Design and Statistical Inference of Response Adaptive Clinical Trials

BY

Yanqing Yi

A Thesis Submitted to the Faculty of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

Department of Statistics
University of Manitoba
Winnipeg, Manitoba

Design and statistical inference of responde adaptive clinical trials

BY

Yanqing Yi

A Thesis/Practicum submitted to the Faculty of Graduate Studies of The University of Manitoba in partial fulfillment of the requirement of the degree

Of

Doctor of Philosophy

Contents

Abstract iii

Acknowledgements v

1 Introduction 1

1.1 Adaptive Designs 1

1.2 Response Adaptive Designs 6

1.2.1 Ethical Issues 7

1.2.2 Types of Response Randomization Procedures 11

1.3 My Research 15

1.3.1 Issues and Literature Review 16

1.3.2 Main Results 20

2 Formulation of Response Adaptive Designs 24

2.1 Formulation of Adaptive Designs with General Responses 24

2.2 Markov Decision Models for Response Adaptive Designs with Dichotomous Responses 26

3 Likelihood Ratio Test and Goodness-of-fit Test 30

3.1 Likelihood ratio test 30

3.2 Goodness-of-fit test for contingency tables with dependent data 34

4 Efficient Estimation **39**

4.1 Asymptotic efficiency of the maximum likelihood estimator 41

4.2 Asymptotically efficient estimation of treatment effects 46

5 Response Adaptive Designs with a Variance-Penalized Criterion **49**

5.1 The variance-penalized criterion 49

5.2 Randomization procedures 52

5.3 A comparison of designs 54

6 Conclusion and Further Research **62**

Bibliography **89**

Abstract

Much attention has been given to response adaptive designs recently because of their ethics advantages. However, the adaptation of the treatment allocation creates a dependence structure in the collected data and raises concerns about the validity of statistical analysis, the loss of power for testing hypotheses and experimental bias. My thesis focuses on the development of statistical inference methodologies and the investigation of optimality properties for response adaptive designs.

The issue of statistical inference for response adaptive clinical trials has been both important and challenging. Due to the dependency in data collected from response adaptive designs, traditional statistical inference methods cannot be applied without modification to analyze data from adaptive clinical trials. I study the treatment randomization processes of response adaptive clinical trials. The information gathering process in the trial is formulated as a stochastic process, in particular a Markov process for dichotomous responses. Then the logarithm of the likelihood ratio test and the goodness-of-fit test are extended to analyze dependent data from the adaptive trial. I also examine the issue of asymptotic efficiency of estimation in response adaptive designs of clinical trials. The asymptotical lower bound of exponential rates for consistent estimators is established and the maximum likelihood estimator of the treatment effect is shown to be asymptotically efficient in the Bahadur sense for response adaptive clinical trials.

Besides the exploration on the statistical inference for response adaptive designs,

I investigate the optimality properties of the designs and explore the evaluation of response adaptive designs using the variance-penalized criterion. It is shown that this criterion evaluates the performance of a response adaptive design based on both the expected number of patients assigned to the better treatment and the power of the statistical test. A new proportion of treatment allocation is proposed and simulation studies are conducted to compare the proposed design with the existing designs. The proposed design has the advantage of assigning more patients to the potentially better treatment with less loss in power of the statistical test in common clinical trial conditions.

However, the optimal treatment allocation under the variance-penalized criterion is deterministic, and hence is vulnerable to selection bias in clinical trials. Searching for an optimal randomization allocation is still under study. Constrained dynamic programming techniques may be employed and algorithms will developed to search for an optimal adaptive allocation rule in my further research.

Acknowledgements

I am very grateful to my dissertation advisor Dr. Xikui Wang of the University of Manitoba for his encouragement, patience and his discussion and very careful reading of the manuscript which led to many improvements. To the faculty and staff of the Department of Statistics, I wish to express my gratitude for their rich teaching and kind support throughout my education. I would like to express special thanks to Dr. James C. Fu for his helpful discussions. I would also like to thank committee members Dr. Liqun Wang and Dr. Alex Leblanc of the Department of Statistics, Dr. Lisa Lix of the Department of Department of Community Health Sciences at the University of Manitoba and the external examiner Dr. Peter Song of the Biostatistics School of Public Health at the University of Michigan.

