


Outcome-dependent randomisation schemes for clinical trials
with fluctuations in patient characteristics



D. Stephen Coad
Pembroke College

A thesis submitted to the Faculty of Mathematical Sciences for the degree of
Doctor of Philosophy of the University of Oxford

Department of Statistics
1 South Parks Road

Trinity Term, 1989

Outcome-dependent randomisation schemes for clinical trials

with fluctuations in patient characteristics

D. Stephen Coad

D. Phil. thesis

Pembroke College, Oxford

Trinity Term, 1989

Abstract

A clinical trial is considered in which two treatments are to be compared. Treatment allocation schemes are usually designed to assign approximately equal numbers of patients to each treatment. The purpose of this thesis is to investigate the efficiency of estimation and the effect of instability in the response variable for allocation schemes which are aimed at reducing the number of patients who receive the inferior treatment.

The general background to outcome-dependent allocation schemes is described in Chapter 1. A discussion of ethical and practical problems associated with these methods is presented together with brief details of actual trials conducted.

In Chapter 2, the response to treatment is Bernoulli and the trial size is fixed. A simple method for estimating the treatment difference is proposed. Simulation results for a selection of allocation schemes indicate that the effect of instability upon the performance of the schemes can sometimes be substantial.

A decision-theory approach is taken in Chapter 3. The trial is conducted in a number of stages and the interests of both the patients in the trial and those who will be treated after the end of the trial are taken into account. Using results for conditional normal distributions, analytical results are derived for estimation of the treatment difference for both a stable and an unstable normal response variable for three allocation schemes. Some results for estimation are also given for other responses.

The problem of sequential testing is addressed in Chapter 4. With instability in the response variable, it is shown that the error probabilities for the test for a stable response variable can be approximately preserved by using a modified test statistic with appropriately-widened stopping boundaries. In addition, some recent results for estimation following sequential tests are outlined.

Finally, the main conclusions of the thesis are highlighted in Chapter 5.

To Mum and Dad

Acknowledgements

I am especially grateful to my supervisor Professor P. Armitage for his guidance and patience during the course of this work.

A special thanks to the organisers of the Sequential Methods in Statistics Research Workshop of July 1988 for giving me the opportunity to attend. Valuable discussions with Dr. M.H. Gail and Professor M. Woodroffe led subsequently to several extensions to some parts of this thesis.

Thanks are also extended to Dr. D.N. Geary for his careful reading of drafts of parts of this thesis; to Dr. J.C. Gittins for the loan of the "DAI" program used in Chapters 2 and 4; to You-Gan Wang for useful discussions on some parts of Chapter 4; to other members and the staff of the Department of Statistics for their help and advice; and lastly, but by no means least, to Guillermina for her continual warmth and encouragement.

This work was supported by a Research Studentship from the Medical Research Council, which is gratefully acknowledged.

CONTENTS

Abstract	ii
Acknowledgements	iv
Contents	v
1. Introduction	1
1.1 The different types of outcome-dependent allocation	1
1.2 Ethical problems	2
1.3 Early literature	3
1.3.1 Data-dependent allocation rules	4
1.3.2 Decision-theoretic models	7
1.4 Modifications of practical relevance	9
1.5 Outcome-dependent allocation in practice	10
1.5.1 Examples	10
1.5.2 Why are these methods not used more often?	12
1.6 Direction of thesis	14
2. A comparative study of some data-dependent allocation rules for Bernoulli data	16
2.1 Allocation rules considered	16
2.1.1 Equal allocation rule	17
2.1.2 Play-the-winner rule	17
2.1.3 Fox's rule	17
2.1.4 One-step-ahead rule	18
2.1.5 Bather's rule	18
2.1.6 Rules based on Gittins indices	19
2.2 Review of recent comparisons	22
2.2.1 Definitions	22
2.2.2 Summary of results	22
2.3 Details of simulation study	24
2.3.1 Computation	24
2.3.2 Criteria for comparison of rules	25
2.3.3 Estimation of treatment difference	26
2.4 Stable response variable	28
2.4.1 Analytical results	28
2.4.2 Simulation results	30
2.4.3 Differences between the two estimates of variance	39

2.5	Modelling instability in the response variable	41
2.5.1	Linear time-trends	41
2.5.2	Time-trends and fluctuations on a logit scale	41
2.6	Linear time-trends in the response variable	43
2.6.1	Analytical results	43
2.6.2	Simulation results	46
2.7	Time-trends and fluctuations in the response variable	53
2.8	Discussion	66
3.	Estimation for decision-theoretic models	68
3.1	Models considered	68
3.1.1	Preliminaries	68
3.1.2	Formulation of models	69
3.1.3	Optimisation	71
3.1.4	Numerical results	73
3.2	Estimation for a stable normal response variable	74
3.3	Modelling instability in the response variable	81
3.3.1	Linear time-trends	82
3.3.2	Fluctuations	82
3.4	Estimation for an unstable normal response variable	83
3.4.1	Linear time-trends	83
3.4.2	Fluctuations	85
3.5	Numerical and simulation results	89
3.5.1	Details of simulation study	89
3.5.2	Simulated versus analytical results	89
3.5.3	Analytical results	105
3.6	Estimation of instability parameters	105
3.6.1	Estimation of trend	105
3.6.2	Estimation of serial correlation	106
3.7	A three-stage model	107
3.7.1	Formulation of model	108
3.7.2	Optimisation	108
3.7.3	Numerical results	109
3.8	Other response variables	110
3.8.1	Bernoulli response (exact)	111
3.8.2	Exponential response (exact)	112
3.8.3	Normal approximations	114
3.8.4	Numerical results	115
3.8.5	Future work	117

4. Sequential tests for an unstable response variable	119
4.1 Sequential tests for a stable normal response variable	119
4.1.1 Stopping rule and terminal decision rule	120
4.1.2 Allocation rules considered	121
4.1.3 Estimation	123
4.2 Simulation results	124
4.2.1 Details of simulation study	124
4.2.2 Comparison of allocation rules	125
4.2.3 Estimation	126
4.3 Sequential tests for a normal response variable with linear time-trends	128
4.3.1 Derivation of tests	128
4.3.2 Adjusting the SPRT boundaries	130
4.3.3 Estimation	131
4.4 Simulation results	132
4.4.1 Comparison of allocation rules	132
4.4.2 Estimation	134
4.5 Approximate sequential tests for a Bernoulli response variable with linear time-trends	136
4.5.1 Derivation of tests	137
4.5.2 Adjusting the SPRT boundaries	138
4.5.3 Application to Bernoulli responses	138
4.5.4 Allocation rules considered	139
4.5.5 Estimation	140
4.6 Simulation results	140
4.6.1 Stable Bernoulli response variable	140
4.6.2 Linear time-trends in a Bernoulli response variable	143
4.7 Sequential tests for a normal response variable with fluctuations	146
4.7.1 Derivation of tests	146
4.7.2 Maximum likelihood estimation and adjusting the SPRT boundaries	148
4.7.3 Simulation results when the observed response is serially correlated	152
4.8 Discussion	154
5. Conclusions	157
Appendix 1. Histograms and plots for §2.4.2 when $n = 150$	160

Appendix 2.	Derivation of the $\partial E/\partial p_i$ ($i=1,2,3$) for the three-stage model in §3.7.2	165
Appendix 3.	Results relating to the estimation procedure in §§3.8.1 and 3.8.2	167
Appendix 4.	Results relating to maximum likelihood estimation in §4.7.2	168
Appendix 5.	Results for estimation for the Fixed Model when there are unequal numbers of patients in the first stage (cf. p.158)	169
References		170

Inferiors revolt in order that they may be
equal and equals that they may be superior

– Aristotle

CHAPTER 1. INTRODUCTION

Suppose that we wish to conduct a clinical trial to compare two treatments which are available for a particular illness or disease. In order to reduce the number of patients receiving the inferior treatment, previous observed responses and allocations in the trial are taken into account. Treatment allocation schemes which are designed in this way are termed *outcome-dependent*.

Since the early 1960's, this type of treatment allocation has been studied extensively in the theoretical literature. However, it has received little attention in practice: it appears that only four documented trials have been conducted. This gap between theory and practice has recently been discussed by Armitage (1985), who suggested that there should be "*more collaboration between the proponents of these models and those statisticians concerned directly with the design and analysis of trials*" (p.15). There is some indication of progress in this respect with the documentation of two recent trials conducted in the United States which implemented an outcome-dependent allocation procedure – as we shall see in §1.5, these trials demonstrated the possible merits of this approach.

In this chapter, we begin by distinguishing between the different types of outcome-dependent allocation in §1.1. In §1.2 a short discussion is given of the main ethical problems related to these. Early work in the area is reviewed in §1.3 and, following this, in §1.4 we discuss a selection of publications which address particular practical problems of these methods. Next, in §1.5 we briefly describe the few trials which have been conducted using these methods and outline the main reasons for their neglect. Lastly, some comments on the direction of the thesis are made in §1.6.

1.1 The different types of outcome-dependent allocation

There are essentially two main approaches, which are not exclusive. The first approach involves the use of a *data-dependent allocation rule* which uses the information on the two treatments to allocate proportionately more patients to the currently better treatment. A distinction should be made in this case between those rules that are

deterministic and those that are *probabilistic*. If the allocation rule is deterministic, the next allocation is wholly determined by the previous responses. Such a rule may be based on the sign of the difference between the sample means for the two treatments. On the other hand, if the allocation rule is probabilistic, the allocation probabilities for the next allocation depend on the previous responses. As an example, if we allocate according to whether the standardised value of the estimated mean treatment difference exceeds two in magnitude, say, the appropriate allocation probabilities may be 1/3 and 2/3, etc. (this particular allocation rule is one of the rules considered in Chapter 4). Data-dependent allocation rules are examined in Chapters 2 and 4 in the context of a fixed trial size and sequential testing, respectively.

The second approach involves defining a *patient horizon* – the total number of patients to be treated, including those outside of the trial – and determining the optimal trial size according to a particular loss (or gain) function. If the trial comprises a number of stages, we seek the optimal size of each stage. This formulation is often referred to as a *decision-theoretic* or horizon model, and is the subject of Chapter 3.

Both of these approaches are versions of the *two-armed bandit problem* (Berry & Fristedt, 1985), where we seek a sequence of pulls of the two arms of a slot machine in such a way as to maximise the reward. In the context of a clinical trial, the treatments correspond to the arms, the patients to the pulls of the arms and the treatment response to the reward. Although this type of problem is difficult, and often intractable analytically, solutions to certain bandit problems have been obtained (most notably, Gittins & Jones, 1974). While we will not be dealing specifically with the general theory of bandit problems in this thesis, some of the results which are presented may provide some insight into these.

1.2 Ethical problems

The outcome-dependent allocation approach is aptly summarised by Clayton (1982, p.477) who writes: "*Rather ironically, adaptive designs, although designed specifically with ethical considerations in mind, in a sense heighten the ethical difficulty*". In this section, some of these difficulties are outlined.

Suppose, for instance, that a data-dependent allocation rule is applied in a trial. Although initially patients may be randomised equally to the two treatments, at some later point in the trial the allocation probabilities may become very uneven (e.g. 9:1). Clearly, clinicians may be very reluctant to allocate 10% of the patients to one treatment when the other treatment appears to be significantly more promising. Furthermore, in such a case it may be more ethical to stop the trial and conclude the superiority of one of the treatments. A good example of such a trial is described by Bartlett *et al.* (1985).

The formulation of decision-theoretic models raises a different issue. In the optimisation the interests of all the patients are assumed equal, and thus "individual ethics" are sacrificed for "collective ethics" (Lellouch & Schwartz, 1971). This contrasts with the currently accepted view that the interests of the immediate patient are of prime concern, as is the case for so-called *myopic* allocation rules (see Definition 2.2). An intermediate approach is possible by discounting the interests of the future patients relative to the current patient (Oudin & Lellouch, 1972). Although certain models can be simplified by choosing a particular discounting (see §2.1.6), this approach may not prove convincing to clinicians.

More fundamentally, Peto (1985, p.32) points out that in practice the difference between the two treatments is unlikely to be *large*. Consequently, since the allocation probabilities often fluctuate initially, the probability of assigning most of the patients to the inferior treatment increases. Moreover, the total sample size can be larger than for conventional approaches, in order to achieve a desired level of precision. Such an increase clearly violates the ethical requirement of avoiding an unnecessarily long trial. However, it is stressed by Louis (1985, p.30) that although this is true for normal response variables with equal variances, counterexamples exist when the response variables are exponential or Bernoulli. Results of Hayre & Gittins (1981) also indicate that the case of normal response variables with unequal variances provides a further counterexample. Some of the results in Chapter 4 will highlight this behaviour.

1.3 Early literature

In this section we begin by discussing the early work on the Bernoulli two-armed

bandit problem; since most of this work was not originally formulated in a clinical trial setting, we retain the bandit nomenclature.

1.3.1 Data-dependent allocation rules

(a) Bernoulli response

The first bandit problem was considered by Thompson (1933) who constructed a two-armed bandit with independent Bernoulli arms with parameters p_A and p_B which have beta prior distributions. The following randomised allocation rule is suggested: allocate the next patient to treatment A with probability equal to the current Bayesian probability that $p_A > p_B$, etc. Thompson evaluates this rule using simulation, and continues this work in Thompson (1935). The limited results obtained indicate that this rule performs quite well, though it has not received much attention since (Berry & Fristedt, 1985, pp.252–253).

It was not until 1952 that the bandit problem received significant attention again, this time from H. Robbins. Robbins (1952) discusses the connection of this problem with the Wald (1947) theory of sequential analysis, and poses the problem of choosing sequentially from one of two populations in order to maximise the sum of n observations. Robbins examines a Bernoulli two-armed bandit problem and shows that the *play-the-winner rule* (see §2.1.3) uniformly dominates random allocation. Robbins (1956) generalised this rule by proposing an allocation rule that takes into account the results from the previous r trials – this is commonly referred to as the *finite-memory approach*. This approach was further investigated by Isbell (1959), Smith & Pyke (1965) and Samuels (1968), who derived successively better rules.

Bradt *et al.* (1956) examine several aspects of Bernoulli two-armed bandit problems, the most important being the case where the characteristics of one arm are completely specified (the so-called *one-armed bandit*). This problem was also addressed by Bellman (1956), who adopts a *geometric discount sequence* and produces some results which relate to *dynamic allocation indices* (these terms are defined in §2.1.6). Bradt *et al.* also consider a simple example of a Bernoulli two-armed bandit with dependent arms – the solution of which is the *play-the-loser rule*!

A special class of allocation rules for a Bernoulli two-armed bandit problem was proposed by Vogel (1960a): until some random time $2k$, each arm is chosen alternatively;

at time $2k$, a decision is made using a sequential probability ratio test (SPRT) as to which arm to continue with for the remainder of the experiment. [This approach is similar to the decision-theoretic formulation adopted, for example, by Colton (1963), and which is further described in §1.3.2.] In follow-up work, Vogel (1960b) obtained asymptotic bounds for the minimax value of the *expected successes lost* (see Definition 2.4); further comments on these bounds are included in §2.4.

An important Bernoulli two-armed bandit problem whose solution eluded Bradt *et al.* (1956) was solved by Feldman (1962): if the two arms have probabilities of success equal to either (a,b) or (b,a) , where a and b are known, myopic allocation rules maximise the expected number of successes. Several authors have generalised this work: Kelley (1974) investigates the case when the two arms have probabilities of success equal to either (a,b) or (c,d) , where c and d are known; Zaborskis (1976) and Rodman (1978) both consider a *multi-armed bandit problem* – the former concentrating on Bernoulli arms and maximising the expected number of successes, and the latter on arbitrary distributions with geometric discounting.

Since these early publications, the Bernoulli two-armed bandit problem has gained enormous popularity amongst theoreticians, and consequently there is now an extensive literature on the area. An excellent annotated bibliography is provided by Berry & Fristedt (1985, pp.207–261) – publications of particular interest since 1985 will be mentioned in due course.

(b) Non-Bernoulli response

The use of a data-dependent allocation rule when the response variable is non-Bernoulli was first considered by Flehinger & Louis (1971). The authors investigate a trial with two treatments which have an exponential response variable, and the patients' survival time is of interest. Three hypotheses are under consideration: the death rates for the two treatments are the same, or one or other of the treatments has a higher death rate. A *stopping rule* based on likelihood ratios is proposed and it is shown using simulation that the operating characteristic of the proposed sequential test is insensitive to the allocation rule. The allocation rules considered vary from alternation of treatments to those that allocate patients to the treatment which has currently incurred fewer deaths.

The latter class of allocation rules significantly reduce the number of patients receiving the inferior treatment.

In follow-up work, Flehinger & Louis (1972) consider a trial with two treatments which have a normal response variable with unknown means and variances unity. The response can either be immediate or delayed, and interest lies in the difference in the means. The authors demonstrate, using a similar approach to Flehinger & Louis (1971) and extensive simulation, that the use of a data-dependent allocation rule "*can achieve the stated objective of reducing ITN* [the expected number of patients assigned to the inferior treatment] *without affecting the error probabilities of the clinical trial, and, in fact, accomplish this without sizeable increases in ASN* [the expected number of patients in the trial]" (p.47). Some comments are made regarding the use of randomised versions of the allocation rules proposed. A generalisation of this work is provided by Flehinger *et al.* (1972); in addition, these authors give some mathematical explanation for the simulation results of Flehinger & Louis (1971, 1972).

A special case of the problem posed by Flehinger & Louis (1972) is studied by Robbins & Siegmund (1974): attention is restricted to testing two hypotheses – the hypothesis that there is no significant difference between the two treatments is not considered. Although we examine this problem further in Chapter 4, we note here two important results obtained by Robbins & Siegmund (1974): firstly, the expected number of patients on each treatment must exceed *one-half* of the expected number under pairwise allocation; secondly, the expected total number of patients in the trial is approximately minimised by pairwise allocation (cf. comments in §1.2). In addition, Robbins and Siegmund solve a simplified version of the above problem which, in turn, leads to a better lower bound for the expected number of patients allocated to the inferior treatment than the one-half of pairwise allocation. Some asymptotic results are also proved for the expected sample sizes on the two treatments when using a particular data-dependent allocation rule; details of this rule are given in §4.1.2.

Louis (1975) deals with the same problem as Robbins & Siegmund (1974) but incorporates an additional cost function, namely, $ASN + (\gamma - 1) ITN$, ($\gamma \geq 1$). Louis minimises the expected value of this cost and uses continuous-time results to derive an

asymptotically optimal rule in discrete time. Simulation results indicate that, compared with previous rules, a significant reduction in the expected trial size can be achieved at the cost of only a slight increase in the expected number of patients allocated to the inferior treatment. A general procedure for constructing optimal rules is provided by Louis (1977); this involves defining a measure of information gained in a trial, which is similar to the *Kullback–Liebler number* (Kullback, 1959). To illustrate this procedure, Louis considers exponential response variables and the corresponding simulation results yield conclusions similar to those of Louis (1975).

Unlike the Bernoulli two-armed bandit problem, there has subsequently been comparatively little work carried out using the above sequential testing approach for non-Bernoulli responses. The most valuable is that due to Hayre (1979) who generalises the cost function of Louis (1975) to any *weighted average* of the sample sizes on the two treatments, with the weights being allowed to depend on the true mean treatment difference. The sequential test considered is a two-population analogue of the test proposed by Armitage (1957), which permits the acceptance of a third hypothesis that there is no significant difference between the two treatments. For this test, the error probabilities are shown to be approximately independent of the allocation rule used. An interesting discussion of some of the drawbacks of the data-dependent allocation rule proposed by Hayre is presented by Siegmund (1983).

1.3.2 Decision-theoretic models

One of the earliest papers which deals with loss functions in connection with sequential sampling is Maurice (1957): a loss function is adopted to choose between two normal populations with known variances on the basis of their means. Maurice implements Wald's (1950) *minimax procedure* to determine a sequential allocation rule. A comparison with a fixed-sample allocation rule, and a *double-sampling* allocation rule due to Grundy *et al.* (1956) shows that, in general, the sequential allocation rule has a smaller expected loss. Using a similar approach, Maurice (1959) examines the situation where the two populations are randomised in the first part of an experiment and then a decision is taken as to which population the remaining observations are to be sampled from.

A similar model to that of Maurice (1959) was studied by Colton (1963) in the

context of clinical trials. Colton considers both fixed-sample and sequential allocation rules and the approaches taken include a Bayesian formulation in which a normal prior distribution is placed on the true mean treatment difference. An independent investigation was carried out by Anscombe (1963) for a similar model. Anscombe seeks a sequential allocation rule such that the trial is terminated at that stage after which it is anticipated that the expected posterior loss will increase (the expectation is with respect to the current posterior distribution of the sample mean treatment difference). Several of the findings of Colton and Anscombe are in excellent agreement; in particular, Colton's results indicate that the sequential allocation rule can attain up to about a 25% smaller loss than fixed-sample allocation. A discussion of some of the drawbacks of Anscombe's model is provided by Armitage (1963).

Since this early work, a number of authors have either generalised the approach described above or considered other distributions for the response variable. The generalisations include *two-stage models* (Colton, 1965; Day, 1969) and *r-stage models* (Cornfield *et al.*, 1969), while Bernoulli and exponential response variables have received attention from, among others, Zelen (1969) and Canner (1970), and Mehta (1981), respectively. There is now considerable literature dealing with decision-theoretic models – the comparative study of Simons (1986) provides a useful summary of recent work.

It is worth mentioning two publications which may be of practical interest. Firstly, Mendoza & Iglewicz (1978) extend Colton's (1963) model to take into account the fact that only a small fraction of the patients available for the trial (i.e. the first stage) are entered – most of the remaining patients are usually allocated concurrently to the standard treatment. It is shown that the inclusion of these latter patients leads to a significant reduction in the trial size. Mendoza and Iglewicz also mention an actual trial described by Schneiderman (1975) for which Colton's model may have been appropriate. Secondly, Chernoff & Petkau (1985) modify Anscombe's model to include an *ethical cost* that is incurred whenever the physician applies a treatment he believes to be inferior. Chernoff & Petkau believe this modification to be a compromise to meet Armitage's criticism of the original model; however, Berry & Fristedt (1985, pp.219–220) are somewhat more sceptical and prefer to discount the interests of the future patients.

1.4 Modifications of practical relevance

In the previous sections, we have assumed that the proposed clinical trial takes no account of any external information. In practice, such an assumption is sometimes unrealistic – for instance, there may be differences in the prognoses of the patients being entered. Furthermore, inferences concerning the true treatment difference upon termination of the trial may be seriously affected by ignoring such information. In this section, we comment briefly on three types of information which can be incorporated.

Firstly, covariates such as age or severity of disease play an important role in the design and analysis of clinical trials, and their incorporation can help to eliminate bias and increase precision in the estimation of treatment difference. The first example of a procedure for incorporating covariates in a data-dependent allocation rule is due to Simon *et al.* (1977); the authors consider a trial with two treatments having a Bernoulli response variable and generalise Bartlett's (1946) procedure to derive a SPRT (details of Bartlett's procedure are given in §4.3). Simulation results show that *"the use of covariate information can result in important reductions in the expected sample size for specified error probabilities, and that the use of covariate information is essential for the elimination of bias when adaptive assignment rules are employed"* (Simon *et al.*, 1977, p.777). More recently, Woodroffe (1979) considered a one-armed bandit problem in which the effect of the new treatment (unknown arm) is supposed to depend on a covariate. Although certain assumptions are made concerning the distribution of the response variable and of the covariate, Woodroffe notes that the results obtained indicate that incorporation of covariate information may simplify the solution to the problem. An extension to this work is provided by Woodroffe (1982) and a similar problem has been investigated by Clayton (1988).

Secondly, it may be beneficial to include the results of a previous trial which has been carried out with the same two treatments on a similar population of patients. Donner (1977) attempts to deal with this problem by extending Colton's (1963) model by assuming a bivariate normal prior distribution for the joint distribution of the true treatment differences for the two populations of patients. Donner hopes that this approach may alleviate *"the problem of duplication of effort and wasted time in the*

design of clinical trials" (p.306).

Lastly, for some types of response variables, the observed response may be not immediate but delayed so that some patients' responses are obtained after a decision to stop observation is made. Anderson (1964) first considered this problem in connection with SPRT's and used a normal response variable as an illustration; Choi & Clark (1970) have subsequently applied Anderson's work to a Bernoulli response variable. Although Zelen (1969, p.146) was the first to mention delayed responses in connection with data-dependent allocation rules, the first quantitative results were given by Hoel *et al.* (1975a) who considered a Bernoulli response variable. Langenberg & Srinivasan (1980) propose two simple non-sequential procedures for dealing with delayed observations in Colton's (1963) model. More recently, Eick (1988) studies the one-armed bandit problem with geometric discounting when some of the responses are delayed and, for a similar setting, Eick (1989) investigates the two-armed bandit problem.

1.5 The use of outcome-dependent allocation in practice

So far we have highlighted ethical problems related to outcome-dependent allocation – clearly, there may also be practical difficulties in implementing these methods in an actual clinical trial. A discussion of these is presented in §1.5.2 but firstly we briefly describe actual trials which have used some form of outcome-dependent allocation.

1.5.1 Examples

During the course of the research for this thesis, I have found reports of only five trials which have used a data-dependent allocation rule. To my knowledge, no use has been made of decision-theoretic models.

Example 1.1. A variant of the play-the-winner rule was used in a lung cancer trial by M. Zelen (see Iglewicz, 1984, p.329). Details of this trial have not been reported.

Example 1.2. Zaborskis (1976) quotes a trial which deals with the treatment of disturbances of cardiac rhythm (*cardiac arrhythmias*). The treatments were three doses of isoptin (5.0 mg., 7.5 mg. and 10.0 mg.) and the allocation rule was a generalised version of Feldman's (1962) myopic rule. Of the 21 patients entered into the trial, 16 were

treated successfully. It was concluded that 10.0 mg. was the most effective dose with 12 of the 14 patients receiving this dose being treated successfully. However, Berry & Fristedt (1985, pp.259–260) feel that this approach is inappropriate and suggest that a dose–response relationship be sought instead.

Example 1.3. A trial for the treatment of respiratory distress in newborn infants, recently carried out at the University of Michigan, was the first to receive significant attention in both statistical and medical journals (e.g. Bartlett *et al.*, 1985; Cornell *et al.*, 1986). The two treatments were extracorporeal membrane oxygenation (ECMO) and conventional medical therapy (CMT), and the response variable was Bernoulli – the two outcomes being death or lung recovery, which are observed within a few days of treatment. The allocation rule used was a randomised version of the play-the-winner rule (e.g. Wei & Durham, 1978). The "success" of this trial (of the 12 infants admitted, 1 died – the only one receiving CMT!) subsequently raised a number of questions concerning the appropriateness of the allocation rule used (Paneth & Wallenstein, 1985; Ware & Epstein, 1985).

Example 1.4. An unpublished trial was recently conducted at the University of Newcastle by D. Appleton for the treatment of peptic ulcers. There were three treatments under consideration – two standard forms of surgery and a new technique known as electrocoagulation – and the response variable was Bernoulli. Initially, it was suggested that Bather's (1985) allocation rule (see §2.1.6) be used in this trial; however, after some discussion, one of the standard treatments was dropped and the allocation rule that was adopted randomised to the two remaining treatments in proportions according to the size of the estimated odds ratio. Of the 60 patients recorded only 15 were entered into the trial, the first 14 of which were randomised! The main reason for the small trial size was the lack of cooperation between the three hospitals involved in the study. However, it is reassuring to know that all the patients entered were treated successfully.

Example 1.5. A follow-up trial to the one in Example 1.3 is described by Ware (1989). The allocation rule adopted initially randomises patients to the two treatments until there is a prespecified number of deaths with one of the treatments (in this case, 4

with CMT), at which point a switch is made to the other treatment (ECMO). The documentation by Ware is particularly valuable since some limitations of the proposed allocation rule for the two treatments under study are highlighted; a lengthy discussion also accompanies this work.

1.5.2 Why are these methods not used more often?

This question has been asked by a number of authors; these include Weinstein (1974), Hoel *et al.* (1975b), Bailar (1976), Simon (1977) and Armitage (1985, 1988). In this section, some of the main issues are discussed.

It is natural to look at the relative merits of randomisation versus outcome-dependent allocation. Statistically, randomisation is the only way to ensure that there is no *systematic bias* between the treatment groups, but to control for variables which affect response there are a number of devices which may be more efficient than randomisation such as matching or blocking. From §1.4, it is clear that little work has been carried out to examine the effect of systematic bias upon the performance of outcome-dependent allocation. In particular, Simon (1977), Friedman *et al.* (1981, pp.51–52) and Armitage (1985) have remarked that few authors have considered the problem of *instability* over time in the population of patients. It is hoped that the results in this thesis will provide some guidance in this direction and amplify the recent remarks by Armitage & Coad (1989).

Weinstein (1974) distinguishes between the approaches required if one is only concerned about future patients and if one takes a myopic view. For the former, the use of more randomised clinical trials may be suitable in order to gain more reliable information on the treatments, while for the latter one would be more subjective. Alternatively, if one is concerned about both the immediate and the future patients, then an intermediate approach is sought – in fact, of the adaptive approaches designed for this situation, those that bring about the best results, in terms of the observed responses, are deterministic.

Most of the literature dealing with outcome-dependent allocation does not model the clinical trial setting realistically. For instance, most of the formulations assume implicitly that there is a *single decision maker* – in many trials, there may a number of

clinicians/statisticians involved in the organisation of a trial, and therefore the allocation rule must be modified to take into account each of their subjective beliefs. On the other hand, the accumulated information on the treatments, gathered either prior to the trial or during its course, should not be ignored when new patients are being treated. However, if the trial is multi-centre, then even the simplest adaptive rule may prove difficult to implement, as is evident from Example 1.4 in §1.5.1.

Simon (1977) states as one of the reasons why data-dependent allocation rules are not often used that of *selection bias* (Blackwell & Hodges, 1957) on the part of the physician. This type of bias is due to the knowledge that the physician may have regarding the treatment assignment of the next patient in the trial. As Armitage (1985, p.21) points out, this may result in the admittance of patients with certain characteristics – and subsequently produce misleading treatment comparisons. This particular problem could have dramatic effects when the allocation rule is deterministic but may be partially obviated if there is a randomisation element in the rule, as in Example 1.3 in §1.5.1. It should be noted that whereas complete randomisation eliminates the potential of selection bias, for some forms of randomisation such as permuted block designs, selection bias can be substantial.

A number of other issues are raised by Simon (1977): these include the problem of delayed responses and the need for a simple adaptive rule. As mentioned in §1.4, some work has been carried out for dealing with delayed responses (see also the results of Robinson, 1983) – however, more work is needed on the robustness properties of the adaptive methods. Furthermore, we seek an allocation rule which is reasonably easy to understand by non-statisticians and which is convincing in terms of the validity of the results. If this is not the case, the success or otherwise of a particular trial may have significantly less impact upon medical practice. Moreover, the actual allocation procedure must be straightforward to conduct – but as Bailar (1976, p.203) points out, "*Complicated assignment algorithms may, however, be acceptable if they are reduced to simple rules or, perhaps an easily used computer program*". It is worth noting that certain designs whose objective is to produce a balanced trial with respect to prognostic covariates (e.g. Pocock & Simon, 1975) have been used in practice, although these designs are sometimes more

complicated than corresponding adaptive ones.

Two final points which have been raised (e.g. Armitage, 1985, 1988) are specification of the patient horizon, N , in decision-theoretic models and the effect of unequal allocation proportions upon the efficiency of estimation of treatment effects. The effect of a poor choice of the patient horizon has been investigated by several authors (e.g. Colton, 1973; Upton & Lee, 1976) and the general recommendation is that it is better to *overestimate* N . Anscombe (1963, p.374) suggests that a reasonable choice is one-tenth of the first value thought of for N . However, recent authors (e.g. Simons, 1987) assume a geometric distribution for N . Berry (1978, pp.344–345) remarks that this assumption is similar to the situation where the patients are discounted geometrically. From a practical point of view, a patient horizon is both conceptually difficult and unrealistic (Armitage, 1985, p.20) – especially if one realises that most patients are not in *any* clinical trial.

Clearly, upon termination of a trial we would like reliable information about the comparative effects of the treatments and, in particular, how the precision of these estimates compares with, say, equal allocation. With the exception of Hayre & Turnbull (1981a, 1981b) and Hayre (1983), who consider fixed-width interval estimation and point estimation with a particular loss structure, little work has been carried out concerning estimation for outcome-dependent allocation. We would not expect these methods to perform as well as randomisation procedures due to the unequal sized treatment groups – one of the principal aims of this thesis has been to investigate this problem more thoroughly.

1.6 Direction of thesis

There are two underlying themes in this thesis: (i) instability in the response variable, and (ii) estimation of treatment effects. Although in each chapter we initially focus our attention on a stable response variable, two forms of instability are subsequently modelled and their effects examined.

The first form of instability studied is a *deterministic linear time-trend*, where the trend is the same for both treatments. Secondly, a model for *fluctuations* in the data (with or without a trend) is formulated. The second form of instability is intended to

model the effect of "*clusters of severely affected patients interspersed amongst clusters of patients with better prognosis*" (Armitage, 1975, p.156). Although this model remains essentially the same throughout the thesis, a slightly different approach is taken in each chapter. It will be seen later that these two forms of instability can be used to model a wide range of situations which may arise in practice – as an example, an interesting discussion on the presence of time-trends in clinical trials is provided by Altman & Royston (1988).

In this thesis, three clinical trial settings are examined. Firstly, in Chapter 2 we suppose that the trial size is *fixed* and that the response variable is Bernoulli; for this setting, we study using simulation the efficiency of estimation for a number of data-dependent allocation rules. This work is essentially an extension of the work by Robinson (1983) and Bather (1985). Secondly, in Chapter 3 we derive analytically a number of new results for estimation for the decision-theoretic models of Colton (1963, 1965). Lastly, in Chapter 4 the problem of sequential testing is investigated; we show how the approach of Robbins & Siegmund (1974) can be modified to take into account instability in the response variable and we provide further simulation results for estimation and compare these with the asymptotic results of Woodroffe (1989). A summary of the main conclusions of the thesis together with further comments is given in Chapter 5.

CHAPTER 2. A COMPARATIVE STUDY OF SOME DATA-DEPENDENT ALLOCATION RULES FOR BERNOULLI DATA

We discuss and extend the recent work by Robinson (1983) and Bather (1985) on data-dependent allocation rules designed for Bernoulli responses. It is hoped that §§2.1 and 2.2, with §1.3.1, will provide a balanced summary of the principal work regarding these rules. Corresponding allocation rules for other responses will be examined in Chapter 4.

In §2.1, a selection of allocation rules are described; these include the play-the-winner rule and a rule based on Gittins indices. Simulation results from recent comparisons are summarised in §2.2.

We rely heavily upon simulation in this chapter. A computer program written for this purpose is described in §2.3 together with a method for estimating the true mean treatment difference. Following this, simulation results for a stable response variable are presented in §2.4 and, for two allocation rules, there is a comparison with analytical results.

In §2.5, two forms of instability in the response variable are modelled. The corresponding simulation results are presented in §§2.6 and 2.7; §2.6 also includes a comparison with analytical results for two allocation rules. Lastly, in §2.8, there is a short discussion of possible directions for future work.

2.1 Allocation rules considered

A clinical trial is conducted in which there are n patients. Each patient can be allocated to one of two treatments A and B using a data-dependent allocation rule. The main objective is to reduce the number of patients who receive the inferior treatment.

We initially assume that the response variable for treatment i at time j , X_{ij} ($i=A,B$; $j \in [0, n-1]$), has a Bernoulli distribution with parameter p_i – the unknown probability of success for treatment i . Further, each patient's response is independent of the responses from previous patients in the trial and depends only on the treatment allocated.

2.1.1 Equal allocation rule

This rule corresponds to randomisation in which patients are allocated equally to the two treatments. In this chapter, we have adopted the *random permuted blocks method* (e.g. Pocock, 1983, pp.76–79): the trial is split into a specified number of *blocks*, k , say, which have equal numbers of patients, and for each block we have a random ordering of the treatment allocations.

2.1.2 Play-the-winner rule

This rule was originally introduced by Robbins (1952) and has subsequently been investigated by a number of authors, including Zelen (1969) (though he uses the term *modified play-the-winner rule*). The rule is to randomise the first patient, and then to proceed in the following manner: if a success is observed with treatment i , allocate the next patient to the same treatment; if a failure is observed, allocate the next patient to the other treatment.

Several variants of the original play-the-winner rule have been investigated in the literature; these include randomised versions of the rule (e.g. Wei & Durham, 1978). The randomised rule might be more attractive to clinicians (see §1.5.2) due to the incorporation of randomisation and due to the simplicity of the rule's operation, even when the response is delayed. For these reasons, it is perhaps not surprising that this rule has been adopted in one of the few clinical trials conducted using a data-dependent allocation rule (cf. §1.5.1).

2.1.3 Fox's rule

Definition 2.1 (Robbins, 1952). An allocation rule is *asymptotically optimal* if, with probability one,

$$\lim_{n \rightarrow \infty} r/n = \max(p_A, p_B) ,$$

where r is the total number of successes at the end of the trial.

It is known (Robbins, 1952) that the play-the-winner rule in §2.1.2 does not possess the above property but a simple rule (suggested by Fox, 1974, and a special case of a more general rule due to Robbins, 1952) has been shown to be asymptotically optimal. The rule is to allocate treatment A at trials $2^i - 1$ ($i=1,2,\dots$), and treatment B at trials

2^i ($i=1,2,\dots$); at other trials, allocate the next patient to the treatment currently having the higher average number of successes, ties being decided randomly.

2.1.4 One-step-ahead rule

The strictly *optimal allocation rule* which at each trial maximises the total expected response from the future patients based on the previous responses can, in theory, be obtained by dynamic programming. Although this rule is four-dimensional and a large computer memory is required, a program using the order of $n^3/6$ storage locations is possible (Berry, 1972, p.872). Furthermore, it has been shown that if the responses are discounted in a certain way, and if n is not too small, the optimal allocation rule can be characterised by special indices (Gittins & Jones, 1974). These indices will be discussed in §2.1.6.

A considerable reduction in computer memory can however be achieved if one considers myopic allocation rules.

Definition 2.2. An allocation rule is *myopic* if each patient is allocated to the treatment with the greater expected immediate response.

One such rule, known as one-step-ahead, has been considered by Jones (1975): the p_i ($i=A,B$) are assumed to have independent beta distributions, and each patient is allocated to the treatment which currently has the higher posterior probability of success, ties being decided randomly. As a comparison, Jones (1975) shows numerically that for $n \leq 15$ this rule is close to the optimal rule. More recently, Jones & Kandeel (1983) presented a modification of the rule which for $n \leq 50$ is closer to the optimum; these authors also note that a further modification can make the one-step-ahead rule asymptotically optimal.

2.1.5 Bather's rule

A class of allocation rules, based on *randomised allocation indices* (RAI's), was suggested by Bather (1980); these indices are perturbed values of the current proportions of successes for each treatment. This class of rules was shown by Bather to be asymptotically optimal. A rule from this class, henceforth referred to as Bather's rule, was further investigated in two further papers (Bather, 1981, 1985). This rule is as

follows: allocate treatment i ($i=A,B$) to the $(j+1)^{\text{th}}$ patient if and only if

$$Q_i(j) = \max \{Q_A(j), Q_B(j)\} ,$$

where

$$Q_k(j) = r_k/n_k + \lambda(n_k)X_k(j)$$

in which r_k and n_k are the numbers of successes and patients, respectively, on treatment k ($k=A,B$) after $j = n_A + n_B$ patients; further, Bather (1981) defines

$$\lambda(\ell) = \frac{4 + \sqrt{\ell}}{15\ell} , \quad (\ell=1,2,\dots)$$

and $X_k(j) = 2 + Y_k(j)$, where $Y_k(j)$ are independent and identically distributed as exponential with mean unity; ties are decided randomly.

2.1.6 Rules based on Gittins indices

Definition 2.3. (a) A *discount sequence* is a weighting vector $(a_1, a_2, \dots)^T$, $a_j \geq 0 \forall j$, of the responses for the different patients (e.g. for the uniform case, $a_j = 1 \forall j$, and for the geometric case, $a_j = a^{j-1}$ ($j=1,2,\dots$), where a is called the *discount factor*); (b) the above discount sequence is *regular* if

$$\frac{a_{j+2}}{a_{j+1}} \leq \frac{a_{j+1}}{a_j} , \quad (j=1,2,\dots).$$

Suppose we wish to determine an optimal allocation rule as defined in §2.1.4. Then the *total discounted response* for treatment i is defined as

$$\sum_{j=1}^{n_i} a_j x_{ij} ,$$

where the a_j are as above and x_{ij} is the observed response on treatment i at time j . The optimal allocation rule maximises the total *expected* discounted response (Berry & Fristedt, 1985, p.1).

Gittins & Jones (1974) proved that for a geometric discount sequence an optimal allocation rule is obtained if one allocates each patient to the treatment with the higher *dynamic allocation index* (named the Gittins index by Whittle, 1980), and randomise in the case of ties. In this chapter this allocation rule will be subsequently referred to as

the Gittins rule.

Definition 2.4. The *Gittins index* for a treatment i with response variable X_i is

$$v_i = \sup_{\tau > 0} \frac{E \left\{ \sum_{j=1}^{\tau} a_j X_{i,j} \right\}}{E \left\{ \sum_{j=1}^{\tau} a_j \right\}}, \quad (2.1)$$

where the discount sequence is regular and τ is the (infinite) set of stopping times.

From (2.1), the Gittins index for a treatment may be regarded as a measure of expected discounted response per unit of discounted time. Each index depends only on the treatment in question, and on the discount sequence.

The original Gittins–Jones result depends on the discount sequence being geometric. Recent work by Berry & Fristedt (1985, pp.145–149) has shown that a necessary and sufficient condition for an allocation rule specified by Gittins indices to be optimal is that the discount sequence be geometric. In view of this result, we note a comment made by Wetherill & Glazebrook (1986, p.223) when discussing the use of geometric discounting: *'The use of a discounted cost structure is mainly for mathematical tractability – though it does usefully correspond to an ethical position of the kind espoused by Anscombe (1963), who comments: "a doctor's responsibilities to his future patients and to all patients everywhere should no doubt be discounted in weight in comparison with his immediate responsibility, but not denied"'*. In other words, a discount sequence can be used to reflect the relative importance of individual versus collective ethics (see §1.2).

Concerning the computation of the Gittins index, Gittins & Jones (1979, p.562) point out that *'Although not required by the general theory of dynamic allocation indices, the assumption of a beta distribution [for the prior distributions for the unknown success probabilities] is essential for their calculation to be tractable. It allows an arbitrary specification of the mean and variance and, as pointed out by Raiffa and Schlaifer (1961) [pp.43–44], this is for many purposes quite adequate'*. The priors assumed here and for the one-step-ahead rule in §2.1.4 are uniform on (0,1). Gittins & Jones (1979) also give sample values for the indices when the discount sequence is geometric with

discount factor $a = 0.75$ and, for a slightly different problem, Glazebrook (1978) considers a geometric discount sequence with $a = 0.5$.

