectively, we obtain

$$\begin{aligned} \theta(n_E, m_E, s_E, s'_E) &= \frac{1}{C_\theta} \sum_{i=0}^{n_E - s_E} (-1)^i \binom{n_E - s_E}{i} \frac{1}{s_E + i + 2} \left\{ \frac{\Gamma(s'_E + 1)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + 2)} \right. \\ &\quad \left. - \frac{\Gamma(s_E + s'_E + i + 3)\Gamma(m_E - s'_E + 1)}{\Gamma(m_E + s_E + i + 4))} \right\} \end{aligned} \tag{4.28}$$

and

$$\begin{aligned} \pi(n_S, m_S, s_S, s'_S) &= \frac{1}{C_\pi} \sum_{i=0}^{n_S - s_S} (-1)^i \binom{n_S - s_S}{i} \frac{1}{s_S + i + 1} \left\{ \frac{\Gamma(s'_S + 1)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + 2)} \right. \\ &\quad \left. - \frac{\Gamma(s_S + s'_S + i + 3)\Gamma(m_S - s'_S + 1)}{\Gamma(m_S + s_S + i + 4))} \right\}. \end{aligned} \tag{4.29}$$

We then evaluate the double integral in (4.10) to obtain the marginal posterior predictive distributions of the future observations x_{n+1} and x'_{m+1} under the treatments t_{n+1} and t'_{m+1}

respectively, as Bernoulli distributions, respectively, where

$$\begin{aligned} & Pr(x_{n+1} = 1 \mid t_{n+1}, n_E, s_E, n_S, s_S, m_E, s_E, m_S, s'_S) \\ = \;& 1 - Pr(x_{n+1} = 0 \mid t_{n+1}, n_E, s_E, n_S, s_S, m_E, s_E, m_S, s'_S) \\ = \;& \begin{cases} \theta(n_E, m_E, s_E, s'_E) & \text{if } t_{n+1} = E; \\ \pi(n_S, m_S, s_S, s'_S) & \text{if } t_{n+1} = S \end{cases} \end{aligned}$$

and

$$\begin{aligned} & Pr(x'_{m+1} = 1 \mid t'_{m+1}, n_E, s_E, n_S, s_S, m_E, s_E, m_S, s'_S) \\ = \;& 1 - Pr(x'_{m+1} = 0 \mid t'_{m+1}, n_E, s_E, n_S, s_S, m_E, s_E, m_S, s'_S) \\ = \;& \begin{cases} \theta'(n_E, m_E, s_E, s'_E) & \text{if } t'_{m+1} = E; \\ \pi'(n_S, m_S, s_S, s'_S) & \text{if } t'_{m+1} = S. \end{cases} \end{aligned} \tag{4.30}$$

These marginal predictive distributions are required for computing the expected utilities of stopping at any stages n and m. Suppose $n = m = 0$. Constructing the proposed stopping rules τ^p and τ'^p as outlined in the previous section, based on the assumed prior in (4.19) and the assumed utility function, we obtained the values $u^* = 0.5757$ and $u'^* = 0.5728$ as the respective estimates of the expected utilities corresponding to the pairs of decisions $(0, 1)$ and $(1, 0)$. At the initial stage of the related trial design with no accumulated data the predictive probability of success is equal to $\frac{2}{3}$ for any given treatment in trial 1 and $\frac{1}{3}$ for any given treatment in trial 2. Therefore, by equation (4.3), the maximum expected utility of stopping both trials is obtained as 0.5200. Given that $\min(u^*, u'^*) > 0.5200$, both τ^p and τ'^p indicate the decision to continue. In other words, the pair of decisions $(0, 0)$ is the estimate of the optimal sequential decisions. We note that there are no exact values or past estimation of the expected utilities with which we can compare our estimates. We justify our estimation by evaluating our trial design's operating characteristics in the next section.

4.7 Operating Characteristics

We evaluate our proposed sequential trial design for related trials by investigating its operating characteristics under fixed values of success probabilities θ, π and θ', π' of the experimental and standard treatments in trials 1 and 2 respectively. We study our trial design under six different scenarios and summarize the results obtained in each scenario by reporting, for each trial, the fixed success probabilities of treatments and the average number of patients assigned to each treatment in a table. Under each scenario we draw 1000 simulations of the possible outcomes of the entire trial and compute the average number of patients assigned to each treatment in both groups. In all scenarios, we let the utilities v and v' be as defined in the previous example, the joint priors $p(\theta, \theta')$ and $p(\pi, \pi')$ be as given 4.19 and set the maximum number of patients per trial at $N = 30$.

Scenario 1 represents a situation where both treatments are equally efficacious i.e. the two treatments do not differ. The trial stopped with averages of 25.49 and 21.31 patients in trials 1 and 2 respectively and according to column P_t the two treatments were approximately equally preferred. In the second scenario the standard treatment is superior to the experimental treatment in both trials. On average 22.13 and 18.18 patients were observed in trials 1 and 2 respectively and the standard treatment in both trials had a high probability of preference (0.87 and 0.84 respectively). Scenario 3 represents a situation where the experimental treatment is more effective than the standard in both trials. We observed 23.74 and 14.19 patients in trials 1 and 2 with more preference given to the experimental treatment. Scenarios 4 and 5 correspond to cases where the experimental treatment is superior to the standard treatment in trial 1 and vice versa in trial 2. In scenario 4 averages of 25.40 and 21.20 patients were observed while in scenario 5 we have 24.08 and 21.53 patients observed in trials 1 and 2 respectively with the higher average of patients assigned to the superior treatment. Scenario 6 represents a situation where the two treatments are equally efficacious in trial 1 while in trial 2 the experimental treatment is superior. The two trials stopped with

averages of 23.07 and 20.22 patients respectively.

Scenario 1

	Trial 1					**Trial 2**			
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	13.24	10.37	0.6	0.53	E	10.97	9.19	0.3	0.50
S	12.24	9.91	0.6	0.47	S	10.34	8.55	0.3	0.50
$\tilde{n}$	25.49	4.27				21.31	5.77		

Scenario 2

	Trial 1					**Trial 2**			
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	6.64	8.21	0.5	0.13	E	6.02	7.69	0.2	0.16
S	15.78	7.33	0.7	0.87	S	12.05	5.97	0.4	0.84
$\tilde{n}$	22.13	5.89				18.18	5.92		

Table 4.1: Operating characteristics for related trials. In all tables under each trial, E and S denote treatments, the second column $\tilde{n}_t$ indicates the average number of patients assigned to each treatment, the third column St. dev. represents the corresponding standard deviation, θ_t is the assumed true success probability and P_t is the probability of preferring a given treatment to the other on stopping. See Appendix B for a relevant detailed interpretation of operating characteristics.

Scenario 3

Trial 1					Trial 2				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	21.96	5.99	0.7	0.98	E	11.69	4.92	0.6	0.95
S	1.77	3.22	0.3	0.02	S	2.50	5.65	0.2	0.05
$\tilde{n}$	23.74	5.22				14.19	5.22		

Scenario 4

Trial 1					Trial 2				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	15.79	9.29	0.7	0.69	E	7.58	8.46	0.2	0.22
S	9.61	10.27	0.5	0.31	S	13.62	7.78	0.4	0.78
$\tilde{n}$	25.40	4.16				21.20	5.79		

Scenario 5

Trial 1					Trial 2				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	10.29	10.28	0.5	0.28	E	14.15	8.13	0.4	0.78
S	13.80	8.33	0.7	0.72	S	7.35	8.27	0.2	0.22
$\tilde{n}$	24.08	5.18				21.53	5.48		

Table 4.2: Continuation of table 4.1, scenarios 2, 3, 4 and 5. See Appendix B.

Scenario 6

Trial 1					Trial 2				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	14.19	9.83	0.7	0.60	E	13.71	7.83	0.4	0.81
S	8.88	9.37	0.7	0.40	S	6.52	8.04	0.2	0.19
$\tilde{n}$	23.07	5.07				20.22	5.87		

Table 4.3: Continuation of table 4.1, scenario 6. See Appendix B.

We next evaluate the operating characteristics of our related trial design to investigate if at some given stages continuing one of the two trials contributes to learning about the other. We let 0.7, 0.4 and 0.2 be the assumed true values of the success probabilities θ, θ' and π' of the experimental treatment in both trials and that of the standard treatment in trial 2 respectively. Then, for each of the values $0.3, 0.4, 0.5, 0.7$ and 0.8, taken as the assumed true success probability of the standard treatment in trial 1, we simulate $1,000$ possible observations of the entire trial and give a similar tabular report. This scenario corresponds to the case where no more patients are admitted to trial 2 while trial 1 continues with patients allocated to the standard treatment S. We observe that for the first three values taken in ascending order and close to $\pi' = 0.2$, the assumed true success probability of the standard treatment in trial 2, the average number of patients assigned to the standard treatment increases in both trials and is higher in trial 2. This trend does not continue for the remaining two values which are distant from $\pi' = 0.2$. This is because, when the assumed values of success probabilities of the standard treatment in both trials are taken close to each other, the resulting marginal posterior densities are close and overlap with each other at the tails and hence depend on each other. However, when the two values of the assumed true success probabilities of the standard treatment in both trials differ by a big margin the resulting marginal posterior densities will be too apart to overlap at the tails and hence, apart from satisfying the initial condition that $\pi' < \pi$, do not depend on each other. In this case changing the assumed true value of success probability in one trial does not influence

the other and the two trials would continue as if they were two separate trials. Consequently, if the joint priors for the success probabilities $p(\theta, \theta)$ and $p(\pi, \pi')$ are selected in such away that at any stages of the trials the respective posterior densities are close enough that they overlap at the tails, then by continuing one of the two trials the proposed trial design permits us to learn about the other.

$\pi = 0.3$

	Trial 1					**Trial 2**			
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	20.51	6.06	0.7	0.98	E	13.46	6.72	0.4	0.86
S	2.30	4.24	0.3	0.02	S	5.15	7.46	0.2	.14
$\tilde{n}$	22.81	5.03				18.61	5.92		

$\pi = 0.4$

	Trial 1					**Trial 2**			
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	19.77	6.56	0.7	0.96	E	13.13	6.86	0.4	0.85
S	3.54	5.71	0.4	0.04	S	5.81	8.03	0.2	0.15
$\tilde{n}$	23.31	4.94				18.94	6.03		

Table 4.4: Operating characteristics. In all tables the assumed true values of θ, θ' and π' are fixed to $0.7, 0.4$ and 0.2 respectively. The assumed true value of π is varied as shown in the five tables. Note how the number of patients assigned treatment S in trial 2 increases as π departs from $\pi' = 0.2$.

$\pi = 0.5$

Trial 1					**Trial 2**				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	18.52	7.67	0.7	0.89	E	13.65	7.13	0.4	0.83
S	5.00	7.66	0.5	0.12	S	5.60	7.50	0.2	0.17
$\tilde{n}$	23.52	4.96				19.24	5.84		

$\pi = 0.7$

Trial 1					**Trial 2**				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	14.19	9.83	0.7	0.60	E	13.71	7.83	0.4	0.81
S	8.88	9.37	0.7	0.40	S	6.46	7.79	0.2	0.19
$\tilde{n}$	23.07	5.07				20.22	5.87		

$\pi = 0.8$

Trial 1					**Trial 2**				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	12.29	10.35	0.7	0.45	E	14.10	8.27	0.4	0.8
S	9.58	8.55	0.8	0.55	S	7.09	8.47	0.2	0.20
$\tilde{n}$	21.87	5.30				21.17	5.96		

Table 4.5: Continuation of table 4.4 for $\pi = 0.5$, 0.7 and 0.8

We now investigate more about our two related trials by evaluating the operating characteristics of two separate trial designs. Suppose that the success probabilities of any treatment in the two trials are independent apriori, and have a common uniform distribution over over the interval $[0, 1]$. We fixed the assumed true success probabilities in each trial as in scenario 3 in table 4.2 and obtained the results in the following table.

Trial 1					**Trial 2**				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	13.86	5.90	0.7	0.94	E	14.86	5.97	0.6	0.96
S	3.41	6.16	0.3	0.06	S	3.10	4.93	0.2	0.04
$\tilde{n}$	17.27	5.35				17.97	5.34		

Table 4.6: Operating characteristics: Two separate trials

Compared to the results in scenario 3 in table 4.2, we observe that in the related trials design fewer patients are assigned to the treatment with lower efficacy rate with a high precision. In scenario 3, an average of 1.85 patients were assigned to treatment S with a standard deviation of 3.35 while when the trials are separate 3.41 are assigned to S with a standard deviation of 6.16. Also we observe that when the two trials are considered separate almost equal number of patients from the two groups are used while in the related trials design more patients from the group with a milder condition of the disease are used.

4.8 Discussion

In this chapter, we have considered approximating the optimal stopping rules of two related trials by proposing two novel stopping rules, τ^p and τ'^p, whose construction is based on the optimal stopping rule of a single-arm trial obtained by Backward Induction. Just as in the previous chapter, the single-arm conversion provides a means of selecting some of the branches of the tree diagram in figure 4.1, along which the approximate expected utility

is calculated beginning at the end of the trial and working backwards. The study of the operating characteristics in section 4.4 shows that our proposed trial design allows early stopping, conforms with the assumed order of superiority of the treatments in both trials and allows learning about one trial when patient entry is confined to the other trial. Suppose at some paired stages both trials continue according τ^p and τ'^p, and that at any future stages there will be patients available for admission. We expect that data that we observe from this trial that continues will make the other trial stop without any admission of more patients and in the event that only one of the two trials stops, then the other will be conducted with an optimal stopping rule.

Our proposed trial design is not restricted to the utility function that we have used here and may be implemented with any other utility function that is deemed suitable. We made a change in the utility function in (4.1) such that $v(t,0) = v'(t,0) = 0.3$ and $v(t,1) = v'(t,1) = 1$ and a repeat of scenario 3 in table 4.2 gave the results as in table 4.7.

Trial 1					**Trial 2**				
Trt	$\tilde{n_t}$	St.dev	θ_t	P_t	Trt	$\tilde{n_t}$	St.dev	θ_t	P_t
E	21.75	6.22	0.7	0.98	E	11.76	4.85	0.6	0.96
S	1.82	3.51	0.3	0.02	S	2.33	5.22	0.2	0.04
$\tilde{n}$	23.56	5.45				14.09	5.04		

Table 4.7: Scenario 3: Sensitivity analysis

We also repeated scenarios 1 and 5 each with there different changes of the utility function and our results changed minimally, indicating that our proposed design is not sensitive to changes in the utility function. We, therefore, conclude that our proposed stopping rules have attractive properties as approximation of the optimal stopping rules for the related trials.

We remark that, we have used the family of joint prior densities in (4.16) purposely because

it meets our objective of the study; consequently, the ordering restriction may be removed and any other joint densities may be assigned to the unknown success probabilities in some way.

In each scenario or in the case of sensitivity analysis, the table that we report is a summary of the $1,000$ simulations of the entire trial and a more detailed analysis may be done to understand them more. Simulations for scenarios 1, 3 and 5 and a software written in R language used to generate them are available in $http://www.cimat.mx/\sim jac/material/thesisOrawo.zip$.

Chapter 5

One-arm Clinical Trial with Continuous Response

5.1 Introduction

Clinical trials where patient response is continuous are common in practice. In such trials a continuous response would represent, for instance, time to recovery of a patient, time to occurrence of some adverse event, tumour size, blood pressure or concentration of some chemical in the blood or urine, after being treated by any two competing treatments. If the objective of a trial is to compare the mean responses of two treatments, then the normal model is the most popular. Whitehead and Jones (1979), Freedman, Lowe, and Macaskill (1984), Carlin *et al.* (1998) and Brockwell *et al.* (2003) are a few examples of single-arm or double-arm clinical trials where patient response have been considered normal.