I wish to thank Dr. Xikui Wang for his financial support from his research grants. I thank the Faculty of Graduate Studies for the University Manitoba Graduate Fellowship and the International Graduate Student Scholarship, and thank the UMSU for the UMSU Scholarship. I am also grateful to the UM Alumni Association, the Faculty of Graduate Studies, the Faculty of Science and the Department of Statistics for their financial support for me to attend several national and international conferences.

To my family and my parents, thanks for their love and spiritual support.

Chapter 1

Introduction

Much attention has been given to adaptive designs of clinical trials in recent years because of their efficiency and ethical advantages. However, the adaptation of an ongoing trial raises concerns about the validity of statistical conclusions, the logistical issues and experimental biases. This chapter introduces the background on adaptive designs, the motivation as well as the main results of my research.

1.1 Adaptive Designs

This section briefly introduces the definition of adaptive designs, various classes of adaptive designs and the major issues for the designs.

An adaptive design is defined by Chow et al. (2005) as a design that allows adaptations or modifications to some aspects of a clinical trial after its initiation without undermining the validity and integrity of the trial. This definition is consistent with that given in an executive summary of the PhRMA working group (Gallo et al., 2006). One kind of modification during the course of a trial could be the adaptation of randomization procedures. An adaptive designs with this kind of modifications is called a response adaptive design. In a response adaptive design, the probability of allocating the next patient to a particular treatment is adapted based on the accumulated in-

formation such as the responses of previously treated patients and previous treatment allocations. The treatment allocation is deliberately biased to assign more patients to the potentially better treatment. The modification could also be other changes of the design or statistical analysis procedures based on the observed interim results and/or any information from outside the trial. This kind of modification includes sample size re-estimating, early stopping due to efficacy or futility, dropping inferior treatment groups, modifying statistical hypotheses, changing inclusion/exclusion criteria, adjusting study dose of drugs and adapting endpoints during the course of a trial. This type of adaptive design is said to be a sequential adaptive design in which interim data analysis is a characteristic. No matter what kind of modifications is applied to a trial, the goal of adaptive designs is to learn from the accumulating information and to apply the learned knowledge to benefit the patients within the trial or to speed up the development of efficient drugs.

Adaptive designs are of great advantages including ethical and efficiency advantages. Firstly, in response adaptive designs, the motivation of the adaptation of treatment randomization is to modify the randomization procedure and to assign more patients within the trial to the potentially better treatment. The future patients within the trial are assigned to the better treatment with large probabilities. In sequential adaptive designs, patients in a trial may benefit from the drop of the inferior treatment and early stopping of a trial. A trial may be stopped if the experimental treatment is clearly better or worse than the control, or the futility stopping criterion is met in the interim analysis. These modifications of an on-going trial based on interim results of the trial not only potentially prevent exposing patients to the inferior treatment in the trial, but also reduce the number of experimental units required in the trial. In addition, interim information helps recalculate sample sizes and this may reduce the possible overly large size of the trial, thus limiting the exposure of patients

to the inferior treatment and saving resources. In adaptive designs, the accrued data in a trial are used to modify the on-going trial for economic consideration and/or concerns about the welfare of the patients within the trial.

Adaptive designs have attracted great interest due to their potential efficiency and/or ethical advantages. However, the modification of an on-going trial raises great concerns about the validity of statistical inference, logistical issues and procedural challenges when applying adaptive designs in practice.

The modification of an on-going trial presents statistical challenges to draw conclusions on medical questions at the end of the trial. In response adaptive designs, the modification of treatment allocation is based on the accumulated information in the trial. This adaptation creates a dependency structure in the collected data and the traditional statistical analysis methods cannot be applied directly without justification. In sequential adaptive designs, the modification of the trial is based on the interim analysis, thus resulting in the use of "non-standard" test statistics in the overall data analysis. It is of concern that the Type I error rate may have been inflated (Chang et al., 2006, Bauer and Kohne, 1994, Proschan and Hunsberger, 1995, Posch and Bauer, 2000). "In adaptive designs often test statistics diverging from the conventional test statistics may have to be used for the test decision" (citing Bauer and Einfalt, 2006). Jennison and Turnbull (2006) argued that "the final analysis of data from an adaptive trial design typically involves the use of unfamiliar test statistics that do not satisfy the sufficiency principle". Burman and Sonesson (2006) questioned whether analysis based on non sufficient statistics can be deemed "valid". Chow et al. (2006) noticed that there is a high risk that a clinical trial using adaptive designs/methods may fail in terms of drawing valid statistical conclusions and /or fail to provide useful information with desired power, especially when the sizes of the trials are relatively small and there are a number of protocol amendments.