Several algorithms for evaluating the Gittins index for a Bernoulli response have been described in the literature: these include those by Gittins (1979) and, more recently, Robinson (1982) and Katehakis & Derman (1986). Since exact computation of these indices requires considerable computer time, particularly as $a \rightarrow 1$, Gittins (1983, 1989) has shown that the indices are well approximated by an equation of the form

$$\nu_a(r, n) = rn^{-1} + [A_a(rn^{-1}) + nB_a(rn^{-1}) + n^{-1}C_a(rn^{-1})]^{-1}, \quad (2.2)$$

where r is the number of successes and A_a , B_a and C_a are obtained by weighted least squares [N.B. For simplicity, A_a , B_a and C_a correspond to an obvious standardisation of A , B and C in equation (7.16) of Gittins (1989).] A simpler version of (2.2) is given by Gittins & Jones (1979).

Using (2.2), Gittins (1989) gives approximate values for the indices for $a = 0.5, 0.6, \dots, 0.9, 0.95$ and 0.99 , and $rn^{-1} = 0.025, 0.05, 0.1, 0.2, \dots, 0.9, 0.95$ and 0.975 . For other values of rn^{-1} , interpolation based on the assumption that $\nu_a(r, n) - rn^{-1}$ is linear in $[rn^{-1}(1-rn^{-1})]^{\frac{1}{2}}$ has been shown to work well: the discrepancy between the approximate and exact values of the indices is generally at most one in the fourth significant figure, the main exceptions being for values of rn^{-1} outside the range $(0.025, 0.975)$.

Other important contributions to the theory of Gittins indices for a Bernoulli response have been made by: Glazebrook (1980) who considers the so-called *randomised dynamic allocation indices* (RDAI's) – these are perturbed values of those defined in (2.1), the perturbations being similar to those used in Bather's RAI's in §2.1.5 (these RDAI's are investigated in §§4.5 and 4.6 in a sequential testing setting) – and shows that rules based on RDAI's are asymptotically optimal and also (almost) Bayes with respect to the discounted response criterion; Whittle (1980, 1981a) who gives an alternative proof of the Gittins–Jones result and extends it; Kelly (1981) who shows that as $a \rightarrow 1$, the Gittins rule tends to the *least failures rule* (a variation on the play-the-winner rule) which allocates the next patient to the treatment which has incurred the least number of failures, or if this does not define a treatment uniquely, allocate (from among these treatments) the

treatment with the largest number of successes; and lastly, Whittle (1988) who investigates the Gittins index when the characteristics of the treatments are changing, whether or not they are in use (it should be noted that this work is in the context of industrial production rather than clinical trials).

This completes the description of the allocation rules which will be investigated in §§2.4, 2.6 and 2.7. We now give a brief summary of the results of some recent comparisons of these and other rules.

2.2 Review of recent comparisons

2.2.1 Definitions

Many of the recent comparisons of data-dependent allocation rules for a Bernoulli response variable have been based on two criteria: error probability and expected successes lost.

Definition 2.5. (a) The *error probability* (EP) for an allocation rule is the probability that the inferior treatment has the higher proportion of successes at the end of the experiment. (b) The *expected successes lost* (ESL) is the difference between the expected number of successes (not discounted) achieved in n trials using only the better treatment, that is, $n \max(p_A, p_B)$, and the expected number of successes achieved in n trials using the allocation rule.

In §2.3.2, additional criteria will be defined on which the comparisons in §§2.4, 2.6 and 2.7 will be based.

2.2.2 Summary of results

One of the most recent comparisons is due to Robinson (1983). This work gives simulation results for five allocation rules – those in §2.1 with the exception of the play-the-winner rule. Robinson considers the setting described in §2.1 and tabulates results for expected successes lost and error probability for a range of pairs of values for (p_A, p_B) ; limited results are also given when there are delayed responses (see §1.4).

From his results, Robinson makes a number of remarks; these include the following:

(i) for the Gittins rule, as a is increased, the expected successes lost usually increases and

the error probability decreases; (ii) for $a = 0.9999$, the Gittins rule and Bather's rule are similar in their performance; (iii) in general, the one-step-ahead rule and Fox's rule have larger error probabilities than Bather's rule; (iv) the Gittins rule with $a = 0.9999$ has considerably less expected successes lost than the equal allocation rule, but has slightly larger error probabilities.

The above remarks by Robinson were not supported by any explanation or justification as to what results might be expected. The following elucidation, intended as a basis for interpretation of the results in §2.4, is informative.

The behaviour in (i) might be expected due to the values of a considered – only values very close to one (0.99, 0.995 and 0.9999) are used. It is known (Kelly, 1981) that as $a \rightarrow 1$, the Gittins rule tends to the least failures rule. Since this rule is a variation on the play-the-winner rule (Kelly, 1981, p.297), we might intuitively expect the expected successes lost to increase and the error probability to decrease as $a \rightarrow 1$ due to the rule becoming less adaptive.

Berry & Fristedt (1985, pp.249–250) express concern over the values chosen for the discount factor; this is echoed by J.C. Gittins (in a private communication, 1987). For the values of n considered here (50 and 100), use of $a = 0.9999$ for example achieves negligible discounting; a better choice of a might be obtained from $a^n \simeq \frac{1}{2}$ or $a^{n-j} \simeq \frac{1}{2}$, where j is the number of patients so far. For the second choice, the value of a at any time depends on the number of patients remaining in the trial; however, this raises problems akin to the choice of a patient horizon as discussed in §1.5.2.

Result (iii) might be expected since Bather's rule is randomised whereas the one-step-ahead rule and Fox's rule are deterministic except when there are ties. Lastly, concerning (iv): because the Gittins rule with $a = 0.9999$ has properties similar to the play-the-winner rule, we might expect the expected successes lost for this rule to be substantially smaller than for equal allocation when $p_A + p_B > 1$ (see Zelen, 1969, pp.140–141).

Other recent comparisons are provided by Bather (1981, 1985). The former includes simulation results for expected successes lost for the play-the-winner rule and Bather's rule. Although Bather finds that the play-the-winner rule does not perform well, his rule

performs well over a wide range of values for (p_A, p_B) . Several other allocation rules are included in the comparison: these include special cases of the class of allocation rules described in §2.1.5 (i.e. different expressions are used for $\lambda(\ell)$) and a myopic rule due to D. Feldman, which was discussed briefly in §1.3.1.

In addition, Bather (1981) compares his rule with one based on a sequential probability ratio test (SPRT) proposed by Vogel (1960a) (see §1.3.1) using the error probability. It is again shown that Bather's rule performs well: there is clear preference for this rule when $n = 1000$ and the error probabilities for the two rules are comparable (the latter result is encouraging since the error probabilities for SPRT's are *minimal*). Whereas the first observation may reflect the asymptotic optimality property of Bather's rule, the latter indicates that incorporation of the randomised element can strike a balance between the loss of information and the increased number of patients on the better treatment.

Lastly, Bather (1985) carries out a comparison based on expected successes lost and error probabilities between a fixed-sample allocation rule and several sequential allocation rules, one of which is Bather's rule. The other rules are:

- (a) an example of the *repeated significance tests* examined by Armitage (1975);
- (b) a symmetric version of the *triangular tests* detailed by Whitehead (1983, §4.3);
- (c) a *truncated* version of an SPRT considered by Bather (1981).

The author ends by noting a number of reservations which one should bear in mind before drawing any definite conclusions as to the advantages of the various rules considered.

2.3 Details of the simulation study

2.3.1 Computation

A FORTRAN program, ALLOCATE, was written to simulate the six allocation rules in §2.1 for both the stable setting and for more general settings which will be described in §2.5. For each set of parameters considered, 4000 simulations were carried out; this number was chosen in order that comparisons could readily be made with the results of Robinson (1983).

Before any data are generated, the NAG routine G05CBF is called to set the generator to a repeatable state. Following this, three NAG routines are called: (i) G05DDF – this simulates a normal random variate (for use when the response variable is unstable); (ii) G05CAF – this simulates a continuous uniform random variate on (0,1) (used in deciding whether a success or failure is observed with a treatment); and (iii) G05DBF – this simulates an exponential random variate (for use with Bather's rule).

Values for the Gittins index for $n \leq 50$ were calculated using a FORTRAN program, DAI (written by D.M. Jones); for larger values of n , a FORTRAN program, APPROXDAI, was written to implement the approximation given by (2.2). Values for the coefficients A_a , B_a and C_a are tabulated by Gittins (1989, pp.241–242).

2.3.2 Criteria for comparison of rules

We have used five criteria to compare the six allocation rules:

- (a) error probability;
- (b) expected successes lost;
- (c) histograms of the number of patients on the better treatment;
- (d) plots of allocation rates to the better treatment;
- (e) efficiency of estimation of the true mean treatment difference.

Criteria (a) and (b) were defined in §2.2.1. The histograms in (c) are constructed using the 4000 simulations of the trial and hence some impression of the distribution of the number of patients on the better treatment is gained. These distributions provide further insight into the behaviour of the allocation rules – for example, although, on the average, an allocation rule may assign more patients to the better treatment, it would be interesting to have some idea of the nature of the variability about the "average". These histograms were obtained using the procedure GCHART in SAS/GRAPH.

Further insight can also be obtained from (d), the rate at which the proportion of patients on the better treatment increases over time; plots are given for the average number of patients (over the simulations) on the better treatment as the trial proceeds. Clearly, the more adaptive the allocation rule, the larger the gradient of the plot. These plots were obtained using the procedure GPLOT in SAS/GRAPH.

Lastly, using (e) we attempt to assess the magnitude of the loss in efficiency of

estimation due to the adaptive nature of the allocation rules. Suitable measures of the bias and the variance of the estimated true mean treatment difference are proposed in §2.3.3.

2.3.3 Estimation of treatment difference

It was suggested by Armitage (1985, p.21) that if there is instability in the response variable, then it may be possible to alleviate some of its possible effects by splitting the trial into a number of equal-length *periods* (or blocks), and analysing the estimated treatment difference *within* these periods. We now describe an empirical and a theoretical estimate of the variance of a suitable weighted estimate.

Denoting the number of simulations by NSIM and the number of periods by PER and using Armitage & Berry (1987, p.382), an appropriate weighted estimate of $\delta = p_B - p_A$ ($p_B > p_A$) for each simulation k is

$$\bar{d}_k' = \frac{\sum_{j=1}^{PER} w_j \bar{d}_j}{\sum_{j=1}^{PER} w_j} \quad (2.3)$$

with

$$w_j = \frac{n_{Aj} n_{Bj}}{n_{Aj} + n_{Bj}}, \quad \bar{d}_j = \hat{p}_{Bj} - \hat{p}_{Aj}, \quad \hat{p}_{ij} = \frac{r_{ij}}{n_{ij}},$$

where r_{ij} and n_{ij} are, respectively, the number of successes and the number of trials on treatment i ($i=A,B$) in period j ($j \in [1, PER]$).

If we denote the estimated bias by BIAS, then

$$BIAS = \bar{\bar{d}} - \delta,$$

where using obvious notation,

$$\bar{\bar{d}} = \sum_{k=1}^{NSIM} \bar{d}_k' / NSIM.$$

An *empirical* estimate of the variance of $\bar{d}_k'$ is given by

$$VARD = \sum_{k=1}^{NSIM} (\bar{d}_k' - \bar{\bar{d}})^2 / (NSIM - 1) \quad (2.4)$$

and a corresponding *theoretical* estimate is

$$\text{MVAR} = \frac{\sum_{k=1}^{\text{NSIM}} \text{Var}(\bar{d}_k')}{\text{NSIM}}, \quad (2.5)$$

where, for example, assuming that the n_{ij} are *fixed* (except for equal allocation this is, of course, an oversimplification but it permits further interpretation which would otherwise be difficult),

$$\text{Var}(\bar{d}') = \frac{\sum_{j=1}^{\text{PER}} w_j^2 \left\{ \frac{\hat{p}_{Aj}(1-\hat{p}_{Aj})}{n_{Aj}} + \frac{\hat{p}_{Bj}(1-\hat{p}_{Bj})}{n_{Bj}} \right\}}{\left\{ \sum_{j=1}^{\text{PER}} w_j \right\}^2}.$$

Thus, from (2.4) and (2.5),

$$E(\text{VARD}) = \frac{\sum_{j=1}^{\text{PER}} w_j^2 \left\{ \frac{p_A(1-p_A)}{n_{Aj}} + \frac{p_B(1-p_B)}{n_{Bj}} \right\}}{\left\{ \sum_{j=1}^{\text{PER}} w_j \right\}^2} \quad (2.6)$$

and

$$E(\text{MVAR}) = \frac{\sum_{j=1}^{\text{PER}} w_j^2 \left\{ \frac{E[\hat{p}_{Aj}(1-\hat{p}_{Aj})]}{n_{Aj}} + \frac{E[\hat{p}_{Bj}(1-\hat{p}_{Bj})]}{n_{Bj}} \right\}}{\left\{ \sum_{j=1}^{\text{PER}} w_j \right\}^2}. \quad (2.7)$$

Furthermore, since

$$E[\hat{p}_{ij}(1-\hat{p}_{ij})] = p_i(1-p_i)[1 - 1/n_{ij}],$$

it follows from (2.6) and (2.7) that

$$E(\text{MVAR}) < E(\text{VARD}).$$

Although this inequality has been deduced under the assumption that the n_{ij} are fixed, it provides a rough guide as to the behaviour of MVAR and VARD. For instance, suppose $p_B > p_A$ and that the difference between each term of the sum in the numerator of $E(\text{VARD}) - E(\text{MVAR})$ is denoted by c_j ($j \in [1, \text{PER}]$). Firstly, if $p_i > \frac{1}{2}$ ($i=A, B$), then since n_{Aj} is generally much smaller than n_{Bj} , the contribution to the c_j from the term in p_A will be greater than that for the term in p_B . If on the other hand $p_i < \frac{1}{2}$ ($i=A, B$), then the converse is true. Consequently, c_j will generally be larger for $p_i > \frac{1}{2}$ than for $p_i < \frac{1}{2}$ ($i=A, B$). Thus, we might expect MVAR and VARD to differ more for $p_i > \frac{1}{2}$

than for $p_i < \frac{1}{2}$ ($i=A,B$). We shall see in §§2.4.2, 2.6.2 and 2.7 that the corresponding simulation results support this argument. However, a more detailed study of the simulation values for VARD - MVAR will be presented in §2.4.3.

2.4 Stable response variable

We begin by considering the setting described in §2.1, where the p_i are stable over time. Some analytical results are given in §2.4.1 for the equal allocation rule and the play-the-winner rule; for equal allocation, exact results are available while for the play-the-winner rule we rely on asymptotic theory. [Of the more adaptive allocation rules, limited analytical results are available only for Bather's rule: when one of the p_i is known, values for the exact expected successes lost are known (Bather, 1980, pp.180–181), while for the more general case, Bather (1983) has derived bounds on the minimax value of the expected successes lost, based on earlier work of Vogel (1960b) and Fabius & van Zwet (1970). The bounds obtained by Bather are compared with simulated values by Robinson (1983).] Following a comparison with the corresponding simulated values, further simulation results are presented in §2.4.2 for the criteria discussed in §2.3.2.

2.4.1. Analytical results

(a) Equal allocation rule

If we assume $p_B > p_A$, then by Definition 2.4,

$$ESL = \frac{1}{2}n(p_B - p_A) \quad (2.8)$$

and

$$EP = \Pr\{r_A > r_B\} \quad (2.9)$$

where r_i is the number of successes on treatment i ($i=A,B$) upon termination of the trial.

Since $r_i \sim \text{Bin}(\frac{1}{2}n, p_i)$, the distribution of r_i is easily computed, and hence exact values for (2.9) can be obtained. Furthermore, we can use a normal approximation to give

$$EP \approx 1 - \Phi\left\{ \frac{\sqrt{n} (p_B - p_A)}{\sqrt{2\{p_A(1-p_A) + p_B(1-p_B)\}}} \right\} \quad (2.10)$$

where Φ denotes the standard normal distribution function. A numerical comparison of the normal approximation and exact values for the error probability showed that the

discrepancy is at most two in the third decimal place.

(b) Play-the-winner rule

The play-the-winner rule can be regarded as a two-state Markov chain, with one state being success with a treatment and the other failure (Zelen, 1969, p.139). The exact distribution of n_A is rather cumbersome. However, the approximation based on the asymptotic result

$$n_A \sim N \left[\frac{n(1-p_B)}{(2-p_A-p_B)}, \frac{n(1-p_A)(1-p_B)(p_A+p_B)}{(2-p_A-p_B)^3} \right] \quad (2.11)$$

works well for values of $p_A + p_B$ not too far from one and for n large (Gabriel, 1959). A similar result holds for n_B .

Thus, by Definition 2.4, we can write

$$ESL \approx \frac{n(1-p_B)(p_B-p_A)}{(2-p_A-p_B)} \quad (2.12)$$

and

$$EP \approx \Pr\{n_A > n_B\}$$

if we assume that the numbers of failures upon termination on the two treatments are the same. Hence, using (2.11), we obtain the normal approximation

$$EP \approx 1 - \Phi \left\{ \frac{(p_B-p_A) \sqrt{n(2-p_A-p_B)}}{2 \sqrt{(1-p_A)(1-p_B)(p_A+p_B)}} \right\} \quad (2.13)$$

(cf. equation (22) of Zelen, 1969, p.140).

If we now compare the results (2.12) and (2.13) with those obtained by simulation, we see that from Table 2.1 that the discrepancy between the approximate and simulated values for expected successes lost is at most two in the first decimal place. The corresponding results for the error probability are generally in good agreement though there are some large discrepancies when $p_i > \frac{1}{2}$ ($i=A,B$). It appears therefore that even for modest values of n , the normal approximation (2.11) works well.

Table 2.1: *Simulated and approximate values for expected successes lost and error probability for the play-the-winner rule*

n	p_A	p_B	ESL		EP	
			Approx.	Simul.	Approx.	Simul.
50	0.3	0.4	2.3	2.4	0.229	0.233
	0.1	0.3	4.4	4.5	0.037	0.033
	0.7	0.9	2.5	2.6	0.021	0.037
	0.1	0.4	6.0	6.1	0.006	0.006
	0.35	0.65	5.3	5.3	0.013	0.012
	0.6	0.9	3.0	3.1	0.001	0.005
	0.3	0.7	6.0	6.0	0.001	0.002
	0.5	0.55	1.2	1.2	0.362	0.375
100	0.3	0.4	4.6	4.7	0.147	0.146
	0.1	0.3	8.8	8.8	0.006	0.006
	0.7	0.9	5.0	5.0	0.002	0.005
	0.1	0.4	12.0	12.0	0.000	0.001
	0.35	0.65	10.6	10.5	0.001	0.001
	0.6	0.9	6.0	6.1	0.000	0.000
	0.3	0.7	12.0	12.2	0.000	0.000
	0.5	0.55	2.4	2.3	0.308	0.300

2.4.2 Simulation results

We now present results from the simulation study described in §2.3. Three trial sizes $n = 50, 100$ and 150 are considered and results are given in Tables 2.2, 2.3 and 2.4. For $n = 50$ and 100 , results for the Gittins rule are tabulated for $a = 0.9, 0.95$ and 0.99 while for $n = 150$, results are only given for $a = 0.9$ and 0.99 [when $n = 150$, the values of A_a , B_a and C_a used in (2.2) for $a = 0.95$ are not reliable (Gittins, 1989, p.242)]. The number of periods, PER, and the number of random permuted blocks are both five.

A selection of histograms of the number of patients on the better treatment and plots of allocation rates to the better treatment for $n = 50$ are displayed, respectively, in Figures 2.1 – 2.3 and Figure 2.4; further histograms and plots for $n = 150$ are included in Appendix 1.

Several of the conclusions from Tables 2.2 – 2.4 are similar to those of Robinson (1983), and are summarised in §2.2.2. We restrict our attention here to any new features of the allocation rules highlighted by our tables. Firstly, the play-the-winner rule has error probabilities similar to those for equal allocation and, in some cases, they are

Table 2.2: *Simulation results for six allocation rules when $n = 50$ (stable response variable)*

Rule	p_A p_B	0.3 0.4	0.1 0.3	0.7 0.9	0.1 0.4	0.35 0.65	0.6 0.9	0.3 0.7	0.5 0.55
Equal	ESL	2.6	5.1	4.9	7.4	7.5	7.5	10.0	1.3
	EP	0.233	0.033	0.043	0.007	0.017	0.004	0.001	0.369
	BIAS	-0.001	0.002	-0.002	0.001	-0.001	0.002	-0.001	-0.002
	VARD	0.018	0.012	0.012	0.013	0.019	0.013	0.017	0.020
	MVAR	0.014	0.010	0.010	0.011	0.015	0.011	0.013	0.016
Bather	ESL	1.8	2.3	2.1	2.4	2.6	2.2	2.5	1.1
	EP	0.237	0.048	0.037	0.007	0.021	0.007	0.004	0.369
	BIAS	0.021	-0.003	0.073	-0.005	0.046	0.096	0.039	0.021
	VARD	0.042	0.019	0.050	0.024	0.053	0.061	0.050	0.067
	MVAR	0.018	0.011	0.022	0.015	0.026	0.025	0.027	0.024
Gittins a=0.9	ESL	1.7	1.8	2.2	1.6	2.1	1.8	1.5	1.0
	EP	0.267	0.051	0.177	0.014	0.066	0.068	0.015	0.397
	BIAS	0.027	-0.005	0.005	-0.008	0.093	-0.006	0.064	0.034
	VARD	0.072	0.023	0.193	0.031	0.108	0.182	0.091	0.194
	MVAR	0.023	0.014	0.030	0.021	0.039	0.029	0.038	0.036
Gittins a=0.95	ESL	1.7	1.9	1.8	1.8	1.9	1.5	1.4	1.0
	EP	0.260	0.044	0.125	0.011	0.043	0.050	0.013	0.393
	BIAS	0.026	-0.010	0.029	-0.010	0.107	0.009	0.092	0.025
	VARD	0.065	0.020	0.188	0.028	0.098	0.172	0.084	0.162
	MVAR	0.021	0.013	0.034	0.019	0.038	0.033	0.038	0.034
Gittins a=0.99	ESL	1.8	2.1	1.6	2.0	2.0	1.2	1.6	1.0
	EP	0.239	0.047	0.083	0.009	0.028	0.024	0.006	0.375
	BIAS	0.011	-0.021	0.058	-0.027	0.073	0.062	0.076	0.025
	VARD	0.044	0.016	0.168	0.022	0.075	0.150	0.071	0.103
	MVAR	0.019	0.011	0.041	0.017	0.036	0.040	0.037	0.030
Fox	ESL	1.9	2.3	1.9	2.5	2.7	2.2	2.8	1.1
	EP	0.277	0.106	0.066	0.042	0.055	0.020	0.019	0.394
	BIAS	0.011	-0.000	0.031	-0.004	0.015	0.023	0.014	0.003
	VARD	0.052	0.028	0.039	0.031	0.051	0.043	0.048	0.064
	MVAR	0.020	0.013	0.022	0.015	0.024	0.022	0.023	0.026
One- step- ahead	ESL	1.8	1.8	3.2	1.6	2.5	3.2	1.9	1.1
	EP	0.295	0.071	0.306	0.025	0.121	0.189	0.051	0.421
	BIAS	0.045	0.012	-0.051	0.021	0.015	-0.062	-0.027	0.007
	VARD	0.108	0.030	0.165	0.039	0.148	0.186	0.129	0.209
	MVAR	0.029	0.019	0.016	0.026	0.032	0.019	0.030	0.036
Play- the- winner	ESL	2.4	4.5	2.6	6.1	5.3	3.1	6.0	1.2
	EP	0.233	0.033	0.037	0.006	0.012	0.005	0.002	0.375
	BIAS	-0.009	-0.018	0.005	-0.028	-0.018	0.011	-0.026	-0.007
	VARD	0.016	0.010	0.037	0.011	0.020	0.041	0.021	0.021
	MVAR	0.014	0.008	0.024	0.010	0.016	0.026	0.016	0.017

Table 2.3: *Simulation results for six allocation rules when $n = 100$ (stable response variable)*

Rule	p_A p_B	0.3 0.4	0.1 0.3	0.7 0.9	0.1 0.4	0.35 0.65	0.6 0.9	0.3 0.7	0.5 0.55
Equal	ESL	5.1	10.1	9.9	15.0	14.9	14.9	20.1	2.5
	EP	0.146	0.006	0.005	0.000	0.001	0.000	0.000	0.319
	BIAS	0.000	0.000	-0.001	0.002	-0.001	0.000	0.003	-0.002
	VARD	0.009	0.006	0.006	0.007	0.009	0.007	0.008	0.010
	MVAR	0.008	0.005	0.005	0.006	0.008	0.006	0.008	0.009
Bather	ESL	3.0	2.9	2.8	2.9	3.1	2.7	3.0	2.0
	EP	0.176	0.008	0.009	0.001	0.006	0.001	0.001	0.337
	BIAS	0.022	0.006	0.075	0.004	0.048	0.084	0.051	0.019
	VARD	0.029	0.012	0.039	0.015	0.038	0.046	0.036	0.047
	MVAR	0.013	0.008	0.019	0.011	0.021	0.022	0.021	0.017
Gittins a=0.9	ESL	3.0	2.1	3.8	1.8	2.5	2.9	1.9	2.0
	EP	0.230	0.020	0.155	0.004	0.038	0.065	0.012	0.372
	BIAS	0.043	0.015	0.010	0.020	0.144	-0.008	0.102	0.056
	VARD	0.064	0.016	0.178	0.020	0.079	0.174	0.071	0.167
	MVAR	0.018	0.011	0.026	0.016	0.029	0.027	0.029	0.027
Gittins a=0.95	ESL	2.8	2.3	3.0	2.1	2.6	2.1	1.8	2.1
	EP	0.207	0.017	0.106	0.003	0.032	0.037	0.006	0.374
	BIAS	0.043	0.005	0.030	0.009	0.129	0.023	0.124	0.045
	VARD	0.055	0.015	0.174	0.019	0.077	0.158	0.061	0.141
	MVAR	0.018	0.011	0.031	0.015	0.029	0.030	0.030	0.026
Gittins a=0.99	ESL	3.0	2.6	2.1	2.3	2.3	1.4	1.8	1.9
	EP	0.183	0.014	0.054	0.001	0.009	0.010	0.001	0.342
	BIAS	0.030	-0.002	0.091	-0.002	0.092	0.092	0.103	0.028
	VARD	0.038	0.012	0.157	0.017	0.057	0.133	0.054	0.090
	MVAR	0.015	0.009	0.035	0.013	0.029	0.037	0.029	0.022
Fox	ESL	3.3	3.3	2.7	3.3	3.4	2.6	3.5	2.0
	EP	0.262	0.073	0.044	0.026	0.029	0.008	0.010	0.373
	BIAS	0.007	0.002	0.029	-0.001	0.019	0.023	0.016	0.011
	VARD	0.040	0.020	0.030	0.022	0.037	0.034	0.035	0.049
	MVAR	0.017	0.010	0.020	0.011	0.022	0.022	0.020	0.022
One- step- ahead	ESL	3.3	2.3	6.3	1.9	4.4	6.0	3.0	2.1
	EP	0.295	0.036	0.311	0.013	0.117	0.189	0.042	0.422
	BIAS	0.046	0.027	-0.048	0.037	0.021	-0.065	0.002	0.010
	VARD	0.099	0.021	0.160	0.024	0.135	0.178	0.114	0.200
	MVAR	0.019	0.012	0.013	0.016	0.022	0.016	0.020	0.024
Play- the- winner	ESL	4.7	8.8	5.0	12.0	10.5	6.1	12.2	2.3
	EP	0.146	0.006	0.005	0.001	0.001	0.000	0.000	0.300
	BIAS	-0.003	-0.010	0.007	-0.014	-0.008	0.005	-0.010	-0.002
	VARD	0.009	0.006	0.013	0.006	0.010	0.016	0.010	0.010
	MVAR	0.008	0.005	0.011	0.005	0.009	0.013	0.009	0.009

Table 2.4: *Simulation results for six allocation rules when $n = 150$ (stable response variable)*

Rule	p_A p_B	0.3 0.4	0.1 0.3	0.7 0.9	0.1 0.4	0.35 0.65	0.6 0.9	0.3 0.7	0.5 0.55
Equal	ESL	7.3	15.0	14.9	22.5	22.5	22.4	30.1	3.7
	EP	0.096	0.000	0.001	0.000	0.000	0.000	0.000	0.270
	BIAS	0.002	0.000	0.000	-0.001	-0.000	-0.001	0.002	0.001
	VARD	0.006	0.004	0.004	0.005	0.006	0.004	0.006	0.007
	MVAR	0.006	0.004	0.004	0.004	0.006	0.004	0.005	0.006
Bather	ESL	3.7	3.3	3.2	3.1	3.3	3.1	3.2	2.8
	EP	0.136	0.005	0.004	0.001	0.001	0.001	0.000	0.304
	BIAS	0.029	0.004	0.067	0.009	0.054	0.078	0.048	0.021
	VARD	0.024	0.010	0.033	0.013	0.032	0.041	0.032	0.039
	MVAR	0.011	0.007	0.017	0.009	0.019	0.021	0.019	0.014
Gittins a=0.9	ESL	3.8	2.4	5.4	1.9	3.3	3.7	1.8	2.9
	EP	0.202	0.016	0.151	0.004	0.039	0.060	0.007	0.377
	BIAS	0.054	0.022	-0.002	0.031	0.146	0.001	0.120	0.042
	VARD	0.059	0.014	0.181	0.017	0.073	0.173	0.061	0.164
	MVAR	0.016	0.010	0.023	0.013	0.026	0.024	0.026	0.023
Gittins a=0.99	ESL	3.6	2.8	2.6	2.3	2.4	1.6	1.8	2.7
	EP	0.146	0.004	0.047	0.001	0.007	0.008	0.001	0.309
	BIAS	0.036	0.004	0.107	0.009	0.106	0.099	0.112	0.044
	VARD	0.035	0.011	0.153	0.014	0.050	0.128	0.049	0.083
	MVAR	0.013	0.008	0.032	0.011	0.025	0.035	0.026	0.019
Fox	ESL	4.2	4.1	3.2	3.9	4.4	3.2	3.9	2.9
	EP	0.232	0.056	0.030	0.019	0.026	0.008	0.005	0.351
	BIAS	0.008	-0.004	0.024	0.001	0.010	0.021	0.012	0.008
	VARD	0.030	0.015	0.024	0.017	0.030	0.028	0.029	0.037
	MVAR	0.012	0.006	0.015	0.007	0.015	0.016	0.014	0.016
One- step- ahead	ESL	4.4	2.6	9.4	1.9	5.4	9.4	4.0	3.1
	EP	0.276	0.030	0.312	0.009	0.093	0.200	0.043	0.419
	BIAS	0.062	0.033	-0.041	0.049	0.043	-0.076	-0.001	0.008
	VARD	0.088	0.017	0.169	0.019	0.118	0.180	0.112	0.194
	MVAR	0.015	0.010	0.012	0.012	0.019	0.015	0.018	0.020
Play- the- winner	ESL	6.7	13.2	7.5	18.0	15.7	9.1	18.0	3.5
	EP	0.097	0.001	0.001	0.000	0.000	0.000	0.000	0.273
	BIAS	-0.003	-0.006	0.002	-0.007	-0.004	0.004	-0.008	0.000
	VARD	0.006	0.004	0.008	0.004	0.007	0.010	0.006	0.007
	MVAR	0.006	0.004	0.007	0.004	0.006	0.009	0.006	0.006




Figure 2.1: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.1$, $p_B = 0.3$)




Figure 2.2: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.7$, $p_B = 0.9$)






Figure 2.3: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.35$, $p_B = 0.65$)

PA=0.1 and PB=0.3



PA=0.7 and PB=0.9



PA=0.35 and PB=0.65




Figure 2.4: Plots of allocation rates to the better treatment ($n = 50$)

smaller; the play-the-winner rule has values for expected successes lost which are about 50% of those for equal allocation when $p_i > \frac{1}{2}$ ($i=A,B$) and about 60–70% when $p_A < \frac{1}{2} < p_B$ – this behaviour is suggested by the results of Zelen (1969, p.141).

Secondly, we see that as a increases for the Gittins rule, both the error probability and the variance of the estimated treatment difference generally decrease. The behaviour of expected successes lost as a increases appears to depend on the magnitude of the p_i ($i=A,B$): for the most part, when $p_A, p_B < \frac{1}{2}$, the expected successes lost increases, otherwise it decreases (cf. remark (i) in §2.2.2); the bias in estimation exhibits the opposite behaviour.

Thirdly, the magnitude of the increase in variance due to the adaptive nature of the allocation rules is particularly noticeable for the Gittins rule when $p_i > \frac{1}{2}$ ($i=A,B$) – the empirical variance tends to be about ten-fold the corresponding value for equal allocation. This is, of course, one of the more extreme cases – an example of the opposite extreme is when the p_i 's are small. Bather's rule appears to be the better rule in balancing the loss in efficiency of estimation with the reduction in the number of patients on the inferior treatment.

This last point is further emphasised in Figure 2.4. Although the Gittins rule with $a=0.99$ generally achieves a larger reduction in the number of patients on the inferior treatment, this reduction is small relative to the resulting loss in efficiency when compared with Bather's rule. The corresponding reduction for Bather's rule over the play-the-winner rule is substantial in several cases. Further, unlike the play-the-winner rule and the Gittins rule, Bather's rule has the asymptotic optimality property which seems to be reflected in Figure A1.4 in Appendix 1. Note that the "discontinuities" in Figures 2.4 and A1.4 for Fox's rule are due to the use of the inferior treatment at prespecified trials (i.e. trial numbers $2^i - 1$, $i = 1, 2, \dots$).

Finally, a few remarks should be made regarding Figures 2.1 – 2.3. Firstly, the distribution of the number of patients on the better treatment for the play-the-winner rule tends to be symmetric with relatively short tails (hence, the error probability is smaller), and centred to the right of $\frac{1}{2}n$; this property is not too surprising since from §2.4.1 the distributions of the treatment group sizes are asymptotically normal. Secondly,

Fox's rule, the one-step-ahead rule and the Gittins rule tend to perform very well for moderate to large differences ($\delta \geq 0.3$), with usually over 85% of the patients being assigned to the better treatment; for small differences ($\delta \leq 0.2$), these rules do not perform as well and, in some cases, there is a significant proportion of the distribution in the lower range of the scale, which explains some of the larger error probabilities in Tables 2.2 - 2.4. Thirdly, although for moderate to large differences the distribution for Bather's rule is concentrated in the upper range of the scale, a "tailing-off" effect is evident for small differences, and hence a smaller error probability is generally attained.

2.4.3 Differences between the two estimates of variance

In §2.3.3, a tentative explanation was provided for the behaviour of the difference between the empirical and the theoretical variances of the estimated treatment difference. We now examine these differences more closely by comparing the differences VARD-MVAR in Tables 2.2 and 2.3 with $E(\text{VARD}) - E(\text{MVAR})$ with the n_{ij} replaced with their average values (over the simulations). The results are summarised in Table 2.5; SIM indicates the values for VARD - MVAR in Tables 2.2 and 2.3 and AV the values for $E(\text{VARD}) - E(\text{MVAR})$ using average values for the n_{ij} .

The two values for the difference in the estimates of variance are in good agreement for equal allocation and the play-the-winner rule and, to a lesser extent, for Fox's rule and Bather's rule. However, for the Gittins rule and, particularly for the one-step-ahead rule, the discrepancy is more apparent when $p_i > \frac{1}{2}$ ($i=A,B$); there are also large discrepancies for the one-step-ahead rule when $p_A + p_B = 1$. Thus, one of the conclusions from the above results is that, with the exception of equal allocation, the variance of the n_{ij} can be substantial and consequently the theoretical variance (2.5) may severely underestimate the empirical variance (2.4). Further simulation results indicated that the variance of the n_{ij} for the one-step-ahead rule can be up to twice that for the Gittins rule. Some further remarks will be made on this problem in §2.8.

Table 2.5: *Simulation results for the difference between the two estimates of variance*

(a) n = 50

Rule	p _A p _B	0.3	0.1	0.7	0.1	0.35	0.6	0.3	0.5
		0.4	0.3	0.9	0.4	0.65	0.9	0.7	0.55
Equal	SIM	0.004	0.002	0.002	0.002	0.004	0.002	0.004	0.004
	AV	0.004	0.002	0.002	0.003	0.004	0.003	0.003	0.004
Bather	SIM	0.024	0.008	0.028	0.009	0.027	0.036	0.023	0.043
	AV	0.005	0.005	0.011	0.010	0.020	0.026	0.034	0.004
Gittins a=0.99	SIM	0.025	0.005	0.127	0.005	0.039	0.110	0.034	0.073
	AV	0.005	0.006	0.019	0.013	0.032	0.076	0.079	0.004
Fox	SIM	0.032	0.015	0.017	0.016	0.027	0.021	0.025	0.038
	AV	0.005	0.005	0.014	0.010	0.018	0.031	0.030	0.004
One- step- ahead	SIM	0.079	0.011	0.149	0.013	0.116	0.167	0.099	0.173
	AV	0.005	0.008	0.005	0.022	0.017	0.011	0.047	0.004
Play- the- winner	SIM	0.002	0.002	0.013	0.001	0.004	0.015	0.005	0.004
	AV	0.004	0.002	0.006	0.003	0.005	0.011	0.005	0.004

(b) n = 100

Rule	p _A p _B	0.3	0.1	0.7	0.1	0.35	0.6	0.3	0.5
		0.4	0.3	0.9	0.4	0.65	0.9	0.7	0.55
Equal	SIM	0.001	0.001	0.001	0.001	0.001	0.001	0.000	0.001
	AV	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Bather	SIM	0.016	0.004	0.020	0.004	0.017	0.024	0.015	0.030
	AV	0.002	0.003	0.007	0.006	0.012	0.017	0.024	0.001
Gittins a=0.99	SIM	0.023	0.003	0.122	0.004	0.028	0.096	0.025	0.068
	AV	0.002	0.004	0.010	0.010	0.023	0.056	0.069	0.001
Fox	SIM	0.023	0.010	0.010	0.011	0.015	0.012	0.015	0.027
	AV	0.001	0.002	0.007	0.005	0.011	0.020	0.018	0.001
One- step- ahead	SIM	0.080	0.009	0.147	0.008	0.113	0.162	0.094	0.176
	AV	0.001	0.005	0.001	0.014	0.006	0.003	0.021	0.001
Play- the- winner	SIM	0.001	0.001	0.002	0.001	0.001	0.003	0.001	0.001
	AV	0.001	0.001	0.002	0.001	0.001	0.003	0.001	0.001

2.5 Modelling instability in the response variable

We consider two types of instability: (i) time-trends; (ii) fluctuations. Reasons for considering these two types of instability were discussed in §1.6. In this section we show how to model these for a Bernoulli response variable and an indication of their anticipated effect is given.

2.5.1 Linear time-trends

If we denote as in §2.1 the response variable for treatment i at time j by X_{ij} ($i=A,B$; $j \in [0, n-1]$), then until now we have assumed that $X_{ij} \sim \text{Bin}(1, p_i) \forall j$. Suppose instead that the response variable is subject to a linear time-trend, then clearly we can model this by assuming $X_{ij} \sim \text{Bin}(1, p_{ij})$ where $p_{ij} = p_i + tj$ and t is the linear time-trend per unit time (the same for both treatments).

We would intuitively expect the efficiency of estimation for the adaptive allocation rules to be affected by the trend. In particular, the estimates of the treatment difference should be more biased, and the corresponding variances may also be affected.

2.5.2 Time-trends and fluctuations on a logit scale

A drawback with modelling time-trends on a linear scale is that the values for t are restricted since $p_{ij} \in (0,1) \forall i,j$. A more flexible approach is to model $\text{logit } p_{ij}$ which is defined over $(-\infty, \infty)$. Thus, if we wished to preserve a constant treatment difference $\delta = p_B - p_A$ we would use the model

$$\left. \begin{aligned} \text{logit } p_{Bj} &= a + tj, \quad a = \text{logit } p_B \\ p_{Aj} &= p_{Bj} - \delta \end{aligned} \right\}. \quad (2.14)$$

If, in addition, we wished to model fluctuations in the data, then a simple model for this is (2.14) with a first-order autoregressive process superimposed, namely,

$$\text{logit } p_{Bj} = a + tj + Z_j, \quad (2.15)$$

where

$$Z_j = \begin{cases} \epsilon_0 & , \quad (j = 0) \\ \rho Z_{j-1} + \epsilon_j & , \quad (j \geq 1; |\rho| < 1) \end{cases}$$

and

$$\epsilon_j \sim \begin{cases} N(0, \sigma_\epsilon^2 / (1 - \rho^2)) & , \quad (j = 0) \\ N(0, \sigma_\epsilon^2) & , \quad (j \geq 1) \end{cases}.$$

For simulation purposes, two steps were taken: (i) to ensure that $p_{Aj} > 0$ ($j \in [0, n-1]$), a constraint was incorporated in the program ALLOCATE (see §2.3.1) so that if for any realisation, $p_{Aj} \leq 0$, further values for the ϵ_j are generated until $p_{Aj} > 0$; we would not expect this constraint to significantly affect the results; (ii) a constraint was imposed on the variance of logit p_{ij} - the variance when $t = 0$ is equal to the *total* variance (due to trend and fluctuations) when $t \neq 0$; if we assume that p_{ij} has a certain beta distribution, this enables us to obtain an approximate expression for σ_ϵ^2 in terms of ρ and a given value of $p_{B, n-1}$. We now derive this expression.

Firstly, the δ -technique yields, to first-order terms,

$$\text{Var}\{\text{logit } p_{Bj}\} \approx \frac{\text{Var}(p_{Bj})}{[E(p_{Bj})\{1-E(p_{Bj})\}]^2} \quad (2.16)$$

and, if we assume $p_{ij} \sim B(\alpha, \beta)$, say, (2.16) becomes

$$\text{Var}\{\text{logit } p_{Bt}\} \approx \frac{(\alpha + \beta)^2}{\alpha\beta(\alpha + \beta + 1)} \quad (2.17)$$

Secondly, considering the right-hand-side of (2.15), we can write

$$\text{Var}\{\text{logit } p_{Bt}\} = \frac{h^2}{12} + \frac{\sigma_\epsilon^2}{(1-\rho^2)} \quad (2.18)$$

where

$$h = \log \left\{ \frac{p_{B, n-1}(1-p_B)}{p_B(1-p_{B, n-1})} \right\}. \quad (2.19)$$

The term involving h in (2.18) follows from the fact that the trend component is uniformly distributed over an interval of length h .