In this chapter we use a numerical method known as 'gridding' to implement a Bayesian sequential design for a one-arm clinical trial comparing an experimental treatment E with

a standard treatment S, where patient response is normal. We note that the gridding approximation that we do here can applied to a double-arm clinical trial except for the complex case when the variances of the responses of the two treatments are unknown and unequal. Here we use the exact predictive probabilities to obtain the gridding approximation unlike in Brockwell *et al.* (2003) where a simulation method is used.

In section 2, we describe the trial design. Section 3 presents our gridding approximation. Illustrative examples are considered in section 4 and we end the chapter by a discussion in section 5.

5.2 The trial design

We describe a Bayesian sequential design for a one-arm clinical trial comparing an experimental treatment E with a standard treatment S. We assume that under treatments E and S the patient responses are $x|t = E \sim N(\mu_E, \lambda_E)$ and $x|t = S \sim N(\mu_S, \lambda_S)$ respectively, where μ_S and λ_S are known. Suppose a and b are known constants, we define the utility, denoted by $v(t, x)$, of making an observation x under treatment t as

$$v(x, t) = a + bx. \tag{5.1}$$

We assume that a large patient response under any treatment in the course of the trial suggests superiority of the treatment. Therefore we assume that $b > 0$. As in the previous chapters, we let N denote the maximum number of patients enrolled in the trial and allow continuous perusal of the accumulating data at N interim stages. At stage $n < N$ of the trial we have to decide whether to continue to next stage $n+1$ or not. Suppose from the first n patients we have the sequence of observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$. If the trial stops at stage n, we propose the utility function $u(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_{N+1}, v)$ for the entire trial given

in (2.2) with $t_i = E$ for all $i = 1, 2, \ldots, n$. i.e

$$u(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_{N+1}, v) = \alpha v(t_{n+1}, x_{N+1}) + \frac{(1-\alpha)}{N}\left\{\sum_{i=1}^{n} v(t_i, x_i) + \sum_{i=n+1}^{N} v(t_{n+1}, x_i)\right\}.$$

The objective is to obtain the optimal sequential decision at any stage of the trial by maximizing the expected utility. The expected utility, denoted by $U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n)$, of stopping the trial with treatment t_{n+1} is found as

$$U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n) = E\{u(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_{N+1}, v)\}, \tag{5.2}$$

where the above expectation is with respect to future observations $x_{n+1}, x_{n+2}, \ldots, x_{N+1}$. Hence the more effective treatment, denoted by t^*_{n+1}, is obtained by maximizing the above expected utility over $\{t_{n+1} = S, E\}$. If we assume that the future observations are independent and identically distributed random variables, then analogous to the expected utility in (2.4) the above expected utility is evaluated as

$$\begin{aligned} U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n) &= (1-\alpha)\sum_{i=1}^{n} v(t_i = E, x_i) + \\ &\quad \times \int (\alpha + (1-\alpha)\frac{N-n}{N}) v(t_{n+1}, x) p(x \mid, t_{n+1}, \mathbf{x}_n) dx. \end{aligned} \tag{5.3}$$

If, as before, we let $U_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n)$ denote the expected utility of continuing, then the maximum expected utility is given by

$$U^*_n(t_{n+1}, \mathbf{t}_n, \mathbf{x}_n) = \max_{d_n=0,1}\{\max_{t_{n+1}=S,E}\{U_n(d_n, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n)\}\}. \tag{5.4}$$

According to the backward induction method the expected utility of continuing the trial is

$$U_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \mathbf{x}_n) = \int U^*_{n+1}(\mathbf{t}_{n+1}, \mathbf{x}_{n+1}) p(x_{n+1} \mid \mathbf{t}_{n+1}, \mathbf{x}_n) dx \tag{5.5}$$

where $p(x_{n+1} \mid \mathbf{t}_{n+1}, \mathbf{x}_n)$ is the predictive density of x_{n+1} given the observed data $\mathbf{x}_n$. We note that the implementation of backward induction as in the previous chapters is not

possible since the number of possible future sequences of the continuous observations will be infinite. We therefore develop, in the following section, a numerical approximation to backward induction based on a discretization of the sufficient statistics for the unknown parameters.

5.3 Numerical approximation of the optimal decision

We present a numerical procedure known as "gridding"(see Brockwell and Kadane, 2003 and Berry and Chi-Hsiang, 1988) for approximating the optimal sequential decision at any interim stage n of the clinical trial described above. We consider two cases; first where the patient response under the experimental treatment E is normal with precision λ_E assumed known and second when λ_E is unknown.

5.3.1 The case of known precision

Suppose we have the observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$ from the first n patients, with $x_i \mid t = E \sim N(\mu_E, \lambda_E)$. We assume that λ_E is known and that the prior is $\mu_E \sim N(\mu_0, \lambda_0)$. The posterior distribution of μ_E conditioned on the observed data $\mathbf{x}_n$ is normal, with mean $\mu_n = \frac{\mu_0\lambda_0 + n\lambda_E\bar{x}_n}{\lambda_0 + n\lambda_E}$ and precision $\lambda_n = \lambda_0 + n\lambda_E$. We note that the sample mean $\bar{x}_n$ is the sufficient statistic for the posterior distribution of μ_E. By evaluating the integral $\int f(x_{n+1} \mid \mu_E)p(\mu_E \mid \mathbf{t}_n, \mathbf{x}_n)d\mu_E$ we obtain the posterior predictive distribution of a future observation x_{n+1} as normal, with mean μ_n and precision $\frac{\lambda_E\lambda_n}{\lambda_n + \lambda_E}$. To approximate the optimal sequential decision at any stage n we proceed as follows.

We construct a grid in the real line over which the summary statistic takes values. We choose two values as the lower and the upper grid bounds q_L and q_U respectively. We

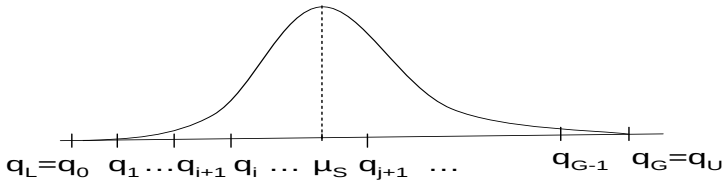

Figure 5.1: The grid points on the real line.

then make G subdivisions $(q_i, q_{i+1}), i = 0, 1, \ldots, G$ of the interval (q_L, q_U) using quantiles of the normal distribution $N(\mu_S, \lambda_S)$, including the mean μ_S with $q_0 = q_L$ and $q_G = q_U$ as indicated in figure 5.1. Suppose that μ_S corresponds to the quantile marked q_j. We use Q to denote the set $Q = \{q_0, q_1, \ldots, q_G\}$ of grid points. Since our interest is to investigate the effectiveness of treatment E relative to treatment S by determining if $\mu_E > \mu_S$, we approximate the distribution of the continuous sufficient statistic at any future stage by a discrete distribution defined on a grid of points as follows. Suppose the trial is at stage n. We note that for any given value $\bar{x}_n \in Q$, the value of the sample mean $\bar{x}_{n+1}$ at the next stage may fall in any of the $G+2$ subdivisions of the real line. Suppose that $\bar{x}_n = q_k$ and that $y_{(n+1,k)}$ denote the discrete random variable whose distribution approximates the continuous distribution of $\bar{x}_{n+1}$. Then $y_{(n+1,k)}$ assumes values over the grid as follows.

$$\begin{aligned}
\text{If} \quad & \bar{x}_{n+1} \leq q_L, && y_{(n+1,k)} = q_L; \\
\text{else if} \quad & \bar{x}_{n+1} \in [q_i, q_{i+1}] \text{ and } i \leq j-1, && y_{(n+1,k)} = q_{i+1}; \\
\text{else if} \quad & \bar{x}_{n+1} \in [q_{j-1}, q_{j+1}], && y_{(n+1,k)} = \mu_S; \\
\text{else if} \quad & \bar{x}_{n+1} \in [q_i, q_{i+1}] \text{ and } i \geq j+1, && y_{(n+1,k)} = q_i; \\
& \text{else} && y_{(n+1,k)} = q_U,
\end{aligned} \tag{5.6}$$

where $i = 0, 1, \ldots, j-2, j+1, \ldots, G$. Suppose that, for each $z = 0, 1, \ldots, G$, the probability that $y_{(n+1,k)} = q_z$ denoted by $p_{n+1}(k, z)$ is computed using the normal predictive probability of a future observation x_{n+1} at stage $n+1$ mentioned above as follows:

$$p_{n+1}(k,z) = \begin{cases}
Pr[\bar{x}_{n+1} \in (q_{z-1}, q_z) \mid \bar{x}_n = q_k] & \text{if } z < j \text{ and } z \neq 0; \\
Pr[\bar{x}_{n+1} \in (q_z, q_{z+1}) \mid \bar{x}_n = q_k] & \text{if } z > j \text{ and } z \neq G; \\
Pr[\bar{x}_{n+1} \in (q_{j-1}, q_{j+1}) \mid \bar{x}_n = q_k] & \text{if } z = j; \\
Pr[\bar{x}_{n+1} \leq q_z) \mid \bar{x}_n = q_k] & \text{if } z = 0; \\
Pr[\bar{x}_{n+1} \geq q_z) \mid \bar{x}_n = q_k] & \text{otherwise.}
\end{cases} \tag{5.7}$$

If we represent the data by the sufficient statistic $\bar{x}_n = \frac{1}{n}\sum_{i=1}^{n} x_i$, we can evaluate the expected utility of stopping, $U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \bar{x}_n = q_k)$, in (5.3) for any grid point $q_k \in Q$ at any stage n. Using the discrete distribution outlined above we approximate the expected utility of continuing in 5.5 as

$$\widehat{U}_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \bar{x}_n = q_k) = \sum_{z=0}^{G} U^*_{n+1}(\mathbf{t}_{n+1}, \bar{x}_{n+1} = q_z) p_{n+1}(k, z). \tag{5.8}$$

We note that at the last stage N, where the trial must stop, we have for each $q_k \in Q, \quad k = 0, 1, \ldots, G$

$$U^*_N(\mathbf{t}_N, \bar{x}_N = q_k) = \max_{t=S,E}\{U_N(d_N = 1, t_{n+1}, \mathbf{t}_N, \bar{x}_N = q_k)\}. \tag{5.9}$$

Then, moving one step backwards we apply (5.8) to estimate $\widehat{U}_{N-1}(d_{N-1} = 0, \mathbf{t}_{N-1}, \bar{x}_{N-1} = q_k)$. Working backward in that manner yields $\widehat{U}_n(d_n = 0, \mathbf{t}_n, \bar{x}_n = q_k)$ for earlier stages $n = N-2, \ldots, 1, 0$.

5.3.2 The case of unknown precision

We now outline the above numerical approximation of the optimal sequential decision at any stage n considering patient observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$ which are $N(\mu_E, \lambda_E)$ with μ_E and λ_E both unknown. Suppose we assume

$$\lambda_E \sim \Gamma(\alpha_0, \beta_0) \quad \text{and} \quad \mu_E \mid \lambda_E \sim N(\mu_0, k_0\lambda_E) \tag{5.10}$$

where $\alpha_0 > 0$, $\beta_0 > 0$, μ_0 and $k_0 > 0$ are constants. The joint prior density $p(\mu_E, \lambda_E)$ corresponds to the normal-gamma density, denoted by $Ng(\mu_E, \lambda_E|\mu_0, k_0, \alpha_0, \beta_0)$, which is a conjugate prior for the normal likelihood

$$l(\mu_E, \lambda_E|\mathbf{x}_n) \propto \lambda_E^{n/2}\exp\{-\frac{\lambda_E}{2}[ns^2 + n(\bar{x}_n - \mu_E)^2]\} \tag{5.11}$$

where $ns^2 = \sum_{i=1}^n (x_i - \bar{x}_n)^2$. Hence the joint posterior distribution $p(\mu_E, \lambda_E|\mathbf{x}_n)$ of μ_E and λ_E also has a normal-gamma density, $Ng(\mu_E, \lambda_E|\mu_n, k_n, \alpha_n, \beta_n)$, where

$$\begin{aligned}
\mu_n &= \frac{k_0\mu_0 + n\bar{x}_n}{k_0 + n} \\
k_n &= k_0 + n \\
\alpha_n &= \alpha_0 + \frac{n}{2} \\
\beta_n &= \beta_0 + \frac{1}{2}ns^2 + \frac{1}{2}(k_0 + n)^{-1}k_0 n(\mu_0 - \bar{x}_n)^2.
\end{aligned} \tag{5.12}$$

It is clear that the two dimensional summary statistic $(\bar{x}_n, s^2)$ is sufficient for the parameter vector (μ_E, λ_E) given the data $\mathbf{x}_n$. Given that $x_{n+1} \sim N(\mu_E, \lambda_E)$, the posterior predictive distribution of a future patient observation x_{n+1} at stage $n + 1$ is obtained, by evaluating the integral $\int\int f(x_{n+1}|\mu_E, \lambda_E)p(\mu_E, \lambda_E|\mathbf{x}_n)d\mu_E d\lambda_E$. The result is a student's t distribution, denoted by $\text{st}(\mu_n, \phi_n, w_n)$, where

$$\mu_n = \frac{k_0\mu_0 + n\bar{x}_n}{k_0 + n}, \quad \phi_n = (k_0 + n)(k_0 + n + 1)^{-1}(\alpha_0 + \frac{n}{2})\beta_n^{-1} \quad \text{and} \quad w_n = 2\alpha_0 + n. \tag{5.13}$$

Suppose we let $s_n = \frac{\sum_{i=1}^n x_i^2}{n}$. Since $s^2 = s_n - \bar{x}_n^2$, we may use $(\bar{x}_n, s_n)$ as the joint sufficient statistics for (μ_E, λ_E). We prefer to use the sample mean sum of squares instead of sample

variance, s^2, because it is easier to evaluate a probability statement about the future mean sum of squares s_{n+1} using the predictive posterior distribution of the future observation x_{n+1} given the observed sample mean $\bar{x}_n$. Besides that, it is simpler than s^2. We approximate the optimal sequential decision based on a two dimensional sufficient statistic as follows. We construct a grid in the space $\mathbf{R} \times \mathbf{R}^+$ over which the sufficient statistic, $(\bar{x}_n, s_n)$, takes values. As in the previous section, we start by choosing the lower and the upper grid bounds, denoted by q_L^h and q_U^h with $q_L^h < q_U^h$ for $h = 1, 2$ for the values of $\bar{x}_n$ and s_n respectively. Let $Q^h = \{q_0^h, q_1^h, \ldots, q_G^h\}$ be sets of values corresponding to the quantiles of the distributions of the random variables $X|t = S$ and $X^2|t = S$ respectively. We note that values in Q^1 are chosen in the same way we choose grid points in Q in the previous section and while for the set Q^2 we may do the same or replace the q_i^2's by the mid points of every subdivision (q_i^2, q_{i+1}^2). The set of grid points $(q_i^1, q_l^2) \in Q^1 \times Q^2, \quad i, l = 0, 1, \ldots, G$ is the support of the discrete distribution approximating the predictive distribution of continuous sufficient statistic at any future stage.