The estimation problem for treatment effects is another concern. The executive summary of the PhRMA working group (Gallo et al., 2006) argued that issues with the estimation of treatment effect have not been fully resolved for some adaptive designs in the frequentist paradigm. This point of view is supported by Wassmer (2006). The research by Bauer and Einfalt (2006) found that, in the published literature, a problem exists with estimation in adaptive designs. Bauer and Einfalt (2006) claimed that "mid-trial design modification may have a negative impact on the persuasiveness and perception of the results." They further suggested that more research on the properties of suitable estimates following design adaptations is important for applying adaptive designs in practice in the future.

In addition, the modification of an on-going trial raises concerns on logistical or/and procedural issues in adaptive designs. The adaptations of an on-going trial are based either on cumulated information on treatment effects or interim analysis. This requires rapid data collection, effective communication between patients/investigators and the randomization center, and appropriate management of interim information to maintain the validity of a clinical trial, thus bringing great challenges to reduce selection bias when implementing adaptive designs in practice.

How to reduce selection bias is a common issue in clinical trials including traditional and adaptive trials. The issue is more specific in adaptive trials because the modification of an on-going trial requires conveying the information on responses of previous patients or interim analysis during the trial. This leads to the difficulty of reducing selection bias. In adaptive designs, the selection bias could come from the investigator (trial personnel and/or sponsor representatives) or from the patients that are involved or will be involved in the trial. A trial personnel could guess the probability of treatment allocation for the next patient according to the adaption rule from previous responses of patients. A sponsor for a pharmaceutical company may

manipulate the available interim results to affect the adaptation of the trial at the anxiety to see the company's latest pharmaceutical product succeed. As a result, the subconscious preference or deliberate dishonesty of the investigator could be included in the data. The selection bias could also come from the patients' aspect. A patient in the trial may choose to quit the treatment if he or she feels it is not effective or requires to transfer to another treatment group provided that he or she guesses the other treatment is more effective based on the responses of previous patients. In some adaptive designs, patients are informed of the nature of the adaptation of the trial for ethical consideration or in order to be compliant with regulations on good clinical trial practice. In the screening phase of a trial, except for emergencies, patients may prefer to be recruited later so that the chance for her or him to get the better treatment becomes larger. This behavior of patients results in the problem of data missing not at random or casts the doubts on the randomness of treatment allocation. Consequently, it leads to some challenges and some problems for statistical analysis.

Selection bias raises serious concerns about the validity of the analytical results or creates problems in statistical analysis. The best way to reduce selection bias is to mask the interim information of the trial from the investigators and patients. However, the design feature of response adaptive trials needs to convey the accumulated information for the modification of the trial or it is difficult to mask the study in some situations. As pointed out by Rosenberger and Lachin (2002), "it is not unusual for patients to be unmasked during the course of the trial due to either adverse events known to be highly associated with one of the treatments, life-threatening emergencies requiring unmasking, or distinguishing features of the masked treatment, such as taste". Recently, the potential impact of conveying interim analysis information to investigators or other people were discussed by many researchers. Lokhnygina (2006) noticed that a potential problem of unblinding the interim data and the resulting op-

erational bias have long been the source of concerns for many researchers considering the use of adaptive designs. Bauer and Brannath (2004) argued that applying such adaptive designs requires new tools of statistical monitoring. The executive summary of the PhRMA working group (Gallo et al., 2006) recommended expanding the responsibility of the independent Data Monitoring Committees, limiting the extent of sponsor involvement and withholding the details of the adaptive procedure to a separate document in order to reduce the negative impact of leaking interim results to investigators or other persons involved in the trial. Lokhnygina (2006) strongly supported this recommendation and claimed that it could make the implementation of the adaptive designs in practice more plausible.

In summary, adaptive designs of clinical trials are potentially efficient and/or have ethical advantages. But the modification of an on-going trial presents some of the greatest challenges in statistical analysis and in the reduction of selection bias. My dissertation focuses on the design and statistical inference of response adaptive clinical trials.