Finally, equating (2.17) and (2.18) yields

$$\sigma_\epsilon^2 \approx \left\{ \frac{(\alpha + \beta)^2}{\alpha\beta(\alpha + \beta + 1)} - \frac{h^2}{12} \right\} (1-\rho^2)$$

and if $\alpha = \beta = 1$, this reduces to

$$\sigma_\epsilon^2 \approx \frac{4}{3} \left\{ 1 - \frac{h^2}{16} \right\} (1-\rho^2). \quad (2.20)$$

Using (2.20), values for the ϵ_j can be generated and consequently the p_{Bj} modelled by (2.15). We would intuitively expect the variance of the estimated treatment difference to increase with ρ and perhaps also the error probability.

2.6 Linear time-trends in the response variable

We now consider the setting described in §2.5.1. Analytical results are derived in §2.6.1 for equal allocation and the play-the-winner rule; for equal allocation, these results are approximate, except in a special case, while for the play-the-winner rule, an exact result is given. Following a comparison with the corresponding simulated values, further simulation results are presented in §2.6.2.

2.6.1 Analytical results

(a) Equal allocation rule

It is clear that since a constant treatment difference is preserved, the expected successes lost is unaffected by the trend, and so is given by (2.8). To obtain the error probability (2.9) in the presence of linear time-trends, we require the probability distribution for each treatment of the number of successes in $\frac{1}{2}n$ independent trials, when the probability of success *varies* from trial to trial. This distribution is known as the *Poisson binomial* (e.g. Kotz & Johnson, 1985, Vol.7, p.19), which is a particular case of the more general *Lexian distribution* (e.g. Kotz & Johnson, 1985, Vol.4, p.622).

If we choose k random permuted blocks as in §2.1.1, and let $X_{ij}^{(k)} \sim \text{Bin}(1, \bar{p}_{ij}^{(k)})$, where $\bar{p}_{ij}^{(k)}$ denotes the mean of the p_{ij} 's which occur in the same block as the $(j+1)^{\text{th}}$ trial on the i^{th} treatment, then we see that the distribution of

$$X_i^{(k)} = \sum_{j=0}^{\frac{1}{2}n-1} X_{ij}^{(k)}, \quad (i=A, B)$$

is Poisson binomial, which reduces to binomial for $k = 1$. When $k = \frac{1}{2}n$, the distribution of the number of successes with treatment i in $\frac{1}{2}n$ trials is given *exactly* by $X_i^{(k)}$; this is easily seen by writing down the distributions for $n = 4$, say, when $k = 1$ and 2 . Furthermore, it will be seen below (Table 2.6) that for $k < \frac{1}{2}n$, $X_i^{(k)}$ provides a reasonable approximation to the required distribution.

The distribution of $X_i^{(k)}$ can be computed using a generalisation of the convolution equation (2.1) of Feller (1968, p.266), and hence the error probability deduced. A FORTRAN program, ERRPROB, was written for this purpose. A normal approximation to the error probability is also available, details of which are now given.

Following, for example, Kotz & Johnson (1985, Vol.4, p.622) it can be shown that

$$E(X_i^{(k)}) = \frac{1}{2}n\bar{p}_i^{(k)} \tag{2.21}$$

and

$$\text{Var}(X_i^{(k)}) = \frac{1}{2}n\bar{p}_i^{(k)} [1-\bar{p}_i^{(k)}] - \sum_{j=0}^{\frac{1}{2}n-1} [\bar{p}_{ij}^{(k)} - \bar{p}_i^{(k)}]^2, \tag{2.22}$$

where

$$\bar{p}_i^{(k)} = \sum_{j=0}^{\frac{1}{2}n-1} \bar{p}_{ij}^{(k)} / \frac{1}{2}n, \quad (i=A,B).$$

Hence, assuming $p_B > p_A$, we can write

$$EP \approx 1 - \Phi\left\{ \frac{E(X_B^{(k)}) - E(X_A^{(k)})}{\sqrt{[\text{Var}(X_A^{(k)}) + \text{Var}(X_B^{(k)})]}} \right\}. \tag{2.23}$$

The exact/approximate, the normal approximation (2.23) and the simulated values for the error probability are presented in Table 2.6. The results indicate that, except when the p_i 's are both small, the error probability based on the distribution of $X_i^{(5)}$ is a reasonable approximation to the simulated value for $k = 5$. Further, the discrepancy between the normal approximation and the exact/approximate values for the error probability is at most two in the third decimal place.

Table 2.6: Exact, approximate and simulated values for error probability for equal allocation

n	p_A	p_B	t	Exact	$k = \frac{1}{2}n$		$k = 5$		
					Normal	Simul.	Approx.	Normal	Simul.
50	0.3	0.4	0.010	0.238	0.236	0.226	0.233	0.231	0.233
	0.1	0.3	0.008	0.057	0.055	0.066	0.040	0.038	0.069
	0.7	0.9	0.002	0.027	0.027	0.020	0.034	0.033	0.021
	0.1	0.4	0.008	0.011	0.009	0.014	0.006	0.005	0.011
	0.35	0.65	0.004	0.015	0.013	0.014	0.015	0.013	0.013
	0.6	0.9	0.002	0.004	0.003	0.002	0.005	0.004	0.002
	0.3	0.7	0.004	0.002	0.001	0.001	0.002	0.001	0.001
	0.5	0.55	0.006	0.359	0.359	0.361	0.362	0.361	0.361
100	0.3	0.4	0.002	0.153	0.152	0.162	0.149	0.147	0.162
	0.1	0.3	0.004	0.013	0.012	0.014	0.007	0.006	0.019
	0.7	0.9	0.001	0.003	0.003	0.002	0.005	0.005	0.002
	0.1	0.4	0.004	0.001	0.000	0.001	0.000	0.000	0.001
	0.35	0.65	0.002	0.001	0.001	0.001	0.001	0.001	0.001
	0.6	0.9	0.001	0.000	0.000	0.000	0.000	0.000	0.000
	0.3	0.7	0.002	0.000	0.000	0.000	0.000	0.000	0.000
	0.5	0.55	0.003	0.305	0.304	0.292	0.308	0.308	0.293

(b) Play-the-winner rule

Recall from §2.4.1 that the play-the-winner rule can be formulated as a two-state Markov chain. Then denoting the transition matrix at time j by P_j ($j \in [0, n-1]$), we can write

$$P_j = P_0 + t_j Q, \quad (2.24)$$

where

$$P_0 = \begin{bmatrix} p_A & q_A \\ q_B & p_B \end{bmatrix}, \quad Q = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \quad q_i = 1 - p_i \quad (i=A, B).$$

If we wish to compute the expected successes lost, then we require the probabilities of allocating treatment i at time j . Clearly, these probabilities may be computed from the matrix $P_0 P_1 P_2 \dots P_j$. Thus, if we define

$$T = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix},$$

then a little algebra yields

$$T^{-1} P_j T = \begin{bmatrix} 1 & p_B - p_A \\ 0 & p_A - q_B + 2t_j \end{bmatrix},$$

and hence

$$P_0 P_1 P_2 \dots P_j = \frac{1}{2} \begin{bmatrix} 1 - c_j + d_j & 1 + c_j - d_j \\ 1 - c_j - d_j & 1 + c_j + d_j \end{bmatrix}, \quad (2.25)$$

in which

$$c_j = \begin{cases} p_B - p_A & , \quad (j=0) \\ (p_B - p_A) \left[1 + \sum_{\ell=1}^j \prod_{k=1}^{\ell} \{ p_A - q_B + 2t(j+1-k) \} \right] & , \quad (j \geq 1) \end{cases}$$

and

$$d_j = \prod_{k=1}^j \{ p_A - q_B + 2tk \}.$$

It is now straightforward to compute the expected number of successes on each treatment using (2.25). Assuming $p_B > p_A$, we obtain

$$ESL = \frac{1}{2} (p_B - p_A) \left\{ n - (p_B - p_A) \left[n + \sum_{j=1}^{n-1} \sum_{\ell=1}^j \prod_{k=1}^{\ell} \{ p_A - q_B + 2t(j+1-k) \} \right] \right\}. \quad (2.26)$$

It is easily checked that for $t = 0$ and as $n \rightarrow \infty$, (2.26) reduces to (2.12), as required. Evaluation of the corresponding error probability appears difficult and has not been attempted here.

To examine the effect of the trend upon the play-the-winner rule, exact values for expected successes lost using (2.26) are presented in Table 2.7 for both a stable response variable and in the presence of a linear time-trend. For a stable response variable, a comparison of the exact and approximate values for expected successes lost (Table 2.1) shows that little accuracy is lost in using approximation (2.12). Additionally, we see that one of the effects of the trend is to reduce the expected successes lost; this effect is more evident when the p_i are large which might be expected, since from (2.11) $E(n_B)$ increases more if the p_i are large.

Table 2.7: *A comparison of exact values for expected successes lost in the absence and presence of a linear time-trend for the play-the-winner rule*

n	p_A	p_B	t	ESL	
				Stable	Trend
50	0.3	0.4	0.010	2.3	2.2
	0.1	0.3	0.008	4.4	4.1
	0.7	0.9	0.002	2.6	1.8
	0.1	0.4	0.008	6.0	5.4
	0.35	0.65	0.004	5.3	4.7
	0.6	0.9	0.002	3.1	2.0
	0.3	0.7	0.004	6.0	4.9
	0.5	0.55	0.006	1.2	1.2
100	0.3	0.4	0.002	4.6	4.5
	0.1	0.3	0.004	8.8	8.3
	0.7	0.9	0.001	5.1	3.3
	0.1	0.4	0.004	12.0	10.7
	0.35	0.65	0.002	10.5	9.3
	0.6	0.9	0.001	6.1	3.7
	0.3	0.7	0.002	12.0	9.8
	0.5	0.55	0.003	2.4	2.3

2.6.2 Simulation results

The results are a natural extension of those in §2.4.2 and are given in Tables 2.8 and 2.9 for $n = 50$ and 100 , respectively. All other parameters are as before. A selection of histograms of the number of patients on the better treatment for $n = 50$ is

displayed in Figures 2.5 - 2.7. Further plots of allocation rates to the better treatment are not included since the effect of the trend upon these will become clear from those results presented.

Since several of the conclusions are similar to those discussed in §2.4.2, we restrict our attention here to studying the effect of the linear time-trend. Firstly, the play-the-winner rule for the most part has similar error probabilities to equal allocation. It is also clear that for $p_i > \frac{1}{2}$ ($i=A,B$), the play-the-winner rule allocates a much larger proportion of patients to the better treatment than for a stable response variable (cf. remark in §2.6.1), and hence the variance of the estimated treatment difference is increased, as is the bias. This effect is particularly evident in Figures 2.5 - 2.7, where there is a "shift" to the right in the distribution (cf. Figures 2.1 - 2.3).

Secondly, with the exception of a slight positive bias in estimation, the trend has little effect upon the performance of Fox's rule, whereas both the expected successes lost and the error probability increase for the one-step-ahead rule, with the effect being more pronounced when $p_A + p_B \leq \frac{1}{2}$ (cf. Tables 2.2 and 2.3). The "robustness" property of Fox's rule is perhaps not surprising since this rule allocates patients to a particular treatment at a number of specified times, and hence there is a "balancing-out" effect. However, an explanation for the corresponding behaviour of the one-step-ahead rule is less clear.

Thirdly, the expected successes lost tends to decrease in the presence of a trend for Bather's rule and increase for the Gittins rule. For both rules, the error probability increases when $p_A + p_B \leq \frac{1}{2}$, while for $p_A + p_B > \frac{1}{2}$ there is some evidence that it decreases for Bather's rule and increases for the Gittins rule. Also, when $p_i > \frac{1}{2}$ ($i=A,B$) Bather's rule tends to allocate more patients to the better treatment - this effect is again clear from Figures 2.5 - 2.7 where there is a slight "shift" to the right in the distribution.

In conclusion, it appears that of the more adaptive allocation rules considered, the Gittins rule is the best in terms of expected successes lost but Fox's rule and, to a lesser extent, Bather's rule are reasonably "robust" with respect to both expected successes lost and error probability over a wide range of values for (p_A, p_B) . However, Fox's rule, like

Table 2.8: *Simulation results for six allocation rules when $n = 50$ (linear time-trends in response variable)*

Rule	p_A	0.3	0.1	0.7	0.1	0.35	0.6	0.3	0.5
	p_B	0.4	0.3	0.9	0.4	0.65	0.9	0.7	0.55
	t	0.010	0.008	0.002	0.008	0.004	0.002	0.004	0.006
Equal	ESL	2.6	5.0	4.9	7.5	7.5	7.5	10.0	1.3
	EP	0.233	0.069	0.021	0.011	0.013	0.002	0.001	0.361
	BIAS	-0.001	0.000	-0.002	0.000	0.003	-0.002	-0.004	-0.001
	VARD	0.018	0.017	0.010	0.017	0.017	0.011	0.016	0.017
	MVAR	0.014	0.014	0.008	0.014	0.014	0.009	0.013	0.014
Bather	ESL	1.9	2.3	1.9	2.3	2.4	2.2	2.5	1.1
	EP	0.269	0.095	0.021	0.025	0.022	0.003	0.004	0.366
	BIAS	0.024	0.017	0.086	0.025	0.054	0.085	0.051	0.023
	VARD	0.066	0.036	0.052	0.038	0.059	0.057	0.056	0.075
	MVAR	0.026	0.019	0.022	0.022	0.029	0.025	0.029	0.028
Gittins a=0.9	ESL	1.8	2.0	2.3	1.9	1.9	1.9	1.5	1.1
	EP	0.324	0.116	0.200	0.053	0.068	0.083	0.026	0.418
	BIAS	0.042	0.029	-0.004	0.034	0.102	-0.019	0.047	0.022
	VARD	0.145	0.051	0.189	0.055	0.123	0.190	0.111	0.240
	MVAR	0.038	0.024	0.028	0.029	0.040	0.029	0.036	0.044
Gittins a=0.95	ESL	1.8	1.9	2.0	2.0	2.0	1.7	1.4	1.1
	EP	0.317	0.115	0.156	0.047	0.061	0.064	0.020	0.409
	BIAS	0.046	0.026	0.016	0.026	0.110	0.005	0.080	0.020
	VARD	0.134	0.048	0.196	0.054	0.117	0.183	0.104	0.225
	MVAR	0.035	0.023	0.034	0.028	0.041	0.034	0.039	0.043
Gittins a=0.99	ESL	1.8	2.2	1.6	2.0	1.8	1.3	1.5	1.1
	EP	0.290	0.108	0.097	0.034	0.036	0.029	0.006	0.378
	BIAS	0.032	0.001	0.044	0.015	0.118	0.046	0.109	0.036
	VARD	0.098	0.037	0.189	0.043	0.093	0.162	0.081	0.156
	MVAR	0.035	0.022	0.038	0.027	0.041	0.037	0.042	0.041
Fox	ESL	2.0	2.4	1.9	2.6	2.8	2.1	2.8	1.1
	EP	0.319	0.139	0.060	0.061	0.066	0.016	0.019	0.398
	BIAS	0.018	0.014	0.036	0.015	0.023	0.028	0.021	0.009
	VARD	0.064	0.045	0.036	0.043	0.055	0.041	0.050	0.067
	MVAR	0.023	0.016	0.021	0.017	0.024	0.021	0.023	0.027
One- step- ahead	ESL	1.9	2.1	3.3	2.0	2.7	3.3	2.2	1.1
	EP	0.351	0.154	0.327	0.076	0.146	0.197	0.070	0.429
	BIAS	0.024	0.032	-0.059	0.031	-0.000	-0.056	-0.027	0.007
	VARD	0.166	0.061	0.174	0.062	0.174	0.191	0.151	0.224
	MVAR	0.037	0.027	0.015	0.030	0.031	0.018	0.028	0.035
Play- the- winner	ESL	2.3	4.2	1.8	5.4	4.7	2.1	5.1	1.2
	EP	0.230	0.073	0.014	0.015	0.014	0.003	0.002	0.361
	BIAS	-0.003	-0.012	0.038	-0.019	-0.017	0.037	-0.021	-0.003
	VARD	0.021	0.016	0.063	0.018	0.023	0.069	0.026	0.022
	MVAR	0.017	0.014	0.034	0.014	0.018	0.035	0.019	0.018

Table 2.9: *Simulation results for six allocation rules when $n = 100$ (linear time-trends in response variable)*

Rule	p_A	0.3	0.1	0.7	0.1	0.35	0.6	0.3	0.5
	p_B	0.4	0.3	0.9	0.4	0.65	0.9	0.7	0.55
	t	0.002	0.004	0.001	0.004	0.002	0.001	0.002	0.003
Equal	ESL	5.0	10.0	10.0	14.9	15.0	15.0	20.1	2.5
	EP	0.162	0.019	0.002	0.001	0.001	0.000	0.000	0.293
	BIAS	-0.000	-0.003	-0.001	0.001	-0.001	0.001	0.000	-0.001
	VARD	0.010	0.009	0.005	0.008	0.008	0.006	0.008	0.009
	MVAR	0.009	0.008	0.004	0.008	0.008	0.005	0.007	0.008
Bather	ESL	2.9	3.0	2.6	2.7	3.0	2.6	2.9	2.0
	EP	0.209	0.048	0.004	0.007	0.006	0.000	0.001	0.339
	BIAS	0.024	0.028	0.078	0.036	0.059	0.084	0.061	0.020
	VARD	0.038	0.023	0.038	0.025	0.042	0.046	0.043	0.056
	MVAR	0.017	0.014	0.019	0.016	0.024	0.023	0.024	0.021
Gittins a=0.9	ESL	3.2	2.7	4.1	2.1	2.8	2.9	1.9	1.9
	EP	0.274	0.084	0.180	0.025	0.055	0.067	0.018	0.405
	BIAS	0.040	0.043	-0.006	0.054	0.123	-0.001	0.093	0.034
	VARD	0.089	0.034	0.181	0.034	0.099	0.178	0.088	0.221
	MVAR	0.022	0.017	0.025	0.020	0.031	0.026	0.028	0.031
Gittins a=0.95	ESL	2.9	2.7	3.3	2.3	2.7	2.4	1.9	2.0
	EP	0.245	0.078	0.127	0.026	0.049	0.050	0.013	0.393
	BIAS	0.042	0.040	0.023	0.048	0.134	0.019	0.120	0.036
	VARD	0.077	0.032	0.180	0.034	0.093	0.172	0.075	0.194
	MVAR	0.023	0.017	0.030	0.019	0.032	0.030	0.031	0.032
Gittins a=0.99	ESL	2.8	2.8	2.1	2.1	2.3	1.5	1.7	2.0
	EP	0.214	0.059	0.057	0.012	0.018	0.013	0.004	0.370
	BIAS	0.036	0.030	0.077	0.046	0.123	0.071	0.133	0.034
	VARD	0.056	0.028	0.158	0.029	0.073	0.141	0.064	0.138
	MVAR	0.020	0.016	0.034	0.019	0.033	0.034	0.031	0.031
Fox	ESL	3.1	3.7	2.7	3.4	3.7	2.6	3.4	2.1
	EP	0.267	0.117	0.044	0.038	0.044	0.009	0.010	0.395
	BIAS	0.012	0.011	0.029	0.024	0.024	0.025	0.021	0.005
	VARD	0.044	0.032	0.029	0.030	0.040	0.033	0.037	0.051
	MVAR	0.018	0.012	0.020	0.013	0.022	0.021	0.020	0.024
One- step- ahead	ESL	3.2	3.1	6.3	2.4	4.7	6.3	3.4	2.1
	EP	0.312	0.113	0.313	0.050	0.133	0.199	0.064	0.431
	BIAS	0.048	0.046	-0.046	0.061	0.014	-0.064	0.006	0.005
	VARD	0.119	0.043	0.161	0.041	0.154	0.183	0.136	0.212
	MVAR	0.022	0.016	0.013	0.018	0.022	0.015	0.020	0.024
Play- the- winner	ESL	4.4	8.3	3.4	10.7	9.2	3.8	10.0	2.3
	EP	0.162	0.020	0.002	0.002	0.001	0.000	0.000	0.297
	BIAS	-0.005	-0.007	0.034	-0.008	-0.009	0.039	-0.007	0.002
	VARD	0.010	0.008	0.027	0.009	0.011	0.033	0.012	0.010
	MVAR	0.009	0.008	0.017	0.008	0.010	0.020	0.011	0.009




Figure 2.5: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.1$, $p_B = 0.3$, $t = 0.008$)




Figure 2.6: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.7$, $p_B = 0.9$, $t = 0.002$)




Figure 2.7: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.35$, $p_B = 0.65$, $t = 0.004$)

the play-the-winner rule, is not particularly ethical (Fox, 1974, p.965) and, in its present form, the Gittins rule is not randomised, and so it is particularly susceptible to selection bias (cf. remarks in §1.5.2). Therefore, of the data-dependent allocation rules considered, Bather's rule seems to be the better in the present context.

2.7 Time-trends and fluctuations in the response variable

We now consider the setting described in §2.5.2. Simulation results are presented separately for the six allocation rules in Tables 2.10 – 2.15 for $n = 50$. These results are tabulated as functions of ρ and $p_{B,n-1}$; only positive values for ρ were considered since negative values did not seem appropriate for modelling between-patient variation. A selection of histograms of the number of patients on the better treatment are displayed in Figures 2.8 – 2.13.

On comparing Tables 2.10 – 2.15 with Table 2.8, it is clear that the general effect of a time-trend on a logit scale is the same as for a linear time-trend, and hence the remarks in §2.6.2 also apply here. It should, of course, be noted that whereas in §2.6 a deterministic model was assumed for the time-trend, here a random component is incorporated.

Let us consider the effect of fluctuations in the response variable. Firstly, it is clear that the play-the-winner rule is affected differently to the other allocation rules. This may be due to the finite-memory nature of the play-the-winner rule – only the most recent observed response is used. A good example of this conflicting behaviour is when $p_A = 0.35$, $p_B = 0.65$ and $p_{B,n-1} = 0.95$: whereas the variance of the estimated treatment difference increases for the play-the-winner rule as ρ increases, there is generally a decrease for the other allocation rules (this effect is also evident from Figures 2.8 – 2.13). The increase in variance might be expected since the closer ρ is to one, the more reliable is the previous response, and hence the more adaptive the rule. Furthermore, since the treatment group sizes for the other allocation rules are less affected by the fluctuations than for the play-the-winner rule, we might expect the resulting variance of the estimated treatment difference to decrease with ρ . This is best seen by considering the difference in two sample means when the observations are correlated – for fixed

Table 2.11: Simulation results for Bather's rule when $n = 50$ (time-trends and/or fluctuations in response variable)

PA	PB	PB, n-1	ρ	ESL	EP	BIAS	VARD	MVAR
0.1	0.3	0.3	0.0	2.4	0.075	0.021	0.036	0.019
			0.6	0.0	2.4	0.094	0.028	0.022
		0.6	0.3	2.4	0.096	0.027	0.048	0.022
			0.7	2.3	0.102	0.032	0.047	0.021
			0.95	2.4	0.108	0.029	0.046	0.020
			0.99	2.3	0.095	0.026	0.046	0.021
			0.0	2.3	0.122	0.031	0.051	0.025
			0.3	2.3	0.128	0.033	0.053	0.025
			0.7	2.3	0.121	0.029	0.055	0.025
			0.95	2.4	0.121	0.029	0.053	0.024
			0.99	2.2	0.123	0.036	0.054	0.025
0.7	0.9	0.9	0.0	2.3	0.054	0.070	0.060	0.024
			0.99	0.0	2.0	0.025	0.087	0.023
		0.99	0.3	2.0	0.021	0.075	0.056	0.023
			0.7	2.0	0.027	0.082	0.057	0.022
			0.95	1.9	0.026	0.083	0.054	0.022
			0.99	2.0	0.024	0.084	0.054	0.022
0.35	0.65	0.65	0.0	2.5	0.024	0.051	0.054	0.026
			0.8	0.0	2.5	0.021	0.058	0.028
		0.8	0.3	2.5	0.027	0.056	0.061	0.028
			0.7	2.5	0.021	0.055	0.057	0.027
			0.95	2.5	0.026	0.057	0.055	0.025
			0.99	2.4	0.019	0.059	0.055	0.025
			0.0	2.3	0.020	0.066	0.062	0.031
		0.95	0.3	2.4	0.027	0.060	0.064	0.030
			0.7	2.3	0.026	0.073	0.063	0.029
			0.95	2.3	0.024	0.065	0.060	0.029
			0.99	2.3	0.021	0.066	0.061	0.028
0.5	0.55	0.55	0.0	1.1	0.368	0.016	0.068	0.023
			0.75	0.0	1.1	0.370	0.017	0.070
		0.75	0.3	1.0	0.377	0.018	0.073	0.026
			0.7	1.0	0.374	0.017	0.071	0.025
			0.95	1.0	0.378	0.013	0.063	0.022
			0.99	1.0	0.363	0.010	0.061	0.021
			0.0	1.0	0.359	0.024	0.082	0.029
		0.95	0.3	0.9	0.365	0.023	0.083	0.029
			0.7	1.0	0.363	0.023	0.084	0.029
			0.95	1.0	0.356	0.024	0.077	0.026
			0.99	0.9	0.351	0.024	0.072	0.026

Table 2.13: *Simulation results for Fox's rule when $n = 50$ (time-trends and/or fluctuations in response variable)*

P_A	P_B	$P_{B, n-1}$	ρ	ESL	EP	BIAS	VARD	MVAR
0.1	0.3	0.3	0.0	2.5	0.127	0.011	0.045	0.020
			0.6	0.0	2.6	0.147	0.017	0.020
			0.3	2.4	0.138	0.021	0.051	0.021
			0.7	2.4	0.138	0.025	0.051	0.021
			0.95	2.5	0.137	0.017	0.048	0.019
		0.9	0.99	2.4	0.136	0.012	0.047	0.019
			0.0	2.5	0.168	0.018	0.055	0.019
			0.3	2.4	0.152	0.023	0.054	0.019
			0.7	2.4	0.153	0.027	0.054	0.019
			0.95	2.5	0.150	0.025	0.051	0.019
			0.99	2.4	0.155	0.028	0.053	0.019
0.7	0.9	0.9	0.0	2.2	0.086	0.030	0.045	0.024
			0.99	0.0	2.0	0.073	0.027	0.022
		0.99	0.3	2.0	0.064	0.034	0.039	0.021
			0.7	2.0	0.066	0.034	0.037	0.021
			0.95	1.9	0.066	0.030	0.036	0.021
			0.99	2.0	0.071	0.036	0.037	0.021
0.35	0.65	0.65	0.0	2.7	0.055	0.019	0.053	0.025
			0.8	0.0	2.7	0.061	0.015	0.025
			0.3	2.7	0.059	0.024	0.052	0.024
			0.7	2.7	0.053	0.025	0.051	0.024
			0.95	2.6	0.056	0.019	0.048	0.022
		0.95	0.99	2.7	0.056	0.017	0.050	0.022
			0.0	2.7	0.065	0.024	0.057	0.024
			0.3	2.7	0.057	0.029	0.055	0.024
			0.7	2.5	0.052	0.036	0.053	0.023
			0.95	2.6	0.058	0.029	0.051	0.022
			0.99	2.6	0.060	0.030	0.051	0.022
0.5	0.55	0.55	0.0	1.1	0.396	0.006	0.066	0.026
			0.75	0.0	1.1	0.400	0.011	0.027
			0.3	1.0	0.385	0.009	0.064	0.026
			0.7	1.1	0.410	0.009	0.064	0.025
			0.95	1.0	0.376	0.012	0.054	0.023
		0.95	0.99	1.0	0.363	0.016	0.054	0.022
			0.0	1.0	0.402	0.013	0.063	0.027
			0.3	0.9	0.388	0.018	0.062	0.027
			0.7	1.0	0.404	0.008	0.063	0.026
			0.95	1.0	0.368	0.024	0.059	0.024
			0.99	0.9	0.372	0.021	0.055	0.024




Figure 2.8: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.1$, $p_B = 0.3$, $\rho = 0.3$, $p_{B,n-1} = 0.6$)




Figure 2.9: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.1$, $p_B = 0.3$, $\rho = 0.95$, $p_{B,n-1} = 0.6$)




Figure 2.10: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.7$, $p_B = 0.9$, $\rho = 0.3$, $p_{B,n-1} = 0.99$)




Figure 2.11: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.7$, $p_B = 0.9$, $\rho = 0.95$, $p_{B,n-1} = 0.99$)




Figure 2.12: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.35$, $p_B = 0.65$, $\rho = 0.3$, $p_{B,n-1} = 0.8$)




Figure 2.13: Histograms of the number of patients on the better treatment ($n = 50$, $p_A = 0.35$, $p_B = 0.65$, $\rho = 0.95$, $p_{B,n-1} = 0.8$)

sample sizes and constant variances, the variance of this difference decreases as ρ increases for positive ρ .

Secondly, as in §2.6.2, Bather's rule and Fox's rule are reasonably "robust" with respect to changes in ρ but, as expected, Bather's rule has the smaller error probabilities and a larger positive bias in estimation. As before, the one-step-ahead rule is affected by the magnitude of the p_i : for small values, there is a positive bias in estimation while for large values, there is a negative bias.

Thirdly, the expected successes lost for the play-the-winner rule generally decreases as ρ increases; this behaviour might be expected with reference to the first point above. However, the variation in the expected successes lost for the other rules shows no definite trends.

In conclusion, the performance of the various allocation rules is not severely affected by the presence of fluctuations. The results for the play-the-winner rule may indicate that in order to take into account possible fluctuations, the "total" responses for each treatment in the adaptive allocation rules should be replaced by "recent" totals. This approach was mentioned briefly by Whittle (1981b, p.289), when in a discussion of Bather's rule, he suggested that in such a setting $\lambda(\ell)$ should tend to a positive rather than a zero limit. This aspect of the problem deserves further attention.

2.8 Discussion

In this chapter, we proposed a method for estimating the difference between two unknown success probabilities following a trial in which a data-dependent allocation rule is used. The estimation involves splitting the trial into a number of periods and estimating the difference within these periods; a weighted estimate is then constructed based on these estimated differences. From the results in §2.4.3, it is clear that further theoretical work is needed regarding the variances of the sample sizes n_{ij} within each period. Progress in this direction may lead to a theoretical estimate of variance which is closer to the empirical variance (2.4). This problem, however, has not been pursued in this thesis.

The simulation results presented in §2.6.2 indicated that for some allocation rules such as the play-the-winner rule and Bather's rule, one of the effects of a linear

time-trend is to allocate more patients to the better treatment. This might not be surprising since there will be less uncertainty as to which is the better treatment. The effect of a trend upon the other adaptive allocation rules was less clear. In §2.6.1 we provided some theoretical justification for the behaviour of the play-the-winner rule in this setting but a corresponding justification for, say, Bather's rule has not been investigated. Further insight into the Gittins rule may also be possible by studying its limiting case, the least failures rule (see §2.1.6).

The use of randomisation theory has been recently studied by several authors in the context of restricted randomisation schemes (e.g. Mehta *et al.*, 1988; Wei, Smythe & Mehta, 1989). However, it appears that only Wei (1988) has investigated this problem in the context of data-dependent allocation. Wei constructs an exact permutation test for testing the equality of two treatments when the randomised play-the-winner rule is adopted. As an illustration, Wei applies this test to the trial described in Example 1.3 of §1.5.1, and shows that the allocation rule cannot be ignored in the analysis. It would be interesting to know whether such a permutation test can be derived for other randomised data-dependent allocation rules such as Bather's rule. Moreover, how is the permutation test affected by the presence of linear time-trends in the response variable? Although the problem of linear time-trends has been examined for some restricted randomisation schemes (e.g. Mehta *et al.*, 1988, pp.300–301), it remains open in a data-dependent allocation setting.

CHAPTER 3. ESTIMATION FOR DECISION-THEORETIC MODELS

We now return to the one- and two-stage models of Colton (1963, 1965) which were briefly mentioned in §1.3.2.

In §3.1, the notation and objectives for these models are defined and the optimal values for the parameters determined for each model.

A procedure for estimating the true mean treatment difference is described in §3.2 and then applied to the models; an alternative procedure is also suggested.

In Chapter 2, two forms of instability in the response variable were considered; these are briefly mentioned again in §3.3 and, following this, their effect upon estimation is investigated in §3.4. Numerical and simulation results for estimation are presented in §3.5. Least squares estimation of the parameters characterising the instability is briefly outlined in §3.6 and appropriate adjustments in the estimation formulae of §§3.2 and 3.4 are indicated.

In §3.7, a natural extension of the models is examined – a three-stage model. Some results are presented to indicate the potential gain from using a more adaptive model. Lastly, in §3.8, some comments and results are given for Bernoulli and exponential responses.

3.1 Models considered

3.1.1 Preliminaries

Each model has associated with it an *experimental* group and a *postexperimental* group of patients. The patients in the experimental group can be allocated to one of two treatments A and B in a number of *stages*, in each of which patients can either be randomised equally to the two treatments or all allocated to the currently better treatment. The postexperimental group consists of all the remaining patients and these are allocated to the currently better treatment.

We make the following assumptions:

- (i) there are a total of N patients to be treated (the patient horizon);
- (ii) the response variable, X_i , for treatment i is normally distributed with mean μ_i and

standard deviation σ ;

(iii) the only cost involved is that of treating the patient;

(iv) the cost is proportional to the true mean treatment difference, δ , with $\delta = \mu_A - \mu_B$.

[N.B. Colton (1965) assumed a normal prior for δ . This assumption is not made here; instead, we standardise δ by σ and examine the performance of the models for different values of δ/σ .]

3.1.2 Formulation of models

In this section, we will consider three models: the first is a one-stage model (called the Fixed Model by Colton, 1963); the second and third are two-stage models (called Model I and Model II by Colton, 1965).

a) Fixed model

Stage 1: Perform a clinical trial on $2Np_1$ patients, Np_1 on each treatment. Compute the observed difference in means, $\bar{d}_1 = \bar{x}_{1A} - \bar{x}_{1B}$.

If $\bar{d}_1 > 0$, use treatment A on the remaining $N(1-2p_1)$ patients;

if $\bar{d}_1 < 0$, use treatment B on the remaining $N(1-2p_1)$ patients.

b) Model I

Stage 1: Perform a clinical trial on $2Np_1$ patients, Np_1 on each treatment.

If $\bar{d}_1 > D$, use treatment A on the remaining $N(1-2p_1)$ patients;

if $\bar{d}_1 < -D$, use treatment B on the remaining $N(1-2p_1)$ patients;

if $-D < \bar{d}_1 < D$, then

Stage 2: Continue the clinical trial on $2Np_2$ additional patients, Np_2 on each treatment.

Compute $\bar{d}_2 = \bar{x}_{2A} - \bar{x}_{2B}$ and $y = (p_1\bar{d}_1 + p_2\bar{d}_2)/(p_1 + p_2)$.

If $y > 0$, use treatment A on the remaining $N(1-2p_1-2p_2)$ patients;

if $y < 0$, use treatment B on the remaining $N(1-2p_1-2p_2)$ patients.

Model I allows early termination of the trial if the observed difference is sufficiently large at the end of the first stage. Otherwise, a further randomised stage is used.

c) Model II

Stage 1: Perform a clinical trial on $2Np_1$ patients, Np_1 on each treatment.

If $\bar{d}_1 > 0$, use treatment A in the second stage;

if $\bar{d}_1 < 0$, use treatment B in the second stage.

Stage 2: Use one treatment only on Np_2 patients.

If treatment A is used, calculate

$$y_A = \frac{p_1 \bar{x}_{1A} + p_2 \bar{x}_{2A}}{p_1 + p_2} - \bar{x}_{1B} .$$

If $y_A > 0$, continue with A on the remaining $N(1-2p_1-p_2)$ patients;

if $y_A < 0$, switch to B on the remaining $N(1-2p_1-p_2)$ patients.

If treatment B is used, calculate

$$y_B = \bar{x}_{1A} - \frac{p_1 \bar{x}_{1B} + p_2 \bar{x}_{2B}}{p_1 + p_2} .$$

If $y_B > 0$, switch to A on the remaining $N(1-2p_1-p_2)$ patients;

if $y_B < 0$, continue with B on the remaining $N(1-2p_1-p_2)$ patients.

With Model II, a temporary decision is made at the end of the first stage to use what appears to be the better treatment. After using this treatment for a further Np_2 patients, there is an opportunity to switch to the other treatment.

Objectives

For each model we wish to construct the appropriate cost or net gain function using the assumptions and then to determine the parameters (p_1 for Fixed Model; D , p_1 and p_2 for Model I; p_1 and p_2 for Model II) such that the overall cost is minimised or overall net gain is maximised. Using these optimal parameters, we can then proceed to compare the three models in their efficiency in estimating δ .

Net gain formulation

'The net gain formulation states that if a patient is treated with the superior treatment a positive gain proportional to the true difference is scored, while if treated with the inferior treatment a negative gain proportional to the true difference is scored' (Colton, 1965, p.170).

Using this formulation, we can write

$$\begin{aligned} E(\text{Net gain}) &= NG\delta \{E(\text{Proportion on A}) - E(\text{Proportion on B})\} \\ &= NG\delta \{1 - 2E(\text{Proportion on B})\} , \end{aligned}$$

where G is a constant of proportionality.

If we assume $\delta > 0$, then B is the inferior treatment. Therefore, maximising

$E(\text{Netgain})$ is equivalent to minimising $E(\text{Proportion on } B)$. In §3.1.3, we minimise $E(\text{Proportion on } B)$ for the three models.

3.1.3 Optimisation

a) Fixed Model

Defining $E = E(\text{Proportion on } B)$, we may write

$$E = p_1 + (1-2p_1)\Phi(h),$$

where $h = -\delta \vee (\frac{1}{2}Np_1)/\sigma$ and $\Phi(x)$ denotes the standard normal distribution function.

Setting $\partial E/\partial p_1 = 0$ gives equation (6) of Colton (1963),

$$\frac{2\Phi(h)-1}{h\varphi(h)} = \frac{1-2p_1}{2p_1},$$

where $\varphi(x)$ denotes the standard normal density function. Given N and δ/σ , this equation can be solved numerically for p_1 .

b) Model I

Using equation (6) of Colton (1965), we may write

$$E = p_1 + p_2 \Phi(h_U) + (1-2p_1-p_2)\Phi(h_L) + (1-2p_1-2p_2) \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u,v) du dv,$$

where

$$h_U = (D-\delta) \vee (\frac{1}{2}Np_1)/\sigma, \quad h_L = (D+\delta) \vee (\frac{1}{2}Np_1)/\sigma, \quad k = -\delta \vee (\frac{1}{2}N(p_1+p_2))/\sigma$$

and $\varphi_2(u,v)$ is a standard bivariate normal density with correlation coefficient

$$\rho_{1y} = \vee \{p_1/(p_1+p_2)\}.$$

Now

$$\begin{aligned} \frac{\partial E}{\partial D} &= p_2 \varphi(h_U) \frac{\partial h_U}{\partial D} + (1-2p_1-p_2) \varphi(h_L) \frac{\partial h_L}{\partial D} \\ &+ (1-2p_1-2p_2) \left\{ \frac{\partial h_U}{\partial D} \int_{-\infty}^k \varphi_2(u, h_U) du - \frac{\partial h_L}{\partial D} \int_{-\infty}^k \varphi_2(u, h_L) du \right\}, \end{aligned}$$

and since

$$\int_{-\infty}^k \varphi_2(u, h_i) du = \varphi(h_i) \Phi(k_i), \quad k_i = (k - h_i \rho_{1y}) / \vee (1 - \rho_{1y}^2), \quad (i=U, L), \quad (3.1)$$

setting $\partial E/\partial D = 0$ gives

$$e^{Np_1 D \delta / \sigma} \left\{ p_2 + (1-2p_1-2p_2) \Phi(k_U) \right\} + \left\{ (1-2p_1-2p_2) \Phi(k_L) - (1-2p_1-p_2) \right\} = 0. \quad (3.2)$$

Similarly, we have

$$\begin{aligned} \frac{\partial E}{\partial p_1} &= 1 + p_2 \varphi(h_U) \frac{\partial h_U}{\partial p_1} + (1-2p_1-p_2) \varphi(h_L) \frac{\partial h_L}{\partial p_1} - 2 \Phi(h_L) \\ &+ (1-2p_1-2p_2) \frac{\partial}{\partial p_1} \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u, v) du dv - 2 \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u, v) du dv, \end{aligned}$$

where

$$\begin{aligned} \frac{\partial}{\partial p_1} \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u, v) du dv &= \frac{\partial k}{\partial p_1} \int_{h_L}^{h_U} \varphi_2(k, v) dv + \frac{\partial h_U}{\partial p_1} \int_{-\infty}^k \varphi_2(u, h_U) du \\ &- \frac{\partial h_L}{\partial p_1} \int_{-\infty}^k \varphi_2(u, h_L) du + \int_{h_L}^{h_U} \int_{-\infty}^k \frac{\partial}{\partial p_1} \varphi_2(u, v) du dv \end{aligned}$$

using the chain rule for differentiation, and

$$\frac{\partial}{\partial p_1} \varphi_2(u, v) = \varphi_2(u, v) \frac{\partial \rho_{1y}}{\partial p_1} \left\{ \frac{\rho_{1y}(1-\rho_{1y}^2) + (1+\rho_{1y}^2)uv - \rho_{1y}(u^2+v^2)}{(1-r^2)^2} \right\}$$

with

$$\frac{\partial \rho_{1y}}{\partial p_1} = \frac{1}{2} \frac{p_2}{\sqrt{p_1(p_1+p_2)^3}}.$$

Setting $\partial E / \partial p_1 = 0$ gives after some algebra

$$\begin{aligned} 1 - 2 \left\{ \Phi(h_L) + \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u, v) du dv \right\} &+ \frac{1}{2} \varphi(h_U) \frac{h_U}{p_1} \left\{ p_2 + (1-2p_1-2p_2) \Phi(k_U) \right\} \\ &+ \frac{1}{2} \varphi(h_L) \frac{h_U}{p_1} \left\{ (1-2p_1-p_2) - (1-2p_1-2p_2) \Phi(k_L) \right\} \\ &+ \frac{(1-2p_1-2p_2)}{(p_1+p_2)} \frac{1}{2} k \varphi(k) \left\{ \Phi(h_U') - \Phi(h_L') \right\} \\ &+ (1-2p_1-2p_2) \int_{h_L}^{h_U} \int_{-\infty}^k \frac{\partial}{\partial p_1} \varphi_2(u, v) du dv \\ &= 0, \end{aligned} \quad (3.3)$$

where

$$h_i' = \frac{h_i - k \rho_{1y}}{\sqrt{(1-\rho_{1y}^2)}}, \quad (i=U, L).$$

In the same way, we can obtain $\partial E / \partial p_2$, which upon setting to zero gives

$$\begin{aligned} \Phi(h_U) - \Phi(h_L) - 2 \int_{h_L}^{h_U} \int_{-\infty}^k \varphi_2(u, v) du dv &+ (1-2p_1-2p_2) \int_{h_L}^{h_U} \int_{-\infty}^k \frac{\partial}{\partial p_2} \varphi_2(u, v) du dv \\ &+ \frac{(1-2p_1-2p_2)}{(p_1+p_2)} \frac{1}{2} k \varphi(k) \left\{ \Phi(h_U') - \Phi(h_L') \right\} = 0. \end{aligned} \quad (3.4)$$

Given N and δ/σ , the simultaneous equations (3.2), (3.3) and (3.4) can be solved numerically for p_1 , p_2 and D .

c) Model II

Clearly, we may write

$$E = p_1 + p_2 \Phi(h) + (1-2p_1-p_2) \Phi(k) ,$$

where h is as for the Fixed Model and

$$k = -\delta / \{ N p_1 (p_1 + p_2) / (2p_1 + p_2) \} / \sigma .$$

Proceeding as for Model I, we obtain

$$\frac{\partial E}{\partial p_1} = 1 + p_2 \varphi(h) \frac{\partial h}{\partial p_1} + (1-2p_1-p_2) \varphi(k) \frac{\partial k}{\partial p_1} - 2 \Phi(k) = 0 , \quad (3.5)$$

where

$$\frac{\partial h}{\partial p_1} = \frac{1}{2} \frac{h}{p_1} , \quad \frac{\partial k}{\partial p_1} = \frac{1}{2} k \left\{ \frac{(p_1 + p_2)^2 + p_1^2}{p_1 (p_1 + p_2) (2p_1 + p_2)} \right\} ,$$

and

$$\frac{\partial E}{\partial p_2} = \Phi(h) + (1-2p_1-p_2) \varphi(k) \frac{\partial k}{\partial p_2} - \Phi(k) = 0 \quad (3.6)$$

in which

$$\frac{\partial k}{\partial p_2} = \frac{1}{2} k p_1^2 / \{ p_1 (p_1 + p_2) (2p_1 + p_2) \} .$$

Again, given N and δ/σ , we can set (3.5) and (3.6) to zero and solve numerically for p_1 and p_2 .