Suppose the trial is at stage n. Given that $(\bar{x}_n, s_n) = (q_{i'}^1, q_{l'}^2)$ for some $i', l' = 0, 1, \ldots, G$, we let the pair $(y_{(n+1,i')}^1, y_{(n+1,l')}^2)$ be the bivariate discrete random variable, defined on the grid $Q^1 \times Q^2$, whose distribution approximates that of a future continuous observation $(\bar{x}_{n+1}, s_{n+1})$ at the next stage . This approximation is such that if the future observation $(\bar{x}_{n+1}, s_{n+1})$ lies in any subdivision of $\mathbf{R} \times \mathbf{R}^+$, then the pair of values (a grid point) assumed by $(y_{(n+1,i')}^1, y_{(n+1,l')}^2)$ is found by assigning values to $y_{(n+1,i')}^1$ and $y_{(n+1,l')}^2$ in Q^1 and Q^2 respectively, as in (5.6). Hence we have that to every subdivision there corresponds a grid point, a value of $(y_{(n+1,i')}^1, y_{(n+1,l')}^2)$, or vice versa. The probability that $(y_{(n+1,i')}^1, y_{(n+1,l')}^2)$ assumes a grid point (q_i^1, q_l^2) given that $(\bar{x}_n, s_n) = (q_i'^1, q_l'^2)$, denoted by $p_{n+1}(i, l, i', l')$, is equal to the predictive probability that the future observation $(\bar{x}_{n+1}, s_{n+1})$ lies in the subdivision that corresponds to (q_i^1, q_l^2) and, may be computed using the above stated student's t predictive probability density. Hence for each grid point $(q_i^1, q_l^2) \in Q^1 \times Q^2$ at any stage n we calculate the expected utility of stopping $U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \bar{x}_n = q_i^1, s_n = q_l^2)$ and approximate the

expected utility of continuing as

$$\widehat{U}_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \bar{x}_n = q^1_{i'}, s_n = q^2_{l'}) = \sum_{i=0}^{G}\sum_{l=0}^{G} U^*_{n+1}(\mathbf{t}_{n+1}, \bar{x}_{n+1} = q^1_i, s_{n+1} = q^2_l) p_{n+1}(i, h, i', l'). \quad (5.14)$$

As in the previous subsection, we apply backward induction to compute $\widehat{U}_n$ for earlier stages $n = N-1, N-2, \ldots, 1, 0$.

5.4 Examples

We consider two examples to illustrate our proposed numerical approximation of the optimal sequential decision at any stage of a one-arm clinical trial.

5.4.1 Example 1: the case of known precision

Suppose we consider a one-arm clinical trial design described in section 5.1 where the patient observation is $x|t = E \sim N(\mu_E, \lambda_E)$ with λ_E known. For the purpose of illustration we set $N = 20, a = 1, b = 2, \lambda_E = \lambda_S = 1.2^{-1}$ and $\mu_S = 2.5$. Also we assume that the prior is $\mu_E \sim N(2.7, 1.8^{-1})$. We recall that at any given stage n of the trial the posterior predictive distribution of a future observation x_{n+1} given past data is normal with mean $\mu_n = \frac{n\lambda_E\bar{x}_n + \mu_0\lambda_0}{n\lambda_E + \lambda_0}$ and precision $\frac{\lambda_E\lambda_n}{\lambda_n + \lambda_E}$ where $\lambda_n = n\lambda_E + \lambda_0$. Hence the expected utility of stopping in (5.3) becomes

$$U_n(d_n = 1, t_{n+1}, \mathbf{t}_n, \bar{x}_n) = \frac{1-\alpha}{N} n(a + b\bar{x}_n) + \{\alpha + (1-\alpha)\frac{N-n}{N}\}(a + b\mu), \quad (5.15)$$

where

$$\mu = \begin{cases} \mu_S & \text{if } t_{n+1} = S; \\ \mu_n & \text{if } t_{n+1} = E. \end{cases} \quad (5.16)$$

We let the lower and upper grid bounds q_L and q_U be 5% and 95% quantiles of the distribution $N(2.5, 1.2^{-1})$ respectively. The interval (q_L, q_U) is then subdivided by 89 quantiles of the same normal distribution giving a total of $G = 91$ grid points . Computing the approximate expected utility $\widehat{U}_n(d_n = 0, t_{n+1} = E, \mathbf{t}_n, \bar{x}_n)$ as described in subsection 5.2.1 for every $y_k \in Q$ at any future stage we obtain the approximate stopping rule given in figure 5.2.

5.4.2 Example 2: the case of unknown precision

We now take a second example to illustrate our numerical approximation of the optimal sequential decision in one-arm trial design where $x|t = E \sim N(\mu_E, \lambda_E)$ with both parameters unknown. Again for the purpose of illustration we assume in this case that $N = 12$, $x|t = S \sim N(2.5, 1.2^{-1})$ and assign the prior probability density

$$\lambda_E \sim \Gamma(1.2, 1) \qquad \text{and} \qquad \mu_E|\lambda_E \sim N(2.3, \lambda_E).$$

The expected utility of stopping remains the same as in (5.15) except that μ_n is defined as in subsection 5.2.2 (we choose $k_0 = 1$).

For purpose of this example we construct both the sets Q^1 and Q^2 by the same way we construct Q in subsection 5.2.1, and use the grid parameters

$$\begin{aligned} q_L^1 &= 0.6982 \\ q_U^1 &= 4.3018 \\ q_L^2 &= 0.5063 \\ q_U^2 &= 0.19.0358 \\ G &= 91. \end{aligned}$$

Hence applying backward induction as outlined in subsection (5.2.2) we compute the approximate expected utility $\widehat{U}_n(d_n = 0, t_{n+1}, \mathbf{t}_n, \bar{x}_n, s_n)$ for every grid point in $Q^1 \times Q^2$ at

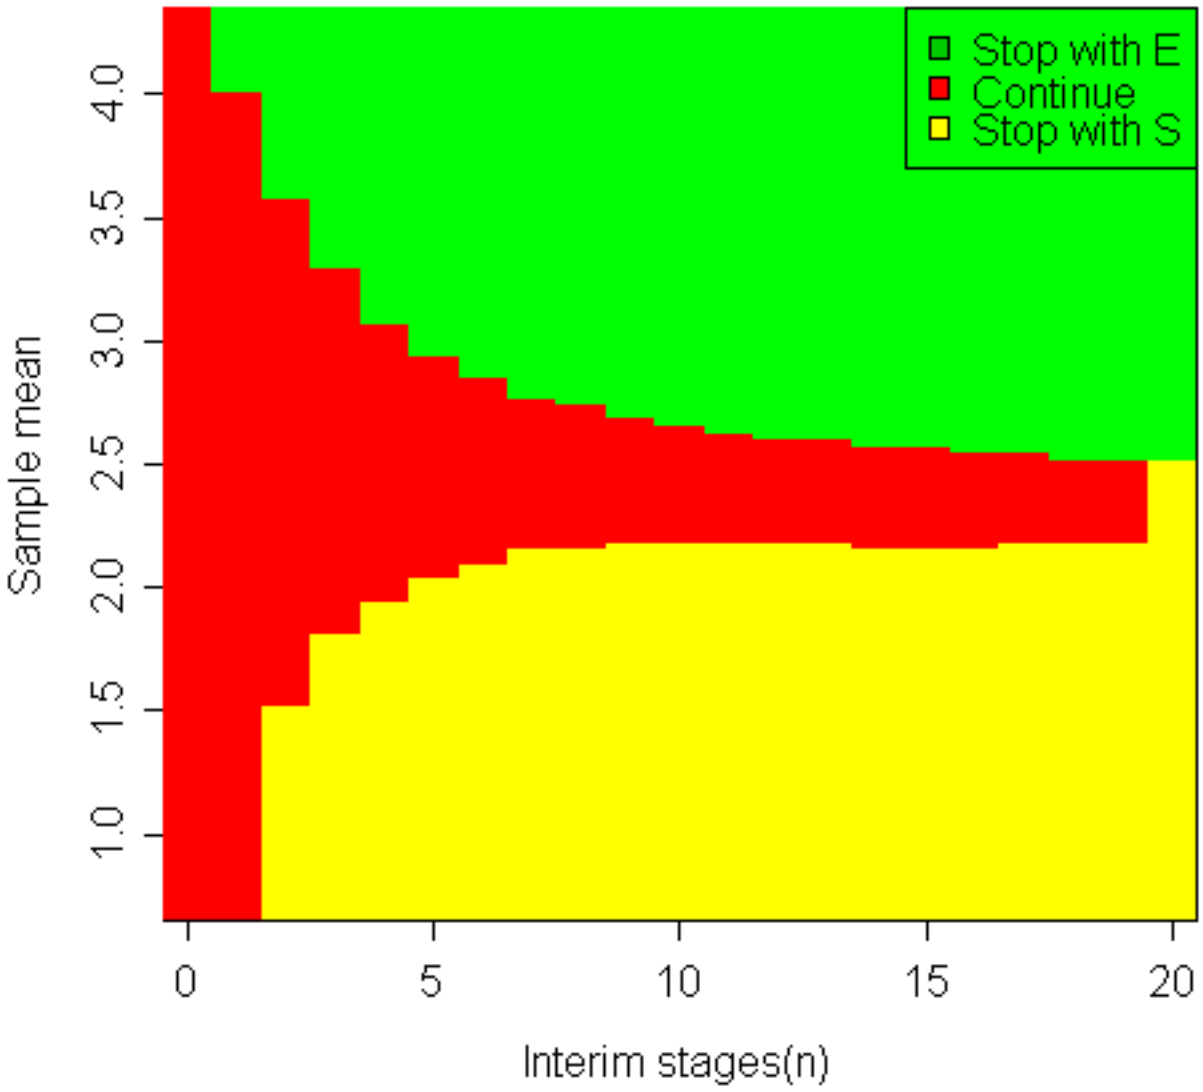

Figure 5.2: Stopping rule: the green, red and yellow regions indicate the decisions to stop with treatment E, continue and stop with treatment S respectively

any future stage. At the initial stage, $n = 0$, the value of the expected utility of stopping in (5.15)is equal to 2.5 and the approximated expected utility of continuing was obtained as 4.828. Hence by our approximation the decision is to continue to the next stage. The resulting approximate stopping rule for $n \geq 1$ is given in the figures 5.3 through 5.14.

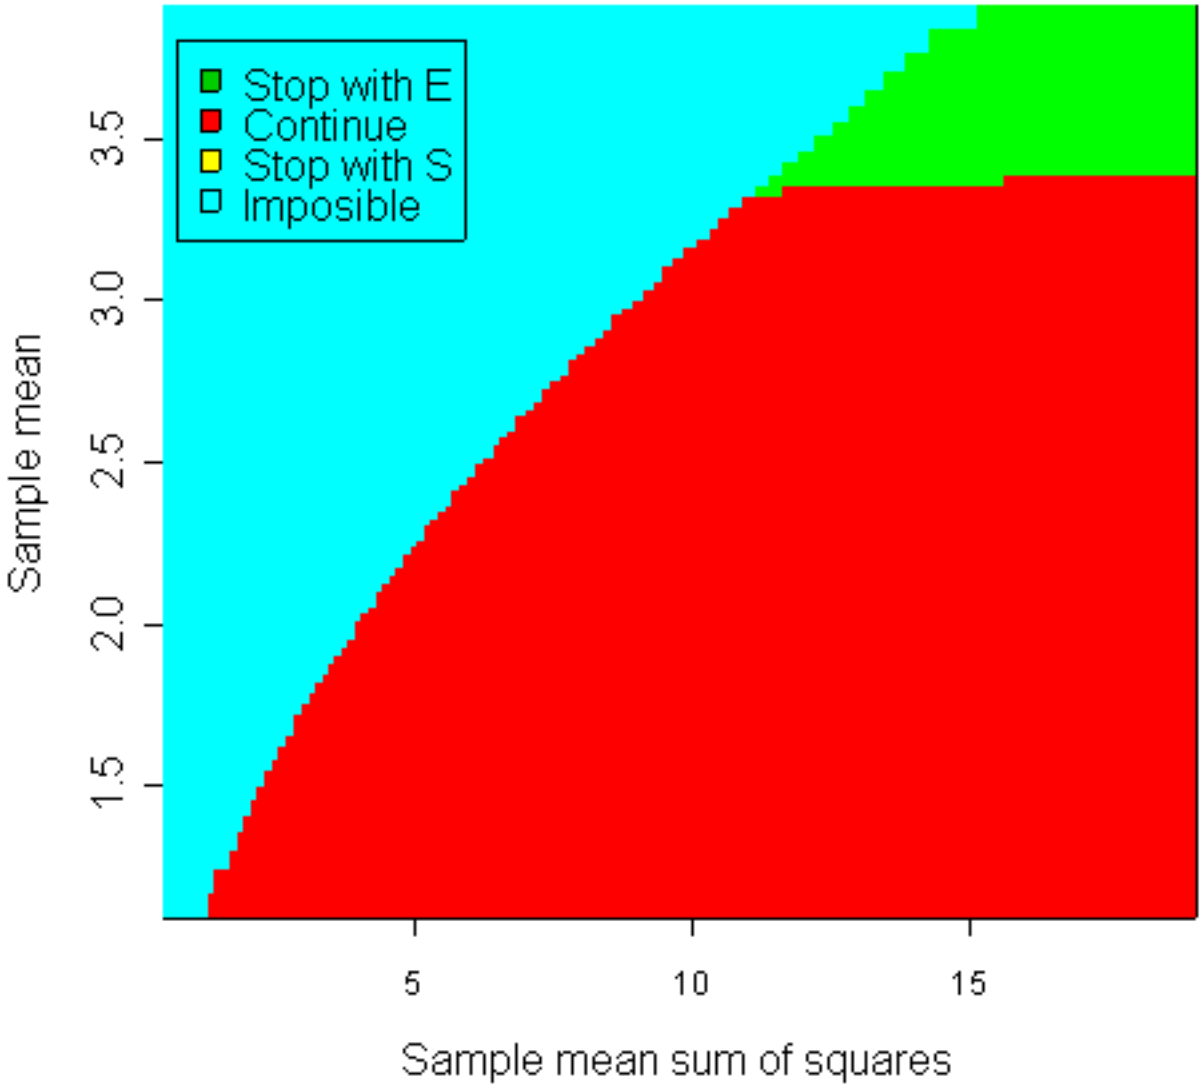

Figure 5.3: Stopping rule: $n = 1$. Stopping rule: In figures 5.3 through 5.14, the green, red and yellow regions indicate the decisions to stop with treatment E, continue and stop with the standard treatment respectively. The light blue region represent impossible case where $\bar{x}_n > s_n$.