1.2 Response Adaptive Designs

With a response adaptive design, the probability of treatment allocation to the next patient is modified based on the cumulating information on previous treatment allocations and responses of previously treated patients in the trial. The purpose of the design is to deliberately bias the treatment allocation in order to assign more patients to the potentially better treatment. Response adaptive designs are developed alternatively for ethical considerations in clinical trials.

1.2.1 Ethical Issues

Traditional randomization provides a powerful method for comparing treatment effects and has many statistical and scientific advantages. The 50-50 randomization is considered as a gold standard in clinical trials. However, this standard is criticized for being unethical when the equipoise on treatment effects is broken, because half of the patients are assigned to the inferior treatment.

Most researchers agree that randomization is the best method for achieving comparability among treatment groups and constitutes the basis of statistical inference. Randomization tends to balance the treatment groups with respect to known or unkown covariates and is used to protect against selection bias from investigators. The probabilities introduced by randomization establish the fundamentals of statistical inference. Randomization guarantees the validity of a statistical conclusion at a significant level. As Byar et al. (1976) claims, "Randomized clinical trials remain the most reliable method for evaluating the efficacy of therapies."

However, randomized clinical trials present a dilemma for investigators between individual and collective ethics when considering the responsibilities of investigators. On the one hand, investigators need to consider the well-being of individual patients within the trial and do what is the best for individual patients. On the other hand, investigators want to gather information about and draw valid statistical conclusions on treatment effects to benefit future patients. Royall (1991) examined the personal care principle and argued "that principle can make it difficult or impossible for a physician to participate in a randomized clinical study." The World Medical Association's Declaration of Helsinki clearly states that in medial research on human subjects, considerations related to the well-being of the human subject should take precedence over the interests of science and society. Palmer and Rosenberger (1999) described the role

of individual and collective ethics in designing clinical trials and claimed that greater attention should be paid to the former. The ethical tension between individual and collective ethics was also noticed by Pullman and Wang (2001).

For instance, the zidovudine trial (AZT) done by the AIDS clinical trial group (ACTG 076) was controversial on ethical grounds. The trial was to investigate the effect of a short term zidovudine (AZT) therapy on reducing the risk of maternal-infant HIV transmission. From April 1991 to December 1993, 477 pregnant women with HIV-infection were recruited into the trial. 239 of the women were given the short term AZT therapy using permuted block randomization and the remaining 238 constituted the placebo group. The first publication of the ACTG 076 study (1994) indicated that the effect of AZT on reducing the HIV transmission from mother to infants is statistically significant (P-value 0.00006) and the reduction is approximately two-thirds. A data analysis in 1996 confirmed the results in 1994 by ACTG 076 and the infection rate was 7.6% in the AZT group and 22.6% in the placebo group. The findings presented by ACTG 076 is a scientific breakthrough, but the randomization employed in the trial is controversial. Lurie and Wolfe (1997) argued that "on the basis of the ACTG 076 data, knowledge about the timing of perinatal transmission, and pharmacokinetic data, the researchers should have had every reason to believe that well-designed shorter regimens would be more effective than placebo. These findings seriously disturb the equipoise (uncertainty over the likely study result) necessary to justify a placebo-controlled trial on ethical grounds." Yao and Wei (1996) criticized the randomization procedure used in the trial and claimed that the procedure had put too many pregnant women on the placebo group and resulted in a large number of HIV-positive infants.

A good clinical trial practice should address both the individual and collective ethics. The 50-50 randomization is criticized for being unethical because individual

ethics (for trial patients) is often sacrificed, especially in desperate medical situations. According to Zelen (1969), "the ethical principle has motivated the development of statistical techniques which attempt to end a trial at the earliest moment when a decision can be made at to which therapy (if any) is most beneficial". Plenty of ethically motivated designs for clinical trials (see Armitage 1960, Anscombe 1963 and Palmer and Shahumyan, 2007), including response adaptive designs pioneered by Zelen (1969), have been proposed to alleviate the tension between the individual and collective ethics. Response adaptive designs modify the treatment allocation probabilities based on the treatment allocations and responses so far accumulated in the trial, and tend to assign more patients to the better treatment. Zelen (1969) found that, as a consequence of his response adaptive procedure (play-the-winner), "the ethical problem posed by an unnecessarily long clinical trial is not as serious as a trial where patients are allocated in equal numbers to the treatments for the length of the clinical trial."