Computation

A FORTRAN program, STANDMINIMISE, was written to carry out the above numerical work. The NAG routine C05NBF was used to solve the equations; this routine uses a modification of the Powell hybrid method for finding a zero of a system of nonlinear equations. Evaluation of the standard normal distribution function was carried out using the routine S15ABF. Lastly, the double integrals for Model I were evaluated using the routines D01BBF, D01FBF, D01BAZ and D01BAX, which use Gaussian quadrature formulae. Some results from this program are now given.

3.1.4 Numerical results

Table 3.1 gives optimal values for the parameters for the three models. [A grid

search similar to Colton's (1965) was considered prior to the minimisation in §3.1.3, and a comparison of results confirmed that the optimal values obtained using the latter are local minima, as required.] The results yield similar conclusions to Colton (1965, pp.177–178): Model II always requires a larger proportion of patients at the first stage than Model I; as expected, for Model II as δ/σ increases a larger proportion of patients are allocated in the second stage relative to the first stage; for both Models I and II, a similar effect is evident as N increases, the effect being more pronounced for Model II.

Table 3.1: *Optimal values for parameters for the three models*

δ/σ	N	Fixed	Model I		D	Model II	
		P ₁	P ₁	P ₂		P ₁	P ₂
0.5	100	0.116	0.080	0.091	0.371	0.096	0.214
	200	0.091	0.060	0.074	0.322	0.073	0.193
	500	0.058	0.036	0.049	0.290	0.044	0.152
	1000	0.038	0.023	0.032	0.279	0.028	0.119
1.5	100	0.041	0.025	0.034	0.841	0.030	0.124
	200	0.026	0.015	0.022	0.823	0.018	0.092
	500	0.013	0.007	0.011	0.815	0.009	0.059
	1000	0.008	0.004	0.006	0.814	0.005	0.041
2.5	100	0.020	0.012	0.017	1.364	0.014	0.079
	200	0.012	0.007	0.010	1.357	0.008	0.056
	500	0.006	0.003	0.005	1.359	0.004	0.034
	1000	0.003	0.002	0.003	1.365	0.002	0.023

3.2 Estimation for a stable normal response variable

We wish to estimate the true mean treatment difference, δ , by the overall estimate

$$\bar{d}^* = \frac{\sum_A x}{n_A} - \frac{\sum_B x}{n_B} ,$$

where $\sum_i x$ denotes the sum of the observed responses for treatment i and n_i is the number of patients on treatment i upon termination of the trial. [This approach is different to that taken in Chapter 2 where we split the trial into a number of blocks and based an estimate on the differences *within* blocks. For the present setting, it might seem more reasonable to estimate the difference using only the randomised stages, but this can lead to a substantial loss of information due to the randomised stages being relatively

small (cf. Table 3.1). Further, the stages are essentially analogous to the blocks in Chapter 2 but now, instead of analysing the difference within these stages, a decision is being made at the end of each stage which, as we shall see, often leads to an *optional stopping effect*. The results in Appendix 5 (which are discussed in Chapter 5, pp.158–159) for the Fixed Model when a data-dependent allocation rule is used in the first stage provide a useful supplement to both the present results and to those in Chapter 2.]

In this section we assume that the response variable is stable. Two methods are suggested to obtain the bias and variance of $\bar{d}^*$. The first method involves integrating $\bar{d}^*$ over the joint distributions of linear combinations of the estimated treatment means at the different stages; the ranges of integration are of course affected by the stopping rules [for the Fixed Model, the joint distributions are bivariate normal whereas for Models I and II, we need to consider the trivariate and even the quadrivariate normal]. For all the models, we can reduce the number of dimensions of integration by using standard results for conditional normal distributions (e.g. Anderson, 1984, Ch.2).

The second method, which was suggested by M.H. Gail (in a private communication, 1988), involves no integration; it makes use of the independence property of the sum and difference of two ^{independent} normal random variables with equal variances. Although in this section we only apply this method to the Fixed Model, it will be seen in §3.8 that some useful results for non-normal responses can also be obtained.

a) Fixed Model

We may write

$$\bar{d}^* = \left\{ \bar{d}_2 + \frac{(1-2p_1)}{(1-p_1)} \bar{x}_A' \right\} Y + \left\{ \bar{d}_3 - \frac{(1-2p_1)}{(1-p_1)} \bar{x}_B' \right\} (1-Y) \quad (3.7)$$

with

$$\bar{d}_2 = \frac{p_1 \bar{x}_{1A} - (1-p_1) \bar{x}_{1B}}{(1-p_1)}, \quad \bar{d}_3 = \frac{(1-p_1) \bar{x}_{1A} - p_1 \bar{x}_{1B}}{(1-p_1)},$$

where the $\bar{x}_{1i}$ are defined as before, $\bar{x}_i'$ is the mean of the responses from the $N(1-2p_1)$ patients after the better treatment has been chosen ($i=A,B$) and Y is an indicator variable with value one when $\bar{d}_1 > 0$, and zero otherwise.

We can derive the expectation of $\bar{d}^*$ by integrating (3.7) over the joint distributions

of the $\bar{d}_i$, ($i \in [1,3]$). Since we are dealing with linear combinations of normal random variables, these joint distributions are bivariate normal. If we define $\mu_i = E(\bar{d}_i)$, $\sigma_i^2 = \text{Var}(\bar{d}_i)$ and $\rho_{ij} = \text{Corr}(\bar{d}_i, \bar{d}_j)$, then the first term of $E(\bar{d}^*)$ is

$$E(\bar{d}_2 Y) = \int_{-\infty}^{\infty} \int_0^{\infty} \bar{d}_2 f_2(\bar{d}_1, \bar{d}_2) d\bar{d}_1 d\bar{d}_2 ,$$

where $f_2(\bar{d}_1, \bar{d}_2)$ is the bivariate normal density function for $\bar{d}_1$ and $\bar{d}_2$, and hence

$$E(\bar{d}_2 Y) = \int_0^{\infty} f(\bar{d}_1) \left\{ \mu_2 + \rho_{12} \frac{\sigma_2}{\sigma_1} (\bar{d}_1 - \mu_1) \right\} d\bar{d}_1$$

in which $f(\bar{d}_1)$ is the normal density function for $\bar{d}_1$. Furthermore, by simple integration

$$\int_0^{\infty} \bar{d}_1 f(\bar{d}_1) d\bar{d}_1 = \sigma_1 \varphi(h) + \delta [1 - \Phi(h)] \quad (3.8)$$

since $\mu_1 = \delta$. Thus,

$$E(\bar{d}_2 Y) = \mu_2 [1 - \Phi(h)] + \rho_{12} \sigma_2 \varphi(h) \quad (3.9)$$

and similarly,

$$E(\bar{d}_3 (1-Y)) = \mu_3 \Phi(h) - \rho_{13} \sigma_3 \varphi(h). \quad (3.10)$$

Lastly, since $\rho_{12} = \rho_{13}$, we can write

$$E(\bar{x}_A' Y) = \mu_A [1 - \Phi(h)] , \quad E(\bar{x}_B' (1-Y)) = \mu_B \Phi(h)$$

and simple algebra yields

$$E(\bar{d}^*) = \delta . \quad (3.11)$$

To derive the variance of (3.7), we firstly write $\bar{d}^{*2}$ as

$$\begin{aligned} \bar{d}^{*2} = & \left\{ \bar{d}_2^2 + 2\bar{d}_2 \frac{(1-2p_1)}{(1-p_1)} \bar{x}_A' + \frac{(1-2p_1)^2}{(1-p_1)^2} \bar{x}_A'^2 \right\} Y^2 \\ & + \left\{ \bar{d}_3^2 - 2\bar{d}_3 \frac{(1-2p_1)}{(1-p_1)} \bar{x}_B' + \frac{(1-2p_1)^2}{(1-p_1)^2} \bar{x}_B'^2 \right\} (1-Y)^2 . \end{aligned} \quad (3.12)$$

Considering the first term in (3.12), we have

$$E(\bar{d}_2^2 Y^2) = \int_0^{\infty} f(\bar{d}_1) \left\{ \sigma_2^2 (1 - \rho_{12}^2) + [\mu_2 + \rho_{12} \frac{\sigma_2}{\sigma_1} (\bar{d}_1 - \delta)]^2 \right\} d\bar{d}_1 , \quad (3.13)$$

where a simple integration by parts gives

$$\int_0^{\infty} \bar{d}_1^2 f(\bar{d}_1) d\bar{d}_1 = (\sigma_1^2 + \delta^2) [1 - \Phi(h)] + \sigma_1 \delta \varphi(h) . \quad (3.14)$$

We can obtain a similar expression to (3.13) for $E\{\bar{d}_3^2(1-Y)^2\}$. Also, we have

$$E(\bar{x}_A'^2 Y^2) = \mu_A^2 [1 - \Phi(h)] \quad , \quad E\{\bar{x}_B'^2 (1-Y)^2\} = \mu_B^2 \Phi(h) \quad ,$$

and thus using the above results, we obtain

$$\text{Var}(\bar{d}^*) = \frac{\sigma^2}{Np_1(1-p_1)} \quad . \quad (3.15)$$

Alternative method

If we define

$$D = Np_1 \bar{d}_1 \quad , \quad S = Np_1 (\bar{x}_{1A} + \bar{x}_{1B})$$

and

$$X_2 = N(1-2p_1) \bar{x}_{1A} \quad , \quad Y_2 = N(1-2p_1) \bar{x}_{1B} \quad ,$$

then we may write our estimate of δ as

$$\bar{d}^* = \frac{[\frac{1}{2}(S+D) + Y X_2]}{N[p_1 + (1-2p_1)Y]} - \frac{[\frac{1}{2}(S-D) + (1-Y) Y_2]}{N[p_1 + (1-2p_1)(1-Y)]} \quad ,$$

where Y is defined as in (3.7). Rearranging yields

$$\bar{d}^* = \frac{S(1-2p_1)(1-2Y) + D + 2p_1[Y(X_2 + Y_2) - Y_2]}{2Np_1(1-p_1)} \quad .$$

Taking expectations and using the fact that the random variables D and S are independent, we obtain (3.11). Using a similar approach, we can also obtain (3.15).

Comments

(i) Note that (3.15) is the usual variance for the sum and difference of two independent sample means based on Np_1 and $N(1-p_1)$ observations. Our result might seem intuitively plausible in view of the symmetry of the stopping rule.

(ii) The alternative method outlined above has only been applied to the Fixed Model – for Models I and II, a generalisation is possible but it was considered unnecessary since the integration method appeared to be more readily adaptable for an unstable response variable.

b) Model I

We may write

$$\bar{d}^* = \left\{ \bar{d}_3 + \frac{(1-2p_1)}{(1-p_1)} \bar{x}_{1A}' \right\} Y_1 + \left\{ \bar{d}_4 - \frac{(1-2p_1)}{(1-p_1)} \bar{x}_{1B}' \right\} Y_2$$

$$\begin{aligned}
& + \left\{ \bar{d}_5 + \bar{d}_6 + \frac{(1-2p_1-2p_2)}{(1-p_1-p_2)} \bar{x}_A' \right\} Y_3 \\
& + \left\{ \bar{d}_7 + \bar{d}_8 - \frac{(1-2p_1-2p_2)}{(1-p_1-p_2)} \bar{x}_B' \right\} (1-Y_1-Y_2-Y_3), \quad (3.16)
\end{aligned}$$

where $\bar{d}_3$ and $\bar{d}_4$ are $\bar{d}_2$ and $\bar{d}_3$, respectively, for the Fixed Model,

$$\bar{d}_5 = \frac{p_1 [(p_1+p_2) \bar{x}_{1A} - (1-p_1-p_2) \bar{x}_{1B}]}{(1-p_1-p_2)(p_1+p_2)}$$

(= $-\bar{d}_7$ with $\bar{x}_{1A}$ and $\bar{x}_{1B}$ interchanged),

$$\bar{d}_6 = \frac{p_2 [(p_1+p_2) \bar{x}_{2A} - (1-p_1-p_2) \bar{x}_{2B}]}{(1-p_1-p_2)(p_1+p_2)}$$

(= $-\bar{d}_8$ with $\bar{x}_{2A}$ and $\bar{x}_{2B}$ interchanged), and the Y_i ($i \in [1,3]$) are the indicator variables

$$(Y_1, Y_2, Y_3) = \begin{cases} (1, 0, 0) & , \quad \bar{d}_1 > D \\ (0, 1, 0) & , \quad \bar{d}_1 < -D \\ (0, 0, 1) & , \quad |\bar{d}_1| \leq D, \quad y > 0 \\ (1, 1, 1) & , \quad |\bar{d}_1| \leq D, \quad y < 0 \end{cases}$$

with y as defined in §3.1.2.

The notation used previously will be extended in an obvious way. To derive $E(\bar{d}^*)$ and $\text{Var}(\bar{d}^*)$ for Model I (and Model II), we firstly require three results for conditional normal distributions: if we let $X_i \sim N(\mu_i, \sigma_i^2)$ and denote by $\rho_{12.34}$ the partial correlation between X_1 and X_2 , keeping X_3 and X_4 fixed, then

$$(i) \quad E\{X_1 | X_2 = x_2, X_3 = x_3\} = \mu_1 + \beta_{12.3} (x_2 - \mu_2) + \beta_{13.2} (x_3 - \mu_3), \quad (3.17)$$

where

$$\beta_{ij.k} = \frac{\sigma_i}{\sigma_j} \rho_{ij.k} \sqrt{\frac{1-\rho_{ik}^2}{1-\rho_{jk}^2}}, \quad (i, j, k \in [1, 4]);$$

$$(ii) \quad \text{Var}\{X_1 | X_2 = x_2, X_3 = x_3\} = \frac{\sigma_1^2 \{1 - \rho_{12}^2 - \rho_{13}^2 - \rho_{23}^2 + 2\rho_{12}\rho_{13}\rho_{23}\}}{(1-\rho_{23}^2)}; \quad (3.18)$$

$$(iii) \quad \rho_{12.34} = \frac{\rho_{12.4} - \rho_{13.4}\rho_{23.4}}{\sqrt{(1-\rho_{13.4}^2)(1-\rho_{23.4}^2)}} \quad (3.19)$$

(e.g. Anderson, 1984, p.34).

We begin by considering the first term in (3.16). In the same way as (3.9),

$$\begin{aligned}
E(\bar{d}_3 Y_1) &= \int_{-\infty}^{\infty} \int_0^{\infty} \bar{d}_3 f_2(\bar{d}_1, \bar{d}_3) d\bar{d}_1 d\bar{d}_3 \\
&= \mu_3 [1 - \Phi(h_U)] + \rho_{13} \sigma_3 \varphi(h_U) .
\end{aligned}$$

Similarly, we obtain

$$E(\bar{d}_4 Y_2) = \mu_4 \Phi(h_L) - \rho_{14} \sigma_4 \varphi(h_L) .$$

Next, we have

$$E(\bar{d}_5 Y_3) = \int_{-\infty}^{\infty} \int_0^{\infty} \int_{-D}^D \bar{d}_5 f_3(\bar{d}_1, y, \bar{d}_5) d\bar{d}_1 dy d\bar{d}_5 ,$$

where $f_3(\bar{d}_1, y, \bar{d}_5)$ is the trivariate normal density function of $\bar{d}_1$, y and $\bar{d}_5$, and hence

$$E(\bar{d}_5 Y_3) = \int_0^{\infty} \int_{-D}^D E\{\bar{d}_5 / \bar{d}_1, y\} f_2(\bar{d}_1, y) d\bar{d}_1 dy .$$

It is straightforward to show that $\rho_{5y} = \rho_{1y}\rho_{15}$, and so using (3.17),

$$E(\bar{d}_5 Y_3) = \mu_5 I_{11} + \beta_{51 \cdot y} \sigma_1 I_{21} , \quad (3.20)$$

where

$$I_{11} = \int_k^{\infty} \int_{h_L}^{h_U} \varphi_2(u, v) du dv , \quad I_{21} = \int_k^{\infty} \int_{h_L}^{h_U} u \varphi_2(u, v) du dv ,$$

with a similar expression for $E\{\bar{d}_7(1-Y_1-Y_2-Y_3)\}$.

Clearly, $\rho_{16} = 0$, and therefore

$$E(\bar{d}_6 Y_3) = \mu_6 I_{11} - \beta_{61 \cdot y} \sigma_1 I_{21} + \frac{\beta_{61 \cdot y}}{\rho_{1y}} \sigma_y I_{31} , \quad (3.21)$$

where

$$I_{31} = \int_k^{\infty} \int_{h_L}^{h_U} v \varphi_2(u, v) du dv ,$$

with a similar expression for $E\{\bar{d}_8(1-Y_1-Y_2-Y_3)\}$.

Therefore, since the $\bar{X}_{1i}'$ and the $\bar{X}_i'$ ($i=A, B$) are both independent of $\bar{d}_1$ and y , and using (3.20) and (3.21), we obtain the result

$$E(\bar{d}^*) = \delta + \frac{1}{2} \frac{\sigma_1 p_2 (1 - 2p_1 - p_2)}{(p_1 + p_2)(1 - p_1 - p_2)(1 - p_1)} [\varphi(h_U) - \varphi(h_L)] . \quad (3.22)$$

To derive $\text{Var}(\bar{d}^*)$, we proceed as for the Fixed Model. The first term of $E(\bar{d}^{*2})$, $E(\bar{d}_3^2 Y_1^2)$, is similar to (3.13) (as is $E(\bar{d}_4^2 Y_2^2)$). Next we consider terms like

$$E(\bar{d}_5^2 Y_3^2) = \int_0^\infty \int_{-D}^D \left\{ \text{Var}\{\bar{d}_5 | \bar{d}_1, y\} + \{E(\bar{d}_5 | \bar{d}_1, y)\}^2 \right\} f_2(\bar{d}_1, y) d\bar{d}_1 dy .$$

Clearly, using (3.17) and (3.18), this can be expressed in terms of double integrals such as I_{11} and I_{21} as defined for (3.20). With a few modifications, the other terms of this type can be treated in the same way. The remaining terms of $E(\bar{d}^{*2})$ are cross-products such as

$$E(\bar{d}_5 \bar{d}_6 Y_3^2) = \int_{-\infty}^\infty \int_{-\infty}^\infty \int_0^\infty \int_{-D}^D \bar{d}_5 \bar{d}_6 f_4(\bar{d}_1, y, \bar{d}_5, \bar{d}_6) d\bar{d}_1 dy d\bar{d}_5 d\bar{d}_6 ,$$

where $f_4(\bar{d}_1, y, \bar{d}_5, \bar{d}_6)$ is the quadrivariate normal density function of $\bar{d}_1$, y , $\bar{d}_5$ and $\bar{d}_6$, and hence

$$E(\bar{d}_5 \bar{d}_6 Y_3^2) = \int_0^\infty \int_{-D}^D \left\{ \text{Cov}(\bar{d}_5, \bar{d}_6 | \bar{d}_1, y) + E(\bar{d}_5 | \bar{d}_1, y) E(\bar{d}_6 | \bar{d}_1, y) \right\} f_2(\bar{d}_1, y) d\bar{d}_1 dy ,$$

where it can be shown that

$$\text{Cov}(\bar{d}_5, \bar{d}_6 | \bar{d}_1, y) = - \rho_{6y} \sigma_6 \sigma_y \beta_{5y.1}$$

using (3.19), and similarly for $E\{\bar{d}_7 \bar{d}_8 (1 - Y_1 - Y_2 - Y_3)\}$.

Obtaining an explicit expression for $E(\bar{d}^{*2})$, and hence for $\text{Var}(\bar{d}^*)$, seemed formidable in terms of algebra. Some simplifications were carried out – however, it was not considered necessary to include the details here. It will be seen in §3.5 that more emphasis is placed on a graphical presentation of the results for $E(\bar{d}^*)$ and $\text{Var}(\bar{d}^*)$.

c) Model II

We may write

$$\begin{aligned} \bar{d}^* = & \left\{ \bar{d}_3 + \frac{p_2}{(1-p_1)} \bar{x}_{2A} + \frac{(1-2p_1-p_2)}{(1-p_1)} \bar{x}_A' \right\} Y_1 \\ & + \left\{ \bar{d}_4 + \frac{p_2}{(p_1+p_2)} \bar{x}_{2A} - \frac{(1-2p_1-p_2)}{(1-p_1-p_2)} \bar{x}_B' \right\} Y_2 \\ & + \left\{ \bar{d}_5 - \frac{p_2}{(p_1+p_2)} \bar{x}_{2B} + \frac{(1-2p_1-p_2)}{(1-p_1-p_2)} \bar{x}_A' \right\} Y_3 \\ & + \left\{ \bar{d}_6 - \frac{p_2}{(1-p_1)} \bar{x}_{2B} - \frac{(1-2p_1-p_2)}{(1-p_1-p_2)} \bar{x}_B' \right\} (1 - Y_1 - Y_2 - Y_3) , \end{aligned} \quad (3.23)$$

where $\bar{d}_3$, $\bar{d}_4$, $\bar{d}_5$ and $\bar{d}_6$ are $\bar{d}_3$, $\bar{d}_7$, $\bar{d}_5$ and $\bar{d}_4$, respectively, for Model I and the Y_i ($i \in [1,3]$) are the indicator variables

$$(Y_1, Y_2, Y_3) = \begin{cases} (1, 0, 0), & \bar{d}_1 > 0, y_A > 0 \\ (0, 1, 0), & \bar{d}_1 > 0, y_A < 0 \\ (0, 0, 1), & \bar{d}_1 < 0, y_B > 0 \\ (0, 0, 0), & \bar{d}_1 < 0, y_B < 0 \end{cases}$$

with y_A and y_B as defined in §3.1.2.

We can proceed as for Model I. For instance, the first term of $E(\bar{d}^*)$ is

$$E(\bar{d}_3 Y_1) = \mu_3 I_{11} + \sigma_1 \beta_{31.y} I_{21} + \sigma_y \beta_{3y.1} I_{31} \quad (3.24)$$

using (3.17), where

$$I_{11} = \int_k^\infty \int_h^\infty \varphi_2(u, v) du dv, \quad I_{21} = \int_k^\infty \int_h^\infty u \varphi_2(u, v) du dv$$

and

$$I_{31} = \int_k^\infty \int_h^\infty v \varphi_2(u, v) du dv.$$

Similar expressions can also be obtained for $E(\bar{d}_4 Y_2)$, $E(\bar{d}_5 Y_3)$ and $E\{\bar{d}_6(1-Y_1-Y_2-Y_3)\}$.

In addition, we have terms like

$$E(\bar{X}_{2A} Y_1) = \mu_A I_{11} + \sigma_1 \beta_{x_2, 1.y} I_{21} + \sigma_y \beta_{x_2, y.1} I_{31} \quad (3.25)$$

in which

$$\beta_{x_2, y.1} = \beta_{x_2, 1.y} / \rho_{1y}^2$$

since $\rho_{x_2, 1} = 0$.

Using (3.24) and (3.25) and the fact that both $\bar{d}_1$ and y_A (and y_B) are independent of the $\bar{x}_i$ ($i=A, B$), we obtain the result

$$\begin{aligned} E(\bar{d}^*) &= \delta + \sigma_1 [A\beta_{x_2, 1.y} + \beta_{31.y} - \beta_{41.y}] [I_{21} + I_{22}] \\ &\quad + \sigma_y [A\beta_{x_2, y.1} + \beta_{3y.1} - \beta_{4y.1}] [I_{31} + I_{32}], \end{aligned} \quad (3.26)$$

where

$$A = - \frac{p_2(1-2p_1-p_2)}{(1-p_1)(p_1+p_2)}, \quad I_{22} = \int_{-\infty}^k \int_{-\infty}^h u \varphi_2(u, v) du dv, \text{ etc.}$$

Considering $\text{Var}(\bar{d}^*)$, the approach is as for Model I. For instance, there are a number of cross-product terms such as

$$E(\bar{d}_3 \bar{X}_{2A} Y_1^2) = \int_0^\infty \int_0^\infty \left\{ \text{Cov}(\bar{d}_3, \bar{X}_{2A} | \bar{d}_1, y_A) + E(\bar{d}_3 | \bar{d}_1, y_A) E(\bar{X}_{2A} | \bar{d}_1, y_A) \right\} f_2(\bar{d}_1, y_A) d\bar{d}_1 dy$$

where it can be shown that

$$\text{Cov}(\bar{d}_3, \bar{X}_{2A} | \bar{d}_1, y_A) = -\sigma_{x_2} \sigma_y \rho_{x_2, y} \beta_{3y.1}$$

using (3.19). Again, some simplification of the terms in $\text{Var}(\bar{d}^*)$ have been carried out.

Computation

A FORTRAN program, STANDEFFICCOLT, was written to incorporate the above results together with more general ones which will be derived in §3.4. The NAG routines for numerical integration described in §3.1 are again implemented. Results from this program will be presented and discussed in §3.5.

3.3 Modelling instability in the response variable

As in Chapter 2, we consider two types of instability: (i) linear time-trends, and (ii) fluctuations. Here we show how to model (i) and (ii) for normal responses and an indication of their anticipated effect is given.

3.3.1 Linear time-trends

If we denote the response variable for treatment i at time j by X_{ij} ($i=A,B$; $j \in [0, N-1]$), then until now we have assumed that $X_{ij} \sim N(\mu_i, \sigma^2) \forall j$. Suppose instead that the response variable is subject to a linear time-trend, then clearly we can model this by assuming that $X_{ij} \sim N(\mu_i + tj, \sigma^2)$ where t is the linear trend per unit time (the same for both treatments).

It is clear that one of the effects of a linear time-trend is to bias the $\bar{d}^*$ for the Fixed Model and to produce a larger bias for Models I and II than for a stable response variable.

3.3.2 Fluctuations

A simple model for fluctuations in the response variable is the stable model with a first-order autoregressive process superimposed. If we also include a linear trend and denote the new response variable by Y_{ij} , then we can write

$$Y_{ij} = X_{ij} + Z_j$$

with X_{ij} as in §3.3.1 and

$$Z_j = \begin{cases} \epsilon_0 & , \quad (j = 0) \\ \rho Z_{j-1} + \epsilon_j & , \quad (j \geq 1; |\rho| < 1) \end{cases} ,$$

where

$$\epsilon_j \sim \begin{cases} N(0, \sigma_\epsilon^2 / (1 - \rho^2)) & , \quad (j = 0) \\ N(0, \sigma_\epsilon^2) & , \quad (j \geq 1) \end{cases}$$

so that

$$Y_{ij} \sim N(\mu_i + t_j, \sigma^2 + \sigma_z^2) \quad , \quad \sigma_z^2 = \sigma_\epsilon^2 / (1 - \rho^2) .$$

Intuitively, if there are fluctuations, we might expect the following: firstly, the bias in $\bar{d}^*$ for the Fixed Model and for Model I are less affected by ρ than Model II; secondly, the variance of $\bar{d}^*$ for all three models increases with ρ , and particularly as $|\rho| \rightarrow 1$; thirdly, derivation of the biases and variances of the $\bar{d}^*$ proves much more cumbersome due to the new correlation structure.

3.4 Estimation for an unstable normal response variable

The approach taken here is the same as in §3.2. Firstly, we consider linear time-trends – only the expectations of the sample means need to be modified. As might be expected, these modifications involve terms which are proportional to arithmetic progressions. Following this, we consider fluctuations – here we need to modify the variances of the sample means. For this case, the modifications involve geometric progressions in ρ and are less straightforward. In fact, exact results for these are difficult to obtain unless certain assumptions are made concerning the randomisation procedure in the trial; to avoid these assumptions, approximations are given which one would expect to be close to the exact results (this seems to be confirmed by the comparison with simulated values in §3.5).

3.4.1 Linear time-trends

a) Fixed Model

By symmetry, $\mu_1 (= E(\bar{d}_1))$ is unaffected by linear time-trends so that $\mu_1 = \delta$. However, both μ_2 and μ_3 are affected. Since the trend is linear, we have an arithmetic progression; thus, taking into account the randomised nature of the first stage, it can be shown that μ_2 has the extra term $-b_1$, and μ_3 the term b_1 (again, by symmetry), where

$$b_1 = \frac{1}{2}t \frac{(2Np_1 - 1)(1 - 2p_1)}{(1 - p_1)} . \quad (3.27)$$

Also,

$$E(\bar{X}_i') = \mu_i + b_4 , \quad b_4 = \frac{1}{2}t [N(1 + 2p_1) - 1] , \quad (i=A, B) . \quad (3.28)$$

Using (3.27) and (3.28) and some simple algebra, we obtain

$$E(\bar{d}^*) = \delta + \frac{1}{2} A^* [1 - 2\Phi(h)] \quad (3.29)$$

and

$$\text{Var}(\bar{d}^*) = \frac{\sigma^2}{Np_1(1 - p_1)} + A^* [2\rho_{12}\sigma_2\varphi(h) + A^*\Phi(h)\{1 - \Phi(h)\}] \quad (3.30)$$

in which

$$A^* = tN(1 - 2p_1)/(1 - p_1) .$$

b) Model I

The approach is the same as for the Fixed Model. Firstly, μ_3 and μ_4 have extra terms given by $-b_1$ and b_1 , respectively, with b_1 defined by (3.27). Similarly, μ_5 and μ_6 have extra terms $-b_2$ and b_2 , respectively, where

$$b_2 = \frac{1}{2}t \frac{p_1(2Np_1 - 1)(1 - 2p_1 - 2p_2)}{(1 - p_1 - p_2)(p_1 + p_2)} \quad (3.31)$$

and μ_6 and μ_8 have extra terms $-b_3$ and b_3 , respectively, with

$$b_3 = \frac{1}{2}t [2N(3p_1 + p_2) - 1](1 - 2p_1 - 2p_2)/(1 - p_1 - p_2) . \quad (3.32)$$

Next, the $E(\bar{X}_{li}')$ are given by (3.28) and, lastly, we have

$$E(\bar{X}_i') = \mu_i + b_5 , \quad b_5 = \frac{1}{2}t [N(1 + 2p_1 + 2p_2) - 1] , \quad (i=A, B) . \quad (3.33)$$

Thus, using (3.31), (3.32) and (3.33) and the method of §3.2, we obtain

$$E(\bar{d}^*) = E(\bar{d}^*)_s + A_1 [1 - \Phi(h_U) - \Phi(h_L)] + A_2 [\Phi(h_U) + \Phi(h_L) + 2I_{11}] , \quad (3.34)$$

where $E(\bar{d}^*)_s$ is given by (3.22), I_{11} is as defined in (3.20) and

$$A_1 = -b_1 + \frac{(1 - 2p_1)}{(1 - p_1)} b_4 , \quad A_2 = -b_2 - b_3 + \frac{(1 - 2p_1 - 2p_2)}{(1 - p_1 - p_2)} b_5 .$$

The terms of $E(\bar{d}^{*2})$ can be similarly modified and hence $\text{Var}(\bar{d}^*)$ obtained.

c) Model II

Corresponding to §3.2, the extra terms for $\bar{d}_3$, $\bar{d}_4$, $\bar{d}_5$ and $\bar{d}_6$ are given by $-b_1$, b_2 ,

$-b_2$ and b_1 , respectively. In addition to these, we have

$$E(\bar{X}_{2i}) = \mu_i + b_4, \quad b_4 = \frac{1}{2}t [N(4p_1 + p_2) - 1], \quad (3.35)$$

and

$$E(\bar{X}_i') = \mu_i + b_3, \quad b_3 = \frac{1}{2}t [N(1+2p_1+p_2) - 1], \quad (i=A, B). \quad (3.36)$$

Modification of the formulae in §3.2 accordingly and some manipulation yields

$$\begin{aligned} E(\bar{d}^*) &= \delta + \sigma_1 [A\beta_{x_2, 1.y} + \beta_{31.y} - \beta_{41.y}] [I_{21} + I_{22}] \\ &+ \sigma_y [A\beta_{x_2, y.1} + \beta_{3y.1} - \beta_{4y.1}] [I_{31} + I_{32}] \\ &+ [-z(A\beta_{x_2, y.1} + \beta_{3y.1} - \beta_{4y.1}) + A_1] [I_{11} - I_{12}] + A_2 [1 - 2\Phi(h)] \end{aligned} \quad (3.37)$$

[cf. (3.26)], where z is the bias of y_A in estimating δ and is given by

$$z = \frac{1}{2}t \frac{Np_2(2p_1 + p_2)}{(p_1 + p_2)};$$

for example,

$$I_{11} = \int_{k_A}^{\infty} \int_h^{\infty} \varphi_2(u, v) du dv, \quad I_{12} = \int_{-\infty}^{k_B} \int_{-\infty}^h \varphi_2(u, v) du dv$$

in which

$$k_A = -(\delta + z)/\sigma_y, \quad k_B = -(\delta - z)/\sigma_y;$$

and

$$A_1 = -b_1 - b_2 + Bb_3 + Ab_4, \quad A_2 = b_2 - \frac{(1 - 2p_1 - p_2)}{(1 - p_1 - p_2)} b_3 + \frac{p_2}{(p_1 + p_2)} b_4$$

with

$$B = \frac{(1 - 2p_1 - p_2)[2(1 - p_1) - p_2]}{(1 - p_1 - p_2)(1 - p_1)}.$$

The terms of $E(\bar{d}^{*2})$ can be modified in the same way to obtain $\text{Var}(\bar{d}^*)$.

3.4.2 Fluctuations

a) Fixed Model

We begin by considering $\text{Var}(\bar{X}_{ij})$ and $\text{Cov}(\bar{X}_{1A}, \bar{X}_{1B})$, and thus $\text{Var}(\bar{d}_1)$. Since $\bar{X}_{ij}$ is based on the responses of a random sample of Np_1 from $2Np_1$ patients, we may write

$$\text{Var}(\bar{X}_{1i}) = \frac{1}{Np_1} \left\{ \sigma^2 + \sigma_z^2 + \frac{1}{Np_1} \sum_{\Omega} I_{\{\omega_{\ell m} \in \Omega_i\}} \omega_{\ell m} \right\}, \quad (3.38)$$

where Ω denotes all $2Np_1(2Np_1 - 1)$ covariances (if there were $2Np_1$ responses from

treatment i), Ω_i the $Np_1(Np_1-1)$ covariances of the Z_j for treatment i and $\omega_{\ell m} = \text{Cov}(Z_\ell, Z_m)$, ($\ell, m \in [0, 2Np_1-1]$). [N.B. It is assumed that p_1 has been rounded to an appropriate number of decimal places.]

From standard results for a first-order autoregressive process,

$$\omega_{\ell m} = \rho^{|\ell-m|} \sigma_z^2.$$

Therefore, we may write

$$\begin{aligned} \sum_{\Omega} \omega_{\ell m} &= 2\sigma_z^2 \rho \sum_{i=1}^{2Np_1-1} \sum_{j=0}^{i-1} \rho^j \\ &= 2\sigma_z^2 \rho^{2Np_1-1} \frac{d}{d(1/\rho)} \left\{ \frac{1-(1/\rho)^{2Np_1}}{1-(1/\rho)} \right\} \\ &= 2\sigma_z^2 \rho [2Np_1(1-\rho) - (1-\rho^{2Np_1})] / (1-\rho)^2. \end{aligned}$$

Since each of the covariances has the same probability, $\frac{1}{2}(Np_1-1)/(2Np_1-1)$, of being one of those associated with treatment i , a "good" approximation to (3.38) would be

$$\text{Var}(\bar{X}_{1i}) \approx \frac{1}{Np_1} \left[\sigma^2 + \sigma_z^2 \left\{ 1 + \frac{(Np_1-1)\rho}{Np_1(2Np_1-1)(1-\rho)^2} [2Np_1(1-\rho) - (1-\rho^{2Np_1})] \right\} \right]. \quad (3.39)$$

out

[As was pointed out earlier, an exact expression for (3.38) is difficult unless we make certain assumptions about the randomisation procedure; one such assumption might be that the patients are randomised in pairs. Although this simplifies the problem, the approximation given by (3.39) was preferred for two reasons: (a) the resulting algebra was still reasonably manageable for the Fixed Model (in obtaining an explicit expression for $\text{Var}(\bar{d}^*)$); (b) the potential improvement in the results obtained using paired observations seemed to be outweighed by the extra algebra.]

Using a similar argument, we can write

$$\text{Cov}(\bar{X}_{1A}, \bar{X}_{1B}) \approx \frac{\sigma_z^2 \rho}{Np_1(2Np_1-1)(1-\rho)^2} [2Np_1(1-\rho) - (1-\rho^{2Np_1})] \quad (3.40)$$

and hence

$$\text{Var}(\bar{d}_1) \approx \frac{2}{Np_1} \left[\sigma^2 + \sigma_z^2 \left\{ 1 - \frac{\rho}{Np_1(2Np_1-1)(1-\rho)^2} [2Np_1(1-\rho) - (1-\rho^{2Np_1})] \right\} \right]. \quad (3.41)$$

Thus, although $\bar{d}^*$ remains unbiased in the absence of a linear time-trend, in the presence of a trend the bias in $\bar{d}^*$ is evident from (3.29), in which h depends on σ_1^2 as given by (3.41).

When we consider $E(\bar{d}^{*2})$, several extra covariance terms need to be included. For example, the $\bar{X}_i'$ ($i=A,B$) are now correlated with $\bar{d}_2$ and $\bar{d}_3$ and, in particular,

$$\text{Cov}(\bar{d}_2, \bar{X}_A') \approx - \frac{\sigma_z^2 \rho}{2[N(1-\rho)]^2 p_1 (1-p_1)} [1-\rho^{2Np_1}] [1-\rho^{N(1-2p_1)}] \quad (3.42)$$

using the product of two geometric progressions. In addition, using (3.39) we have

$$\text{Var}(\bar{X}_i') = \frac{1}{N(1-2p_1)} \left[\sigma^2 + \sigma_z^2 \left\{ 1 + \frac{2\rho}{N(1-2p_1)(1-\rho)^2} [N(1-2p_1)(1-\rho) - (1-\rho^{N(1-2p_1)})] \right\} \right] \quad (3.43)$$

After some further modifications in σ_2^2 and ρ_{12} , the following result was obtained

$$\begin{aligned} \text{Var}(\bar{d}^*) \approx & \text{Var}(\bar{d}^*)_t + \frac{\sigma_z^2}{[N(1-p_1)]^2 p_1 (1-\rho)} \left[-2\rho^{2Np_1+1} \{1-\rho^{N(1-2p_1)}\} p_1 / (1-\rho) \right. \\ & - \frac{(1-\rho^{2Np_1})\rho}{(1-\rho)} \left\{ \frac{1}{(2Np_1-1)} [N(1-2p_1)^2 - \{p_1^2 + (1-p_1)^2\} / p_1] + \{1-\rho^{N(1-2p_1)}\} \right\} \\ & \left. + N(1-2p_1)p_1 [(1+\rho) + 2N(1-2p_1)\rho / (2Np_1-1)] + N[p_1^2 + (1-p_1)^2] [(1-\rho) - 2\rho / (2Np_1-1)] \right] \end{aligned} \quad (3.44)$$

where $\text{Var}(\bar{d}^*)_t$ is given by (3.30).

b) Model I

The approach is as for the Fixed Model; the bias in $\bar{d}^*$ for a stable response variable is evident from (3.22) with σ_1^2 given by (3.41), and in the presence of a linear time-trend from (3.34) with σ_y^2 given by

$$\begin{aligned} \text{Var}(y) \approx & \frac{2}{N(p_1+p_2)} \left[\sigma^2 + \sigma_z^2 \left\{ 1 - \frac{\rho}{N(p_1+p_2)(1-\rho)^2} [2 - \rho(1-\rho^{2Np_1-1}) / (2Np_1-1) \right. \right. \\ & \left. \left. - \rho(1-\rho^{2Np_2-1}) / (2Np_2-1) \right] \right\} \right] \end{aligned} \quad (3.45)$$

using (3.41).