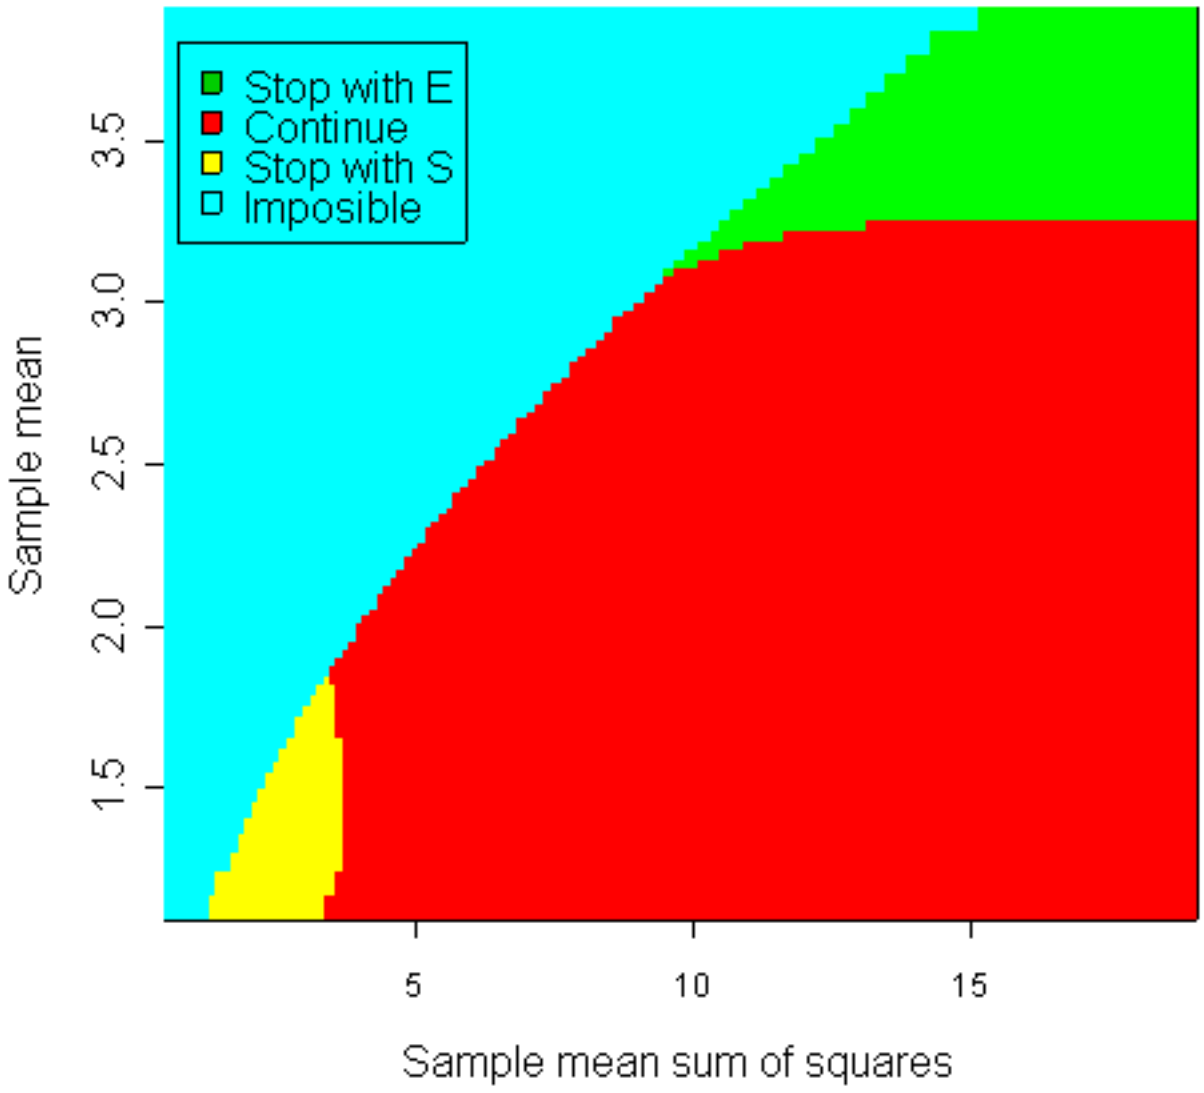

Figure 5.4: Stopping rule: $n = 2$

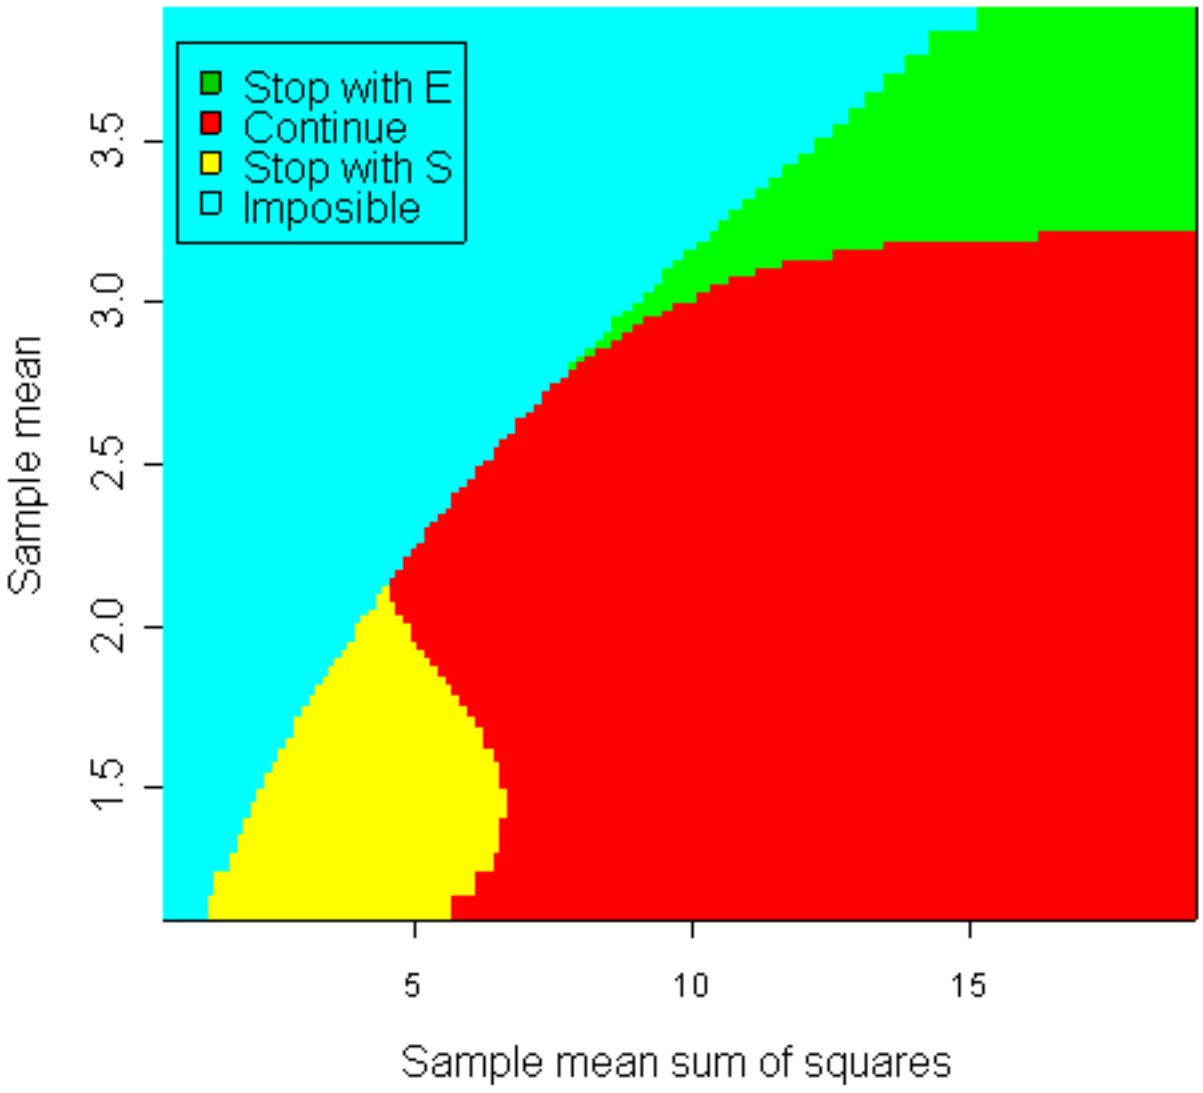

Figure 5.5: Stopping rule: $n = 3$

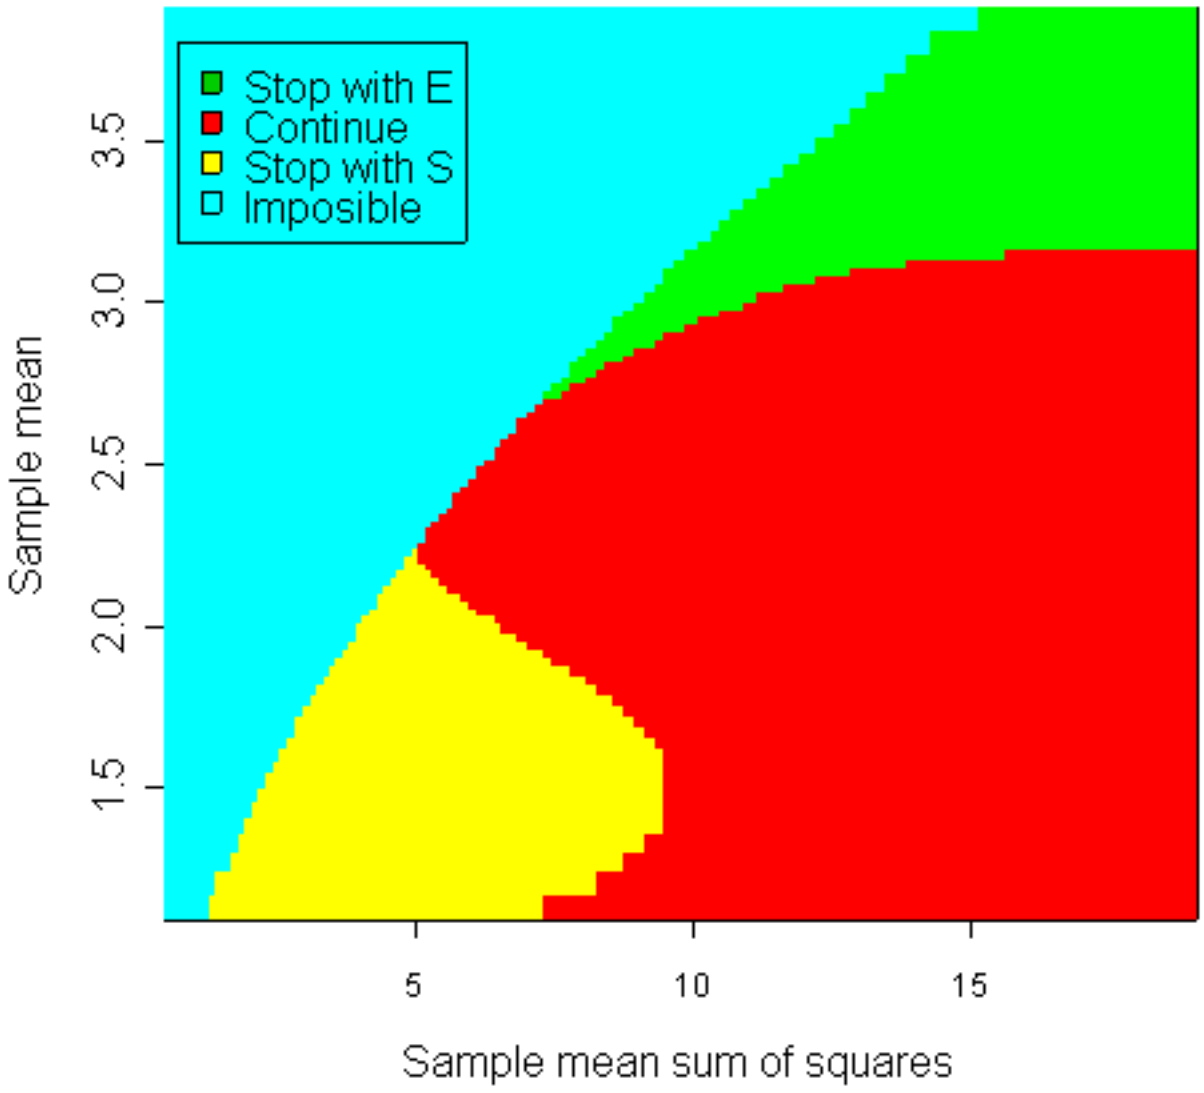

Figure 5.6: Stopping rule: $n = 4$

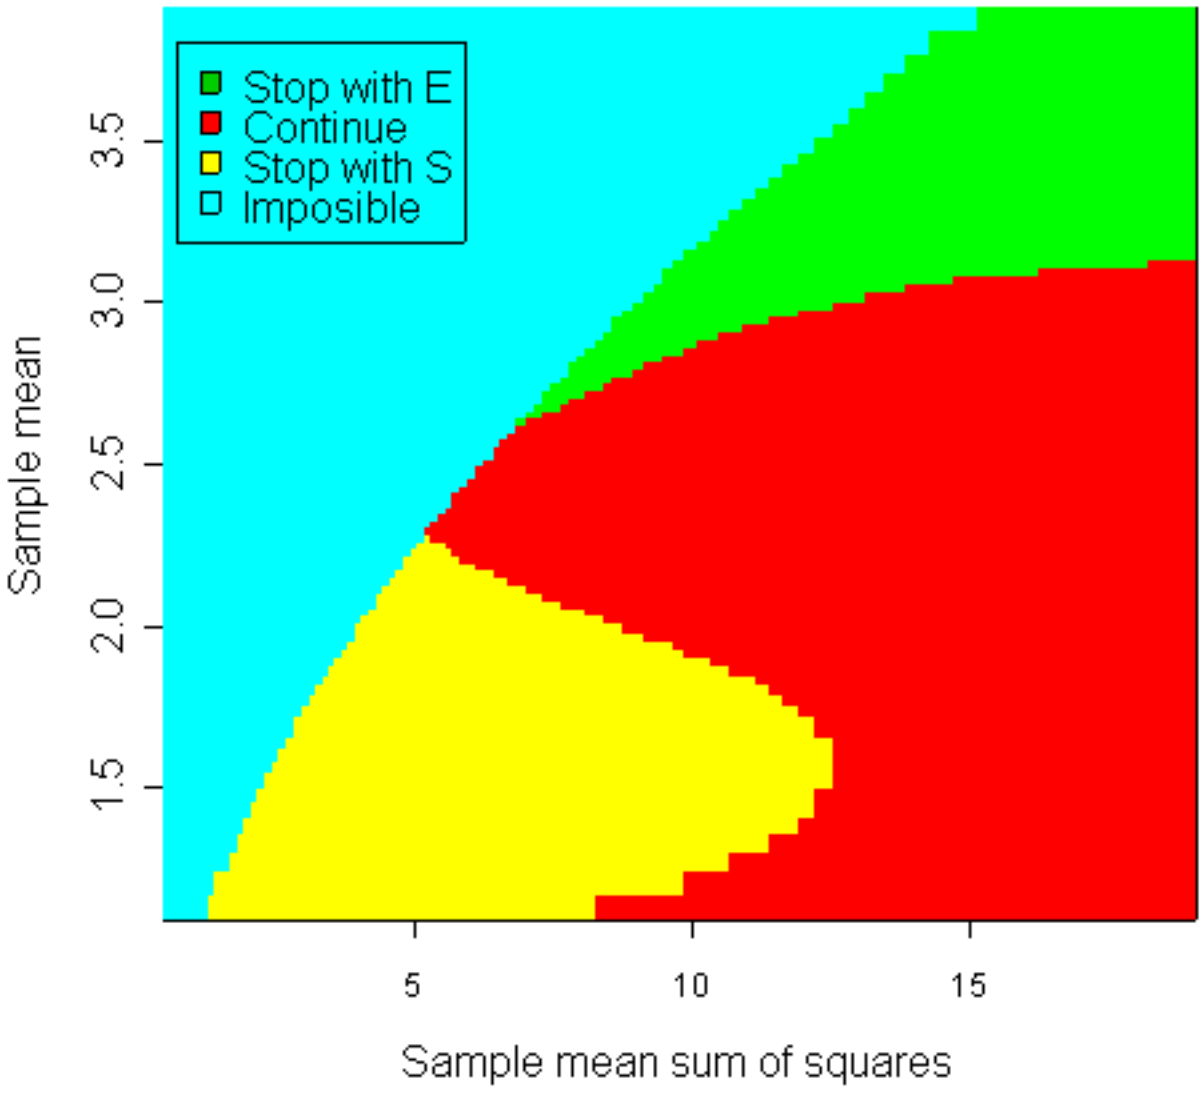

Figure 5.7: Stopping rule: $n = 5$

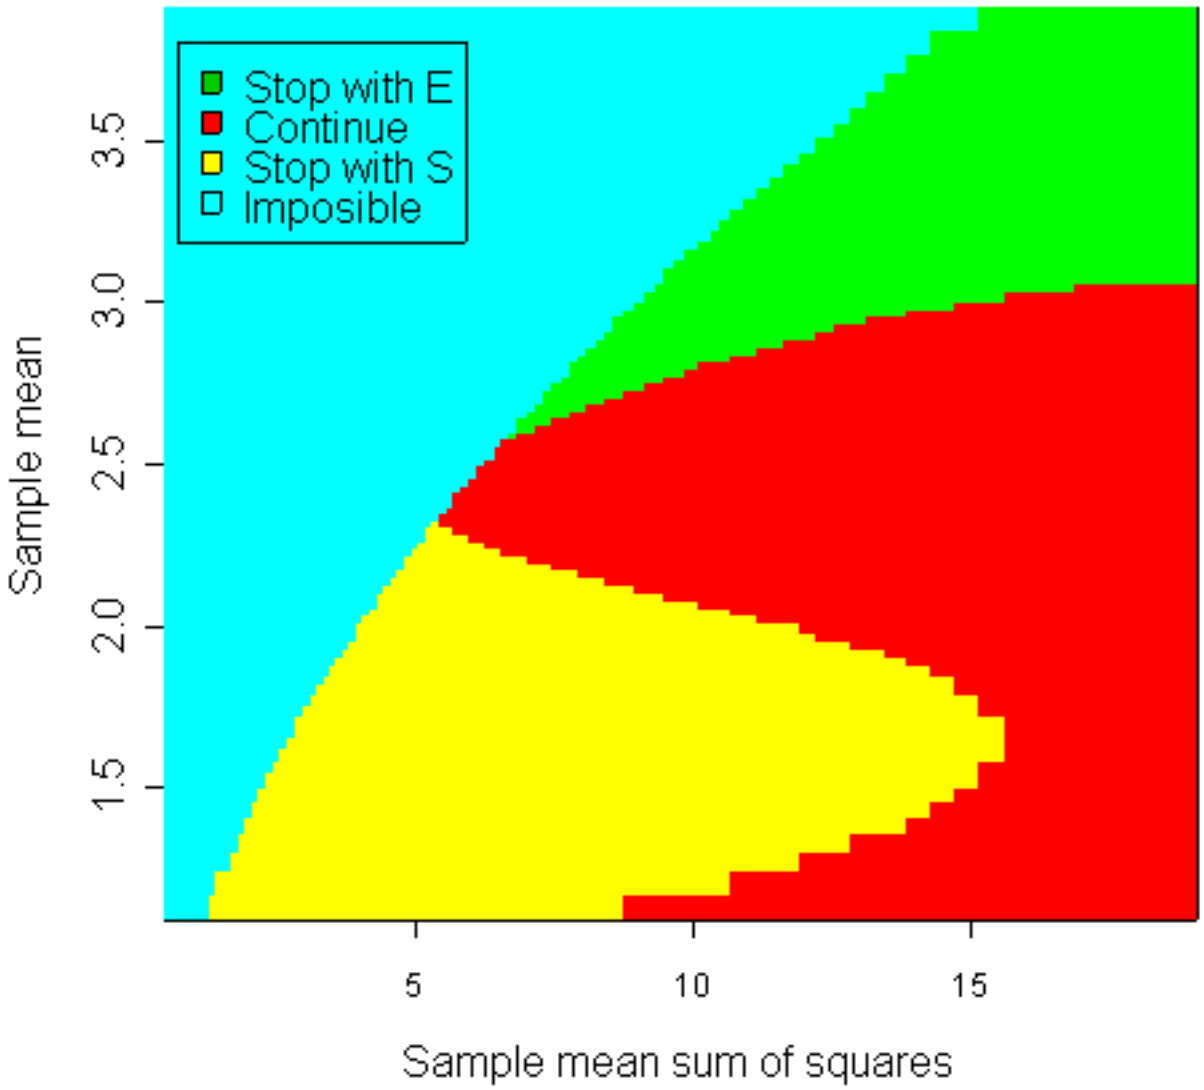

Figure 5.8: Stopping rule: $n = 6$

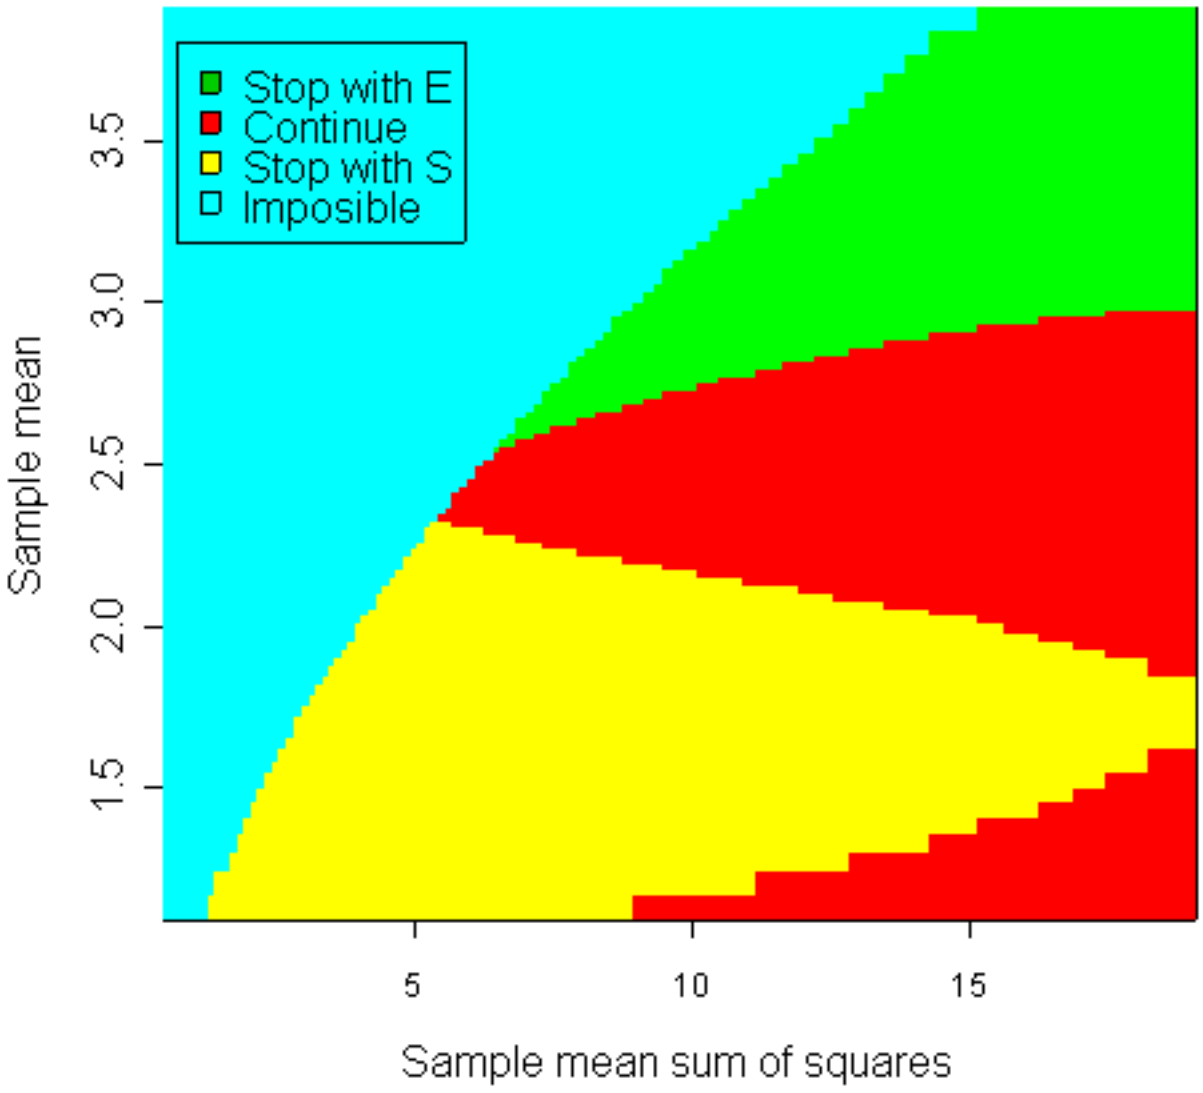

Figure 5.9: Stopping rule: $n = 7$

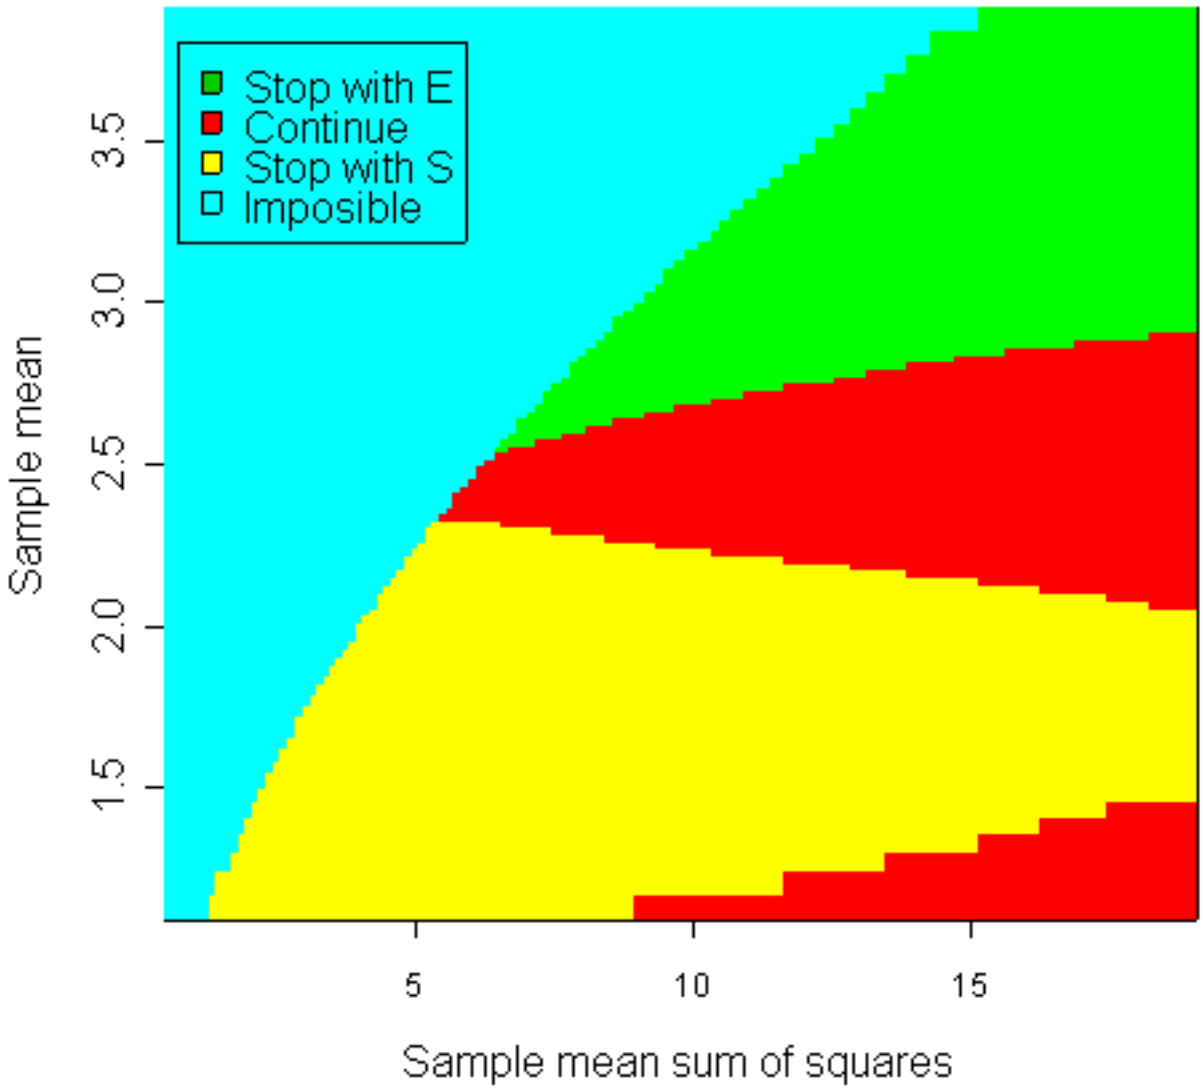

Figure 5.10: Stopping rule: $n = 8$

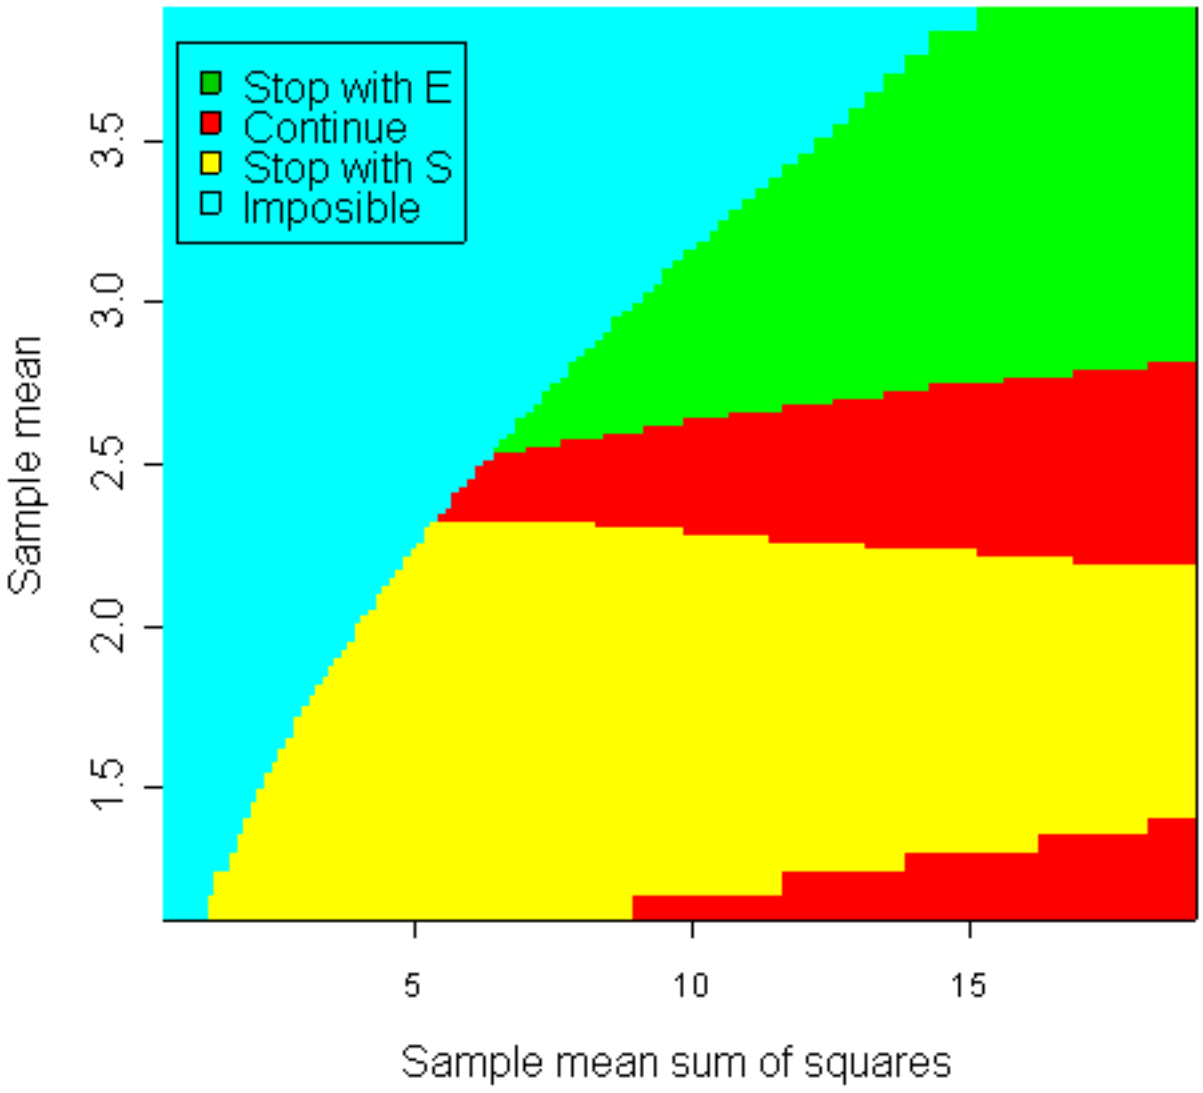

Figure 5.11: Stopping rule: $n = 9$

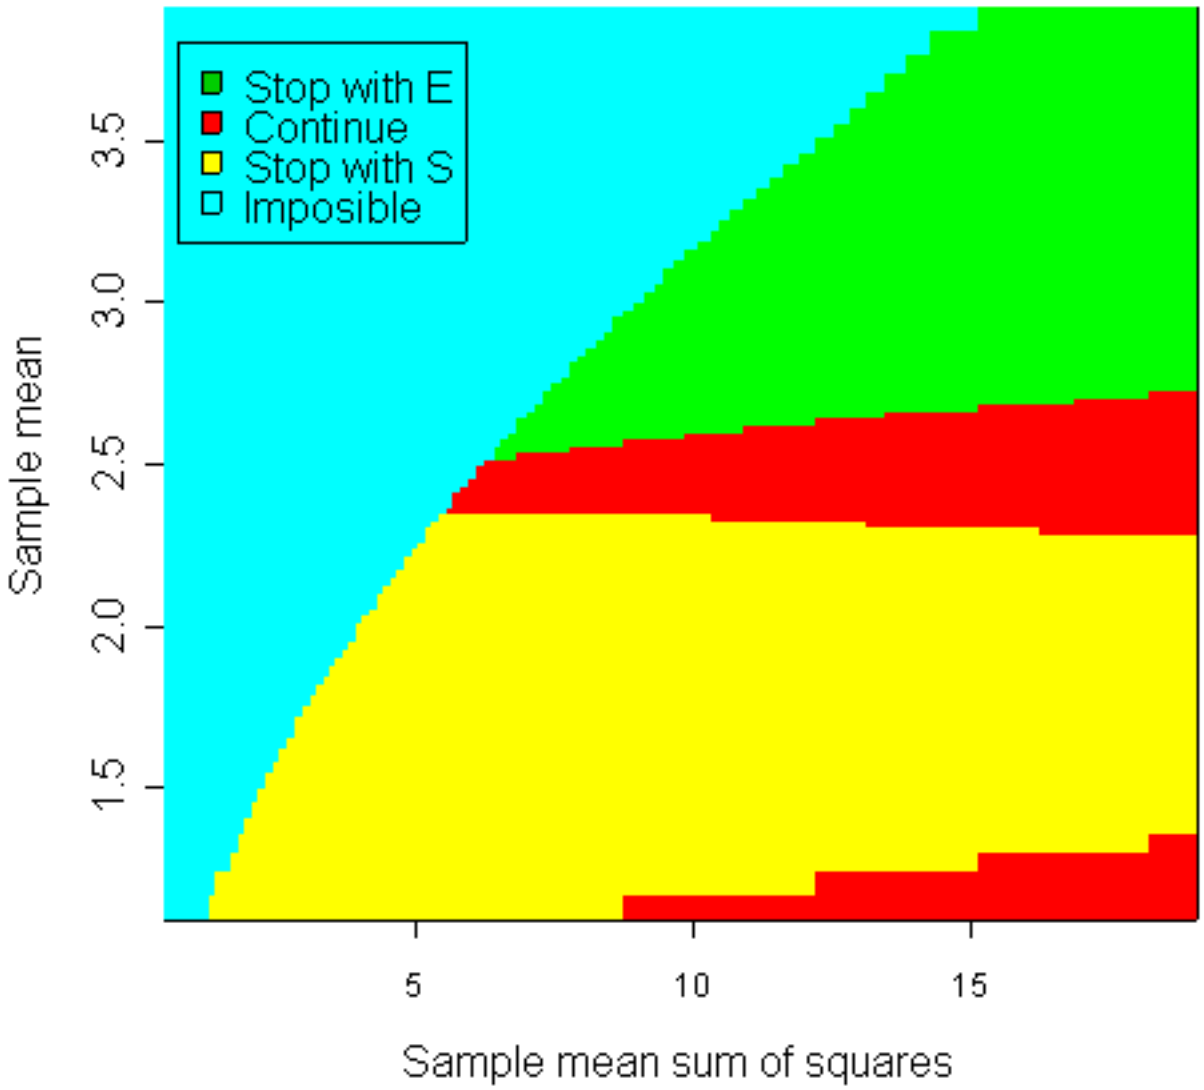

Figure 5.12: Stopping rule: $n = 10$

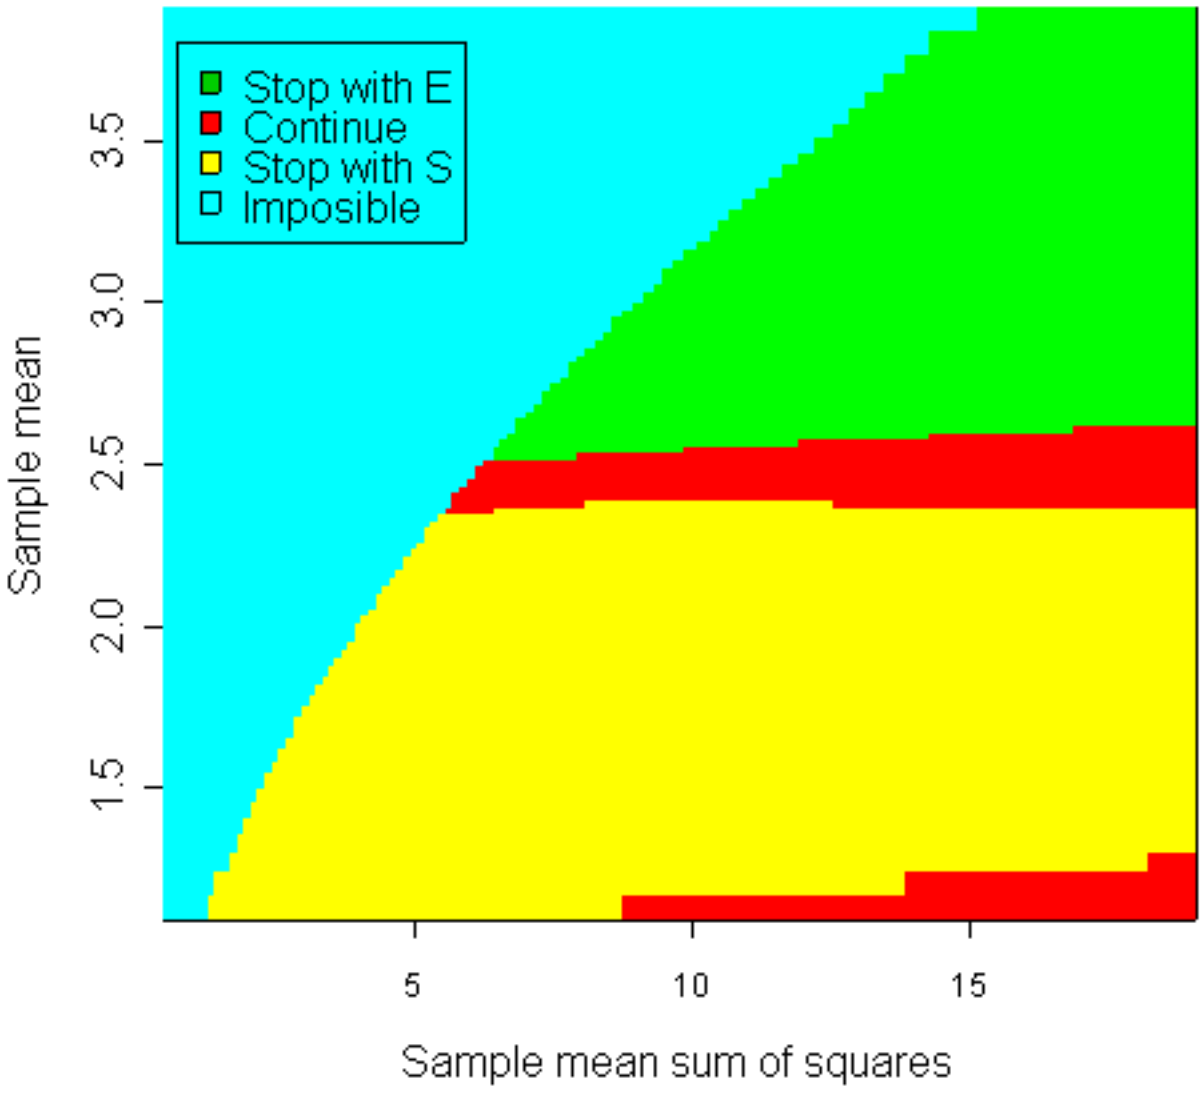

Figure 5.13: Stopping rule: $n = 11$

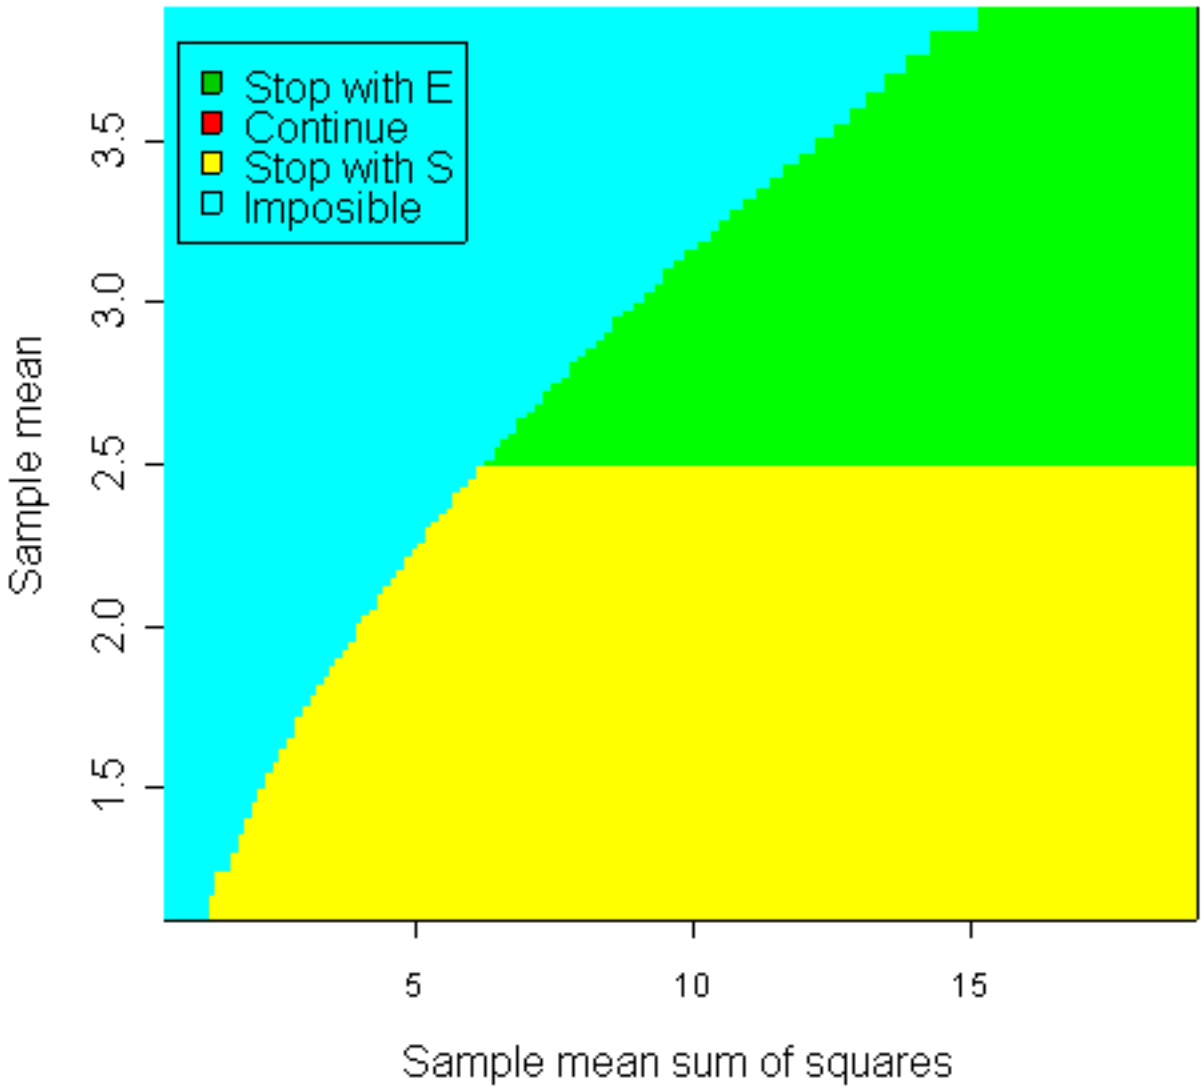

Figure 5.14: Stopping rule: $n = 12$