Zelen's play-the-winner procedure was extended to a randomized play-the-winner (RPW) procedure by Wei and Durham (1978) in the thought of reducing selection bias. This design has been used in the extracorporeal membrane oxygenation (ECMO) trial conducted by Bartlett and his colleagues (1985) at the University of Michigan. The ECMO technique was employed to treat newborns with respiratory failure characterized by persistent pulmonary hypertension (PPH). PPH results in low blood flow through the lungs, thus inadequate oxygenation of the blood. Newborns with PPH are at high risk of death in the first days of life. ECMO is an external system for providing temporary support during respiratory failure. Historically, researches reported 80% survival rate in the use of ECMO to treat newborns with PPH and only 20% or less survival rate in other traditional treatments. Questions were raised on the comparability of the survival rates because of the absence of the concurrent control groups in

those studies. The Michigan trial group recognized the need of a control group to draw a valid conclusion and was also strongly concerned about the ethical issue raised by using 50-50 randomization. The Michigan ECMO trial is the first randomized clinical trial carried out on the effect of ECMO.

The RPW design used in the Michigan ECMO trial can be described by an urn which initially contains a ball labeled with ECMO and a ball with the label CT representing the conventional treatment. When a patient is ready to be treated, a ball is randomly drawn with replacement and the corresponding treatment is applied. The urn is updated when the outcome of an applied treatment is available. An additional ball of the same type is added to the urn if the treatment is successful. Otherwise, an additional ball representing the opposite treatment is added. The allocation probability of the next patient depends on the allocations of previous patients and responses of previous treated patients. In time, the urn is expected to contain a high proportion of balls associated with the more successful treatment. In the Michigan ECMO trial, the first ball drawn from the urn was an ECMO ball and the first baby was assigned to the ECMO treatment. The baby survived and an additional ECMO ball was added to the urn. The second baby was assigned to CT treatment using the urn and the baby died, so another additional ECMO ball was added to the urn. This procedure continued until the next seven babies were assigned to the ECMO treatment and all survived. The randomization ceased since the planned sample size was reached. Later, two more babies were treated with ECMO and both survived. In brief, 10 babies were treated with the ECMO treatment and all survived. Only one was assigned to the CT treatment by the RPW design but died. The data provided encouraging information about the survival rate of infants treated with ECMO, but was not conclusive because only one baby was allocated to the CT treatment.

However, the conclusion in another randomized ECMO trial conducted by the UK

collaborative trial group showed that the RPW design used in Michigan ECMO trial did assign more patients to the better treatment. The UK collaborative ECMO trial enrolled 185 mature infants from 1993 to 1995 and the recruitment was stopped early in November in 1995 on the advise of the Independent Data-monitoring Committee because the accumulated data showed a clear advantage with ECMO. 93 of the recruited infants were randomly assigned to ECMO and 92 to the CT treatment. The survival infants after the treatments were followed up to one year old to observe the morbidity status. The data analysis demonstrated that ECMO reduces the risk of death without a concomitant rise in severe disability. A follow-up study to 7 years (McNally et al., 2006) concluded that the beneficial influence of ECMO is still present. This trial was unethical because many infants were treated with the inferior treatment. But the trial did provide a sound conclusion and justified the routine use of ECMO in medical practice.

In summary, the 50-50 randomization presents ethical difficulties in clinical trials, especially in desperate medical situations. But ethical concerns should not preclude randomization in clinical trials. As one of the alternatives, response adaptive designs integrate randomization with the ethical consideration and provide a better way to alleviate the tension between the individual and collective ethics.

1.2.2 Types of Response Randomization Procedures

Response adaptive designs modify treatment allocations based on the accumulated responses of previously treated patients in the trial and deliberately bias treatment allocations to assign more patients to the potentially better treatment. The designs apply what is learned from the accumulated information on treatment effects to benefit future patients within the trial. Response adaptive designs have been studied for decades and many adaptive randomization procedures have been developed, including

the procedures based on urn models, sequential estimation procedures and decision-analytic procedures.