When we consider $E(\bar{d}^{*2})$, several extra covariance terms need to be considered: firstly, $\bar{d}_3$ is correlated with the $\bar{X}_{ii}'$ ($i=A,B$) with $\text{Cov}(\bar{d}_3, \bar{X}_{ii}')$ given by (3.42); secondly,

$\text{Var}(\bar{X}_{1i}')$ is modified using (3.43) (with a similar expression for $\text{Var}(\bar{X}_i')$); thirdly, $\bar{X}_A'$ is now correlated with $\bar{d}_5$ and $\bar{d}_6$ and, in particular,

$$\text{Cov}(\bar{X}_A', \bar{d}_5) \approx -\frac{1}{2} \frac{\sigma_z^2 \rho^{2Np_2+1}}{[N(1-\rho)]^2 (1-p_1-p_2)(p_1+p_2)} (1-\rho^{2Np_1}) [1-\rho^{N(1-2p_1-2p_2)}]$$

and

$$\text{Cov}(\bar{X}_B', \bar{d}_6) \approx -\frac{1}{2} \frac{\sigma_z^2 \rho^{2Np_1+1}}{[N(1-\rho)]^2 (1-p_1-p_2)(p_1+p_2)} (1-\rho^{2Np_2}) [1-\rho^{N(1-2p_1-2p_2)}];$$

lastly, it can be shown that

$$\text{Cov}(\bar{d}_5, \bar{d}_6 | \bar{d}_1, y) = \text{Cov}(\bar{d}_5, \bar{d}_6) - \rho_{6y} \sigma_6 \sigma_y \beta_{5y.1},$$

where

$$\text{Cov}(\bar{d}_5, \bar{d}_6) \approx \frac{(1-2p_1-2p_2)^2 \sigma_z^2 \rho}{[2N(p_1+p_2)(1-\rho)]^2 (1-p_1-p_2)} (1-\rho^{2Np_1})(1-\rho^{2Np_2}).$$

c) Model II

Unlike the previous two models, the bias for Model II is affected by its imbalanced second stage. For example, one of the main modifications to $E(\bar{d}^*)$ centres on the four terms like $E(\bar{X}_A' Y_1)$ since $\bar{X}_A'$ and y_A are now correlated. In fact, using (3.17) we obtain

$$E(\bar{X}_A' Y_1) = \mu_{x_A} I_{11} + \sigma_1 \beta_{x_A, 1.y} I_{21} + \sigma_y \beta_{x_A, y.1} I_{31}, \quad (3.46)$$

where

$$\beta_{x_A, y.1} = \beta_{x_A, 1.y} \sigma_1 / (\rho_{1y} \sigma_y)$$

since $\rho_{xA,1} = 0$.

It is not difficult to show that, with these modifications, we have

$$\begin{aligned} E(\bar{d}^*) &= \delta + \sigma_1 [A\beta_{x_2, 1.y} + B\beta_{x_A, 1.y} + \beta_{31.y} - \beta_{41.y}] [I_{21} + I_{22}] \\ &\quad + \sigma_y [A\beta_{x_2, y.1} + B\beta_{x_A, y.1} + \beta_{3y.1} - \beta_{4y.1}] [I_{31} + I_{32}] \\ &\quad + [-Z\{A\beta_{x_2, y.1} + B\beta_{x_A, y.1} + \beta_{3y.1} - \beta_{4y.1}\} + A_1] [I_{11} - I_{12}] + A_2 [1 - 2\Phi(h)] \end{aligned} \quad (3.47)$$

[cf. (3.37)], in which σ_1^2 is given approximately by (3.41), etc.

The above modifications are extended in an obvious way for $E(\bar{d}^{*2})$. In addition to these, we have: firstly, that the $\bar{X}_{2i}$ and the $\bar{X}_i'$ ($i=A, B$) are now both correlated with the $\bar{d}_i$ ($i \in [3, 6]$); for the former, it can be shown that

$$\text{Cov}(\bar{d}_3, \bar{X}_{2A} | \bar{d}_1, y) = \text{Cov}(\bar{d}_3, \bar{X}_{2A}) - \sigma_{x_2} \sigma_y \rho_{x_2, y} \beta_{3y.1},$$

where

$$\text{Cov}(\bar{d}_3, \bar{X}_{2A}) \approx -\frac{1}{2} \frac{(1-2p_1) \sigma_z^2 \rho}{(1-p_1) p_1 p_2 [N(1-\rho)]^2} (1-\rho^{2Np_1}) (1-\rho^{Np_2}),$$

while for the latter, a similar result follows; secondly, since the $\bar{X}_{2i}$ and the $\bar{X}_i'$ are now correlated, we have to modify the cross-product terms involving these quantities such as

$$E(\bar{X}_{2A} \bar{X}_A' Y_1^2) = \int_0^\infty \int_0^\infty \left\{ \text{Cov}(\bar{X}_{2A}, \bar{X}_A' | \bar{d}_1, y_A) + E(\bar{X}_{2A} | \bar{d}_1, y_A) E(\bar{X}_A' | \bar{d}_1, y_A) \right\} f_2(\bar{d}_1, y_A) d\bar{d}_1 dy$$

with

$$\text{Cov}(\bar{X}_{2A}, \bar{X}_A' | \bar{d}_1, y_A) = \text{Cov}(\bar{X}_{2A}, \bar{X}_A') - \frac{\text{Cov}(\bar{X}_{2A}, y_A) \text{Cov}(\bar{X}_A', y_A)}{\sigma_y^2 (1-\rho_{1y}^2)} \quad (3.48)$$

since $\rho_{x_2,1} = \rho_{x_A,1} = 0$, and where

$$\text{Cov}(\bar{X}_{2A}, \bar{X}_A') \approx \frac{\sigma_z^2 \rho}{[N(1-\rho)]^2 p_2 (1-2p_1-p_2)} (1-\rho^{Np_2}) [1-\rho^{N(1-2p_1-p_2)}].$$

This completes the derivation of the analytical results for estimation for the three models. In §3.5, we give some numerical results together with details and results of a simulation study which was carried out to check the analytical results.

3.5 Numerical and simulation results

In §3.1.1, it was emphasised that we are considering the standardised mean treatment difference, δ/σ . This emphasis follows from the fact that the results are unaffected if all the linear dimensions are multiplied by a constant. Therefore, we also divide the bias and trend by σ , and the variance, mean-squared error (MSE) and fluctuation variance by σ^2 . In this way, each set of results will be identified by the *parameter set* $N, \delta/\sigma, t/\sigma, \rho$ and σ_z^2/σ^2 .

3.5.1 Details of the simulation study

A FORTRAN program, STANDSIMULCOLT, was written to simulate the three models for both a stable and an unstable response variable. For each set of parameter values, 1000 simulations were carried out for $N = 100$ [with 30 different combinations of parameter values, an average CPU time of about 15 minutes is required – for STANDEFFICCOLT, the equivalent time is only about 10 seconds].

Before any data are generated, the NAG routine G05CBF is called to set the generator to a repeatable initial state. Following this, two NAG routines are called: (i) G05DDF – this simulates a Normal random variate (for use as the response variable); (ii) G05DYF – this simulates a discrete uniform random variate (for use in the randomised stage(s)).

In order that we have whole numbers of patients in each stage of the trial, the optimal values for p_1 and p_2 obtained using STANDMINIMISE are rounded so that Np_1 and Np_2 are integers (cf. remark in §3.4.2).

3.5.2 Simulated versus analytical results

The analytical results for estimation are compared with those obtained by simulation. The two sets of results are displayed in two-dimensional plots for a wide variety of parameter values. These plots were obtained using the procedure GPLOT in SAS/GRAPH. It will be seen that there is good agreement between the analytical and simulated values, even when there are fluctuations present, and hence the approximations derived in §3.4.2 seem to work well.

a) No trend, no fluctuations ($t/\sigma = 0$; $\sigma_z^2/\sigma^2 = 0$)

Figure 3.1 gives plots of biases, variances and MSE's against δ/σ [= 0.5(0.1)2.5]; for all three models the variance, as expected, increases as δ/σ increases – this is due to more patients being assigned to the better treatment. The bias for the Fixed Model is zero for all values of δ/σ (cf. equation (3.11)). For Model I the bias increases as δ/σ increases and for Model II it initially increases and then decreases – intuitively, to zero.

b) Trend, no fluctuations ($t/\sigma \neq 0$; $\sigma_z^2/\sigma^2 = 0$)

Figure 3.2 gives plots of biases, variances and MSE's against t/σ [= 0.0(0.0005)0.005] for fixed δ/σ [= 0.5]. In each case, the plot is an increasing function of t/σ . This behaviour might be expected because the trend will clearly bias the estimated difference, and moreover it will exaggerate the true difference so that the better treatment is more likely to be chosen at the end of each stage, thus making the treatment group sizes more unequal.

Figures 3.3 and 3.4 give plots of biases, variances and MSE's against δ/σ for fixed t/σ [= 0.0025, 0.005]; the conclusions are similar to those in a).





Figure 3.1: Plots of bias, variance and mean-squared error against δ/σ ($t/\sigma = 0.0$)





Figure 3.2: Plots of bias, variance and mean-squared error against t/σ ($\delta/\sigma = 0.5$)




Figure 3.3: Plots of bias, variance and mean-squared error against δ/σ ($t/\sigma = 0.0025$)





Figure 3.4: Plots of bias, variance and mean-squared error against δ/σ ($t/\sigma = 0.005$)

c) No trend, fluctuations ($t/\sigma = 0$; $\sigma_z^2/\sigma^2 \neq 0$)

Figures 3.5 and 3.6 give plots of biases, variances and MSE's against ρ [= 0.04(0.05)0.99] for fixed δ/σ [= 0.5] and σ_z^2/σ^2 [= 0.05, 0.2]; for each model, the variance increases until $\rho \approx 0.95$ and then decreases. This might be expected due to the constraint $\sigma_z^2/\sigma^2 = 0.05$ or 0.2 [with $\sigma_z^2 = \sigma_\epsilon^2/(1-\rho^2)$]. As an illustration, suppose $\rho = 0.99$, then $1-\rho^2 = 0.0199$. This in turn necessitates σ_ϵ^2 becoming "small". The bias for Model II is affected in the same way. Also, the bias for this model increases as σ_z^2/σ^2 increases.

d) Trend, fluctuations ($t/\sigma \neq 0$; $\sigma_z^2/\sigma^2 \neq 0$)

Figures 3.7 and 3.8 give plots of biases, variances and MSE's against ρ for fixed δ/σ [= 0.5], t/σ [= 0.0025] and σ_z^2/σ^2 [= 0.05, 0.2]. The conclusions are similar to those in c) except that the bias for the Fixed Model and Model I increases slightly as $\rho \rightarrow 1$ for $\sigma_z^2/\sigma^2 = 0.2$.

3.5.3 Analytical results

More detailed impressions of the behaviour of the bias, variance and MSE for the three models are obtained by plotting the results in three dimensions. The procedure G3D in SAS/GRAPH was used for this purpose.

a) Bias, variance and MSE as functions of δ/σ and t/σ ($\sigma_z^2/\sigma^2 = 0$)

Figures 3.9 - 3.11 give three-dimensional plots of biases, variances and MSE's for the three models. [N.B. In this case, in order that the biases, etc. would be "smooth" functions, the optimal values for p_1 and p_2 were not rounded.] The conclusions for the biases are similar to those in a) and b) of §3.5.2. Furthermore, the variances tend to increase rapidly as δ/σ increases except when δ/σ is small (< 0.1) and t/σ is large (> 0.003), when they decrease initially.

b) Bias, variance and MSE as functions of t/σ and ρ ($\sigma_z^2/\sigma^2 \neq 0$)

Figures 3.12 - 3.14 give three-dimensional plots of biases, variances and MSE's for fixed δ/σ [= 0.5] and σ_z^2/σ^2 [= 0.2]. The conclusions are similar to those in c) and d) of §3.5.2.




Figure 3.5: Plots of bias, variance and mean-squared error against ρ ($\delta/\sigma = 0.5$, $t/\sigma = 0.0$, $\sigma_z^2/\sigma^2 = 0.05$)




Figure 3.6: Plots of bias, variance and mean-squared error against ρ ($\delta/\sigma = 0.5$, $t/\sigma = 0.0$, $\sigma_z^2/\sigma^2 = 0.2$)




Figure 3.7: Plots of bias, variance and mean-squared error against ρ ($\delta/\sigma = 0.5$, $t/\sigma = 0.0025$, $\sigma_z^2/\sigma^2 = 0.05$)




Figure 3.8: Plots of bias, variance and mean-squared error against ρ ($\delta/\sigma = 0.5$, $\nu/\sigma = 0.0025$, $\sigma_z^2/\sigma^2 = 0.2$)






Figure 3.9: Plots of bias, variance and mean-squared error for the Fixed Model as functions of δ/σ and t/σ





Figure 3.10: Plots of bias, variance and mean-squared error for Model I as functions of δ/σ and t/σ






Figure 3.11: Plots of bias, variance and mean-squared error for Model II as functions of δ/σ and t/σ



RULE=Fixed Model Stand. difference=0.5 Stand. fluct.=0.2




Figure 3.12: Plots of bias, variance and mean-squared error for the Fixed Model as functions of t/σ and ρ ($\delta/\sigma = 0.5$, $\sigma_z^2/\sigma^2 = 0.2$)




Figure 3.13: Plots of bias, variance and mean-squared error for Model I as functions of t/σ and ρ ($\delta/\sigma = 0.5$, $\sigma_z^2/\sigma^2 = 0.2$)




Figure 3.14: Plots of bias, variance and mean-squared error for Model II as functions of t/σ and ρ ($\delta/\sigma = 0.5$, $\sigma_z^2/\sigma^2 = 0.2$)

3.6 Estimation of instability parameters

In §3.3, we introduced two types of instability in the response variable. In practice, the so-called instability parameters (t/σ , ρ and σ_z^2/σ^2) would be unknown at the outset and consequently they would need to be estimated. If we have only a linear time-trend, then we can use ordinary least squares while if we also have fluctuations, then a non-linear least squares procedure is needed. These two aspects of the problem are addressed in §§3.6.1 and 3.6.2, respectively, together with an indication of possible adjustments in the estimation formulae of §3.4.

3.6.1 Estimation of trend

We assume as in §3.3.1 that $X_{ij} \sim N(\mu_i + tj, \sigma^2)$, where t is the linear trend per unit time. If we standardise as in §3.5, we may write

$$x_j = a + t^*(j - \bar{j}) + \delta^*(r_j - \bar{r}) + Z_j, \quad Z_j \text{ i.i.d. } N(0, 1), \quad (j \in [0, N-1]), \quad (3.49)$$

where r_j is an indicator variable with value one if treatment A is used at time j , $t^* = t/\sigma$ and $\delta^* = \delta/\sigma$.

In matrix form, we can write

$$\underline{x} = Y \underline{\theta} + \underline{Z}, \quad \underline{\theta}^T = (a, t^*, \delta^*).$$

The least squares estimates of a , t^* and δ^* are thus given by

$$\hat{a} = \bar{x}, \quad \hat{t}^* = S_{jx.r} / S_{jj.r}, \quad \hat{\delta}^* = S_{rx.j} / S_{rr.j},$$

where

$$S_{jx.r} = S_{jx} - S_{jr} S_{xr} / S_{rr}, \quad S_{jx} = \sum_{j=0}^{N-1} (j - \bar{j})(x_j - \bar{x}),$$

etc., with variances N^{-1} , $S_{jj.r}^{-1}$ and $S_{rr.j}^{-1}$, respectively.

The least squares estimate of t^* can be used to adjust for the bias in $\bar{d}^*$ for each model. As an illustration, consider the Fixed Model. In this case, $E(\bar{d}^*)$ and $\text{Var}(\bar{d}^*)$ are given, respectively, by (3.29) and (3.30). If we denote the bias in $\bar{d}^*/\sigma$ by $b(\delta^*)$, then a suitably adjusted estimate of δ^* , $\tilde{d}^*$, is given by the solution of

$$\tilde{d}^* = \bar{d}^* / \sigma - b(\tilde{d}^*)$$

with t^* replaced by $\hat{t}^*$.

By considering a Taylor series expansion of $b(\tilde{d}^*)$ about $b(\delta^*)$, it is straightforward to show that to first-order terms $\tilde{d}^*$ is unbiased for δ^* and

$$\text{Var}(\tilde{d}^*) \leq \text{Var}(\bar{d}^*/\sigma) [1 + b'(\delta^*)]^{-2} \quad (3.50)$$

(Whitehead, 1986, p.578), where ' denotes differentiation with respect to δ^* .

Furthermore, since $\text{Var}(\bar{d}^*/\sigma)$ depends on t^* and t^{*2} , we use an unbiased estimate of t^{*2} given by $\hat{t}^{*2} - S_{jj,r}^{-1}$ to estimate (3.50) and replace δ^* by $\tilde{d}^*$. We can adjust $\bar{d}^*/\sigma$ for Model I in the same way. The approach is less straightforward for Model II since some of the limits of integration depend on t^* but we can replace t^* by $\hat{t}^*$ and use an adjustment based on the first-order terms of a Taylor series expansion about t^* .

3.6.2 Estimation of serial correlation

If we have fluctuations in the response variable then the response at time j can be written in a similar form to (3.49) but with the Z_j as defined in §3.3.2. Alternatively, we can write

$$x_j = \rho x_{j-1} + a(1-\rho) + t^* [j(1-\rho) + \rho] + \delta^* (r_j - \rho r_{j-1}) + \epsilon_j, \quad \epsilon_j \text{ i.i.d. } N(0, \sigma_\epsilon^2)$$

if the first response x_0 is omitted [N.B. Beach & MacKinnon (1978) have investigated the loss in efficiency incurred by this omission].

Several methods for simultaneous least squares estimation of the regression coefficients and ρ have been suggested in the literature (e.g. Malinvaud, 1980, Ch.12, §7). One approach would involve firstly obtaining the normal equations and then using the following iterative method: (a) choose an initial value, $\rho^{(0)}$, for ρ (zero, say) and determine the values of the regression coefficients which minimise the sum of squared differences between $x_j - \rho^{(0)} x_{j-1}$ and $a(1-\rho^{(0)}) + t^* [j(1-\rho^{(0)}) + \rho^{(0)}] + \delta^* (r_j - \rho^{(0)} r_{j-1})$; (b) determine a revised value, $\rho^{(1)}$, for ρ using the residuals from fitting (a); (c) repeat (a) using the revised value for ρ ; and so on. [Another approach is to apply the *Kalman filter technique* to this problem (e.g. Harvey & Phillipps, 1979).]

A generalisation of the iterative method above for errors (Z_j) following an ARMA process has been considered by Pierce (1971). Asymptotic normal properties of the parameter estimates are investigated together with small-sample properties via a simulation study for a simple regression model. The latter study includes some results for errors following an AR(1) process. There are three main conclusions: (a) for small samples

(15), the estimated regression coefficients appear to be unbiased but the estimate of ρ , $\hat{\rho}$, is badly biased downward with a variance larger than its asymptotic one; (b) for moderate samples (50), the bias in $\hat{\rho}$ is considerably smaller but still significant; (c) for large samples (200), the bias in $\hat{\rho}$ is negligible. From these results, we might expect this approach to work reasonably well if we had, say, $N = 100$ patients. However, no numerical results have yet been obtained to support this suggestion.

Having computed estimates for t^* , ρ and σ_z^2/σ^2 , we can adjust for the bias in $\bar{d}^*/\sigma$ for the three models using the same approach as for Model II in §3.6.1: replace ρ by $\hat{\rho}$ and use the first-order Taylor series terms to adjust accordingly.

3.7 A three-stage model

In §3.1, we compared the optimal values of the parameters for the Fixed Model with those for Models I and II. It was demonstrated numerically by Colton (1965, pp.177–178) that the optimal two-stage models can attain up to almost 50% of the difference in the overall net gain between the optimal one-stage (i.e. Fixed Model) and a fully sequential model. In this short section, we attempt to answer the question: How much closer to a fully sequential model is an optimal three-stage model in terms of overall net gain? The three-stage model under investigation, henceforth referred to as Model III, is a straightforward extension of Model II.

3.7.1 Formulation of model

Stages 1 and 2: As for Model II in §3.1.2.

If $y_i > 0$, use treatment A in the third stage;

if $y_i < 0$, use treatment B in the third stage ($i=A,B$).

Stage 3: Use one treatment only on Np_3 patients.

If $\bar{d}_1 > 0$ and $y_A > 0$, calculate

$$y_{AA} = \frac{p_1 \bar{x}_{1A} + p_2 \bar{x}_{2A} + p_3 \bar{x}_{3A}}{p_1 + p_2 + p_3} - \bar{x}_{1B} ;$$

if $\bar{d}_1 > 0$ and $y_A < 0$, calculate

$$y_{AB} = \frac{p_1 \bar{x}_{1A} + p_2 \bar{x}_{2A}}{p_1 + p_2} - \frac{p_1 \bar{x}_{1B} + p_3 \bar{x}_{3B}}{p_1 + p_3} ;$$

if $\bar{d}_1 < 0$ and $y_B > 0$, calculate

$$y_{BA} = \frac{p_1 \bar{x}_{1A} + p_3 \bar{x}_{3A}}{p_1 + p_3} - \frac{p_1 \bar{x}_{1B} + p_2 \bar{x}_{2B}}{p_1 + p_2} ;$$

if $\bar{d}_1 < 0$ and $y_B < 0$, calculate

$$y_{BB} = \bar{x}_{1A} - \frac{p_1 \bar{x}_{1B} + p_2 \bar{x}_{2B} + p_3 \bar{x}_{3B}}{p_1 + p_2 + p_3} .$$

If $y_{ij} > 0$, use treatment A on the remaining $N(1-2p_1-p_2-p_3)$ patients;

if $y_{ij} < 0$, use treatment B on the remaining $N(1-2p_1-p_2-p_3)$ patients

(i,j=A,B).

3.7.2 Optimisation

We firstly require an expression for E, the expected proportion of patients on the inferior treatment, i.e. B. By considering all possible combinations of decisions for Model III (9 in total), and using the approach of Colton (1965), we obtain

$$\begin{aligned} E = & p_1 + p_2 \Phi(h) + p_3 \Phi(k) \\ & + (1-2p_1-p_2-p_3) \left[\Phi(\ell_2) - \int_{-\infty}^{\ell_2} \left\{ \int_{-\infty}^k \int_{-\infty}^h + \int_k^{\infty} \int_h^{\infty} \right\} \varphi_3^{(2)}(u,v,w) du dv dw \right. \\ & \left. + \int_{-\infty}^{\ell_1} \left\{ \int_{-\infty}^k \int_{-\infty}^h + \int_k^{\infty} \int_h^{\infty} \right\} \varphi_3^{(1)}(u,v,w) du dv dw \right] , \quad (3.51) \end{aligned}$$

where k and h are defined as for Model II in §3.1.3,

$$\ell_1 = -\delta / \left\{ \frac{Np_1(p_1+p_2+p_3)}{2p_1+p_2+p_3} \right\} / \sigma , \quad \ell_2 = -\delta / \left\{ \frac{N(p_1+p_2)(p_1+p_3)}{2p_1+p_2+p_3} \right\} / \sigma$$

and $\varphi_3^{(j)}(u,v,w)$ is a standard trivariate normal density with correlation coefficients ρ_{1y} , ρ_{1,ℓ_j} and ρ_{y,ℓ_j} ($j=1,2$). Clearly, putting $p_3 = 0$ reduces (3.51) to that for Model II in §3.1.3.

To obtain the optimal values for p_1 , p_2 and p_3 , we require the $\partial E / \partial p_i$ ($i=1,2,3$). The derivation of these involves differentiation of standard trivariate normal densities, and is therefore awkward. Details of the three resulting equations are given in Appendix 2; these equations can be solved numerically for p_1 , p_2 and p_3 . A FORTRAN program, THREESTAGE, was written for this purpose using the various NAG routines described in §3.1.3. Results from this program are now discussed.

3.7.3 Numerical results

Optimal values for the p_i ($i=1,2,3$) are given in Table 3.2 (a similar check to that in §3.1.2 was carried out to confirm that these values corresponded to local minima). A comparison of the optimal net gain for Model III with that of the Fixed Model and Model II is given in Table 3.3, where we tabulate the ratio of the optimal net gain for Models II and III to that for the Fixed Model.

From the conclusions of Colton (1965, pp. 177–178), our results in Table 3.3 indicate that the optimal three-stage model can attain up to almost 75% of the difference in net gain between an optimal one-stage and a fully sequential model. As expected, we have a "law of diminishing returns" setting in after a two- or three-stage model.

No attempt at estimation has been made for Model III. Intuitively, one would expect a larger bias and variance than for Model II due to the more adaptive nature of the model.

Table 3.2: *Optimal values for parameters for Model III*

δ/σ	N	p_1	p_2	p_3
0.5	100	0.085	0.134	0.205
	200	0.064	0.115	0.181
	500	0.038	0.079	0.147
	1000	0.024	0.055	0.126
1.5	100	0.026	0.059	0.129
	200	0.016	0.040	0.111
	500	0.008	0.023	0.089
	1000	0.004	0.015	0.073
2.5	100	0.012	0.033	0.103
	200	0.007	0.022	0.087
	500	0.003	0.012	0.066
	1000	0.002	0.008	0.053

Table 3.3: Values for the ratio of the optimal net gain for Models II and III to that for the Fixed Model

δ/σ	N	Model II	Model III
0.5	100	1.0887	1.1288
	200	1.0698	1.0988
	500	1.0452	1.0609
	1000	1.0303	1.0393
1.5	100	1.0323	1.0421
	200	1.0208	1.0263
	500	1.0111	1.0135
	1000	1.0066	1.0079
2.5	100	1.0167	1.0208
	200	1.0102	1.0125
	500	1.0052	1.0061
	1000	1.0030	1.0035

3.8 Other response variables

Until now we have considered a normal response variable. It is natural at this point to highlight modifications which are needed when the response variable is non-normal. As examples, we examine Bernoulli and exponential responses for the Fixed Model. In §§3.8.1 and 3.8.2, we use Gail's method of §3.2 and the change of variable technique to derive the exact bias and variance of $\bar{d}^*$ for a small first stage (not more than 3 patients on each treatment); for a larger first stage, a normal approximation can be employed, details of which are given in §3.8.3. Finally, in §3.8.4 a comparison of the exact and large-sample results is presented.

3.8.1 Bernoulli response (exact)

Let n ($= Np_i$) be the number of patients on each treatment in the first stage. Further let X_i be a sum of n independent Bernoulli random variables with parameters p_i ($i=A,B$) and $S = X_A + X_B$, $D = X_A - X_B$. Firstly, suppose $n = 1$, then the joint probability function of X_A and X_B is

$$P(x_A, x_B) = p_A^{x_A} (1-p_A)^{1-x_A} p_B^{x_B} (1-p_B)^{1-x_B} \text{ , } (p_A, p_B \in [0, 1]; \text{ } x_A, x_B=0, 1).$$

Following Gail's method of §3.2, we consider the joint probability function of S and D , thus

$$P(s, d) = p_A^b (1-p_A)^{1-b} p_B^c (1-p_B)^{1-c}, \quad (s, d) \in \{(0, 0), (1, 1), (1, -1), (2, 0)\},$$

where $b = \frac{1}{2}(s+d)$ and $c = \frac{1}{2}(s-d)$.

Simple algebra yields the marginal probability function of D as

$$P(d) = \begin{cases} p_A(1-p_B) & , \quad d = 1 \\ 1-p_A-p_B+2p_Ap_B & , \quad d = 0 \\ p_B(1-p_A) & , \quad d = -1 \end{cases}.$$

Next, we have

$$E(Y) = P(D = 1) + \frac{1}{2}P(D = 0) = \frac{1}{2}(1+p_A-p_B), \quad E(SY) = p_A.$$

This in turn yields

$$E(\bar{d}^*) = \delta + \frac{1}{2} (N-2)(p_A-p_B)(p_A+p_B-1)/(N-1) \quad (3.52)$$

and similarly,

$$\begin{aligned} \text{Var}(\bar{d}^*) &= [N^2\{(p_A+p_B)-(p_A^2+p_B^2)\} + (N-2)^2\{(p_A+p_B)(p_A-p_B)^2(2-p_A-p_B) \\ &\quad - 2(p_A^2+p_B^2) + (p_A+p_B+2p_Ap_B)\} + 2N(N-2)\{2p_A(1-p_B)(p_A+p_B) \\ &\quad - (p_A-p_B)[(p_A^2-p_B^2) + 2p_B] - (p_A+p_B-2p_Ap_B)\} \\ &\quad + 2(N-2)\{2p_B(1-p_B) + (p_A-p_B)(1+p_A-p_B)(1-p_A-p_B)\}]/[2(N-1)]^2. \end{aligned} \quad (3.53)$$

It is easily verified that for $p_A \simeq p_B (= p, \text{ say})$, and as $N \rightarrow \infty$, $\text{Var}(\bar{d}^*) \rightarrow 2p^2(1-p)$.

Suppose now that $n = 2$, then the joint probability function of S and D as defined for the Fixed Model in §3.2 is

$$P(s, d) = \begin{bmatrix} 2 \\ b \end{bmatrix} \begin{bmatrix} 2 \\ c \end{bmatrix} p_A^b (1-p_A)^{2-b} p_B^c (1-p_B)^{2-c},$$

where $(s, d) \in \{(0, 0), (1, 1), (1, -1), (2, 0), (2, 2), (2, -2), (3, 1), (3, -1), (4, 0)\}$ and b and c are as above.

We can again obtain the marginal probability distribution of D and hence

$$E(Y) = p_A(1-p_B)(1-p_A+p_Ap_B) + \frac{1}{2}[p_A^2p_B^2+(1-p_A)^2(1-p_B)^2]$$

and also

$$E(SY) = 2p_A[(1-p_A)(1-p_B)(1+p_B) + p_A(1-p_B)(1+2p_B) + p_Ap_B^2]$$

and finally

$$E(\bar{d}^*) = \delta + (N-4) \left[\frac{1}{2}(p_A + p_B) \{ 2p_A(1-p_B)(2-p_A+p_A p_B) + p_A^2 p_B^2 + (1-p_A)^2(1-p_B)^2 \} \right. \\ \left. - p_A(1-p_A)(1-p_B)(1+p_B) - p_A^2(1-p_B)(1+2p_B) - p_A^2 p_B^2 \right] / (N-2). \quad (3.54)$$

The corresponding variance for $n = 2$ has not been evaluated. However, an expression for the bias for $n = 3$ has been derived and is included in Appendix 3.

3.8.2 Exponential response (exact)

Redefine X_i to be a sum of n independent exponential random variables with parameters λ_i ($i=A,B$), and suppose $n = 1$. The joint probability density function (p.d.f.) of X_A and X_B is

$$f(x_A, x_B) = \lambda_A \lambda_B e^{-\lambda_A x_A - \lambda_B x_B}, \quad (\lambda_i, x_i > 0, i=A, B).$$

Using the change of variable technique, it is easily shown that

$$f(s, d) = \frac{1}{2} \lambda_A \lambda_B e^{-\frac{1}{2} \{ (\lambda_A + \lambda_B)s + (\lambda_A - \lambda_B)d \}}, \quad (s > 0, d < s, d > -s).$$

By simple integration, the marginal p.d.f. of D is

$$f(d) = \begin{cases} \frac{\lambda_A \lambda_B}{(\lambda_A + \lambda_B)} e^{\lambda_B d} & , \quad d < 0 \\ \frac{\lambda_A \lambda_B}{(\lambda_A + \lambda_B)} e^{-\lambda_A d} & , \quad d > 0 \end{cases}$$

and hence

$$E(Y) = \frac{\lambda_B}{(\lambda_A + \lambda_B)}.$$

Further, using $f(s, d)$, we obtain

$$E(SY) = \frac{\lambda_B (3\lambda_A + \lambda_B)}{\lambda_A (\lambda_A + \lambda_B)^2}$$

and hence

$$E(\bar{d}^*) = \delta + \frac{(N-2)}{(N-1)} \frac{(\lambda_A - \lambda_B)}{(\lambda_A + \lambda_B)^2}, \quad (3.55)$$

where $\delta = \lambda_A^{-1} - \lambda_B^{-1}$. It can be seen that for $\lambda_A \approx \lambda_B (= \lambda, \text{ say})$, we have

$$E(\bar{d}^*) \approx \frac{(\lambda_B - \lambda_A)}{\lambda^2} \left\{ 1 - \frac{1}{4} \frac{(N-2)}{(N-1)} \right\}. \quad (3.56)$$

After some algebra, the variance of $\bar{d}^*$ can be written

$$\text{Var}(\bar{d}^*) = \left[(N-2)^2 (\lambda_A + \lambda_B)^2 \left\{ (\lambda_A^2 + \lambda_B^2) - \frac{4\lambda_A^2 \lambda_B^2 (\lambda_A - \lambda_B)^2}{(\lambda_A + \lambda_B)^4} \right\} \right]$$

$$\begin{aligned}
& + N^2 (\lambda_A + \lambda_B)^2 (\lambda_A^2 + \lambda_B^2) - 2N(N-2) (\lambda_A^2 + \lambda_B^2)^2 \\
& + 4(N-2) (\lambda_A + \lambda_B)^2 (\lambda_A^2 - \lambda_A \lambda_B + \lambda_B^2) \Big] / \left[2(N-1) \lambda_A \lambda_B (\lambda_A + \lambda_B) \right]^2. \quad (3.57)
\end{aligned}$$

For $\lambda_A \approx \lambda_B = \lambda$, and as $N \rightarrow \infty$, $\text{Var}(\bar{d}^*) \rightarrow 1/2\lambda^2$.

Suppose now that $n = 2$, then

$$f(s, d) = \frac{1}{8} (\lambda_A \lambda_B)^2 (s^2 - d^2) e^{-\frac{1}{2}\{(\lambda_A + \lambda_B)s + (\lambda_A - \lambda_B)d\}}, \quad (s > 0, d < s, d > -s).$$

In the same way as for $n = 1$, we can show that

$$E(Y) = \frac{\lambda_B^2 (3\lambda_A + \lambda_B)}{(\lambda_A + \lambda_B)^3}$$

and

$$E(SY) = \frac{2\lambda_B^2}{\lambda_A (\lambda_A + \lambda_B)^4} [(\lambda_A + \lambda_B)^2 + 3\lambda_A (\lambda_A + \lambda_B) + 6\lambda_A^2].$$

Using these two results, we have

$$E(\bar{d}^*) = \delta + 3\lambda_A \lambda_B \frac{(N-4)}{(N-2)} \frac{(\lambda_A - \lambda_B)}{(\lambda_A + \lambda_B)^4} \quad (3.58)$$

and the expression corresponding to (3.56) is

$$E(\bar{d}^*) \approx \frac{(\lambda_B - \lambda_A)}{\lambda^2} \left\{ 1 - \frac{3}{16} \frac{(N-4)}{(N-2)} \right\}. \quad (3.59)$$

As a large-sample guide, if we compare (3.56) and (3.59), we see that the bias in $\bar{d}^*$ is reduced by one-sixteenth by going from one patient to two patients per treatment in the first stage.

As in §3.8.1, the variance of $\bar{d}^*$ for $n = 2$ has not been evaluated and the bias for $n = 3$ is included in Appendix 3.

3.8.3 Normal approximations

We saw in §§3.8.1 and 3.8.2 that for small values of n exact results for the bias and the variance of $\bar{d}^*$ for Bernoulli and exponential responses can be derived. Clearly, for values of n greater than two or three, the algebra required soon becomes awkward. However, for $n > 10$ or 15, say, we would expect results for normal responses with unequal variances to approximate those for Bernoulli and exponential. [In §3.8.4, a numerical comparison is carried out for small values of n .] Details of the results for

normal responses are now outlined.

Let $S = Y_A + Y_B$ and $D = Y_A - Y_B$ where $Y_i \sim N(n\mu_i, n\sigma_i^2)$, ($i=A,B$). Clearly, the joint p.d.f. of S and D is bivariate normal with means $\mu_S = n(\mu_A + \mu_B)$, $\mu_D = n(\mu_A - \mu_B)$, variances $\sigma_j^2 = n(\sigma_A^2 + \sigma_B^2)$ ($j=S,D$) and correlation coefficient $\rho_{S,D} = (\sigma_A^2 - \sigma_B^2)/(\sigma_A^2 + \sigma_B^2)$. It follows that

$$E(Y) = 1 - \Phi(-\mu_D/\sigma_D)$$

and using standard results for conditional normal distributions,

$$E(SY) = \mu_S [1 - \Phi(\mu_D/\sigma_D)] + \rho_{S,D} \sigma_S \varphi(\mu_D/\sigma_D)$$

and hence

$$E(\bar{d}^*) = \delta - \frac{(N-2)}{n(N-1)} \rho_{S,D} \sigma_S \varphi(\mu_D/\sigma_D), \quad (3.60)$$

where $\delta = \mu_A - \mu_B$.

In a similar way, we can show that

$$\begin{aligned} \text{Var}(\bar{d}^*) = & [N^2 \sigma_D^2 + (N-2n)^2 \sigma_S^2 + 4n^2 (N-2n) \{ \sigma_A^2 - (\sigma_A^2 - \sigma_B^2) \Phi(-\mu_D/\sigma_D) \} \\ & + 2(N-2n) \rho_{S,D} \sigma_S \{ N \sigma_D [2\Phi(-\mu_D/\sigma_D) - 1] - 2(N-2n) \rho_{S,D} \sigma_S \varphi^2(\mu_D/\sigma_D) \} \\ & + 4N(N-2n) \rho_{S,D} \sigma_S \mu_D \varphi(\mu_D/\sigma_D)] / [2N(N-n)]^2, \end{aligned} \quad (3.61)$$

where putting $\sigma_A^2 = \sigma_B^2 = \sigma^2$ reduces (3.61) to (3.15), as required.

We can now obtain approximations for the biases and variances for Bernoulli and exponential responses and compare these with the exact results. Firstly, we consider a Bernoulli response for $n = 1$; making the appropriate substitutions in (3.60), we obtain

$$E(\bar{d}^*) = \delta + \frac{(N-2)}{(N-1)} \frac{(p_A - p_B)(p_A + p_B - 1)}{[p_A(1-p_A) + p_B(1-p_B)]} \varphi[(p_A - p_B)/\sqrt{p_A(1-p_A) + p_B(1-p_B)}] \quad (3.62)$$

Similarly, an expression for the variance can be obtained using (3.61); for $p_A \approx p_B = p$, and as $N \rightarrow \infty$, this becomes $p(1-p)$.

For an exponential response and $n = 1$, we obtain

$$E(\bar{d}^*) = \delta + \frac{(N-2)}{(N-1)} \frac{(\lambda_A^2 - \lambda_B^2)}{\lambda_A \lambda_B \sqrt{(\lambda_A^2 + \lambda_B^2)}} \varphi[(\lambda_A - \lambda_B)/\sqrt{(\lambda_A^2 + \lambda_B^2)}] \quad (3.63)$$

and for $\lambda_A \approx \lambda_B = \lambda$, we have

$$E(\bar{d}^*) \simeq \frac{(\lambda_B - \lambda_A)}{\lambda^2} \left\{ 1 - \frac{1}{\sqrt{\pi}} \frac{(N-2)}{(N-1)} \right\} . \tag{3.64}$$

As a comparison, the bias term for $n = 2$ involves a factor of $1/\sqrt{2\pi}$. Therefore, the ratio of the exact bias for $n = 2$ to that for $n = 1$ is 0.76 compared with 0.71 for the normal approximation. For $n = 3$, it is shown in Appendix 3 that the exact bias involves a factor of $5/32$ which compares with $1/\sqrt{3\pi}$ for the normal approximation; the corresponding ratios of bias for $n = 3$ to those for $n = 2$ are, respectively, 0.84 and 0.8.

Again, an expression for the variance follows straightforwardly and for $\lambda_A \simeq \lambda_B = \lambda$, and as $N \rightarrow \infty$, this becomes $1/\lambda^2$.

A FORTRAN program, NONNORMAL, was written to incorporate the exact results of §§3.8.1 and 3.8.2, and the above normal approximations. A numerical comparison of the results is now discussed.

3.8.4 Numerical results

We firstly consider a Bernoulli response. From the results in Table 3.4, the difference between the exact and the approximate bias only strictly decreases as n increases for $p_A = 0.3$ and $p_B = 0.1$. However, the difference decreases when n increases from 2 to 3 for each case considered. The results in Table 3.5 seem to indicate that the agreement between the exact and the approximate variances is reasonable except when both the p_i 's are small or large - this might be expected by referring to the limiting results obtained for small values of $p_A - p_B$.

Table 3.4: *A comparison of the exact and the normal approximation values for the bias for a Bernoulli response when $N = 100$*

		(p_A, p_B)			
		$(0.45, 0.4)$	$(0.3, 0.1)$	$(0.8, 0.3)$	$(0.7, 0.6)$
Exact	1	-0.0037	-0.0594	0.0247	0.0148
	2	-0.0019	-0.0388	0.0093	0.0079
	3	-0.0014	-0.0275	0.0051	0.0057
Normal	1	-0.0042	-0.0801	0.0229	0.0173
	2	-0.0030	-0.0530	0.0116	0.0121
	3	-0.0024	-0.0400	0.0067	0.0097

Table 3.5: *A comparison of the exact and the normal approximation values for the variance for a Bernoulli response when $N = 100$*

		(p_A, p_B)			
	n	(0.45, 0.4)	(0.3, 0.1)	(0.8, 0.3)	(0.7, 0.6)
Exact	1	0.2096	0.0581	0.2046	0.2945
Normal	1	0.2462	0.1444	0.1892	0.2270

Secondly, we consider an exponential response. It is clear from Table 3.6 that as n increases, the difference between the exact and the approximate bias decreases. As an illustration, suppose that $\lambda_A = 0.01$ and $\lambda_B = 0.1$, then the difference in the biases decreases by about a factor of 3. From Table 3.7, the normal approximation for the variance is very poor but as mentioned earlier, for small values of $\lambda_B - \lambda_A$ the exact value tends to be about half that of the approximate one – the case where $\lambda_A = 0.25$ and $\lambda_B = 0.5$ provides a good example.

Table 3.6: *A comparison of the exact and the normal approximation values for the bias for an exponential response when $N = 100$*

		(λ_A, λ_B)			
	n	(0.01, 0.1)	(0.25, 0.5)	(2.5, 10.0)	(50.0, 75.0)
Exact	1	7.3629	0.4400	0.0475	0.0016
	2	1.8065	0.2902	0.0226	0.0011
	3	0.4923	0.2127	0.0119	0.0009
Normal	1	26.0510	0.9588	0.1103	0.0035
	2	12.2072	0.6071	0.0592	0.0024
	3	6.6029	0.4437	0.0367	0.0018

Table 3.7: *A comparison of the exact and the normal approximation values for the variance for an exponential response when $N = 100$*

		(λ_A, λ_B)			
	n	(0.01, 0.1)	(0.25, 0.5)	(2.5, 10.0)	(50.0, 75.0)
Exact	1	872.4347	4.3725	0.0263	0.0001
Normal	1	3706.1031	9.0680	0.0675	0.0003

3.8.5 Future work

It is clear from the results in §3.8.4 that further work is needed to assess the validity of the normal approximation for small values of n . From the limited numerical results presented, there is some indication that the normal approximation is better for Bernoulli responses. Since the algebra becomes awkward for values of $n > 2$ or 3 , a small simulation study might be appropriate for $n > 3$.

An extension of Gail's method of §3.2 to Models I and II for non-normal responses is possible but, in the case of instability, this approach would be difficult and again a simulation study might be more suitable. For large values of n , it may be possible to modify the results of §§3.2 and 3.4 to allow for unequal variances and hence obtain normal approximations. However, more work is needed.