We describe the process of decision making at a given stage of the trial as indicated by the shape of our stopping rule as follows. We consider figure 5.8 which corresponds to stage $n = 6$ of the trial and represents the true shape of our stopping rule. Let the value of s_n be fixed at around 12, for instance, and increase the value of $\bar{x}_n$ starting from the lowest. We note from the relation $s^2 = s_n - \bar{x}_n^2$ that, with s_n kept fixed, increasing $\bar{x}_n$ decreases the sample variance. For the first interval of small values of $\bar{x}_n$ the trial continues because the variation is big enough to guarantee observing a value $\bar{x}_{n+1}$ close to μ_s at the next stage. As we ascend we encounter a second interval of higher values of $\bar{x}_n$ where the variance is too small to enable reaching anywhere near μ_S and hence the trial stops. Once we are near but below μ_S the trial continues again due to the fact that the variation, although small, may enable observing a value close to μ_S. Passing μ_S, the trial still continues over an interval of values of $\bar{x}_n$ for which the variation is small but may enable observing a value of $\bar{x}_{n+1}$ below μ_S. We finally encounter the last interval where the trial stops because the variation is so small that chances are high that we observe values of the sample mean greater than μ_S for a number of consecutive future interim stages.

5.5 Discussion

We have used exact predictive probabilities to obtain gridding approximation in a one-arm clinical trial with normal response. Our approach looks more reasonable and does not need a huge number of grid points to obtain smooth stopping boundaries as indicated in the figures in the previous section. The trial design that we have considered here is such that interim analysis is carried out after every patient has been treated and assessed. Computationally, this may be cumbersome for the case of unknown precision when the number of patients enrolled in the trial is big. This problem can be avoided by allowing patients to enter in a group at any stage of the trial (see Pocock, 1977).