Lastly, if the response variable considered in §3.8.2 represents the survival time of a patient, some of the data may be *censored*. As pointed out in §1.4, for a related problem – delayed observations – Langenberg & Srinivasen (1981) modified a version of the Fixed Model. A similar modification might circumvent the problem of censoring but it is evident that the complexity of the problem would be increased.

CHAPTER 4. SEQUENTIAL TESTS FOR AN UNSTABLE RESPONSE VARIABLE

We now consider the class of sequential tests which was investigated by, among others, Robbins & Siegmund (1974) and Hayre & Turnbull (1981c).

In §4.1, we describe these tests together with a number of data-dependent allocation rules designed for normal responses with known, equal variances. Some results due to Woodroffe (1989) for estimation following sequential tests are then outlined. A simulation study of the allocation rules is presented in §4.2.

With linear time-trends in the response variable, a new sequential test is derived using maximum likelihood theory in §4.3 using the procedures of Bartlett (1946) and Cox (1963). Appropriate adjustments to the original boundaries are then given. Following some comments regarding estimation, simulation results for these tests are presented in §4.4.

In §4.5, we derive a sequential test for normal responses with known, unequal variances in the presence of linear time-trends. An approximate test is then constructed for Bernoulli responses; simulation results for this test are presented in §4.6 for several data-dependent allocation rules.

A sequential test is derived in §4.7 when there are fluctuations in a normal response variable; a special case of the fluctuation model is used to illustrate the merits of this test. Lastly, there is a short discussion of possible extensions to the present work in §4.8.

4.1 Sequential tests for a stable normal response variable

Consider a clinical trial in which patients can be allocated to one of two treatments A and B. We initially assume that the response variable for treatment i at time j , X_{ij} ($j=0,1,2,\dots$) is normally distributed with mean μ_i and variance unity. We are interested in testing $H_0: \delta < 0$ against $H_1: \delta > 0$, where $\delta = \mu_B - \mu_A$. For this problem, Robbins & Siegmund (1974), following earlier work by Flehinger & Louis (1972) (see §1.3.1), derived a stopping rule and a terminal decision rule of the sequential probability ratio test (SPRT) type for which the error probability is *approximately independent* of the allocation

rule adopted. Use of such a test enables us to seek an allocation rule which reduces the expected number of patients on the inferior treatment.

4.1.1 Stopping rule and terminal decision rule

Defining the X_{ij} as above, and given $\delta^* > 0$, Robbins & Siegmund (1974, p.132) show that a sequential test of $H_0^*: \delta = -\delta^*$ against $H_1^*: \delta = \delta^*$ can be given in terms of λ , the likelihood ratio of the observed responses (m on treatment A and n on treatment B) under these two *simple* hypotheses, where

$$\lambda = \exp\left\{2\delta^* \frac{mn}{(m+n)} (\bar{x}_{Bn} - \bar{x}_{Am})\right\}$$

and $\bar{x}_{Am}$ and $\bar{x}_{Bn}$ are the sample means on treatments A and B, respectively, after $m+n$ patients.

For simplicity, we only consider symmetric tests in this thesis. Thus, if we define $0 < A^{-1} < 1 < A$, then we can test H_0^* against H_1^* sequentially as follows: let $(M, N) =$ first (m, n) such that $z_{m,n} \notin (-a, a)$ and for $M + N < \infty$, accept H_0^* or H_1^* according as $z_{M,N} \leq -a$ or $\geq a$, respectively, where

$$z_{m,n} = \frac{mn}{(m+n)} (\bar{x}_{Bn} - \bar{x}_{Am}) \quad (4.1)$$

and

$$a = \log A / 2\delta^* .$$

It is straightforward (Robbins & Siegmund, 1974, p.133) to show that for $\delta \neq 0$ the error probability of the above test is given by

$$\Pr(\text{error}) \simeq (1 + e^{2a\delta})^{-1} , \quad (4.2)$$

where the approximation is due only to overshoot of the stopping boundaries. Furthermore, using a martingale argument (Robbins & Siegmund, 1974, pp.133-134), it can be shown that if $E\{MN/(M+N)\} < \infty$,

$$E\left\{\frac{MN}{M+N}\right\} \simeq \begin{cases} a^2 , & (\delta=0) \\ \frac{a(e^{2a\delta} - 1)}{\delta(e^{2a\delta} + 1)} , & (\delta \neq 0) \end{cases} . \quad (4.3)$$

Thus, the *information* for this test is approximately independent of the allocation rule.

4.1.2 Allocation rules considered

(a) Robbins & Siegmund rule

A data-dependent rule which for $z_{m,n}$ close to zero approximates pairwise allocation and for $z_{m,n}$ close to one of the stopping boundaries $-a$ or a allocates most of the patients to the better treatment, was proposed by Robbins & Siegmund (1974, p.134). This rule is: initially allocate one patient to each treatment; subsequently, allocate the next patient to treatment B if

$$\frac{(n-m)}{(m+n)} \leq \frac{z_{m,n}}{c}, \quad (c \geq a); \quad (4.4)$$

otherwise allocate to treatment A.

(b) Proportionate randomisation rule

This allocation rule is based on the standardised value of the estimated mean treatment difference

$$\hat{\delta}_s = \sqrt{\left\{ \frac{mn}{m+n} \right\}} (\bar{x}_{Bn} - \bar{x}_{Am}), \quad (m, n \geq 1). \quad (4.5)$$

The rule is: initially allocate one patient to each treatment; subsequently, if $|\hat{\delta}_s| < 2$, then randomise equally to the two treatments; otherwise, randomise in the proportions 1:2 or 2:1 to treatments A and B, respectively, according to whether $\hat{\delta}_s > 2$ or $\hat{\delta}_s < -2$.

(c) The Gittins rule

The use of Gittins indices for normal responses when the variance is known (see Gittins, 1989, §7.2) yields the allocation rule: initially allocate one patient to each treatment; subsequently, allocate each patient to the treatment which currently has the higher Gittins index, randomising in the case of ties. From simulation, this allocation rule does not work well in its present form. In particular, the results indicated that for $\delta \leq 0.25$, the expected sample sizes on the two treatments are comparable with those for the Robbins & Siegmund rule but for larger values of δ , the expected sample size on the better treatment tended to be excessively large (e.g. 856.7 for $\delta = 4.0$!). It is evident from (the proof of) Theorem 4.1 below that for this rule, the expected total sample size is not finite. For this reason, two simple modifications to the rule were considered.

Firstly, we allowed more patients to be allocated to each treatment initially. Although, in general, this resulted in a considerable reduction in the expected sample size

on the better treatment (e.g. allocating 3 patients initially to each treatment gave 4.59 for $\delta = 4.0$), it was felt that the adaptive property of the original rule was not being used to its best effect due to little or no reduction in the number of patients on the inferior treatment. The following modified rule provided a flexible compromise: if $m^r < n$ or $n^r < m$ ($r \geq 1$), allocate the next patient to treatment A or B, respectively; otherwise, follow the original rule. [N.B. When $r = 1$, the rule is simply pairwise allocation and as $r \rightarrow \infty$, we have the original rule.] In this chapter, this allocation rule will be subsequently referred to as the Gittins rule.

Theorem 4.1. For the Gittins rule, for each δ and $r \geq 1$ there exist constants $p \in (0,1)$ and $c > 0$ such that for sufficiently large ν ,

$$\Pr\{M+N > \nu\} \leq c\rho^{\nu^{1/r}}.$$

In particular, $E\{(M+N)^k\} < \infty$ for all $k > 0$ and $r < \infty$.

Proof. It suffices to prove the result for $\delta > 0$. A similar argument applies for $\delta \leq 0$. Thus, if after k patients, m_k and n_k denote the sample sizes on treatments A and B, respectively, then

$$\begin{aligned} \Pr\{M+N > \nu\} &= \Pr\{|z_{m_k, n_k}| \leq a, k=1, 2, \dots, \nu\} \\ &\leq \Pr\{z_{m_\nu, n_\nu}^* \leq x\}, \end{aligned}$$

where

$$z_{m_\nu, n_\nu}^* = [z_{m_\nu, n_\nu} - i_\nu^2 \delta] / i_\nu, \quad x = [-i_\nu^2 \delta + a] / i_\nu$$

and

$$i_\nu^2 = m_\nu n_\nu / \nu.$$

Hence, we can write

$$\Pr\{M+N > \nu\} \leq \Phi(x),$$

where Φ denotes the standard normal distribution function.

Noticing that $x \leq 0$ for sufficiently large ν , and using the fact that $\min\{m_\nu, n_\nu\} > (\frac{1}{2}\nu)^{1/r}$, some simple analysis shows that

$$\Pr\{M+N > \nu\} \leq C_1 e^{-C_2 \nu^{1/r}} \quad (4.6)$$

for some constants $C_1, C_2 > 0$. Since the right-hand side of (4.6) is exponentially decreasing in $p = \exp(-C_2)$ for finite values of r only, the required result follows. $\square$

4.1.3 Estimation

The problem of estimation has already been investigated in Chapters 2 and 3. In Chapter 3, explicit expressions were obtained for the bias and variance of the estimated treatment difference. For the sequential test described in §4.1.1 (using the Robbins & Siegmund rule with $c > a$), Woodroffe (1989) has derived an asymptotic expansion (as $a \rightarrow \infty$) for the distribution function of

$$Z_s = \sqrt{\left\{ \frac{MN}{M+N} \right\}} (\hat{\delta} - \delta), \quad (4.7)$$

where

$$\hat{\delta} = \bar{x}_{BN} - \bar{x}_{AM}.$$

This expansion can be written as

$$\gamma^*(\xi) = \Pr(Z_s \leq \xi) \approx \Phi(\xi) - \frac{\text{sgn}(\delta)}{2\sqrt{a|\delta|}} \varphi(\xi) \quad (4.8)$$

(cf. equation (17) of Woodroffe, 1989), where $\varphi(\xi)$ denotes the standard normal density function. Woodroffe notes that the following is asymptotically equivalent to (4.8)

$$\gamma^{**}(\xi) \approx \Phi\left\{ \xi - \frac{\text{sgn}(\delta)}{2\sqrt{a|\delta|}} \right\}. \quad (4.9)$$

The derivation of (4.8) uses a version of Stein's Identity (see Stein, 1987, Chap.2) instead of a Taylor series expansion; use of the former avoids certain integrability problems.

Remark 4.1. From Woodroffe (1989, §6), we note that (4.8) and (4.9) are valid approximations for other data-dependent allocation rules if as $a \rightarrow \infty$, the final numbers of patients on the two treatments depend only on μ_A and μ_B through their difference, δ – intuitively, the latter property might be expected to hold. Simulation results in §4.2.3 demonstrate the insensitivity of the distribution function of (4.7) to the allocation rule used.

If we denote the bias in $\hat{\delta}$ by $b(\delta)$, then we see from (4.8) and (4.9) that asymptotically,

$$b(\delta) = \sqrt{\left\{ \frac{M+N}{MN} \right\}} \frac{\text{sgn}(\delta)}{2\sqrt{a|\delta|}}. \quad (4.10)$$

Thus, a bias corrected estimate of δ , $\tilde{\delta}$, is obtained by solving the cubic equation

$$\tilde{\delta} = \hat{\delta} - b(\tilde{\delta}) \quad (4.11)$$

(cf. approach in §3.6.1). Following Whitehead (1986, p.578), use of a Taylor series expansion of $b(\tilde{\delta})$ about $b(\delta)$ shows that to first-order terms $\tilde{\delta}$ is asymptotically unbiased for δ and

$$\text{Var}(\tilde{\delta}) \leq \text{Var}(\hat{\delta}) [1 + b'(\delta)]^{-2}, \quad (4.12)$$

where $\text{Var}(\hat{\delta}) = (M+N)/MN$. Simulation results for (4.10) and (4.12) are presented in §4.2.3 to illustrate the effect of the bias correction (4.11). The cubic equation (4.11) was solved numerically using the NAG routine C02AEF; this routine solves a real polynomial using the method of Grant and Hitchins. Since (4.11) has three real roots (e.g. Abramowitz & Stegun, 1964, p.17), $\tilde{\delta}$ is taken to be the largest root which is less than $\hat{\delta}$. Values for (4.12) are obtained by substituting $\tilde{\delta}$ for δ .

4.2 Simulation results

4.2.1 Details of simulation study

A FORTRAN program, SPRTSNEW, was written to simulate the sequential tests described in §4.1, and more general tests which will be derived in §§4.3, 4.5 and 4.7. A comparison of the three allocation rules described in §4.1.2 is carried out in §4.2.2; the results are based on 1000 simulations and so can readily be compared with those of Robbins & Siegmund (1974, pp.134–135). Corresponding results for estimation are presented in §4.2.3, and are based on 10000 simulations, as in Woodroffe (1989, §7).

The above program implements NAG routines as described in §§2.3.1 and 3.5.2; details of these are therefore omitted here. Values for the Gittins index for normal responses with known variance are given by Gittins (1989, pp.218–219) for a selection of values for the sample size; for other sample sizes, approximate values for the indices were obtained by interpolation.

4.2.2 Comparison of allocation rules

Simulation results are presented in Tables 4.1 – 4.3 for a range of values for the true mean treatment difference, δ ; due to the symmetry of the problem, only positive values for δ were considered. All the results are based on $a = 6$, which for $\delta = 0.25$, gives an error probability of approximately 0.05 using (4.2). The value of c for the Robbins & Siegmund rule is 6, and the discount factor and r for the Gittins rule are 0.99 and 1.5, respectively. Values for the approximations (4.2) and (4.3) are also included.

Several conclusions emerge from the results. Firstly, both (4.2) and (4.3) provide good approximations to the simulated values, the discrepancies being mainly due to the effect of overshoot. Secondly, while the error probabilities and the information are

Table 4.1: *Simulation results for the error probability for three allocation rules (stable normal response variable)*

δ	Approx.	Pr(error)		Gittins
		Rob.Sieg.	Prop.Ran.	
0.05	0.354	0.358	0.350	0.354
0.1	0.231	0.221	0.228	0.212
0.17	0.115	0.117	0.103	0.114
0.25	0.047	0.029	0.041	0.025
0.5	0.002	0.001	0.002	0.001
1.0	0.000	0.000	0.000	0.000
2.0	0.000	0.000	0.000	0.000
4.0	0.000	0.000	0.000	0.000

Table 4.2: *Simulation results for the expected sample sizes for three allocation rules (stable normal response variable)*

δ	RS	E(M)		RS	E(N)	
		PR	G		PR	G
0.05	90.0	76.2	94.3	113.8	76.5	116.0
0.1	75.8	69.4	73.4	122.8	69.4	115.7
0.17	51.9	58.5	56.3	122.3	59.0	111.5
0.25	35.7	45.3	39.1	120.1	45.8	103.8
0.5	15.2	25.1	18.3	91.8	26.1	60.9
1.0	7.1	12.0	8.8	63.9	14.3	23.8
2.0	3.6	5.8	4.9	35.6	8.7	9.7
4.0	2.1	3.0	2.9	12.7	4.7	3.8

Table 4.3: *Simulation results for the information for three allocation rules (stable normal response variable)*

δ	$E\{MN/(M+N)\}$			
	Approx.	Rob.Sieg.	Prop.Ran.	Gittins
0.05	35.0	36.1	37.9	38.2
0.1	32.2	34.5	34.5	33.5
0.17	27.2	27.4	29.1	29.4
0.25	21.7	23.3	22.5	24.0
0.5	11.9	12.2	12.5	13.3
1.0	6.0	6.2	6.3	6.3
2.0	3.0	3.2	3.3	3.3
4.0	1.5	1.6	1.7	1.6

approximately independent of the allocation rule used, the corresponding differences in the expected sample sizes for the three rules are striking: whereas for $\delta \leq 0.25$ the Gittins rule and the Robbins & Siegmund rule are comparable, for large δ the total expected sample size for the Gittins rule is significantly smaller at the cost of slightly more patients on the inferior treatment; moreover, for $\delta \geq 2$ the total expected sample size for the Gittins rule is only slightly larger than for proportionate randomisation.

4.2.3 Estimation

Simulation results for the distribution function of (4.7) are presented in Tables 4.4 and 4.5 for $\delta = 0.5$ and 1.0 , respectively, when $a = 6$. Further, since the approximations (4.8) and (4.9) require that $c > a$ for the Robbins & Siegmund rule, we have chosen $c = 12$ as in Woodroffe (1989, §7). Values are also included for a direct normal approximation.

It is clear that the amount by which the direct normal approximation overestimates the distribution can be substantial in the middle part of the distribution. The approximations (4.8) and (4.9) are considerable improvements: both tend to "overcorrect" for $\xi \geq 0$ while for $\xi < -1$, (4.9) tends to "undercorrect" and for $\xi < 0$, (4.8) "overcorrect". Furthermore, whereas (4.8) is the better approximation for $\xi \geq 0$, (4.9) appears better for $\xi < 0$. Results obtained for other values of δ indicate that the approximations work well for $\delta \geq 0.5$, but for $\delta < 0.5$ the distribution function is considerably underestimated when $\xi < 0$.

Table 4.4: *Simulated and approximate values for $Pr(Z_\delta \leq \xi)$ when $\delta = 0.5$ (stable normal response variable)*

ξ	$\Phi(\xi)$	Approximate		RS	Simulated	G
		$\gamma^*(\xi)$	$\gamma^{**}(\xi)$		PR	
-2.00	0.0228	0.0072	0.0110	0.0144	0.0103	0.0105
-1.75	0.0401	0.0152	0.0207	0.0246	0.0189	0.0206
-1.50	0.0668	0.0294	0.0368	0.0393	0.0348	0.0395
-1.25	0.1056	0.0529	0.0619	0.0625	0.0626	0.0639
-1.00	0.1587	0.0888	0.0988	0.0991	0.0976	0.1037
-0.75	0.2266	0.1397	0.1495	0.1511	0.1468	0.1557
-0.50	0.3085	0.2069	0.2152	0.2172	0.2134	0.2243
-0.25	0.4013	0.2897	0.2951	0.2950	0.2973	0.3115
0.00	0.5000	0.3848	0.3864	0.3921	0.3909	0.4064
0.25	0.5987	0.4871	0.4846	0.4981	0.4919	0.5096
0.50	0.6915	0.5898	0.5837	0.6068	0.5905	0.6066
0.75	0.7734	0.6864	0.6777	0.6986	0.6908	0.7022
1.00	0.8413	0.7715	0.7616	0.7840	0.7745	0.7813
1.25	0.8944	0.8416	0.8318	0.8496	0.8453	0.8487
1.50	0.9332	0.8958	0.8871	0.8987	0.8994	0.9039
1.75	0.9599	0.9350	0.9280	0.9387	0.9365	0.9448
2.00	0.9772	0.9617	0.9565	0.9657	0.9619	0.9656

Table 4.5: *Simulated and approximate values for $Pr(Z_\delta \leq \xi)$ when $\delta = 1.0$ (stable normal response variable)*

ξ	$\Phi(\xi)$	Approximate		RS	Simulated	G
		$\gamma^*(\xi)$	$\gamma^{**}(\xi)$		PR	
-2.00	0.0228	0.0117	0.0138	0.0128	0.0135	0.0138
-1.75	0.0401	0.0224	0.0253	0.0230	0.0247	0.0257
-1.50	0.0668	0.0404	0.0442	0.0417	0.0409	0.0434
-1.25	0.1056	0.0684	0.0730	0.0707	0.0697	0.0730
-1.00	0.1587	0.1093	0.1143	0.1141	0.1092	0.1161
-0.75	0.2266	0.1652	0.1700	0.1714	0.1668	0.1725
-0.50	0.3085	0.2367	0.2407	0.2432	0.2386	0.2415
-0.25	0.4013	0.3224	0.3249	0.3275	0.3277	0.3239
0.00	0.5000	0.4186	0.4191	0.4263	0.4196	0.4240
0.25	0.5987	0.5198	0.5183	0.5353	0.5202	0.5364
0.50	0.6915	0.6196	0.6163	0.6332	0.6213	0.6212
0.75	0.7734	0.7119	0.7074	0.7143	0.7127	0.7229
1.00	0.8413	0.7920	0.7869	0.8047	0.7963	0.7896
1.25	0.8944	0.8571	0.8522	0.8589	0.8622	0.8558
1.50	0.9332	0.9068	0.9025	0.9162	0.9113	0.9044
1.75	0.9599	0.9423	0.9389	0.9429	0.9437	0.9460
2.00	0.9772	0.9662	0.9637	0.9685	0.9677	0.9639

On comparing the simulation results for the three allocation rules, we see that the particular adaptive rule has little effect, if any, upon the sampling distribution of (4.7).

Simulation results for the bias and variance of $\hat{\delta}$ and $\tilde{\delta}$ are presented in Table 4.6 for $\delta = 1.0$. Quantities denoted by "Theor." are averages (over the simulations) of the appropriate expressions in §4.1.3 and "Empir." denotes the corresponding empirical values. The results for the biases demonstrate that, although (4.10) is the asymptotic bias in $\hat{\delta}$, an adjustment via (4.11) can yield an estimate of δ which is significantly less biased. Moreover, the corresponding increase in variance is negligible.

Table 4.6: *Simulated and approximate values for the bias and variance of $\hat{\delta}$ and $\tilde{\delta}$ when $\delta = 1.0$ (stable normal response variable)*

			Rob.Sieg.	Prop.Ran.	Gittins
Bias	$\hat{\delta}$	Theor.	0.0853	0.0847	0.0847
		Empir.	0.1476	0.1601	0.1494
Variance	$\tilde{\delta}$	Empir.	0.0622	0.0755	0.0648
		$\hat{\delta}$	Theor.	0.1825	0.1820
	$\hat{\delta}$	Theor.	0.1825	0.1820	0.1803
		Empir.	0.2010	0.2050	0.2050
	$\tilde{\delta}$	Theor.	0.1987	0.1978	0.1962
		Empir.	0.2026	0.2066	0.2066

4.3 Sequential tests for a normal response variable with linear time-trends

We now suppose that the response variable is subject to a linear time-trend; thus, using the notation of §4.1, we assume $X_{ij} \sim N(\mu_i + tj, 1)$, ($i=A,B$; $j=0,1,2,\dots$), where t is the linear trend per unit time (the same for both treatments). If t is known, the sequential test given in §4.1.1 is still appropriate (with a slight modification in the numerator of (4.1) to make $\bar{x}_{Bn} - \bar{x}_{Am}$ an unbiased estimator of δ). However, t would usually be unknown at the outset and consequently the sequential test of §4.1 would need modifying in order to preserve the original error probabilities. This problem is now addressed.

4.3.1 Derivation of tests

Let $\theta = \frac{1}{2}(\mu_A + \mu_B)$ and a_{Aj} be an indicator variable with value one if patient $j+1$ (at time j) is allocated to treatment A, and zero otherwise. Then the log-likelihood of

the observed responses (after m patients have been allocated to treatment A and n to treatment B) can be written as

$$L(\underline{x}, \delta, \theta, t) = -\frac{1}{2}(m+n)\log(2\pi) - \frac{1}{2} \left[\sum_{j=0}^{m+n-1} a_{Aj} \{x_{Aj} - \theta + \frac{1}{2}\delta - t_j\}^2 + \sum_{j=0}^{m+n-1} (1-a_{Aj}) \{x_{Bj} - \theta - \frac{1}{2}\delta - t_j\}^2 \right]. \quad (4.13)$$

Suppose that we wish to test H_0^* against H_1^* as defined in §4.1.1, then since θ and t are *unknown nuisance parameters*, we can apply a combination of the procedures of Bartlett (1946) and Cox (1963). Thus, we consider a Taylor series expansion about (δ, θ, t) of

$$\log \lambda = L(\underline{x}, \delta_1, \hat{\theta}_1, \hat{t}) - L(\underline{x}, \delta_0, \hat{\theta}_0, \hat{t}) \quad (4.14)$$

in which $\hat{\theta}_k$ is the maximum likelihood estimate of θ assuming H_k^* ($k=0,1$) and $\hat{t}$ is the maximum likelihood estimate of t based on $\underline{x}$; $\delta_0 = -\delta^*$ and $\delta_1 = \delta^*$. A little algebra yields the expansion

$$\begin{aligned} \log \lambda = & (\hat{\theta}_1 - \hat{\theta}_0) \frac{\partial L}{\partial \theta} + (\delta_1 - \delta_0) \frac{\partial L}{\partial \delta} + \frac{1}{2}(\hat{\theta}_1 - \hat{\theta}_0)(\hat{\theta}_1 + \hat{\theta}_0 - 2\theta) \frac{\partial^2 L}{\partial \theta^2} \\ & + \frac{1}{2}(\delta_1 - \delta_0)(\delta_1 + \delta_0 - 2\delta) \frac{\partial^2 L}{\partial \delta^2} + (\hat{t} - t)(\hat{\theta}_1 - \hat{\theta}_0) \frac{\partial^2 L}{\partial \theta \partial t} \\ & + [\hat{\theta}_1 \delta_1 - \hat{\theta}_0 \delta_0 - \theta(\delta_1 - \delta_0) - \delta(\hat{\theta}_1 - \hat{\theta}_0)] \frac{\partial^2 L}{\partial \theta \partial \delta} \\ & + (\hat{t} - t)(\delta_1 - \delta_0) \frac{\partial^2 L}{\partial \delta \partial t}, \end{aligned} \quad (4.15)$$

where L is defined by (4.13) and its derivatives are evaluated at (δ, θ, t) .

By considering expansions of $\partial L(\underline{x}, \hat{\delta}, \hat{\theta}, \hat{t})/\partial \theta$ and $\partial L(\underline{x}, \hat{\delta}, \hat{\theta}, \hat{t})/\partial \delta$ about (δ, θ, t) , some tedious algebra yields

$$\lambda = \exp \left\{ 2\delta^* \frac{mn}{(m+n)} \hat{\delta} \right\}, \quad (4.16)$$

where $\hat{\delta}$ is the maximum likelihood estimate of δ based on $\underline{x}$.

Using (4.13), we can derive three simultaneous equations for the maximum likelihood estimates of (δ, θ, t) . Solving for $\hat{\delta}$, we obtain

$$\hat{\delta} = \frac{\{(\bar{x}_{Bn} - \bar{x}_{Am}) + QD\}}{C}, \quad (m, n \geq 2) \quad (4.17)$$

in which

$$Q = \frac{6}{mn} \left[2 \sum_{j=0}^{m+n-1} j a_{Aj} - m(m+n-1) \right],$$

$$C = 1 - mnQ^2 / [12(m+n-1)(m+n+1)]$$

and

$$D = \frac{1}{(m+n-1)(m+n+1)} \left[-\frac{1}{2}(m+n-1)(m\bar{x}_{Am} + n\bar{x}_{Bn}) + \sum_{j=0}^{m+n-1} j \{ a_{Aj} x_{Aj} + (1-a_{Aj}) x_{Bj} \} \right].$$

Note that (4.17) can also be obtained using an analysis of covariance (i.e. treating time as a covariate).

Thus, our modified test statistic is given by

$$z_{m,n} = \frac{mn}{(m+n)} \hat{\delta} \quad (4.18)$$

with $\hat{\delta}$ given by (4.17). However, we now need to widen the SPRT boundaries $(-a, a)$ of §4.1.1 to allow for the variability due to estimation of the trend. This aspect of the problem is now considered.

4.3.2 Adjusting the SPRT boundaries

Using Cox (1963, pp.7-8), we see that the SPRT boundaries for test statistic (4.18) should be $(-a_{m,n}, a_{m,n})$, where

$$a_{m,n} = a \frac{mn}{(m+n)} \text{Var}(\hat{\delta}), \quad (m, n \geq 1). \quad (4.19)$$

Now, since we can write

$$\frac{\partial L}{\partial \underline{\beta}} = I(\underline{\beta}) (\hat{\underline{\beta}} - \underline{\beta}), \quad (4.20)$$

where $\underline{\beta}^T = (\delta, \theta, t)$ and $I(\underline{\beta})$ is the *Fisher information matrix*, it is evident from Silvey (1975, §2.12) that $\text{Var}(\hat{\underline{\beta}})$ attains the Cramér-Rao lower bound. Thus, we can either obtain $\text{Var}(\hat{\delta})$ directly from (4.17) or from $I^{-1}(\underline{\beta})$. In either case, a little algebra gives

$$\text{Var}(\hat{\delta}) = \frac{(m+n)}{mn} \frac{1}{C}, \quad (m, n \geq 2) \quad (4.21)$$

with C as defined for (4.17). Again, (4.21) can also be obtained using an analysis of covariance.

Summarising, we define (M,N) = first (m,n) such that $z_{m,n} \notin (-a_{m,n}, a_{m,n})$ and for $M+N < \infty$, accept H_0^* or H_1^* according as $z_{M,N} \leq -a_{M,N}$ or $z_{M,N} \geq a_{M,N}$, respectively, where $z_{m,n}$ and $a_{m,n}$ are given by (4.18) and (4.19). Using this test, we would expect the error probabilities to be given approximately by (4.2). Furthermore, using a similar argument to Hayre & Gittins (1981, p.698), it can be shown that if $E\{MN/(M+N)\} < \infty$,

$$E\left\{ \frac{MN}{M+N} \right\} \approx \begin{cases} aE(a_{M,N}) & , \quad (\delta=0) \\ \frac{E(a_{M,N})(e^{2a\delta} - 1)}{\delta(e^{2a\delta} + 1)} & , \quad (\delta \neq 0) \end{cases} \quad (4.22)$$

(cf. equation (4.3)).

4.3.3 Estimation

In analogy to §4.1.3, we seek an asymptotic expansion for the distribution function of

$$Z_s(t) = \frac{(\hat{\delta} - \delta)}{\sqrt{\text{Var}(\hat{\delta})}} \quad , \quad (4.23)$$

where $\hat{\delta}$ and $\text{Var}(\hat{\delta})$ are given, respectively, by (4.17) and (4.21) with $(m,n) = (M,N)$. When the SPRT boundaries are unadjusted, Woodroffe's (1989, §9) results for a stable normal response variable setting in which there is an unknown variance σ^2 in the responses led to Conjecture 4.1. However, from power considerations we might expect the distribution function of (4.23) for the modified test in §4.3.2 to be asymptotically the same as for a stable response variable (Conjecture 4.2).

Conjecture 4.1. If the sequential test defined by (4.18) is used but the SPRT boundaries are unadjusted, the distribution function of (4.23) is asymptotically

$$\Pr(Z_s(t) \leq \xi) \approx \Phi(\xi) - \frac{\sigma \text{sgn}(\delta)}{2\sqrt{(a|\delta|)}} \varphi(\xi) \quad (4.24)$$

(cf. equation (4.8)), where $\sigma^2 = E(a_{M,N})/a$.

Conjecture 4.2. For the sequential test defined by (4.18) and (4.19), the distribution function of (4.23) is asymptotically

$$\Pr(Z_s(t) \leq \xi) \approx \Phi(\xi) - \frac{\text{sgn}(\delta)}{2\sqrt{(a|\delta|)}} \varphi(\xi) \quad .$$

Formal proofs of the above two results have not yet been attempted. It would appear that a generalisation of the results of Woodroffe (1989) may be required. We shall see in §4.4.2 that simulation results indicate that the accuracy of the approximation in Conjecture 4.2 is comparable with that for a stable normal response variable in §4.2.3.

4.4 Simulation results

4.4.1 Comparison of allocation rules

Simulation results for the modified sequential test are presented in Tables 4.7 – 4.10 for a range of values for δ with a trend, $t = 0.0005$. For comparison purposes, the corresponding results for the sequential test in §4.1.1 (which, of course, does not adjust for estimation of the trend) are given in parentheses. The tabulation is as before (Tables 4.1 – 4.3) with two exceptions: firstly, as an approximation to the information, we have computed the product of the simulated value for $E(a_{M,N})/a$ and the simulated value for the information given in Table 4.3; and secondly, two patients are randomly allocated to each treatment initially.

On comparing the results in parentheses with those in Tables 4.1 – 4.3, we see that the error probabilities have generally increased as a result of the trend and the expected sample sizes on the two treatments are slightly smaller than in the stable case. This behaviour is expected because the stopping boundaries will be crossed earlier due to the increased variability in the responses. Also, in the presence of a trend, the size of the error probability depends on the particular allocation rule used – the trend has a greater effect upon a more adaptive rule due to the imbalance in the sample sizes.

We see from Table 4.7 that our modified sequential test works well – in many cases, the simulated error probabilities are smaller than those when the trend is known and, furthermore, there is good evidence that the error probabilities are approximately independent of the allocation rule. From Table 4.10, we see that (4.22) provides a reasonable approximation to the simulated values, although there are some rather large discrepancies for some of the rules when δ is small; however, this inaccuracy is comparable with that for the stable setting (see Table 4.3).

Table 4.7: *Simulation results for the error probability for three allocation rules (linear time-trends in normal response variable)*

δ	Approx.	Pr(error)		Prop.Ran.	Gittins
		Rob.Sieg.			
0.05	0.354	0.340 (0.399)		0.345 (0.351)	0.351 (0.358)
0.1	0.231	0.211 (0.240)		0.218 (0.213)	0.232 (0.252)
0.17	0.115	0.116 (0.139)		0.117 (0.100)	0.125 (0.117)
0.25	0.047	0.041 (0.054)		0.039 (0.041)	0.034 (0.044)
0.5	0.002	0.000 (0.003)		0.000 (0.004)	0.000 (0.004)
1.0	0.000	0.000 (0.000)		0.000 (0.000)	0.000 (0.000)
2.0	0.000	0.000 (0.000)		0.000 (0.000)	0.000 (0.000)
4.0	0.000	0.000 (0.000)		0.000 (0.000)	0.000 (0.000)

Table 4.8: *Simulation results for the expected sample size on the inferior treatment for three allocation rules (linear time-trends in normal response variable)*

δ	Rob.Sieg.	E(M)		Gittins
		Prop.Ran.		
0.05	152.6 (78.7)	77.7 (76.2)		134.5 (79.2)
0.1	125.6 (65.2)	70.5 (70.3)		107.4 (68.0)
0.17	86.9 (52.3)	57.2 (59.5)		76.5 (52.1)
0.25	55.2 (35.9)	46.7 (46.5)		43.8 (38.3)
0.5	18.0 (15.6)	25.6 (25.4)		17.9 (17.8)
1.0	7.5 (7.3)	12.6 (12.2)		9.3 (8.7)
2.0	3.9 (3.7)	6.1 (5.8)		5.2 (4.9)
4.0	2.1 (2.1)	3.4 (3.1)		3.2 (3.0)

Table 4.9: *Simulation results for the expected sample size on the better treatment for three allocation rules (linear time-trends in normal response variable)*

δ	Rob.Sieg.	E(N)		Gittins
		Prop.Ran.		
0.05	232.9 (94.3)	77.8 (76.6)		175.2 (98.8)
0.1	247.1 (103.7)	71.0 (70.2)		192.1 (97.5)
0.17	244.4 (105.9)	57.7 (59.7)		171.9 (99.7)
0.25	231.3 (104.2)	47.4 (47.1)		147.3 (89.9)
0.5	188.3 (90.0)	26.9 (26.4)		69.5 (55.9)
1.0	120.9 (62.6)	15.3 (14.5)		26.1 (23.7)
2.0	66.0 (34.4)	8.9 (8.3)		10.0 (9.7)
4.0	27.8 (11.1)	4.8 (4.1)		4.3 (3.9)

Table 4.10: *Simulation results for the information for three allocation rules (linear time-trends in normal response variable)*

δ	Rob.Sieg.		E{MN/(M+N)}				Gittins	
	Approx.	Simul.	Prop.Ran.		Approx.	Simul.	Approx.	Simul.
0.05	43.4	45.5 (30.1)	38.3	38.6 (38.0)			44.0	42.8 (32.1)
0.1	41.2	43.7 (28.9)	34.8	35.1 (34.9)	39.2	41.7 (29.7)		
0.17	32.3	35.4 (26.1)	29.4	28.5 (29.6)	33.8	35.2 (26.2)		
0.25	27.1	27.5 (21.1)	22.8	23.3 (23.2)	26.9	27.6 (21.6)		
0.5	13.7	14.5 (12.2)	12.7	12.9 (12.7)	14.1	13.8 (12.7)		
1.0	6.8	6.9 (6.2)	6.5	6.7 (6.4)	6.5	6.8 (6.3)		
2.0	3.5	3.6 (3.2)	3.4	3.5 (3.2)	3.3	3.4 (3.3)		
4.0	1.9	1.9 (1.6)	1.8	1.9 (1.7)	1.7	1.8 (1.7)		

On comparing the expected sample sizes for the modified sequential test with those in Table 4.2, it is clear that, on the average, the total expected sample size has been doubled for the Robbins & Siegmund rule and for small δ has been increased by about 50% for the Gittins rule; for $\delta \leq 0.25$, there is also a significant increase in the number of patients on the inferior treatment for these rules. As might be expected, the corresponding increase is slight for the proportionate randomisation rule (for $\delta = 0.17$, there is a decrease – an indication of the magnitude of random error involved!).

4.4.2 Estimation

Results for the distribution function of (4.23) when using the modified sequential test are presented in Tables 4.11 and 4.12 for $\delta = 0.5$ and 1.0 , respectively. Corresponding results for the bias and variance of $\hat{\delta}$ and $\tilde{\delta}$ are given in Table 4.13 for $\delta = 1.0$.

As in the stable case, the approximations to the distribution function of (4.23) are considerable improvements over a direct Normal approximation with the results exhibiting a similar pattern to that in Tables 4.4 and 4.5. The sampling distribution of (4.23) is insensitive to the data-dependent allocation rule adopted and thus the results provide strong evidence in support of Conjecture 4.2. Finally, the results in Table 4.13 yield similar conclusions to before.

Table 4.11: *Simulated and approximate values for $Pr(Z_{\xi}^{(t)} \leq \xi)$ when $\delta = 0.5$ (linear time-trends in normal response variable)*

ξ	$\Phi(\xi)$	Approximate		RS	Simulated	G
		$\gamma^*(\xi)$	$\gamma^{**}(\xi)$		PR	
-2.00	0.0228	0.0072	0.0110	0.0119	0.0097	0.0109
-1.75	0.0401	0.0152	0.0207	0.0220	0.0192	0.0190
-1.50	0.0668	0.0294	0.0368	0.0365	0.0342	0.0347
-1.25	0.1056	0.0529	0.0619	0.0607	0.0588	0.0610
-1.00	0.1587	0.0888	0.0988	0.0959	0.0970	0.1015
-0.75	0.2266	0.1397	0.1495	0.1493	0.1488	0.1558
-0.50	0.3085	0.2069	0.2152	0.2173	0.2186	0.2263
-0.25	0.4013	0.2897	0.2951	0.3063	0.3061	0.3084
0.00	0.5000	0.3848	0.3864	0.3992	0.3962	0.4043
0.25	0.5987	0.4871	0.4846	0.5007	0.5001	0.5103
0.50	0.6915	0.5898	0.5837	0.6032	0.6054	0.6046
0.75	0.7734	0.6864	0.6777	0.7040	0.6983	0.7065
1.00	0.8413	0.7715	0.7616	0.7835	0.7816	0.7862
1.25	0.8944	0.8416	0.8318	0.8529	0.8476	0.8525
1.50	0.9332	0.8958	0.8871	0.9045	0.9011	0.9076
1.75	0.9599	0.9350	0.9280	0.9419	0.9384	0.9428
2.00	0.9772	0.9617	0.9565	0.9635	0.9649	0.9629

Table 4.12: *Simulated and approximate values for $Pr(Z_{\xi}^{(t)} \leq \xi)$ when $\delta = 1.0$ (linear time-trends in normal response variable)*

ξ	$\Phi(\xi)$	Approximate		RS	Simulated	G
		$\gamma^*(\xi)$	$\gamma^{**}(\xi)$		PR	
-2.00	0.0228	0.0117	0.0138	0.0156	0.0143	0.0128
-1.75	0.0401	0.0224	0.0253	0.0279	0.0271	0.0252
-1.50	0.0668	0.0404	0.0442	0.0465	0.0453	0.0438
-1.25	0.1056	0.0684	0.0730	0.0773	0.0727	0.0717
-1.00	0.1587	0.1093	0.1143	0.1185	0.1133	0.1118
-0.75	0.2266	0.1652	0.1700	0.1761	0.1710	0.1671
-0.50	0.3085	0.2367	0.2407	0.2488	0.2426	0.2360
-0.25	0.4013	0.3224	0.3249	0.3330	0.3263	0.3211
0.00	0.5000	0.4186	0.4191	0.4334	0.4190	0.4278
0.25	0.5987	0.5198	0.5183	0.5357	0.5168	0.5317
0.50	0.6915	0.6196	0.6163	0.6329	0.6138	0.6130
0.75	0.7734	0.7119	0.7074	0.7261	0.7060	0.7311
1.00	0.8413	0.7920	0.7869	0.7978	0.7912	0.7905
1.25	0.8944	0.8571	0.8522	0.8639	0.8542	0.8671
1.50	0.9332	0.9068	0.9025	0.9079	0.9058	0.9029
1.75	0.9599	0.9423	0.9389	0.9482	0.9421	0.9479
2.00	0.9772	0.9662	0.9637	0.9667	0.9675	0.9691

Table 4.13: Simulated and approximate values for the bias and variance of $\hat{\delta}$ and $\tilde{\delta}$ when $\delta = 1.0$ (linear time-trends in normal response variable)

			Rob.Sieg.	Prop.Ran.	Gittins
Bias	$\hat{\delta}$	Theor.	0.0804	0.0834	0.0822
		Empir.	0.1470	0.1443	0.1576
Variance	$\tilde{\delta}$	Empir.	0.0666	0.0609	0.0754
		Theor.	0.1639	0.1739	0.1726
	$\hat{\delta}$	Empir.	0.1795	0.1904	0.1915
		Theor.	0.1779	0.1891	0.1898
	$\tilde{\delta}$	Theor.	0.1779	0.1891	0.1898
		Empir.	0.1807	0.1922	0.1926

4.5 Approximate sequential tests for a Bernoulli response variable with linear time-trends

Until now we have considered normal responses. Suppose instead we wished to derive a similar sequential test for Bernoulli responses. Robbins (1974) suggested such a test for which the error probability depends approximately only on the odds ratio of the two success probabilities. More recently, Hayre & Turnbull (1981c) proposed a class of (approximate) sequential tests, and applied their tests to the Bernoulli setting. The authors showed by simulation that in order to preserve the error probabilities of the test (to those in the setting of §4.1), the SPRT boundaries can be widened using adjustments of the form proposed for normal responses by Robbins *et al.* (1967, p.1387), to allow for the variability due to estimation of the two Bernoulli variances. [N.B. The problem of adjusting SPRT boundaries for normally distributed observations when the variance is unknown was first considered by Baker (1950) and has subsequently been studied by, among others, Hall (1962), Lai *et al.* (1977) and Hayre & Gittins (1981); unlike Robbins *et al.*, Baker and Hall used an initial sample of observations to estimate the variance rather than modifying the estimate as the sampling proceeds.]