Our approach is restricted to conjugate normal and normal-gamma prior densities for the two cases that we have considered, respectively. Making slight changes of the the prior densities, we computed again the stopping boundaries for both cases and observed minimal differences. We noted that for the case of known precision, a huge decrement in the precision caused minimal changes on the stopping boundaries compared to example 1 is subsection 5.4.1. This was due to that fact that the information contained in the prior changes minimally. While a big increase in the precision causes a big change on the stopping boundaries. This is due to the fact that the predictive density becomes more concentrated about the mean $\mu_S = 2.5$ and needs a very small or big observation in order to move away from it. Figure 5.15 is the result of repeat example 1 is subsection 5.4.1 with a normal prior $\mu_E \sim N(2.6, 0.2^{-1})$ with a higher precision.

Given that the exact implementation of gridding approximation of the optimal stopping rule is through the use of the exact predictive probabilities, our approach, save for cases with extreme computational difficulties, is more efficient than the previous simulation-based approaches. All the stopping boundaries displayed in this chapter we obtained by a software written in R and is available from $http://www.cimat.mx/\sim jac/material/thesisOrawo.zip$.

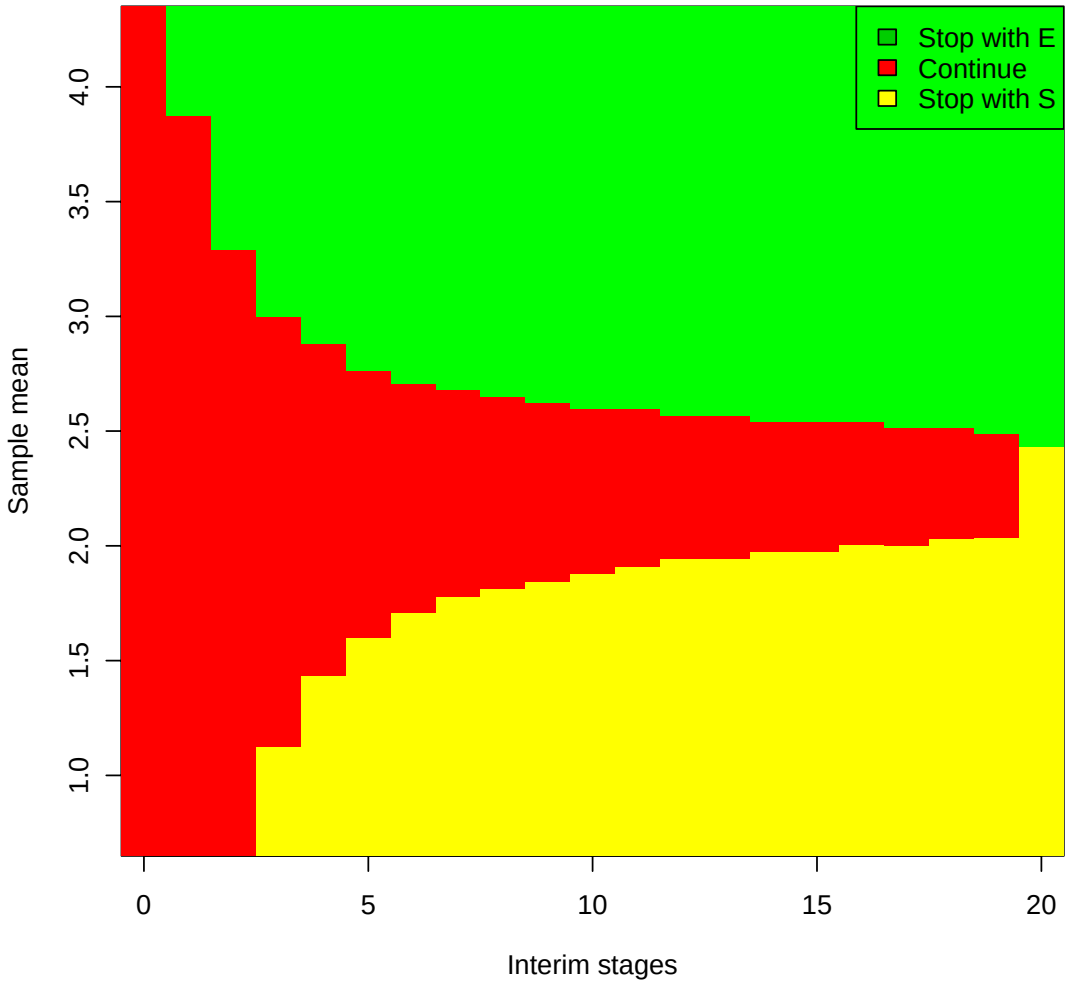
Stop with E
Continue
Stop with S
Sample mean
1.0
1.5
2.0
2.5
3.0
3.5
4.0
0
5
10
15
20
Interim stages

Chapter 6

Discussion

This thesis presents approximation of optimal stopping rules for three different types of clinical trial designs where data are monitored regularly using the Bayesian approach. The three trial designs include a multiple-arm trial design, a trial design for two related trials (both with binary responses) and a one-arm trial design with continuous response, discussed in Chapters 3, 4 and 5 respectively. Our search for approximate stopping rules was motivated by the fact that in general it is impossible to obtain optimal stopping rules by direct implementation of Backward Induction for any of the three trial designs.

In Chapter 3 we have discussed approximation of the optimal stopping rule for a multiple-arm trial design by proposing a novel stopping rule, denoted by τ^p, using the optimal stopping rule of a single-arm trial design obtained by Backward Induction. At any stage of the trial the construction of τ^p consists of making a systematic selection of some the branches of the resulting complex decision tree and then beginning from the end of the trial, working backwards, the estimate of the expected utility of continuing is obtained. We have considered numerical examples to provide evidence that τ^p approximates the optimal stopping rule obtained by Backward Induction with a reasonable accuracy. We fail to conduct exhaustive

evaluation of our proposed trial design when the number of treatment arms exceeds two since the trial design becomes more complex and that there is no previous approximation work under the same design setting available in the literature. In spite of this, we believe that the use of τ^p will enable more efficient comparison of many new treatments with a standard treatment in a single trial as compared to conducting a series of trials with at most two treatment arms. Given that the implementation of Backward Induction for a single-arm clinical trial design with binary outcomes is simple and fast, the construction of our proposed stopping rule is simple and hence requires straight forward programming. Further investigation into the behavior of τ^p may provide more information on its accuracy.

Similarly, we have presented a solution in Chapter 4 to a sequential problem involving two related trials by proposing two stopping rules, denoted by τ^p and τ'^p, one for each trial. We construct the two stopping rules by an original technique such that the value of each stopping rule at any interim stage depends on data from both trials and hence permits learning about one trial while patient admission is confined to the other. Although the assessment of our proposed trial design has been narrowed to studies of the operating characteristics due to the lack of simple examples and previous approximations in the literature, it is possible that the performance achieved may be higher than that of conducting two separate trials. Studies of its operating characteristics with informative joint priors may provide more evidence of its efficiency. Furthermore, assigning other prior distributions that lead to simple expressions for posterior distributions will improve computing time and may allow a more exhaustive study of our proposed stopping rules for the two related trials.

We note that our approach is limited to the assumption that the unknown success probabilities, of treatments within a multiple-arm trial design, are independent; since otherwise the trial design cannot be converted to a single-arm clinical trial. This also applies to each of the two related trials. So far we have justified the effectiveness of our proposed stopping rules using simulated data and the next crucial stage would be to use them to conduct clinical trials involving real patients. Clinical experts, once convinced with the logic behind the develop-

ment of these stopping rules, would require that we quantify the error that we might commit in the course of the trial to see whether it is acceptable or not. For the case a multiple-arm trial design we may use the results of the comparison in table 3.1 in section 3.3 to find out the percentage decrease in the expected utility. Also we may establish the mean estimation bias in all our simulated trials by comparing an estimate of the success probabilities with its true value (this may be done since the final values of the sufficient statistic are provided for each trial simulation in $http://www.cimat.mx/\sim jac/material/thesisOrawo.zip$).

Lastly, we have presented the implementation of Backward Induction using the gridding method in a single-arm clinical trial design where patient response is normal. We note that, given that a computer implementation of Backward Induction is only possible to sequential experiments with discrete outcomes, a gridding method must be employed in a numerical analysis of sequential experiments with continuous outcome. Our improvement over other approaches is due to the fact that we obtain the gridding approximation of the expected utility of continuing using the exact posterior predictive probability distribution of a future observation. We note that evaluating the gridding approximation of expected utility by simulation, as has been done in the previous approaches, is equivalent to approximating the value that we obtain. However, our approach may encounter computational difficulties if the expected utility depends on a high-dimensional sufficient statistic or when N is large and many interim looks are considered. Our approach is the restricted to the family of normal conjugate priors and the existence of a sufficient statistic for the unknown parameters of interest.

Although we have used some specific prior distributions, proper elicitation of prior distributions will require working in collaboration with experts who have the initial knowledge about the unknown success probabilities. Perhaps more importantly, to turn our trial designs into true applicable protocols, the utility value for any outcome under a given treatment, $v(t, r)$, will need to be correctly established. This may be done in terms of quality-adjusted life years (QALY) (see Kaplan, 1995 and Parmigiani, 2002). The value may be determined by asking

a patient the period of time that he or she considers equivalent to one year in a given health state. Other alternative methods have calculated QALY's using preference-based measures of health related quality of life. These methods have been used in the UK and involve eliciting utilities of various health states from a representative sample of the general population. The various heath states defined by each of the above methods are valued using the standard gambling technique. Then based on this data statistical modelling is used to estimate the utilities for all the health states, (see Kharroubi, Brazier, Roberts and O'Hagan, 2005 and reference there in). However, these QALTY of heath states have been done in a handful of countries and may not be directly applicable to other societies like Mexico.

The utility function u, that we have used throughout this thesis, is not limited to any specific form by our algorithms. Hence any other utility functions, leading to similar optimal stopping rules for a single arm trial and is suited to a problem at hand, may be used.

Further investigation may involve studying the approximation of the optimal stopping rules for multiple-arm clinical trial designs when dependence is introduced among the various unknown treatment success probabilities.

We conclude by saying that we have provided simple algorithms that permit finding approximate stopping rules with impressive properties, for each of the three trial designs. Use of more simple examples of these trial designs will enable exhaustive evaluation of these stopping rule. We believe that the evaluation presented in this thesis will be convincing enough for our proposed trial designs to be tested in real clinical trials.

Appendix A

Notation and Basic concepts

A.1 Bayesian framework

Suppose that we are interested in the values of the k unknown parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_k)$ and that $\boldsymbol{\theta}$ itself is a random variable and has a prior probability distribution denoted by $p(\boldsymbol{\theta})$. Now suppose we have n observations $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ whose probability distribution $p(\mathbf{x} \mid \boldsymbol{\theta})$ depends on the values of the k parameters. Then we have

$$\begin{aligned} p(\mathbf{x} \mid \boldsymbol{\theta})p(\boldsymbol{\theta}) &= p(\mathbf{x}, \boldsymbol{\theta}) \\ &= p(\boldsymbol{\theta} \mid \mathbf{x})p(\mathbf{x}). \end{aligned} \tag{A.1}$$

The conditional probability distribution $p(\boldsymbol{\theta} \mid \mathbf{x})$ is called the posterior probability distribution of $\boldsymbol{\theta}$ given the observed data $\mathbf{x}$ and is obtained as

$$p(\boldsymbol{\theta} \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{x})} \tag{A.2}$$

where

$$p(\mathbf{x}) = \begin{cases} \int p(\mathbf{x} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta} & \text{if } \boldsymbol{\theta} \text{ is continuous;} \\ \sum p(\mathbf{x} \mid \boldsymbol{\theta})p(\boldsymbol{\theta}) & \text{if } \boldsymbol{\theta} \text{ is discrete,} \end{cases} \tag{A.3}$$

and the sum is taken over all possible values of $\boldsymbol{\theta}$. The formula in (A.2) is usually referred to as Bayes' Theorem.

Given the observed data $\mathbf{x}$, $p(\mathbf{x} \mid \boldsymbol{\theta})$ is regarded as a function of $\boldsymbol{\theta}$ and is called the likelihood. Noting that $p(\mathbf{x})$ does not depend on $\boldsymbol{\theta}$ and thus for a fixed $\mathbf{x}$ can be considered constant, the Bayes formula can be expressed as

$$p(\boldsymbol{\theta} \mid \mathbf{x}) \propto p(\mathbf{x} \mid \boldsymbol{\theta})p(\boldsymbol{\theta}), \tag{A.4}$$

meaning that the posterior probability distribution of $\boldsymbol{\theta}$ given $\mathbf{x}$ is proportional to the likelihood times the prior probability distribution. Hence the posterior distribution summarizes the prior information and the information provided by the observed data about the unknown parameter $\boldsymbol{\theta}$.

Suppose x_{n+1} is a future observation independent of $\mathbf{x}_n$ given $\boldsymbol{\theta}$. Then the predictive distribution of x_{n+1} is given by

$$p(x_{n+1} \mid \mathbf{x}_n) = \int f(x_{n+1} \mid \boldsymbol{\theta})p(\boldsymbol{\theta} \mid \mathbf{x}_n)d\boldsymbol{\theta}. \tag{A.5}$$

This distribution is usually referred to as the posterior predictive distribution because it summarizes the information about the likely value of a new observation given the prior and the data we have observed so far. If no data have been observed the distribution of a future observation x_{n+1} given in (A.5) now becomes

$$p(x_{n+1}) = \int f(x_{n+1} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}, \tag{A.6}$$

and is called the prior predictive distribution.

Suppose we have a decision problem where a choice has to be made from a set of available actions. Let $a \in \mathcal{A}$ be the class of possible actions. The loss function, denoted by $u(a, \boldsymbol{\theta})$, quantifies the loss incurred when $\boldsymbol{\theta}$ is the true state of nature and we take action a. In

the light of the data $\mathbf{x}$, our knowledge about $\boldsymbol{\theta}$ is represented by the posterior distribution $p(\boldsymbol{\theta} \mid \mathbf{x})$, which combines the prior knowledge of $\boldsymbol{\theta}$ with the information provided by the data. According to a Bayesian approach, the optimal action is that which minimizes the expected posterior loss,

$$\mathcal{U}(a) = \int_{\Theta} u(a, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid \mathbf{x}) \mathbf{d}\boldsymbol{\theta}. \tag{A.7}$$

A.2 Notation

- $t = E_0, E_1, \ldots, E_k$ ($t = E, S$ for $k = 1$): treatments.
- x: patient response.
- $v(t, x)$: utility of observing x after assigning treatment t.
- N: fixed maximum number of patients.
- u: the utility function for entire trial.
- τ: stopping rule.
- $U_n^{(\tau)}$: the expected utility determined by a stopping rule τ.
- U_n: the expected utility.
- $d_n = 0, 1$: the decisions of continuing and stopping respectively.
- U_n^*: the maximum expected utility.
- θ_j: the success probability of treatment E_j, $j = 0, 1, \ldots, k$.
- $\pi_j(\theta_j)$: the prior distribution of θ_j.