We now generalise the approach taken in §4.3 to the case when $X_{ij} \sim N(\mu_i + tj, \sigma_i^2)$ ($i=A,B$; $j=0,1,2,\dots$), where the σ_i^2 are known and unequal. The derived sequential test is then used as an approximation for Bernoulli responses. Of course, the σ_i^2 are actually unknown, but simulation results in §4.6 indicate that our proposed sequential test performs well for a wide range of parameters even when the σ_i^2 are replaced by estimates.

4.5.1 Derivation of tests

Defining θ and a_{Aj} as in §4.3.1, the log-likelihood of the observed responses (after m patients have been allocated to treatment A and n to treatment B) can be written as

$$L(\underline{x}, \delta, \theta, t) = -\frac{1}{2}(m+n)\log(2\pi) - \frac{1}{2} \left[\sum_{j=0}^{m+n-1} a_{Aj} \{x_{Aj} - \theta + \frac{1}{2}\delta - t_j\}^2 / \sigma_A^2 + \sum_{j=0}^{m+n-1} (1-a_{Aj}) \{x_{Bj} - \theta - \frac{1}{2}\delta - t_j\}^2 / \sigma_B^2 \right]. \quad (4.25)$$

Proceeding as in §4.3.1 yields

$$\lambda = \exp \left\{ 2\delta^* \frac{mn}{(m\sigma_B^2 + n\sigma_A^2)} \hat{\delta} \right\} \quad (4.26)$$

(cf. equation (4.16)), and three simultaneous equations for the maximum likelihood estimates of (δ, θ, t) using (4.25). Solving for $\hat{\delta}$, we obtain

$$\hat{\delta} = \frac{\{(\bar{x}_{Bn} - \bar{x}_{Am}) + (m+n)QD^* / (12C^*)\}}{F}, \quad (m, n \geq 2), \quad (4.27)$$

with

$$F = 1 - \frac{mn(m+n)^2 \sigma_A^2 \sigma_B^2 Q^2}{144C^*},$$

where Q is as before,

$$C^* = (m\sigma_B^2 + n\sigma_A^2) \left[(\sigma_B^2 - \sigma_A^2) \sum_{j=0}^{m+n-1} j^2 a_{Aj} + \frac{1}{6} (m+n)(m+n-1)(2m+2n-1)\sigma_A^2 \right] - \left[\frac{1}{2}\sigma_A^2 (m+n)(m+n-1) \right]^2 - (\sigma_B^2 - \sigma_A^2) \sum_{j=0}^{m+n-1} j a_{Aj} \left[(\sigma_B^2 - \sigma_A^2) \sum_{j=0}^{m+n-1} j a_{Aj} + (m+n)(m+n-1)\sigma_A^2 \right]$$

and

$$D^* = (m\sigma_B^2 + n\sigma_A^2) \left[\sigma_B^2 \sum_{j=0}^{m+n-1} j a_{Aj} x_{Aj} + \sigma_A^2 \sum_{j=0}^{m+n-1} j (1-a_{Aj}) x_{Bj} \right] - (m\sigma_B^2 \bar{x}_{Am} + n\sigma_A^2 \bar{x}_{Bn}) \left[(\sigma_B^2 - \sigma_A^2) \sum_{j=0}^{m+n-1} j a_{Aj} + \frac{1}{2}(m+n)(m+n-1)\sigma_A^2 \right].$$

Note that (4.27) can also be obtained using a weighted analysis of covariance, with the weights $1/\sigma_A^2$ and $1/\sigma_B^2$.

Thus, our modified test statistic is given by

$$Z_{m,n} = \frac{mn}{(m\sigma_B^2 + n\sigma_A^2)} \hat{\delta}, \quad (4.28)$$

with $\hat{\delta}$ given by (4.27). We now obtain the corresponding adjustments to the SPRT boundaries.

4.5.2 Adjusting the SPRT boundaries

Analogous to (4.19), we see that the SPRT boundaries for test statistic (4.28) should be $(-a_{m,n}, a_{m,n})$, where

$$a_{m,n} = a \frac{mn}{(m\sigma_B^2 + n\sigma_A^2)} \text{Var}(\hat{\delta}), \quad (m, n \geq 1). \quad (4.29)$$

Proceeding as in §4.3.2, we obtain

$$\text{Var}(\hat{\delta}) = \frac{(m\sigma_B^2 + n\sigma_A^2)}{mn} \frac{1}{F}, \quad (4.30)$$

with F as for (4.27). Again, (4.30) can also be obtained using a weighted analysis of covariance.

Thus, we can form a sequential test as described in §4.3.2 which has error probabilities which are given approximately by (4.2). Furthermore, if $E\{MN/(M\sigma_B^2 + N\sigma_A^2)\} < \infty$ we obtain an analogous expression to (4.22), namely,

$$E\left\{ \frac{MN}{M\sigma_B^2 + N\sigma_A^2} \right\} \simeq \begin{cases} aE(a_{M,N}) & , \quad (\delta=0) \\ \frac{E(a_{M,N})(e^{2a\delta} - 1)}{\delta(e^{2a\delta} + 1)} & , \quad (\delta \neq 0) \end{cases}. \quad (4.31)$$

4.5.3 Application to Bernoulli responses

The results in §§4.5.1 and 4.5.2 can now be adapted for Bernoulli responses by letting $X_{ij} \sim \text{Bin}(1, p_{ij})$, where $p_{ij} = p_i + t_j$ and p_i is the unknown probability of success for treatment i for the first patient. Since the σ_i^2 are unknown, we replace them by the *empirical* estimates

$$\hat{\sigma}_{Am}^2 = \max\{\bar{x}_{Am}(1-\bar{x}_{Am}), m^{-1}\}, \quad \hat{\sigma}_{Bn}^2 = \max\{\bar{x}_{Bn}(1-\bar{x}_{Bn}), n^{-1}\} \quad (4.32)$$

as used by Hayre & Turnbull (1981c, p.2352), where $\bar{x}_{Am}$ and $\bar{x}_{Bn}$ are defined in §4.1.1.

We can apply the sequential test defined by (4.28) and (4.29) with the σ_i^2 replaced by the estimates (4.32). To a normal approximation and neglecting overshoot, the error probabilities for this test are given by (4.2).

4.5.4 Allocation rules considered for a Bernoulli response variable

(a) Robbins & Siegmund rule

This rule is a generalisation of that in §4.1.2: initially randomly allocate two patients to each treatment; subsequently, allocate the next patient to treatment B if

$$\frac{(n\hat{\sigma}_{Am}^2 - m\hat{\sigma}_{Bn}^2)}{(m\hat{\sigma}_{Bn}^2 + n\hat{\sigma}_{Am}^2)} \leq \frac{z_{m,n}}{c}, \quad (c \geq a), \quad (4.33)$$

where $z_{m,n}$ is given by (4.28); otherwise allocate to treatment A.

(b) Proportionate randomisation rule

This rule is as in §4.1.2 but with $\hat{\delta}_s$ defined by

$$\hat{\delta}_s = \sqrt{\left\{ \frac{mn}{m\hat{\sigma}_{Bn}^2 + n\hat{\sigma}_{Am}^2} \right\}} \hat{\delta}, \quad (m, n \geq 1), \quad (4.34)$$

where $\hat{\delta}$ is given by (4.27).

(c) Bather's rule and (d) the Gittins rule

Details of these rules are given in §§2.1.5 and 2.1.6, respectively; the results for the Gittins rule are based on the modification described in §4.1.2.

(e) The randomised Gittins rule

This allocation rule is the same as the Gittins rule except that the indices are now randomised (see Glazebrook, 1980) using the randomised element of Bather's rule. [It should, of course, be noted that although this randomised element was the result of theoretical and empirical work for Bather's rule, it possesses the necessary properties for use with the indices. Since no previous numerical results have been reported for randomised Gittins indices, the simulation results obtained in the present work provide valuable insight into their behaviour.] The randomised rule is initially to randomly allocate two patients to each treatment; subsequently, if $m^r < n$ or $n^r < m$ ($r \geq 1$), allocate the next patient to treatment A or B, respectively; otherwise allocate the next patient to treatment i if

$$G_i = \max\{G_A, G_B\}$$

with

$$G_A = \nu_A + \lambda(m)X_A(m+n), \quad G_B = \nu_B + \lambda(n)X_B(m+n) \quad (4.35)$$

in which ν_i is defined by (2.1) and $\lambda(\cdot)$, $X_i(\cdot)$ are defined in §2.1.5.

4.5.5 Estimation

Suppose that we are interested in the approximate distribution function of (4.23) when the treatment response is Bernoulli. If the σ_i^2 were known, we might expect, for reasons discussed in §4.3.3, that (4.8) and (4.9) would provide reasonable approximations. However, how good are the approximations when the σ_i^2 are estimated?

Simulation results for this setting were obtained for $(p_A, p_B) = \{(0.1, 0.6), (0.3, 0.8)\}$. Although not included here, the results indicated that for $\xi \leq 0$ approximations (4.8) and (4.9) are considerable improvements over a direct Normal approximation, while for $\xi > 0$ the simulated values for the distribution function were rather erratic (for some values of ξ , there was substantial variation between the allocation rules). For smaller values of $\delta = p_B - p_A$, the approximations are slightly better for $\xi > 0$ but poorer for $\xi \leq 0$. One possible explanation for this behaviour is that for smaller values of δ (though not too small), the estimated variances are more accurate and hence the approximations are generally better.

It would be interesting to investigate whether Woodroffe's (1989) results could be improved for this situation. However, this has not been attempted here.

4.6 Simulation results

4.6.1 Stable Bernoulli response variable

We begin by presenting in Tables 4.14 – 4.17 the results obtained using the sequential test in §4.1.1 for a stable Bernoulli response variable. The discount factor and r for the Gittins rules are 0.99 and 1.5, respectively. The tabulation is similar to that in §4.2.2 with values included for approximations (4.2) and (4.31) (with $a_{M,N} = a$).

It is clear from Table 4.14 that the discrepancies between the simulated and approximate error probabilities are substantial in some cases. When one or both of the p_i 's are less than about 0.2, the error probabilities are often smaller than is

Table 4.14: *Simulation results for the error probability for five allocation rules (stable Bernoulli response variable)*

δ	P_A	P_B	Approx.	RS	Pr(error)		G	RG
					PR	B		
0.05	0.1	0.15	0.354	0.330	0.327	0.311	0.304	0.302
	0.5	0.55	0.354	0.370	0.363	0.345	0.372	0.343
0.1	0.1	0.2	0.231	0.186	0.199	0.185	0.177	0.191
	0.4	0.5	0.231	0.247	0.253	0.219	0.268	0.264
0.2	0.2	0.4	0.083	0.087	0.079	0.074	0.087	0.076
	0.4	0.6	0.083	0.100	0.103	0.090	0.075	0.099
0.3	0.1	0.4	0.027	0.020	0.017	0.009	0.016	0.017
	0.4	0.7	0.027	0.030	0.037	0.029	0.036	0.028
0.4	0.1	0.5	0.008	0.006	0.005	0.001	0.004	0.003
	0.3	0.7	0.008	0.007	0.006	0.010	0.005	0.011
0.5	0.1	0.6	0.002	0.000	0.001	0.001	0.000	0.001
	0.3	0.8	0.002	0.002	0.004	0.002	0.000	0.001

indicated by the approximation. However, the results seem to suggest that there is little dependence of the error probability upon the allocation rule.

The results in Tables 4.15 and 4.16 demonstrate that, although the adaptive rules reduce the number of patients on the inferior treatment, the corresponding total expected sample sizes for these rules are in some cases much larger than those for proportionate randomisation. This is most evident for Bather's rule where the total expected sample size is approximately double that for proportionate randomisation when δ is large. In

Table 4.15: *Simulation results for the expected sample sizes on the inferior treatment for five allocation rules (stable Bernoulli response variable)*

δ	P_A	P_B	Rob.Sieg.	Prop.Ran.	E(M)		
					Bather	Gittins	Ran.Gittins
0.05	0.1	0.15	16.6	14.6	15.0	14.7	14.3
	0.5	0.55	19.0	16.8	17.6	18.2	15.3
0.1	0.1	0.2	13.2	13.7	13.0	13.1	12.2
	0.4	0.5	15.4	15.1	15.3	15.6	13.8
0.2	0.2	0.4	9.8	11.4	10.2	9.6	9.9
	0.4	0.6	10.6	12.3	10.9	10.5	10.6
0.3	0.1	0.4	6.4	8.8	6.6	6.6	7.0
	0.4	0.7	6.9	9.5	7.6	7.3	7.8
0.4	0.1	0.5	5.0	7.1	5.3	5.3	5.7
	0.3	0.7	5.2	7.5	5.4	5.6	5.9
0.5	0.1	0.6	4.1	5.9	4.5	4.6	4.7
	0.3	0.8	4.0	6.1	4.4	4.6	4.9

Table 4.16: *Simulation results for the expected sample size on the better treatment for five allocation rules (stable Bernoulli response variable)*

δ	p_A	p_B	Rob.Sieg.	Prop.Ran.	E(N)		
					Bather	Gittins	Ran.Gittins
0.05	0.1	0.15	21.9	14.7	18.2	17.1	15.8
	0.5	0.55	22.5	16.8	21.7	21.2	17.2
0.1	0.1	0.2	22.6	13.8	18.5	17.0	14.6
	0.4	0.5	22.4	15.2	22.6	20.7	16.2
0.2	0.2	0.4	21.9	11.5	20.2	15.3	13.6
	0.4	0.6	22.0	12.5	21.9	18.0	14.8
0.3	0.1	0.4	18.4	8.6	19.3	12.0	10.6
	0.4	0.7	18.0	9.4	20.0	14.7	12.0
0.4	0.1	0.5	16.3	7.3	20.7	10.0	9.1
	0.3	0.7	16.5	7.6	19.7	11.1	9.5
0.5	0.1	0.6	14.0	6.2	22.7	8.6	7.9
	0.3	0.8	12.2	6.3	19.2	8.7	8.0

contrast, the total expected sample sizes for the Gittins rules are only slightly larger than that for proportionate randomisation – the excess diminishing as δ increases.

The approximation (4.31) for the information appears to work reasonably well unless (one or) both of the p_i are smaller than about 0.2; this latter behaviour is mainly due to some estimates of the σ_i^2 being small. The inaccuracy for larger p_i tends to be of about the same magnitude as for a normal response variable (see Table 4.3), which suggests that it is mainly due to the effect of overshoot rather than to the normal approximation.

Table 4.17: *Simulation results for the information for five allocation rules (stable Bernoulli response variable)*

δ	p_A	p_B	$E\{MN/(M\sigma_B^2 + N\sigma_A^2)\}$					
			Approx.	Rob.Sieg.	Prop.Ran.	Bather	Gittins	Ran.Gittins
0.05	0.1	0.15	35.0	72.9	67.9	73.7	72.0	72.1
	0.5	0.55	35.0	34.5	35.6	34.9	37.4	33.7
0.1	0.1	0.2	32.2	57.7	55.6	58.7	60.1	54.1
	0.4	0.5	32.2	32.2	32.6	34.5	35.6	31.9
0.2	0.2	0.4	25.0	29.8	28.6	31.6	30.1	29.6
	0.4	0.6	25.0	26.6	26.8	27.5	27.4	26.7
0.3	0.1	0.4	18.9	23.4	24.1	24.8	23.3	23.1
	0.4	0.7	18.9	20.0	21.1	21.2	21.4	21.3
0.4	0.1	0.5	14.8	17.1	18.2	18.7	16.9	17.8
	0.3	0.7	14.8	16.4	17.1	16.7	16.8	16.7
0.5	0.1	0.6	11.9	13.4	14.1	15.1	13.9	13.9
	0.3	0.8	11.9	12.8	14.2	14.2	14.2	14.1

4.6.2 Linear time-trends in a Bernoulli response variable

We now subject the Bernoulli response variable to a trend, $t = 0.0025$. The simulation results are presented in Tables 4.18 – 4.20 using a similar format to §4.4.1.

Comparing the results in parentheses with those in Tables 4.14 – 4.16, the conclusions are similar to those discussed in §4.4.1.

It is clear from Table 4.18 that our approximate sequential test works well. When $\delta \geq 0.3$, the simulated error probability is generally no larger than the approximation (4.2), and when $\delta < 0.3$ and the p_i 's are small, the simulated error probability tends to be smaller than the approximation, but for larger p_i 's, they are slightly larger. This latter behaviour may be due to large fluctuations in the early estimates of the σ_i^2 . As for a normal response variable, the error probability is relatively insensitive to the allocation rule.

The results for the expected sample sizes on the two treatments for the modified sequential test are particularly revealing. Firstly, for large δ both Bather's rule and the Robbins & Siegmund rule allocate slightly fewer patients to the inferior treatment than proportionate randomisation, but in some cases the total expected sample sizes for these rules are considerably larger. Secondly, whereas for large δ the Gittins rules appear to retain total expected sample sizes which are only slightly larger than that for proportionate randomisation, for smaller values of δ and moderate values for the p_i , the corresponding total expected sample sizes for the Gittins rules are almost 50% larger (cf. conclusions in §4.6.1).

Lastly, the results for the information exhibited a similar behaviour to that in Table 4.17, and were therefore not included here.

Table 4.18: *Simulation results for the error probability for five allocation rules (linear time-trends in Bernoulli response variable)*

δ	P_A	P_B	Approx.	Rob.Sieg.		Pr(error) Prop.Ran.		Bather	
0.05	0.1	0.15	0.354	0.346	(0.365)	0.337	(0.350)	0.345	(0.331)
	0.5	0.55	0.354	0.359	(0.379)	0.338	(0.352)	0.375	(0.376)
0.1	0.1	0.2	0.231	0.217	(0.216)	0.204	(0.207)	0.207	(0.211)
	0.4	0.5	0.231	0.238	(0.285)	0.235	(0.244)	0.255	(0.261)
0.2	0.2	0.4	0.083	0.104	(0.111)	0.076	(0.082)	0.079	(0.093)
	0.4	0.6	0.083	0.098	(0.129)	0.097	(0.104)	0.084	(0.100)
0.3	0.1	0.4	0.027	0.029	(0.034)	0.018	(0.023)	0.018	(0.019)
	0.4	0.7	0.027	0.030	(0.030)	0.023	(0.038)	0.020	(0.028)
0.4	0.1	0.5	0.008	0.006	(0.009)	0.003	(0.004)	0.007	(0.006)
	0.3	0.7	0.008	0.007	(0.011)	0.011	(0.007)	0.005	(0.012)
0.5	0.1	0.6	0.002	0.000	(0.001)	0.000	(0.002)	0.000	(0.001)
	0.3	0.8	0.002	0.004	(0.000)	0.000	(0.004)	0.002	(0.001)

δ	P_A	P_B	Pr(error)	
			Gittins	Ran.Gittins
0.05	0.1	0.15	0.324 (0.327)	0.355 (0.334)
	0.5	0.55	0.359 (0.390)	0.336 (0.350)
0.1	0.1	0.2	0.216 (0.196)	0.200 (0.255)
	0.4	0.5	0.233 (0.273)	0.240 (0.237)
0.2	0.2	0.4	0.090 (0.094)	0.079 (0.096)
	0.4	0.6	0.078 (0.107)	0.079 (0.090)
0.3	0.1	0.4	0.018 (0.024)	0.021 (0.027)
	0.4	0.7	0.020 (0.034)	0.020 (0.023)
0.4	0.1	0.5	0.002 (0.008)	0.005 (0.006)
	0.3	0.7	0.004 (0.008)	0.004 (0.016)
0.5	0.1	0.6	0.000 (0.002)	0.000 (0.002)
	0.3	0.8	0.003 (0.001)	0.002 (0.001)

Table 4.19: *Simulation results for the expected sample size on the inferior treatment for five allocation rules (linear time-trends in Bernoulli response variable)*

δ	p_A	p_B	Rob.Sieg.	E(M)		Bather
				Prop.Ran.		
0.05	0.1	0.15	26.7 (14.8)	14.5 (14.1)		16.2 (14.2)
	0.5	0.55	33.7 (15.2)	18.1 (17.4)		28.5 (16.3)
0.1	0.1	0.2	21.4 (12.2)	13.8 (13.3)		14.3 (11.9)
	0.4	0.5	24.7 (13.9)	17.0 (15.6)		21.8 (14.1)
0.2	0.2	0.4	14.9 (9.6)	12.3 (12.2)		11.2 (9.9)
	0.4	0.6	16.2 (9.4)	12.7 (12.4)		13.0 (10.3)
0.3	0.1	0.4	8.6 (6.5)	9.0 (8.8)		7.1 (6.8)
	0.4	0.7	9.2 (6.6)	9.7 (9.3)		7.7 (7.3)
0.4	0.1	0.5	5.8 (5.0)	7.3 (7.1)		5.6 (5.4)
	0.3	0.7	6.2 (5.0)	7.6 (7.5)		5.6 (5.2)
0.5	0.1	0.6	4.6 (4.0)	6.2 (6.0)		4.6 (4.4)
	0.3	0.8	4.7 (3.9)	6.4 (6.1)		4.5 (4.2)

δ	p_A	p_B	E(M)	
			Gittins	Ran.Gittins
0.05	0.1	0.15	14.9 (13.4)	14.6 (13.3)
	0.5	0.55	25.7 (15.0)	19.6 (14.2)
0.1	0.1	0.2	13.4 (11.6)	13.1 (12.0)
	0.4	0.5	18.8 (12.7)	16.3 (12.5)
0.2	0.2	0.4	10.8 (9.2)	10.4 (9.6)
	0.4	0.6	12.5 (9.7)	11.7 (10.3)
0.3	0.1	0.4	7.2 (6.6)	7.6 (6.9)
	0.4	0.7	8.1 (7.3)	8.1 (7.3)
0.4	0.1	0.5	5.8 (5.3)	6.1 (5.9)
	0.3	0.7	6.3 (5.6)	6.3 (5.8)
0.5	0.1	0.6	5.0 (4.6)	5.0 (4.7)
	0.3	0.8	5.0 (4.5)	5.2 (4.8)

Table 4.20: *Simulation results for the expected sample size on the better treatment for five allocation rules (linear time-trends in Bernoulli response variable)*

δ	P_A	P_B	Rob.Sieg.	E(N)		Bather
				Prop.Ran.		
0.05	0.1	0.15	36.2 (18.5)	14.6 (14.3)		21.2 (16.6)
	0.5	0.55	42.8 (18.9)	18.5 (17.6)		37.2 (19.1)
0.1	0.1	0.2	40.0 (20.0)	13.9 (13.6)		22.5 (17.1)
	0.4	0.5	42.3 (18.8)	17.0 (15.6)		35.8 (20.7)
0.2	0.2	0.4	40.3 (19.5)	12.4 (12.4)		31.4 (20.0)
	0.4	0.6	38.5 (18.5)	12.7 (12.6)		37.9 (21.7)
0.3	0.1	0.4	36.0 (17.3)	9.0 (8.9)		26.7 (19.1)
	0.4	0.7	29.6 (16.4)	9.9 (9.5)		32.5 (19.0)
0.4	0.1	0.5	32.4 (16.0)	7.4 (7.2)		29.0 (20.9)
	0.3	0.7	25.1 (15.0)	7.9 (7.4)		29.7 (19.4)
0.5	0.1	0.6	25.2 (14.3)	6.5 (6.2)		32.8 (21.5)
	0.3	0.8	17.0 (11.6)	6.6 (6.3)		26.3 (20.2)

δ	P_A	P_B	E(N)	
			Gittins	Ran.Gittins
0.05	0.1	0.15	19.1 (15.7)	16.2 (14.5)
	0.5	0.55	33.7 (16.5)	26.5 (15.7)
0.1	0.1	0.2	19.5 (15.5)	16.3 (13.9)
	0.4	0.5	27.7 (17.0)	24.0 (15.3)
0.2	0.2	0.4	21.4 (14.7)	17.6 (13.2)
	0.4	0.6	27.2 (16.2)	21.6 (14.5)
0.3	0.1	0.4	14.7 (11.8)	12.5 (10.5)
	0.4	0.7	19.1 (13.8)	19.1 (11.6)
0.4	0.1	0.5	11.7 (9.9)	10.6 (9.3)
	0.3	0.7	14.1 (10.7)	14.1 (9.3)
0.5	0.1	0.6	9.5 (8.6)	9.5 (7.8)
	0.3	0.8	9.9 (8.6)	9.0 (7.7)

4.7 Sequential tests for a normal response variable with fluctuations

In §§4.3 and 4.5, we derived suitable sequential tests when the response variable is subject to a unknown linear time-trend. We now proceed as in Chapters 2 and 3 and outline a possible approach when there are fluctuations in the response variable. If we denote the new response variable by Y_{ij} ($i=A,B$; $j=0,1,2,\dots$), then a suitable model for the Y_{ij} is detailed in §3.3.2. A sequential test for this setting is now derived.

4.7.1 Derivation of tests

Assume that $\sigma_A^2 = \sigma_B^2 = \sigma^2$ with σ^2 known, and let w_j be the deviation from the

mean of the observed response for patient $j+1$ (at time j), then

$$w_j = a_{Aj} (y_{Aj} - \theta + \frac{1}{2}\delta - t_j) + (1 - a_{Aj}) (y_{Bj} - \theta - \frac{1}{2}\delta - t_j) , \quad (4.36)$$

where θ , δ , t and a_{Aj} are defined as before. The p.d.f. of the vector of responses, $\underline{y}$, is the multivariate normal

$$f_{\underline{Y}}(\underline{y}) = (2\pi)^{-\frac{1}{2}(m+n)} |\Sigma_Y|^{-1} \exp\left\{-\frac{1}{2}\underline{w}^T \Sigma_Y^{-1} \underline{w}\right\} \quad (4.37)$$

for which

$$\Sigma_Y = \sigma^2 I_{m+n} + \sigma_\epsilon^2 V ,$$

where

$$V = \frac{1}{(1-\rho^2)} \begin{bmatrix} 1 & \rho & \rho^2 & \cdot & \cdot & \cdot & \rho^{m+n-1} \\ \rho & 1 & \rho & \cdot & \cdot & \cdot & \rho^{m+n-2} \\ \vdots & \vdots & \cdot & \cdot & \cdot & \cdot & \vdots \\ \rho^{m+n-1} & \rho^{m+n-2} & \rho^{m+n-3} & \cdot & \cdot & \cdot & 1 \end{bmatrix} .$$

Thus, the log-likelihood of the observed responses (after m patients have been allocated to treatment A and n to treatment B) can be written as

$$L(\underline{y}, \delta, \theta, t, \rho, \sigma_\epsilon^2) = -\frac{1}{2}(m+n) \log(2\pi) - \frac{1}{2} \log |\Sigma_Y| - \frac{1}{2} \underline{w}^T \Sigma_Y^{-1} \underline{w} . \quad (4.38)$$

Since we now have *four* unknown nuisance parameters (θ, t, ρ and σ_ϵ^2), in order to derive a sequential test for H_0^* : $\delta = -\delta^*$ against H_1^* : $\delta = \delta^*$, we consider a Taylor series expansion about $(\delta, \theta, t, \rho, \sigma_\epsilon^2)$ of

$$\log \lambda = L(\underline{y}, \delta_1, \hat{\theta}_1, \hat{t}, \hat{\rho}, \hat{\sigma}_\epsilon^2) - L(\underline{y}, \delta_0, \hat{\theta}_0, \hat{t}, \hat{\rho}, \hat{\sigma}_\epsilon^2) \quad (4.39)$$

in which $\hat{\theta}_i$ and δ_i are as defined in §4.3.1; and $\hat{t}$, $\hat{\rho}$ and $\hat{\sigma}_\epsilon^2$ are the maximum likelihood estimates of t , ρ and σ_ϵ^2 based on $\underline{y}$.

Working through some algebra, this expansion yields, to quadratic terms,

$$\begin{aligned} \log \lambda \simeq & - [\hat{\theta} - \frac{1}{2}(\hat{\theta}_0 + \hat{\theta}_1)] \left\{ (\hat{\theta}_1 - \hat{\theta}_0) \frac{\partial^2 L}{\partial \theta^2} + 2\delta^* \frac{\partial^2 L}{\partial \theta \partial \delta} \right\} \\ & - \hat{\delta} \left\{ (\hat{\theta}_1 - \hat{\theta}_0) \frac{\partial^2 L}{\partial \theta \partial \delta} + 2\delta^* \frac{\partial^2 L}{\partial \delta^2} \right\} , \end{aligned} \quad (4.40)$$

where L is defined by (4.38) and its derivatives are evaluated at $(\delta, \theta, t, \rho, \sigma_\epsilon^2)$. Using (4.38), we can write one of the maximum likelihood equations as

$$\frac{\partial L}{\partial \theta} = - \frac{\partial \underline{w}^T}{\partial \theta} \Sigma_Y^{-1} \underline{w} . \quad (4.41)$$

Setting $\partial L / \partial \theta = 0$, it can be shown that

$$(\hat{\theta}_1 - \hat{\theta}_0) \frac{\partial^2 L}{\partial \theta^2} + 2\delta^* \frac{\partial^2 L}{\partial \theta \partial \delta} = 0 \quad (4.42)$$

and, thus, (4.40) becomes

$$\begin{aligned} \log \lambda &\approx -2\hat{\delta}\delta^* \left[\frac{\partial^2 L}{\partial \delta^2} - \left\{ \frac{\partial^2 L}{\partial \theta \partial \delta} \right\}^2 \left\{ \frac{\partial^2 L}{\partial \theta^2} \right\}^{-1} \right] \\ &= 2\hat{\delta}\delta^* / I^{\delta\delta}, \end{aligned} \quad (4.43)$$

say, where $\hat{\delta}$ is the maximum likelihood estimate of δ based on y .

Therefore, our modified test statistic is given by

$$Z_{m,n} = \frac{\hat{\delta}}{I^{\delta\delta}}, \quad (4.44)$$

where, for $\rho = 0$ and $\sigma_\epsilon^2 = 0$, $I^{\delta\delta} = (m+n)/mn$. As in §§4.3.2 and 4.5.2, we widen the SPRT boundaries $(-a, a)$ of §4.1.1 to allow for the variability due to estimation of the trend, t , and the fluctuation characteristics ρ and σ_ϵ^2 . This procedure, together with maximum likelihood estimation of the parameters, is now considered.

4.7.2 Maximum likelihood estimation and adjusting the SPRT boundaries

Clearly in order to implement the test statistic (4.44), we require the maximum likelihood estimates of θ , δ , t , ρ and σ_ϵ^2 together with various asymptotic variance and covariance terms.

From (4.38), we can write down five simultaneous equations for the maximum likelihood estimates of θ , δ , t , ρ and σ_ϵ^2 . Since Σ_Y^{-1} is of full rank ($= m+n$), and using well-known results for matrices, we obtain

$$\left. \begin{aligned} \frac{\partial L}{\partial k} &= -\text{tr} \left\{ \Sigma_Y^{-1} \frac{\partial \underline{w}}{\partial k} \underline{w}^T \right\}, \quad (k = \theta, \delta, t) \\ \frac{\partial L}{\partial \rho} &= -\frac{1}{2} \sigma_\epsilon^2 \left[\text{tr} \left\{ \Sigma_Y^{-1} \frac{\partial \underline{v}}{\partial \rho} \right\} - \text{tr} \left\{ \Sigma_Y^{-1} \frac{\partial \underline{v}}{\partial \rho} \Sigma_Y^{-1} \underline{w} \underline{w}^T \right\} \right] \\ \frac{\partial L}{\partial \sigma_\epsilon^2} &= -\frac{1}{2} \left[\text{tr} \left\{ \Sigma_Y^{-1} \underline{v} \right\} - \text{tr} \left\{ \Sigma_Y^{-1} \underline{v} \Sigma_Y^{-1} \underline{w} \underline{w}^T \right\} \right] \end{aligned} \right\}. \quad (4.45)$$

Furthermore, using equations (13) and (14) of Galbraith & Galbraith (1974), we can write

$$\Sigma_Y^{-1} = \frac{(1-\rho\eta)(\rho-\eta)}{\rho\sigma_\epsilon^2/\sigma^2} M, \quad (4.46)$$

where

$$\eta = p^{-1} \{1 + \sqrt{(1-p^2)}\}, \quad p = \frac{2\rho}{\{\sigma_\epsilon^2/\sigma^2 + (1+\rho^2)\}} \quad (4.47)$$

and M is a matrix with elements

$$\begin{aligned} m_{ii} &= [(1-\rho\eta)^2(\eta-\rho)^2\{1-\eta^{2(i-1)}\}\{1-\eta^{2(m+n-i)}\} \\ &\quad + (1-\eta^2)\{(1-\rho\eta)^2-(\eta-\rho)^2\eta^{2(m+n-1)}\}]/\alpha \\ m_{ij} &= \eta^{j-i-1}(1-\rho\eta)(\eta-\rho)\{(1-\rho\eta)-(\eta-\rho)\eta^{2i-1}\} \\ &\quad \times \{(1-\rho\eta)-(\eta-\rho)\eta^{2(m+n-j)+1}\}/\alpha, \quad (i < j) \end{aligned} \quad (4.48)$$

with

$$\alpha = (1-\eta^2)\{(1-\rho\eta)^2-(\eta-\rho)^2\eta^{2(m+n)}\}.$$

Setting equations (4.45) to zero and solving simultaneously leads to the maximum likelihood estimate of $\underline{\beta}^T = (\delta, \theta, t, \rho, \sigma_\epsilon^2)$. Using these values, we can estimate $\partial^2 L / \partial \delta^2$, etc., and hence compute an approximate value for (4.44). Analogous to the earlier results, the SPRT boundaries for test statistic (4.44) should be $(-a_{m,n}, a_{m,n})$, where

$$a_{m,n} = a \frac{\text{Var}(\hat{\delta})}{I_{\delta\delta}}. \quad (4.49)$$

We can derive the second partial derivatives of (4.38) with respect to each of the parameters in $\underline{\beta}^T$. [Explicit expressions for some of these are included in Appendix 4.] Thus, by inverting the Fisher information matrix and using the maximum likelihood estimates, we can estimate $\text{Var}(\hat{\delta})$.

It is clear that an adequate assessment of the performance of the sequential test defined by (4.44) and (4.49) requires considerable numerical and simulation work. Moreover, in order to obtain "reasonable" convergence of the parameter estimates in the early stages of the trial, several patients need to be randomised initially. It was considered impractical to apply the test strictly sequentially - a "group" sequential approach was taken where each group contained 20 patients [This approach leads on the average to a larger expected sample size].

Using this approach, the above sequential test was incorporated in the program

SPRTSNEW. [The NAG routine C05NBF (for more details, see §3.1.3) was used to solve the equations (4.45).] Although no results are presented here, this limited investigation indicated that even for moderate values of δ such as 1.0, the CPU time required for only 100 simulations can be up to 20 minutes. In many cases the estimate for δ converged relatively quickly whereas the other parameter estimates were rather poor in the early stages of the trial. As a comparison, parameter estimates based on an initial randomised phase of 100 patients were also computed for a few simulations – these clearly indicated the need for a larger randomised phase and also good starting values.

One method for obtaining starting values for the parameters is as follows: use the treatment means to estimate θ and δ , and hence form a single time series accordingly; compute first differences of the resultant time series to give a stationary ARMA(1,1) process with parameters ρ and η (see Galbraith & Galbraith, 1974, pp.68–69); the first two autocorrelations for this process may be estimated and hence initial estimates of ρ and η (and thus σ_ϵ^2) obtained. However, the efficiency of this method has not been studied.

Before we consider a special case of the fluctuation model, it is worth examining the performance of the sequential test in §4.1.1 for the above situation. Values for the error probabilities based on 1000 simulations are presented in Tables 4.21 and 4.22 for $\sigma_z^2/\sigma^2 = 0.05$ and 0.2, respectively, where $\sigma_z^2 = \sigma_\epsilon^2/(1-\rho^2)$ and $t = 0$.

The results demonstrate that the error probabilities for the original sequential test can be significantly exceeded if the dependence structure of the response variable is ignored. A comparison of Tables 4.21 and 4.22 highlights the effect of a greater contribution from the variance component of the fluctuations, σ_z^2 . It should, however, be noted that there is comparatively little effect upon the proportionate randomisation rule while for the more adaptive rules, the effect can be substantial for a modest treatment difference. The behaviour of the error probability for $\rho = 0.99$ may be explained by a parallel argument to that in §3.5.3.

Table 4.21: *Simulation results for the error probability for three allocation rules*
(fluctuations in normal response variable with $\sigma_z^2/\sigma^2 = 0.05$)

δ	ρ	Pr(error)			
		Approx.	Rob.Sieg.	Prop.Ran.	Gittins
0.1	0.0	0.231	0.234	0.227	0.246
	0.3		0.223	0.232	0.250
	0.7		0.234	0.223	0.250
	0.95		0.268	0.221	0.281
	0.99		0.264	0.220	0.267
0.25	0.0	0.047	0.058	0.050	0.035
	0.3		0.052	0.046	0.049
	0.7		0.053	0.052	0.046
	0.95		0.079	0.042	0.075
	0.99		0.066	0.037	0.048
0.5	0.0	0.002	0.005	0.001	0.002
	0.3		0.002	0.000	0.001
	0.7		0.006	0.000	0.001
	0.95		0.004	0.001	0.006
	0.99		0.006	0.000	0.002

Table 4.22: *Simulation results for the error probability for three allocation rules*
(fluctuations in normal response variable with $\sigma_z^2/\sigma^2 = 0.2$)

δ	ρ	Pr(error)			
		Approx.	Rob.Sieg.	Prop.Ran.	Gittins
0.1	0.0	0.231	0.266	0.256	0.257
	0.3		0.286	0.262	0.302
	0.7		0.299	0.255	0.284
	0.95		0.333	0.253	0.306
	0.99		0.312	0.243	0.265
0.25	0.0	0.047	0.058	0.070	0.054
	0.3		0.098	0.058	0.069
	0.7		0.122	0.061	0.105
	0.95		0.126	0.061	0.129
	0.99		0.105	0.049	0.099
0.5	0.0	0.002	0.005	0.002	0.004
	0.3		0.010	0.004	0.009
	0.7		0.008	0.001	0.009
	0.95		0.015	0.002	0.019
	0.99		0.011	0.001	0.012

4.7.3 Simulation results when the observed response is serially correlated

Since the likelihood equations (4.45) for the fluctuation model are rather cumbersome to work with, we now consider a special case of the model, namely, $\sigma^2 = 0$, $\sigma_\epsilon^2 = 1$ and $t = 0$. Clearly, this simplified model may be written

$$Y_{ij} = \mu_i + Z_j, \quad (i=A,B), \quad (4.50)$$

where Z_j is defined in §3.3.2. Defining w_j as in (4.36), and using the results of Kendall & Stuart (1983, pp.497-498), the log-likelihood is

$$L(\underline{y}, \delta, \theta, \rho) = -\frac{1}{2}(m+n)\log(2\pi) + \frac{1}{2}\log(1-\rho^2) - \frac{1}{2}\left\{w_0^2 + w_{m+n-1}^2 + (1+\rho^2)\sum_{j=1}^{m+n-2} w_j^2 - 2\rho\sum_{j=1}^{m+n-1} w_j w_{j-1}\right\}. \quad (4.51)$$

Using (4.51), we can write down $\partial L/\partial\theta$, $\partial L/\partial\delta$ and $\partial L/\partial\rho$. Setting the resulting expressions to zero and solving simultaneously leads to the maximum likelihood estimates of θ , δ and ρ . Using these estimates, the test statistic (4.44) can be implemented and the adjusted SPRT boundaries (4.49) computed using the second partial derivatives of (4.51).

The NAG minimisation routine E04LAF was used as an alternative to solving the three equations; this routine is a modified-Newton algorithm for finding a minimum of a function of several variables, subject to fixed upper and lower bounds on the variables. The routine requires the first and second partial derivatives of the function. The function to be minimised here is the negative log-likelihood; θ and δ are unbounded and ρ is bounded below by -0.99 and bounded above by 0.99 .

The above routine was incorporated in the program SPRTSNEW and simulation results are presented in Tables 4.23 - 4.25. For each set of parameter values considered, 100 simulations were carried out [for $\delta = 0.1$, the CPU time required in some cases was about 35 minutes and so 1000 simulations was not considered admissible]. For comparison purposes, the corresponding results for the sequential test in §4.1.1 (which, of course, does not adjust for estimation of ρ) are given in parentheses. As before, these are based on 1000 simulations. The parameters for the three rules are as in §4.2.2, and an initial 20 patients are randomly allocated to each treatment. The approximation (4.2) for the error

probability is also included.

Table 4.23: *Simulation results for the proportionate randomisation rule (serially correlated responses)*

δ	ρ	Pr(error)		E(M)	E(N)
		Approx.	Simul.		
0.1	0.3	0.231	0.230 (0.254)	65.1 (67.1)	66.6 (66.8)
	0.7		0.240 (0.330)	52.3 (49.4)	52.3 (49.4)
	0.95		0.250 (0.412)	50.1 (28.6)	49.9 (28.4)
	0.99		0.180 (0.407)	41.7 (25.6)	41.6 (25.6)
0.25	0.3	0.047	0.040 (0.045)	42.6 (46.8)	43.1 (46.9)
	0.7		0.020 (0.139)	37.8 (40.7)	37.7 (40.8)
	0.95		0.030 (0.286)	30.5 (27.1)	30.2 (27.2)
	0.99		0.030 (0.316)	30.9 (25.5)	30.8 (25.3)
0.5	0.3	0.002	0.000 (0.001)	27.6 (28.6)	27.7 (28.6)
	0.7		0.000 (0.024)	24.7 (28.3)	24.3 (28.3)
	0.95		0.000 (0.152)	21.8 (25.2)	21.4 (25.1)
	0.99		0.000 (0.207)	21.6 (24.3)	21.4 (24.2)

Table 4.24: *Simulation results for the Robbins & Siegmund rule (serially correlated responses)*

δ	ρ	Pr(error)		E(M)	E(N)
		Approx.	Simul.		
0.1	0.3	0.231	0.220 (0.338)	104.2 (57.2)	180.4 (72.4)
	0.7		0.100 (0.376)	118.9 (34.1)	228.5 (39.8)
	0.95		0.080 (0.400)	135.8 (30.2)	246.4 (32.2)
	0.99		0.110 (0.427)	132.7 (31.5)	226.0 (32.9)
0.25	0.3	0.047	0.040 (0.094)	47.6 (36.0)	147.2 (63.2)
	0.7		0.020 (0.208)	64.7 (30.1)	209.9 (38.7)
	0.95		0.020 (0.297)	63.7 (27.1)	214.9 (30.9)
	0.99		0.030 (0.330)	59.2 (29.4)	223.5 (32.5)
0.5	0.3	0.002	0.000 (0.007)	24.8 (23.1)	74.8 (41.4)
	0.7		0.000 (0.070)	24.4 (23.5)	114.1 (33.2)
	0.95		0.000 (0.145)	21.1 (24.1)	71.8 (28.8)
	0.99		0.000 (0.181)	23.2 (26.3)	88.1 (30.9)

From Tables 4.23 – 4.25, we see that the behaviour of the error probability and the expected sample sizes for the original sequential test is analogous to that in §§4.4.1 and 4.6.2, where there is an unknown linear time-trend. In the present case, the effect is more pronounced with substantial decreases in the expected sample sizes and corresponding increases in the error probability as the serial correlation coefficient, ρ , increases. In

particular, one effect of serial correlation is to force the two sample sizes to become more nearly equal.

Table 4.25: *Simulation results for the Gittins rule (serially correlated responses)*

δ	ρ	Pr(error)		E(M)	E(N)
		Approx.	Simul.		
0.1	0.3	0.231	0.240 (0.303)	77.5 (59.1)	151.1 (79.9)
	0.7		0.100 (0.380)	118.6 (35.4)	127.4 (39.0)
	0.95		0.180 (0.420)	83.0 (30.0)	107.4 (32.0)
	0.99		0.220 (0.446)	89.8 (31.6)	83.8 (33.1)
0.25	0.3	0.047	0.040 (0.101)	46.1 (39.2)	134.5 (71.4)
	0.7		0.050 (0.217)	55.4 (30.0)	94.6 (40.4)
	0.95		0.010 (0.318)	50.1 (27.8)	69.0 (31.7)
	0.99		0.040 (0.297)	48.3 (29.9)	47.0 (33.0)
0.5	0.3	0.002	0.000 (0.011)	23.1 (23.9)	53.1 (46.6)
	0.7		0.000 (0.064)	26.7 (24.2)	35.3 (35.5)
	0.95		0.000 (0.179)	27.0 (25.1)	30.1 (30.9)
	0.99		0.000 (0.163)	30.6 (25.4)	32.0 (29.9)

Our modified sequential test generally works well. In most cases, the error probability for this test is no larger than the approximate value (4.2); in many cases, it is substantially smaller. For the more adaptive rules and, in particular, the Robbins & Siegmund rule, the corresponding sample sizes required can be up to eight-fold those for the original test when δ is small. The fact that in some cases these sample sizes are smaller than for the original test is mainly attributable to random error since the values are averaged over only 100 simulations.

4.8 Discussion

We have attempted in this chapter to generalise some of the recent work on sequential tests involving two treatments when a data-dependent allocation rule is used. Although we adopted a non-rigorous justification of the proposed tests, a more thorough treatment may be possible along lines similar to Breslow (1969). We also need to check formally that the expected sample size, etc. for each of the allocation rules is finite. Although the proofs for this are relatively straightforward for a stable normal response variable (cf. Lemma 2 of Robbins & Siegmund, 1974, and Theorem 4.1 here), corresponding proofs when, say, linear time-trends are present appear difficult.

For a special case of the fluctuation model, the sequential test in §4.7.3 performs well. It would be interesting to extend this test to take into account a linear time-trend. Clearly, the maximum likelihood estimation procedure would necessitate more CPU time; however, since the structure of Σ_Y^{-1} is unaffected, this should not be severely affected.

The results for the sampling distribution of the standardised estimated treatment difference upon termination of the trial demonstrate that Woodroffe's (1989) approximation works well, even when the SPRT boundaries are widened to take into account the variability due to estimation of a linear time-trend. It was also shown that, using this approximation, a significant reduction in the bias of the maximum likelihood estimate, $\hat{\delta}$, of the true mean treatment difference, δ , can be attained. Further refinements in Woodroffe's approximation may improve the debiasing. Also, as we have already remarked in §4.5.5, an extension of Woodroffe's approach to non-normal response variables may be possible.

With the exception of §§4.5.4, 4.5.5 and 4.6, we have only considered normal responses here. As we pointed out in §1.3.1, some results have been reported for other responses such as exponential and Bernoulli (e.g. Hayre & Turnbull, 1981c) but these all assume a stable response variable. Since the parameter of interest, δ , can be transformed by a monotone function and the sequential test performed accordingly (Cox, 1963, p.8), the tests proposed in this chapter can be extended to Hayre and Turnbull's setting. As an example, consider an exponential response variable which might represent patients' survival time. In this case, we may be interested in the ratio of the mean survival times on the two treatments.

We have compared several data-dependent allocation rules with regard to reducing the number of patients receiving the inferior treatment. One of the conclusions is that the Gittins rule, proposed in §4.1.2, performs well for a wide range of parameter values. Since values for the Gittins index are now readily available (Gittins, 1989), this rule is both easy to use and flexible: choosing r close to one produces approximately pairwise allocation while r close to two, say, produces a strongly adaptive rule. The particular form of this rule arose from a discussion with Y-G. Wang. [In the context of a fixed trial size (cf. Chapter 2), Wang (1989, §4) has applied a version of the Gittins rule to

the situation where the distribution of the response is *unknown*; simulation results for a range of distributions for the response variable demonstrate the merits of the rule. It is hoped that at a later stage a similar approach will be developed for the sequential testing setting.]

Finally, the approach taken in this chapter can be applied to more general sequential tests such as those proposed by Armitage (1957), which permit the acceptance of a third hypothesis. We mentioned in §1.3.1 that for a stable response variable some work has been carried out in this area by Hayre (1979). It would appear that the adjustments in the SPRT boundaries derived in the present chapter can be applied directly to the situation investigated by Hayre. However, as yet no results have been obtained for this.

CHAPTER 5. CONCLUSIONS

Two important practical aspects of outcome-dependent allocation have been addressed in this work: (i) the efficiency of estimation of treatment effects, and (ii) the effect of changes in patient characteristics during the course of a clinical trial. If one allocates proportionately more patients to the better treatment, one would expect the efficiency of the estimated treatment difference to decrease. This is shown analytically in Chapter 3 for certain schemes for comparing proportions or normal means, and confirmed by simulation in Chapter 2 for more complicated schemes for comparing two proportions. For a sequential testing setting, Armitage (1985, p.22) conjectured from power considerations that the efficiency of estimation for outcome-dependent allocation schemes would be much the same as for equal allocation. This conjecture is confirmed by simulation in Chapter 4 for a certain class of sequential tests for normal responses.

Our results indicate that in some cases changes in patient characteristics can have a dramatic effect upon estimation. It is shown in Chapter 2 that the effect upon estimation can be reduced by the use of "blocking". However, it will be difficult to allow (completely) for such changes because for some schemes the presence of a trend, say, can significantly affect the distributions of the sample sizes on the two treatments. This effect is highlighted, for example, by the play-the-winner rule. [N.B. Use of "blocking" to allow for linear time-trends in the data has also been examined for restricted randomisation schemes (Mehta *et al.*, 1988; Wei, Smythe & Mehta, 1989).] It would be interesting to investigate the use of "blocking" in the context of sequential testing (cf. Chapter 4). From the results of Chapters 2 and 4, it is clear that data-dependent allocation rules which change the allocation probabilities gradually are significantly less affected by changes in patient characteristics than those leading to sudden changes in allocation probabilities (see also the remarks by Armitage & Coad (1989) in connection with the trial described in Example 1.5 in §1.5.1). This contrast is illustrated in Chapter 4 by the difference in expected sample sizes on the two treatments between the proportionate randomisation rule and the Robbins & Siegmund rule. The analytical results in Chapter 3 demonstrate the effect of changes in patient characteristics on schemes of

varying adaptiveness. These results indicate that even for simple outcome-dependent allocation schemes and a modest trend, the effect upon the bias and variance of the estimated treatment difference can be serious. However, as noted in §3.6 an approximate correction can be made.

In this work we have investigated three different clinical trial settings. Until now trials conducted using outcome-dependent allocation have been mainly for a fixed trial size for the comparison of two proportions (for more details, see §1.5.1). This might not be too surprising in view of the voluminous literature dealing with this setting. Clearly, in order that a corresponding trial involving survival data be possible, considerably more research is needed with regard to the problem of censoring. Some progress has been made in this area by Hayre & Turnbull (1981c). It is hoped that at a later stage some parts of this thesis (i.e. §§3.8 and 4.5) will be extended to address the problem of censoring.

It was pointed out in §2.8 that randomisation theory has been applied by Wei (1988) to the randomised play-the-winner rule, and subsequently applied to the trial described in Example 1.3. [The problem of inference based on *profile likelihood methods* has been studied by Wei, Smythe & Park (1989); this work complements the analysis of a trial described by Ware (1989).] As we remarked in §2.8, it would be interesting to see whether randomisation theory can be applied to other data-dependent allocation rules.

Two further extensions to the present work which deserve attention are: (i) estimation for the decision-theoretic models in Chapter 3 when a data-dependent allocation rule is used in the randomised stage(s), and (ii) termination rules for the Bernoulli setting of Chapter 2. Zelen (1969) has studied the Fixed Model for Bernoulli responses when the play-the-winner rule is used in the first stage, but no quantitative results were given for estimation. Although we would expect the bias and variance of the estimated treatment difference to increase as a result of allowing patients to be allocated adaptively in the first stage, it would be interesting to know the magnitude of this increase. Some guidance in this direction may be obtained by considering normal responses and allowing for unequal numbers of patients in the first stage. Results in Appendix 5 indicate that if one allocates a proportion θ ($\theta \geq \frac{1}{2}$) of the patients in the first stage to the better treatment,

and one moves from a slightly adaptive rule ($\theta = 0.55$) to a strongly adaptive one ($\theta = 0.8$), the bias in estimation can increase almost ten-fold but, more importantly, the corresponding variance increases by about 50%. A similar behaviour would be expected for more complex models.

The problem of sequential testing for Bernoulli data was briefly examined in §§4.5 and 4.6; termination rules were constructed using normal distribution theory. However, there is a large literature due mainly to D.G. Hoel, R. Simon, M. Sobel and G.H. Weiss which addresses the question of termination rules specifically designed for Bernoulli data when treatment allocation is adaptive (for a review of some of this work, see Hoel *et al.*, 1975). The early work deals with termination rules based on the difference in the numbers of successes with each treatment, and more recent work with likelihood-ratio type termination rules derived in an analogous way to those of Flehinger & Louis (1971, 1972). The likelihood ratio tests for Bernoulli data have been modified by Simon *et al.* (1977) to incorporate covariates. It should also be possible to extend these tests to allow for time-trends, etc.

In conclusion, it is hoped that the results presented in this thesis will be of interest to both theorists and practitioners of clinical trials; and also provide quantitative evidence for some of the remarks of Armitage (1985) and Armitage & Coad (1989). As indicated above, several aspects of outcome-dependent allocation still remain unexplored – we hope to pursue some of these in due course.

Appendix 1. Histograms and plots for §2.4.2 when $n = 150$

A selection of histograms of the number of patients on the better treatment and plots of allocation rates to the better treatment for $n = 150$ are displayed, respectively, in Figures A1.1 – A1.3 and Figure A1.4.




Figure A1.1: Histograms of the number of patients on the better treatment ($n = 150$, $p_A = 0.1$, $p_B = 0.3$)




Figure A1.2: Histograms of the number of patients on the better treatment ($n = 150$, $p_A = 0.7$, $p_B = 0.9$)







Figure A1.3: Histograms of the number of patients on the better treatment ($n = 150$, $p_A = 0.35$, $p_B = 0.65$)

PA=0.1 and PB=0.3



PA=0.7 and PB=0.9



PA=0.35 and PB=0.65




Figure A1.4: Plots of allocation rates to the better treatment ($n = 150$)

Appendix 2. Derivation of the $\partial E/\partial p_i$ ($i=1,2,3$) for the three-stage model in §3.7.2

Proceeding as in §3.1.3, straightforward differentiation of (3.51) yields

$$\begin{aligned} \frac{\partial E}{\partial p_1} = & 1 + p_2 \varphi(h) \frac{\partial h}{\partial p_1} + p_3 \varphi(k) \frac{\partial k}{\partial p_1} - 2\{\Phi(\ell_2) + I_1\} \\ & + (1-2p_1-p_2-p_3) \left\{ \varphi(\ell_2) \frac{\partial \ell_2}{\partial p_1} + \frac{\partial I_1}{\partial p_1} \right\}, \end{aligned} \quad (A2.1)$$

$$\begin{aligned} \frac{\partial E}{\partial p_2} = & \Phi(h) + p_3 \varphi(k) \frac{\partial k}{\partial p_2} - \{\Phi(\ell_2) + I_1\} \\ & + (1-2p_1-p_2-p_3) \left\{ \varphi(\ell_2) \frac{\partial \ell_2}{\partial p_2} + \frac{\partial I_1}{\partial p_2} \right\} \end{aligned} \quad (A2.2)$$

and

$$\frac{\partial E}{\partial p_3} = \Phi(k) - \{\Phi(\ell_2) + I_1\} + (1-2p_1-p_2-p_3) \left\{ \varphi(\ell_2) \frac{\partial \ell_2}{\partial p_3} + \frac{\partial I_1}{\partial p_3} \right\}, \quad (A2.3)$$

in which

$$\begin{aligned} I_1 = & - \int_{-\infty}^{\ell_2} \left\{ \int_{-\infty}^k \int_{-\infty}^h + \int_k^{\infty} \int_h^{\infty} \right\} \varphi_3^{(2)}(u, v, w) du dv dw \\ & + \int_{-\infty}^{\ell_1} \left\{ \int_{-\infty}^k \int_{-\infty}^h + \int_k^{\infty} \int_h^{\infty} \right\} \varphi_3^{(1)}(u, v, w) du dv dw, \end{aligned}$$

where the notation is as before.

Explicit expressions for the $\partial I_1/\partial p_i$ ($i=1,2,3$) were obtained using the chain rule for differentiation. These expressions involve derivatives of the standard trivariate normal densities $\varphi_3^{(j)}$ ($j=1,2$). Using Johnson & Kotz (1972, p.85), we can write

$$\begin{aligned} \frac{\partial \varphi_3^{(j)}(u, v, w)}{\partial p_i} = & -\frac{1}{2} \varphi_3^{(j)}(u, v, w) \left\{ \frac{1}{\Delta} \frac{\partial \Delta}{\partial p_i} + \frac{\partial A_{11}}{\partial p_i} u^2 + \frac{\partial A_{yy}}{\partial p_i} v^2 \right. \\ & \left. + \frac{\partial A_{\ell_j}}{\partial p_i} \ell_j w^2 + 2 \frac{\partial A_{1y}}{\partial p_i} uv + 2 \frac{\partial A_{1\ell_j}}{\partial p_i} uw + 2 \frac{\partial A_{y\ell_j}}{\partial p_i} vw \right\}, \end{aligned} \quad (A2.4)$$

where, for example,

$$A_{11} = (1 - \rho_y^2 \ell_j) / \Delta, \quad A_{1y} = -(\rho_{1y} - \rho_1 \ell_j \rho_y \ell_j) / \Delta$$

and

$$\Delta = 1 - \rho_{1y}^2 - \rho_{1\ell_j}^2 - \rho_{y\ell_j}^2 + 2\rho_{1y}\rho_{1\ell_j}\rho_{y\ell_j} \quad .$$

The correlation coefficients ρ_{km} ($k,m = 1,y,\ell_j$) were obtained algebraically. Following some further work, equations (A2.1) – (A2.3) were solved numerically as detailed in §3.7.2.

Appendix 3. Results relating to the estimation procedure in §§3.8.1 and 3.8.2

Exact expressions for $E(\bar{d}^*)$ are derived in §§3.8.1 and 3.8.2 for Bernoulli and exponential responses, respectively, for $n = 1$ and 2. Corresponding expressions for $n = 3$ are given below.

(a) Bernoulli response

Using the approach in §3.8.1, we obtain

$$\begin{aligned}
 E(\bar{d}^*) = & \delta + (N-6) \{ (p_A + p_B) \left[\frac{1}{2} \{ (1-p_A)^3 (1-p_B)^3 + p_A^3 p_B^3 \} \right. \\
 & + 3p_A(1-p_B) \{ (1-p_A)(1-p_B)^2 + p_A p_B \{ p_A p_B + (1-p_B)(3-2p_A) \} \} \\
 & + p_A^3 (1-p_B)^3 + 9p_A p_B (1-p_A)(1-p_B) \{ (1-p_A)(1-p_B) + p_A p_B \} / 2 \\
 & - p_A(1-p_A)(1-p_B) \{ (1-p_A)(1-p_B)(2p_B+1) + 6p_A p_B^2 + 2p_A(1-p_B)^2 \} \\
 & \left. - p_A^2 (1-p_B)^2 [p_A + 9p_B - 6p_A p_B] - p_A^3 p_B^2 (5-4p_B) \right\} / (N-3). \quad (A3.1)
 \end{aligned}$$

(b) Exponential response

Using the approach in §3.8.2, we obtain

$$E(\bar{d}^*) = \delta + 10\lambda_A^2 \lambda_B^2 \frac{(N-6)}{(N-3)} \frac{(\lambda_A - \lambda_B)}{(\lambda_A + \lambda_B)^6}, \quad (A3.2)$$

which for $\lambda_A = \lambda_B = \lambda$, can be written as

$$E(\bar{d}^*) \approx \frac{(\lambda_B - \lambda_A)}{\lambda^2} \left\{ 1 - \frac{5}{32} \frac{(N-6)}{(N-3)} \right\}. \quad (A3.3)$$

Appendix 4. Results relating to maximum likelihood estimation in §4.7.2

Starting from the likelihood equations (4.45), we can write down expressions for $\partial^2 L / \partial k \partial \ell$ ($k, \ell = \delta, \theta, t, \rho, \sigma_\epsilon^2$). With the exception of $\partial^2 L / \partial \rho^2$, $\partial^2 L / \partial (\sigma_\epsilon^2)^2$ and $\partial^2 L / \partial \rho \partial \sigma_\epsilon^2$, these are straightforward and therefore only expressions for these three quantities are included here.

Defining $\Phi = \Sigma_Y^{-1} V \Sigma_Y^{-1} V$, $\Omega = \Sigma_Y^{-1} \partial V / \partial \rho \Sigma_Y^{-1} \partial V / \partial \rho$, $\Lambda = \Sigma_Y^{-1} \partial^2 V / \partial \rho^2$ and $\Psi = \partial V / \partial \rho \Sigma_Y^{-1} V \Sigma_Y^{-1}$, and using well-known results for matrices, these may be written as

$$\frac{\partial^2 L}{\partial \rho^2} = -\frac{1}{2} \sigma_\epsilon^2 \left[\text{tr}(\Lambda - \sigma_\epsilon^2 \Omega) - \text{tr}((\Lambda - 2\Omega \sigma_\epsilon^2) \Sigma_Y^{-1} \underline{w} \underline{w}^T) \right], \quad (\text{A4.1})$$

$$\frac{\partial^2 L}{\partial (\sigma_\epsilon^2)^2} = \frac{1}{2} \text{tr}(\Phi) - \text{tr}(\Phi \Sigma_Y^{-1} \underline{w} \underline{w}^T) \quad (\text{A4.2})$$

and

$$\frac{\partial^2 L}{\partial \rho \partial \sigma_\epsilon^2} = \frac{1}{2} \sigma_\epsilon^2 \left[\text{tr}(\Psi) - 2 \text{tr}(\Sigma_Y^{-1} \Psi \underline{w} \underline{w}^T) \right]. \quad (\text{A4.3})$$

Appendix 5. Results for estimation for the Fixed Model when there are unequal numbers of patients in the first stage (cf. pp.158-159)

Using the notation of Chapter 3, assume that the response variable is stable and that $2Np_1\theta$ and $2Np_1(1-\theta)$ patients are allocated to treatments A and B, respectively, in the first stage of the Fixed Model, where $\theta \in (0,1)$. Then using the same approach as in §3.2, it can be shown that

$$E(\bar{d}^*) = \delta - \frac{c_3}{c_1 c_2} \varphi(\delta/\sigma_1) \tag{A5.1}$$

and

$$\begin{aligned} \text{Var}(\bar{d}^*) = & \frac{\sigma^2}{Nc_1(1-c_1)} + \frac{c_3\varphi(\delta/\sigma_1)}{(c_1 c_2)^2} \left\{ \delta [1-2p_1(1-\theta+\theta^2)] - c_3\varphi(\delta/\sigma_1) \right\} \\ & + \frac{\sigma^2}{2Np_1} \left\{ \frac{[1-2p_1\theta(3-4p_1)]}{c_2^2\theta} - \frac{[1-2p_1(1-\theta)(3-4p_1)]}{c_1^2(1-\theta)} \right\} \Phi(-\delta/\sigma_1), \end{aligned} \tag{A5.2}$$

where

$$c_3 = \frac{(1-2p_1)(1-2\theta)\sigma}{\sqrt{2Np_1\theta(1-\theta)}}, \quad \sigma_1^2 = \frac{\sigma^2}{2Np_1\theta(1-\theta)}$$

and

$$c_1 = 1-2p_1(1-\theta) \quad , \quad c_2 = 1-2p_1\theta \quad .$$

As an example, we suppose $\delta^* = \delta/\sigma = 0.5$ and $N = 100$ so that $p_1 = 0.116$ from Table 3.1. Values for the bias and variance of $\bar{d}^*/\sigma$ using (A5.1) and (A5.2) are presented in Table A5.1. Without loss of generality, we assume that A is the better treatment and $\theta \geq 0.5$.

Table A5.1: Standardised bias and variance of the estimated treatment difference for the Fixed Model with an unequal first stage

θ	Bias	Variance
0.5	0.0	0.098
0.55	0.008	0.101
0.6	0.017	0.105
0.8	0.077	0.152
0.9	0.169	0.235

References

- Abramowitz, M. & Stegun, I.A. (1964). *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables*. New York: Dover Publications.
- Altman, D.G. & Royston, J.P. (1988). The hidden effect of time. *Statist. Med.* 7, 629–637.
- Anderson, T.W. (1964). Sequential analysis with delayed observations. *J. Amer. Statist. Assoc.* 59, 1006–1015.
- Anderson, T.W. (1984). *An Introduction to Multivariate Statistical Analysis*, 2nd edition. New York: Wiley.
- Anscombe, F.J. (1963). Sequential medical trials. *J. Amer. Statist. Assoc.* 58, 365–383.
- Armitage, P. (1957). Restricted sequential procedures. *Biometrika* 44, 9–26.
- Armitage, P. (1963). Sequential medical trials: some comments on F.J. Anscombe's paper. *J. Amer. Statist. Assoc.* 58, 384–387.
- Armitage, P. (1975). *Sequential Medical Trials*, 2nd edition. Oxford: Blackwell.
- Armitage, P. (1985). The search for optimality in clinical trials. *Int. Statist. Rev.* 53, 15–24.
- Armitage, P. (1988). Some aspects of phase-III trials. In *Clinical Trials and Related Topics: Proceedings of an ISI Meeting in Biometry*. Amsterdam: Elsevier.
- Armitage, P. & Berry, G. (1987). *Statistical Methods in Medical Research*, 2nd edition. Oxford: Blackwell.
- Armitage, P. & Coad, D.S. (1989). Discussion on 'Investigating therapies of potentially great benefit: ECMO' by J.H. Ware. Due to appear in *Statist. Sci.*
- Bailar, J.C. (1976). Patient assignment algorithms – an overview. In *Proceedings of the 9th International Biometric Conference*, Vol. I, pp.189–206. Raleigh, North Carolina: The Biometric Society.

- Baker, A.G. (1950). Properties of some tests in sequential analysis. *Biometrika* 37, 334–346.
- Bartlett, M.S. (1946). The large sample theory of sequential tests. *Proc. Camb. Phil. Soc.* 42, 239–244.
- Bartlett, R.H., Roloff, P.W., Cornell, R.G., Andrews, A.F., Dillon, P.W. & Zwischenberger, J.B. (1985). Extracorporeal circulation in neonatal respiratory failure: a prospective randomised trial. *Pediatrics* 76, 479–487.
- Bather, J.A. (1980). Randomised allocation of treatments in sequential trials. *Adv. Appl. Prob.* 12, 174–182.
- Bather, J.A. (1981). Randomised allocation of treatments in sequential experiments (with discussion). *J. R. Statist. Soc. B* 43, 265–292.
- Bather, J.A. (1983). The minimax risk for the two-armed bandit problem. In *Mathematical Learning Models – Theory & Algorithms*, Eds. U. Herkenrath, D. Kalin & W. Vogel, pp. 1–11. New York: Springer-Verlag.
- Bather, J.A. (1985). On the allocation of treatments in sequential medical trials. *Int. Statist. Rev.* 53, 1–13.
- Beach, C.M. & Mackinnon, J.G. (1978). A maximum likelihood procedure for regression with autocorrelated errors. *Econometrica* 46, 51–58.
- Bellman, R. (1956). A problem in the sequential design of experiments. *Sankhyā A* 16, 221–229.
- Berry, D.A. (1972). A Bernoulli two-armed bandit. *Ann. Math. Statist.* 43, 871–897.
- Berry, D.A. (1978). Modified two-armed bandit strategies for certain clinical trials. *J. Amer. Statist. Assoc.* 73, 339–345.
- Berry, D.A. & Fristedt, B. (1985). *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman & Hall.
- Blackwell, D. & Hodges, J.L. (1957). Design for the control of selection bias. *Ann. Math. Statist.* 28, 449–460.

- Bradt, R.N., Johnson, S.M. & Karlin, S. (1956). On sequential designs for maximising the sum of n observations. *Ann. Math. Statist.* **27**, 1060–1074.
- Breslow, N. (1969). On large sample sequential analysis with applications to survivorship data. *J. Appl. Prob.* **6**, 261–274.
- Canner, P.L. (1970). Selecting one of two treatments when the responses are dichotomous. *J. Amer. Statist. Assoc.* **65**, 291–306.
- Chernoff, H. & Petkau, A.J. (1985). Sequential medical trials with ethical cost. In *Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer*, Vol. 2, Eds. L.M. LeCam & R.A. Olshen, pp. 521–537. Monterey, California: Wadsworth.
- Choi, S.C. & Clark, V.A. (1970). Sequential decision for a binomial parameter with delayed observations. *Biometrics* **26**, 411–420.
- Clayton, D.G. (1982). Ethically optimised designs. *Br. J. Clin. Pharmac.* **13**, 469–480.
- Clayton, M.K. (1988). Covariate models for Bernoulli bandits. Paper read at 206th I.M.S. meeting, Fort Collins, Colorado.
- Colton, T. (1963). A model for selecting one of two medical treatments. *J. Amer. Statist. Assoc.* **58**, 388–400.
- Colton, T. (1965). A two-stage model for selecting one of two treatments. *Biometrics* **21**, 169–180.
- Colton, T. (1973). Determination of the patient horizon in a decision-theoretic approach to the design of clinical trials. In *I.S.I. Vienna Conference Proceedings*, pp. 283–288.
- Cornell, R.G., Landenberger, B.D. & Bartlett, R.H. (1986). Randomised play-the-winner clinical trials. *Comm. Stat. – Theor. Meth.* **15**, 159–178.
- Cornfield, J., Halperin, M. & Greenhouse, S.W. (1969). An adaptive procedure for sequential clinical trials. *J. Amer. Statist. Assoc.* **64**, 759–770.
- Cox, D.R. (1963). Large sample sequential tests for composite hypotheses. *Sankhyā* **25**, 5–12.

- Day, N.E. (1969). Two-stage designs for clinical trials. *Biometrics* 25, 111-118.
- Donner, A. (1977). The use of auxiliary information in the design of a clinical trial. *Biometrics* 33, 305-314.
- Eick, S.G. (1988). Gittins procedures for bandits with delayed responses. *J. R. Statist. Soc. B* 50, 125-132.
- Eick, S.G. (1989). The two-armed bandit with delayed responses. Due to appear in *Ann. Statist.*
- Fabius, J. & van Zwet, W.R. (1970). Some remarks on the two-armed bandit. *Ann. Math. Statist.* 41, 1906-1916.
- Feldman, D. (1962). Contributions to the 'two-armed bandit problem'. *Ann. Math. Statist.* 33, 847-856.
- Feller, W. (1968). *An Introduction to Probability Theory and its Applications*, Vol. 1, 3rd edition. New York: Wiley.
- Flehinger, B.J. & Louis, T.A. (1971). Sequential treatment allocation in clinical trials. *Biometrika* 58, 419-426.
- Flehinger, B.J. & Louis, T.A. (1972). Sequential medical trials with data-dependent treatment allocation. *Sixth Berkeley Symp. of Math. Statist. & Prob.* 4, 43-51.
- Flehinger, B.L., Louis, T.A., Robbins, H. & Singer, B.H. (1972). Reducing the number of inferior treatments in clinical trials. *Proc. Nat. Acad. Sci. U.S.A.* 69, 2993-2994.
- Fox, B.L. (1974). Finite horizon behaviour of policies for two-armed bandits. *J. Amer. Statist. Assoc.* 69, 963-965.
- Friedman, L.M., Furberg, C.D. & Demets, D.L. (1981). *Fundamentals of Clinical Trials*. Boston: John Wright, PSG Inc.
- Gabriel, K.R. (1959). The distribution of the number of successes in a sequence of dependent trials. *Biometrika* 46, 454-460.
- Galbraith, R.F. & Galbraith, J.I. (1974). On the inverses of some patterned matrices arising in the theory of stationary time series. *J. Appl. Prob.* 11, 63-71.

- Gittins, J.C. (1979). Bandit processes and dynamic allocation indices (with discussion). *J. R. Statist. Soc B* 41, 148–177.
- Gittins, J.C. (1983). Dynamic allocation indices for Bayesian bandits. In *Mathematical Learning Models - Theory and Algorithms*, Eds. U. Herkenrath, D. Kalin & W. Vogel, pp. 50–67. New York: Springer-Verlag.
- Gittins, J.C. (1989). *Multi-armed bandit allocation indices*. Chichester: Wiley.
- Gittins, J.C. & Jones, D.M. (1974). A dynamic allocation index for the sequential design of experiments. In *Progress in Statistics*, Vol. 1, Eds. J. Gani, K. Sarkadi & I. Vincze, pp. 241–266. Amsterdam: North-Holland.
- Gittins, J.C. & Jones, D.M. (1979). A dynamic allocation index for the discounted multi-armed bandit problem. *Biometrika* 66, 561–565.
- Glazebrook, K.D. (1978). On the optimal allocation of two or more treatments in a controlled clinical trial. *Biometrika* 65, 335–340.
- Glazebrook, K.D. (1980). On randomised dynamic allocation indices for the sequential design of experiments. *J. R. Statist. Soc B* 42, 342–346.
- Grundy, P.M., Healy, M.J.R. & Rees, D.H. (1956). Economic choice of the amount of experimentation. *J. R. Statist. Soc. B* 19, 32–55.
- Hall, W.J. (1962). Some sequential analogs of Stein's two-stage test. *Biometrika* 49, 367–378.
- Harvey, A.C. & Phillips, G.D.A. (1979). Maximum likelihood estimation of regression models with autoregressive-moving average disturbances. *Biometrika* 66, 49–58.
- Hayre, L.S. (1979). Two population sequential tests with three hypotheses. *Biometrika* 66, 465–474.
- Hayre, L.S. (1983). Sequential estimation of the difference between the means of two normal populations. *Metrika* 30, 101–107.
- Hayre, L.S. & Gittins, J.C. (1981). Sequential selection of the larger of two normal means when the variances are unknown. *J. Amer. Statist. Assoc.* 76, 696–700.

- Hayre, L.S. & Turnbull, B.W. (1981a). Sequential estimation in two-armed exponential clinical trials. *Biometrika* 68, 411-416.
- Hayre, L.S. & Turnbull, B.W. (1981b). Estimation of the odds ratio in the two-armed bandit problem. *Biometrika* 68, 661-668.
- Hayre, L.S. & Turnbull, B.W. (1981c). A class of simple approximate sequential tests for adaptive comparison of two treatments. *Comm. Statist. - Theor. Meth.* 10, 2339-2360.
- Hoel, D.G., Sobel, M. & Weiss, G.H. (1975a). Comparison of sampling methods for choosing the best binomial population with delayed observations. *J. Statist. Comput. Simul.* 3, 299-313.
- Hoel, D.G., Sobel, M. & Weiss, G.H. (1975b). A survey of adaptive sampling for clinical trials. In *Perspectives in Biometrics I*, Ed. R. Elashoff, pp.29-61. New York: Academic Press.
- Iglewicz, B. (1984). Alternative designs: sequential, multi-stage, decision theory and adaptive designs. In *Cancer Clinical Trials: Methods and Practice*, Eds. M.E. Buyse, J. Staquet & R.J. Sylvester, pp. 312-334. Oxford University Press.
- Isbell, J.R. (1959). On a problem of Robbins. *Ann. Math. Statist.* 30, 606-610.
- Johnson, N.L. & Kotz, S. (1972). *Distributions in Statistics: Continuous Multivariate Distributions*. New York: Wiley.
- Jones, P.W. (1975). The two-armed bandit. *Biometrika* 62, 523-524.
- Jones, P.W. & Kandeel, H.A. (1983). Numerical investigation of the two-armed bandit. In *Mathematical Learning Models - Theory and Algorithms*, Eds. U. Herkenrath, D. Kalin & W. Vogel, pp.101-107. New York: Springer-Verlag.
- Katehakis, M.N. & Derman, C. (1986). Computing optimal sequential allocation rules in clinical trials. In *Adaptive Statistical Procedures and Related Topics*, Ed. J.V. Ryzin, pp.29-39. Hayward, California: I.M.S.
- Kelley, T.A. (1974). A note on the Bernoulli two-armed bandit problem. *Ann. Statist.* 2, 1056-1062.

- Kelly, F.P. (1981). Multi-armed bandits with discount factor near one: the Bernoulli case. *Ann. Statist.* 9, 987-1001.
- Kendall, M.G. & Stuart, A. (1983). *The Advanced Theory of Statistics*, Vol. 3, 4th edition. London: Griffin.
- Kotz, S. & Johnson, N.L. (Eds.) (1985). *Encyclopedia of Statistical Sciences*, Vols. 4 & 7. New York: Wiley.
- Kullback, S. (1959). *Information Theory and Statistics*. New York: Wiley.
- Lai, T.L., Robbins, H. & Siegmund, D.O. (1977). Sequential decision about a normal mean. In *Statistical Decision Theory and Related Topics II*, Eds. S.S. Gupta & D.S. Moore, pp.213-222. New York: Academic Press.
- Langenberg, P. & Srinivasan, R. (1981). On the Colton model for clinical trials with delayed observations - normally-distributed responses. *Biometrics* 37, 143-148.
- Lellouch, J. & Schwartz, D. (1971). L'essai thérapeutique: éthique individuelle ou éthique collective? *Rev. Inst. Int. Statist.* 39, 127-136.
- Louis, T.A. (1975). Optimal allocation in sequential tests comparing the means of two Gaussian populations. *Biometrika* 62, 359-369.
- Louis, T.A. (1977). Sequential allocation in clinical trials comparing two exponential survival curves. *Biometrics* 33, 627-634.
- Louis, T.A. (1985). Discussion on 'On the allocation of treatments in sequential medical trials' by J.A. Bather and 'The search for optimality in clinical trials' by P. Armitage. *Int. Statist. Rev.* 53, 28-31.
- Malinvaud, E. (1980). *Statistical Methods of Econometrics*, 3rd edition. Oxford: Elsevier.
- Maurice, R.J. (1957). A minimax procedure for choosing between two populations using sequential sampling. *J. R. Statist. Soc. B* 19, 255-261.
- Maurice, R.J. (1959). A different loss function for the choice between two populations. *J. R. Statist. Soc. B* 21, 203-213.

- Mehta, C.R. (1981). Sequential comparison of two exponential distributions with censored survival data. *Biometrika* 68, 669-675.
- Mehta, C.R., Patel, N.R. & Wei, L.J. (1988). Constructing exact significance tests with restricted randomisation rules. *Biometrika* 75, 295-302.
- Mendoza, G. & Iglewicz, B. (1978). An extension of Colton's model for comparing two medical treatments. *J. Amer. Statist. Assoc.* 73, 646-649.
- Oudin, C. & Lellouch, J. (1972). La comparaison de quelques stratégies dans la conduite des essais thérapeutiques. *Rev. Statist. Appl.* 20, 5-21.
- Paneth, N. & Wallenstein, S. (1985). Extracorporeal membrane oxygenation and the play-the-winner rule. *Pediatrics* 76, 622-623.
- Peto, R. (1985). Discussion on 'On the allocation of treatments in sequential medical trials' by J.A. Bather and 'The search for optimality in clinical trials' by P. Armitage. *Int. Statist. Rev.* 53, 31-34.
- Pierce, D.A. (1971). Least squares estimation in the regression model with autoregressive-moving average errors. *Biometrika* 58, 299-312.
- Pocock, S.J. (1983). *Clinical Trials: A Practical Approach*. New York: Wiley.
- Pocock, S.J. & Simon, R. (1975). Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics* 31, 103-115.
- Raiffa, H. & Schlaifer, R. (1961). *Applied Statistical Decision Theory*. Cambridge, Massachusetts: Harvard University Press.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* 58, 527-535.
- Robbins, H. (1956). A sequential decision problem with a finite memory. *Proc. Nat. Acad. Sci. U.S.A.* 42, 920-923.
- Robbins, H. (1974). A sequential test for two binomial populations. *Proc. Nat. Acad. Sci. U.S.A.* 71, 4435-4436.

- Robbins, H. & Siegmund, D.O. (1974). Sequential tests involving two populations. *J. Amer. Statist. Assoc.* **69**, 132-139.
- Robbins, H., Simons, G. & Starr, N. (1967). A sequential analogue of the Behrens-Fisher problem. *Ann. Math. Statist.* **38**, 1384-1391.
- Robinson, D.R. (1982). Algorithms for evaluating the dynamic allocation index. *Oper. Res. Letters* **1**, 72-74.
- Robinson, D.R. (1983). A comparison of sequential treatment allocation rules. *Biometrika* **70**, 492-495.
- Rodman, L. (1978). On the many-armed bandit problem. *Ann. Prob.* **6**, 491-498.
- Samuels, S.M. (1968). Randomised rules for the two-armed bandit problem with finite memory. *Ann. Math. Statist.* **39**, 2103-2107.
- Schneiderman, M.A. (1975). How do you know you've done any better? *Cancer* **35**, 64-69.
- Siegmund, D.O. (1983). Allocation rules for sequential clinical trials. In *Mathematical Learning Models - Theory and Algorithms*, Eds. U. Herkenrath, D. Kalin & W. Vogel, pp.203-212. New York: Springer-Verlag.
- Silvey, S.D. (1975). *Statistical Inference*. London: Chapman & Hall.
- Simon, R. (1977). Adaptive treatment assignment methods and clinical trials. *Biometrics* **33**, 743-749.
- Simon, R., Hoel, D.G. & Weiss, G.H. (1977). The use of covariate information in the sequential analysis of dichotomous response experiments. *Comm. Statist. - Theor. Meth.* **8**, 777-788.
- Simons, G.D. (1986). A comparison of seven allocation rules for a clinical trial model. *Sequential Anal.* **3**, 193-221.
- Simons, G.D. (1987). A random horizon model for sequential clinical trials. Paper read at 200th I.M.S. meeting, Blacksburg, Virginia.

- Smith, C.V. & Pyke, R. (1965). The Robbins-Isbell two-armed bandit problem with finite memory. *Ann. Math. Statist.* **36**, 1375-1386.
- Stein, C. (1987). *Approximate Computation of Expectations*. Hayward, California: I.M.S.
- Thompson, W.R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25**, 285-294.
- Thompson, W.R. (1935). On the theory of apportionment. *Amer. J. Math.* **57**, 450-456.
- Upton, G.J.G. & Lee, R.D. (1976). The importance of the patient horizon in the sequential analysis of binomial clinical trials. *Biometrika* **63**, 335-342.
- Vogel, W. (1960a). A sequential design for the two-armed bandit. *Ann. Math. Statist.* **31**, 430-443.
- Vogel, W. (1960b). An asymptotic minimax theorem for the two-armed bandit problem. *Ann. Math. Statist.* **31**, 444-451.
- Wald, A. (1947). *Sequential Analysis*. New York: Wiley.
- Wald, A. (1950). *Statistical Decision Functions*. New York: Wiley.
- Wang, Y-G. (1989). Contributions to the theory of dynamic allocation indices. University of Oxford.
- Ware, J.H. (1989). Investigating therapies of potentially great benefit: ECMO (with discussion). Due to appear in *Statist. Sci.*
- Ware, J.H. & Epstein, M.F. (1985). Extracorporeal circulation in neonatal respiratory failure: a prospective randomised study. *Pediatrics* **76**, 850-851.
- Wei, L.J. (1988). Exact two-sample permutation tests based on the randomised play-the-winner rule. *Biometrika* **75**, 603-606.
- Wei, L.J. & Durham, S. (1978). The randomised play-the-winner rule in medical trials. *J. Amer. Statist. Soc.* **73**, 840-843.

- Wei, L.J., Smythe, R.T. & Mehta, C.R. (1989). Interval estimation with restricted randomisation rules. *Biometrika* **76**, 363–368.
- Wei, L.J., Smythe, R.T. & Park, T.S. (1989). Statistical inferences with data-dependent treatment allocation rules. To appear.
- Weinstein, M.C. (1974). Allocation of subjects in medical experiments. *New Eng. J. Med.* **291**, 1278–1285.
- Wetherill, G.B. & Glazebrook, K.D. (1986). *Sequential Methods in Statistics*, 3rd edition. London: Chapman & Hall.
- Whitehead, J. (1983). *The Design and Analysis of Sequential Clinical Trials*. Chichester: Horwood.
- Whitehead, J. (1986). On the bias of maximum likelihood estimation following a sequential test. *Biometrika* **73**, 573–581.
- Whittle, P. (1980). Multi-armed bandits and the Gittins index. *J. R. Statist. Soc. B* **42**, 143–149.
- Whittle, P. (1981a). Arm-acquiring bandits. *Ann. Prob.* **9**, 284–292.
- Whittle, P. (1981b). Discussion on 'Randomised allocation of treatments in sequential experiments' by J.A. Bather. *J. R. Statist. Soc. B* **43**, 288–289.
- Whittle, P. (1988). Restless bandits: activity allocation in a changing world. *J. Appl. Prob.* **25A**, 287–298.
- Woodroffe, M. (1979). A one-armed bandit problem with a concomitant variable. *J. Amer. Statist. Assoc.* **74**, 799–806.
- Woodroffe, M. (1982). Sequential allocation with covariates. *Sankhyā A* **44**, 403–414.
- Woodroffe, M. (1989). Very weak expansions for sequentially designed experiments: linear models. Due to appear in *Ann. Statist.*
- Zaborskis, A.A. (1976). Sequential Bayesian plan for choosing the best method of medical treatment. *Automation Remote Control* **37**, 1750–1757.

Zelen, M. (1969). Play the winner rule and the controlled clinical trial. *J. Amer. Statist. Assoc.* **64**, 131-146.