A.3 Models for clinical trials with binary responses

Suppose we have a sequential clinical trial design where k experimental treatments $E_1, E_2, \ldots, E_k$ are compared with a standard treatment E_0. Assume that the success probabilities of the $k+1$ treatments, θ_0, θ_1, $\ldots$, θ_k, have the prior probability distributions $\pi_0(\theta_0)$, $\pi_1(\theta_1)$, $\ldots$, $\pi_k(\theta_k)$ respectively. Suppose we assume that patient outcome is binary. Let t_i and x_i denote the assigned treatment and the the observed response for the i^{th} patient, $i = 1, 2, \ldots, N$. At any stage n of the trial, we have the sequence of binary observations $\mathbf{x}_n = (x_1, x_2, \ldots, x_n)$ from the first n patients arising from applying treatments $\mathbf{t}_n = (t_1, t_2, \ldots, t_n)$. Suppose that we summarize the information contained in $(\mathbf{t}_n, \mathbf{x}_n)$ by the sufficient statistic $(s_{(k)}, n_{(k)})$ with $n_{(k)} = (n_0, n_1, ..., n_k)$ and $s_{(k)} = (s_0, s_1, ..., s_k)$. Suppose for each j we let $A_j = \{i : t_i = E_j, \quad i = 1, 2, \ldots, n\}$ be the set indicating the patients assigned to treatment E_j, the respective j^{th} components of $n_{(k)}$ and $s(k)$ are given by $n_j =\mid A_j \mid$ and $s_j = \sum_{i \in A_j} x_i$ respectively. If we assume that the random variables $\theta_0, \theta_1, \ldots, \theta_k$ are independent apriori, then the posterior distribution is given by

$$\pi(\theta_0, \theta_1, \ldots, \theta_k \mid s_{(k)}, n_{(k)}) \propto f(s_{(k)} \mid n_{(k)}, \theta_0, \theta_1, \ldots, \theta_k) \prod_{j=0}^{k} \pi_j(\theta_j), \tag{A.8}$$

where $f(s_{(k)} \mid n_{(k)}, \theta_0, \theta_1, \ldots, \theta_k) = \prod_{j=0}^{k} \theta_j^{s_j}(1-\theta_j)^{n_j - s_j}$ is the likelihood function. Consequently, the marginal posterior density function for θ_j, $j = 0, 1, \ldots, k$ is obtained as

$$\pi_j(\theta_j \mid s_j, n_j) \propto \theta_j^{s_j}(1-\theta_j)^{n_j - s_j}\pi_j(\theta_j). \tag{A.9}$$

Suppose $\pi_j(\theta_j)$ is a beta probability distribution with parameters α_j and β_j. We obtain the marginal posterior density $\pi(\theta_j \mid n_j, s_j)$ in (A.9) as a beta probability density with parameters $\alpha_j + s_j$ and $\beta_j + n_j - s_j$. Then the posterior predictive distribution of a future observation x_{n+1} at the next stage of the trial under treatment $t_{n+1} = E_j$ is given by

$$p(x_{n+1} \mid s_j, t_{n+1} = E_j) = \int_0^1 p(x_{n+1} \mid \theta_j)\pi_j(\theta_j \mid s_j, n_j)d\theta_j. \tag{A.10}$$

This reduces to

$$\begin{aligned} p(1 \mid s_j, t_{n+1} = E_j) &= \int_0^1 \theta_j \pi_j(\theta_j \mid s_j, n_j)d\theta_j \\ &= E[\theta_j \mid n_j, s_j] \\ &= \frac{\alpha_j + s_j}{n_j + \alpha_j + \beta_j} \end{aligned}$$

and

$$p(0 \mid s_j, t_{n+1} = E_j) = 1 - \frac{\alpha_j + s_j}{n_j + \alpha_j + \beta_j}. \tag{A.11}$$

We note that, at any stage n, the posterior predictive success probability of any treatment $t_{n+1} = E_j$ is equal to the mean of the marginal posterior distribution of θ_j.

Appendix B

Some Considerations on Operating Characteristics

We note that the tables appearing in sections 3.4, 3.5, 4.6 and 4.7 are summaries of the full simulations and hence many other things may be done to understand fully the operating characteristics of our proposed trial designs. For instance, drawing a histogram of the number of patients assigned to each of the three treatments may help understand more on how consistent our proposed trial designs are in obeying the assumed order of superiority in each scenario and allowing for early stopping, over the $1,000$ simulations. The histograms in figure B.1 correspond to scenario 3 in table 3.2 of section 3.4, where a non informative beta prior $beta(0.5, 0.5)$ was assigned to each of the unknown treatment success probabilities. We observe from the first bar of these histograms that the trial stopped with at most five patients assigned to the less effective treatments E_0 and E_1 approximately 80% and 70% of the times, respectively. This is happens with the most effective treatment approximately 30% of the time. According to the order of the assumed true success probabilities of the three treatments it is easier to get a failure with treatments E_0 and E_1 than with treatment

E_2, while simulating the trial. Hence the less effective treatments easily loose to the most effective treatment. Furthermore, there is a possibility that the trial design allocates patients to only one of the three treatments until stopping without assigning any patient to the other two; the superior treatment will always be preferred most of the time. This also explains the high variation reported in the column marked St.dev in other tables and in table B.1. The histograms also indicate that most of the time more patients were assigned to the superior treatment.

The histograms in figure B.2 report the simulations obtained from a repeat of scenario 3 in table 3.2 with informative beta priors in (3.3). Compared to histograms in figure B.1, we observe a remarkable decrease in the number of simulations for the entire trial where more than five patients were assigned to the less effective treatments; for all the 1000 simulations not more than five patients were assigned to the least effective treatment E_0. Our proposed trial design assigns more patients, most of the times, to the most effective treatment with a high precision as the the corresponding histogram is peaked and concentrated around the average number of patients assigned to it. Also, we observe by comparing the last bars in the histograms corresponding to treatment E_1 in figures B.1 and B.2 that in this case the trial stops early most of the time. Similar detailed interpretation may also be done using the results of the operating characteristics in section 4.7.

Scenario 3

Trt	$\tilde{n_t}$	St.dev.	θ_t	P_t
E_0	5.82	11.64	0.3	0.102
E_1	19.47	16.36	0.6	0.689
E_2	9.02	14.24	0.4	0.209
$\tilde{n}$	34.31	13.10		

Table B.1: Operating characteristics: scenario 3

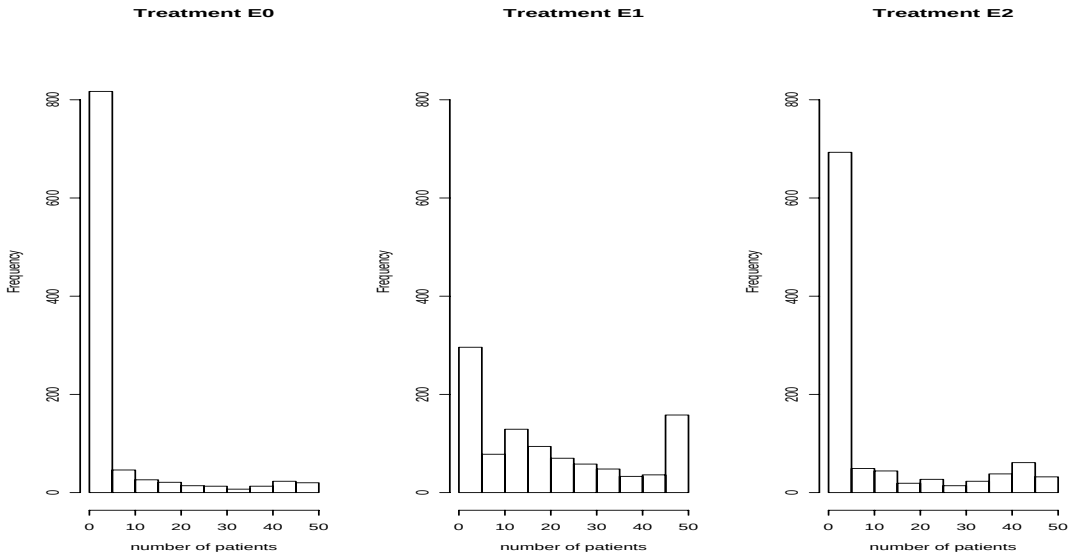

Figure B.1: Operating characteristics: histogram of the number of patients assigned to each treatment in Scenario 3.

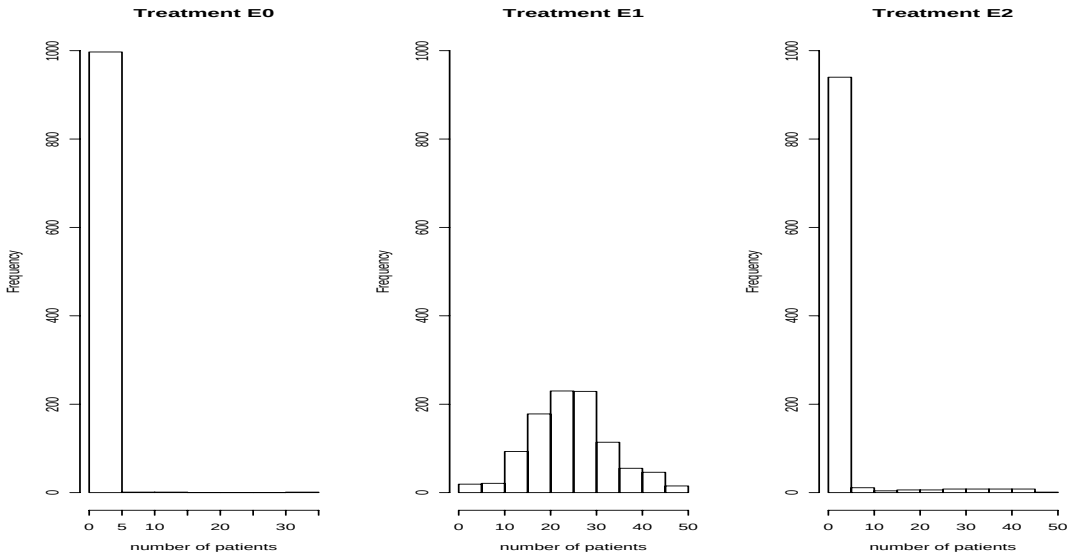

Figure B.2: Operating characteristics: repeat of scenario 3 with informative beta priors, $beta(1.5, 3.5)$, $beta(3, 2)$ and $beta(2, 3)$.

Bibliography

[1] Anscombe F.J. (1963). "Sequential medical trials". *Journal of the American Statistical Association*, **58**, 365 – 383.

[2] Armitage P. (1975). *Sequential Medical Trials*, Oxford: Blackwell,

[3] Berger, J.O. (1985). *Statistical Decision Theory and Bayesian Analysis*, Second Edition. Springer-Verlag: NY.

[4] Berry D.A. (1985). "Interim analyses in clinical trials: classical vs. Bayesian approaches". *Statistics in medicine*, **4**, 521 – 526.

[5] Berry D.A. (1987). "Statistical inference, designing clinical trials, and pharmaceutical company decisions". *The Statistician*, **36**, 181 – 189.

[6] Berry D.A. (1993). "A case of Bayesianism in clinical trials". *Statistics in Medicine*, **12**, 1377-1393.

[7] Berry D.A. and Chih-Hsiang H. (1988) "One-sided sequential stopping boundaries for clinical trials: A decision-theoretic approach". *Biometrics*, **44**, 219-227.

[8] Brockwell, A.E.and Kadane, J.B. (2003). "A gridding method for Bayesian sequential decision problems", *Journal of Computational and Graphical Statistics*, **12**, 566-584

[9] Carlin, B.P., Kadane, J.B. and Gelfand, A.E. (1998). "Approaches for optimal sequential decision analysis in clinical trials". *Biometrics*, **54**, 964 − 975

[10] Chaloner K. (1996). "Elicitation of prior distributions". In *Bayesian Biostatistics*(D.A. Berry and D.K. Stangl, Eds.). Marcel Dekker, New York, pp. 261 − 277.

[11] Choi S.C. and Pepple P.A. (1989). "Monitoring clinical trials based on the predictive probability of signicance". *Biometrics*, *45*, 317 − 323.

[12] Christen, J.A., Müller, P., Kyle, W. and Wolf, J. (2004). " A Bayesian randomized clinical trial: A decision theoretic sequential design". *The Canadian Journal of Statistics*,**4**, 387 − 402.

[13] Christen, J.A. and Nakamura, M. (2003). " Sequential stopping rules for species accumulation". *Journal of Agricultural, Biological and Environmental Statistics*, **8**, 184 − 195.

[14] Cornfield J. (1966). "A Bayesian test of some classical hypotheses-with applications to sequential clinical trials". *Journal of the American Statistical Association*, **61**, 577−594.

[15] DeGroot, M.H. (1970). *Optimal Statistical Decisions*. McGraw-Hill:NY.

[16] Demets D.L. and Ware J.H. (1980). "Group sequential methods for clinical trials with a one-sided hypothesis ". *Biometrika*, **67**, 651 − 660

[17] Elfring G.L. and Schultz J.R. (1973). "Group sequential designs for clinical trials". *Biometrics*, **29**, 471 − 477.

[18] Freedman L.S., Lowe D. and Macaskill P. (1984). "Stopping rules for clinical trials incorporating clinical opinion". *Biometrics*, **40**, 575 − 586.

[19] Geisser S. (1984). "On prior distributions for binary trials". *The American Statistician*, **4**, 244 − 248.

[20] Kadane J.B. and Wolf L.J. (1998). "Experiences in elicitation". *The Statistician*, **47**, 3 − 19.

[21] Kaplan R.M. (1995). "Utility assessment for estimating quality-adjusted life years". In *Valuing Health Care* (Sloan F.A.,Ed.), Cambridge University Press, New York, pp. $31 - 60$.

[22] Kharroubi S.A., Brazier J.E., Roberts J. and O'Hagan A. (2005). "Modelling SF-6D health state preference data using a nonparametric Bayesian method". (to be submited)

[23] Lewis R.J. and Berry D.A. (1994). " Group sequential clinical trials: A classical evaluation of a Bayesian decision-theoretic designs". *Journal of the American Statistical Association*, **428**,$1528 - 1534$.

[24] Parmigiani G. (2002). *Modeling in Medical Decision Making: A Bayesian Approach.* Wiley: Chichester.

[25] Petkau A.J. (1978). " Sequential Medical Trials for comparing an experimental with a standard treatment". *Journal of the American Statistical Association*, **362**, $328 - 338$.

[26] Pocock S.J. (1977). "Group sequential methods in the design and analysis of clinical trials". *Biometrika*, **64**, $191 - 199$.

[27] Spiegelhalter D.J., Abrams K.R. and Myles, J.P. (2004) *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*, Willey: NY.

[28] Spiegelhalter D.J., Freedman L.S. and Parmar M.K.B. (1994). " Bayesian approaches to randomized trials". In *Bayesian Biostatistics* (D.A. Berry and D.K. Stangl, Eds.). Marcel Dekker, New York, pp. $67 - 108$.

[29] Stallard N., Thall P.F., and Whitehead J. (1999). " Decision-theoretic designs for phase II clinical trials with multiple outcomes". *Biometrics*, **55**, $971 - 977$.

[30] Stallard N. (2003). "Decision-theoretic designs for phase II clinical trials allowing for competing studies". *Biometrics*, **59**, $402 - 409$.

[31] Thall P.F. and Simon R. (1994). "Practical Bayesian guidlines for phase IIB clinical trials". *Biometrics*, **50**, $337 - 349$.

[32] Thall P.F., Simon R.M. and Estey E.H. (1995). " Bayesian sequential monitoring for single-arm clinical trials with multiple outcomes". *Statistics in medicine*, **14**, $357 - 379$.

[33] Wathen K. and Christen, J.A. (2004)." Implementation of Backward Induction for Sequential Adaptive Clinical Trials", (submitted).

[34] Whitehead J. (1997). *Design and analysis of sequential clinical trials*, second edition. Wiley: Chichester.

[35] Whitehead J. and Jones A.D. (1979). "The analysis of sequential clinical trials". *Biometrika*, **66**, $443 - 452$.