The treatment randomization based on an urn model is based on the urn composition process, which is modified according the responses of previous patients. This adapted urn composition process automatically represents the information gathering process on treatment effects in the trial.

An important contribution on urn models is the randomized play-the-winner (RPW) procedure proposed by Wei and Durham in 1978. This randomization procedure was used in the ECMO trial by Bartlett et al. in 1985 and in the clinical trial on fluoxetine versus placebo for depressive disorder (Tamura et al., 1994). Urn models originated from the Pólya urn scheme presented by Eggenberger and Pólya (1923). The Pólya urn was generalized by Friedman in 1949. The idea of using urn models for response adaptive randomization can be traced back to the researches in 1960s (Athreya and Karilin, 1967, 1968, and Zelen, 1969). Zelen (1969) introduced the play-the-winner procedure (PW) to alleviate the ethical tension between individual and collective ethics presented by the 50-50 randomization. Zelen's PW was generalized to the randomized play-the-winner design (RPW) by Wei and Durham in 1978, which inherits the spirit of Zelen's procedure that tends to assign more patients to the better treatment, but is less vulnerable to experimental bias because of the randomized allocation of treatments. In general, a RPW(α, β, γ) design is described by an urn that contains initially α balls of type A and β balls of type B representing treatments A and B respectively. When a patient is ready to be treated, a ball is drawn at random and the corresponding treatment is applied. Then the ball is returned to the urn and the urn is updated based on the patient's responses. An additional γ balls of the same type are added to the urn if the response of the patient is a success. For a failure, an additional γ balls of the opposite type are added. In time the urn is expected to

contain a high proportion of balls associated with the more successful treatment, thus future patients within the trial have a large probability to be assigned to the better treatment.

Since the works of Zelen (1969) and Wei and Durham (1978), many urn models have been developed using different adaptation rules on the composition of an urn. The updating rule employed by Durham and Yu (1990) is that a ball of the same type is added if there is a success and the urn composition remains the same for a failure. Ivanova et al. (2000) proposed the birth and death urn in which a ball is added to the urn for a success response and a ball is removed from the urn for a failure. The drop-the-loser rule (DL) developed by Ivanova (2003) adapts the urn composition by removing a ball if there is a failure and keeping the urn composition unchanged if there is a success. Furthermore, Ivanova and Flournoy (2001) generalized the binary response urn models to a ternary urn model with three outcomes. Other classes of urn models include the generalized Friedman's urn model (also called generalized Pólya urn model, GPU for short) (Wei, 1979, Rosenberger et al., 1997, Bai et al., 2002), the randomized Pólya urn model (Durham et al., 1998) and the sequential estimation-adjusted urn models (Zhang et al., 2006). Urn models were reviewed in Dirienzo (2000) and Rosenberger (2002). Recently, the drop-the-loser urn model has been extended by Sun et al. (2007) and Zhang (2007).

Another class of treatment randomization procedures in response adaptive designs is the sequential estimation procedures. This type of randomization is to target a certain proportion of treatment allocation. The target proportion is pre-specified according to the objective of the design and involves the unknown parameters of the treatments in the trial. The unknown parameters are sequentially estimated and the updated estimates are used in the randomization procedure to achieve the target proportion. Major sequential randomization procedures include the doubly-adaptive

biased coin design (DBCD)(Eisels, 1994, Eisele and Woodroofe, 1995, Hu et al., 2003, Hu and Zhang, 2004), the doubly adaptive weighted difference design procedure (Geraldes et al., 2006) and the sequential maximum likelihood procedure (Melfi et al., 2001, Roseberger et al., 2001, Baldi et al., 2005).

Among the sequential randomization procedures, the doubly-adaptive biased coin design is very flexible in that it can target any proportion. Moreover, this procedure is shown to be asymptotically less variable in proportions of treatment allocation than the maximum likelihood procedure (Hu et al., 2003). The idea of the doubly-adaptive biased coin design can be traced back to Efron's biased coin design which is used to balance treatment assignments. But the ethical concern in a clinical trial requires to bias the treatment allocation to assign more patients to the better treatment. To achieve a desired allocation proportion, Eisele (1994) and Eisele and Wooodroofe (1995) proposed the doubly-adaptive biased coin design procedure, where an allocation function was defined and the conditions on the allocation function were given. As Melfi et al. (2001) pointed out, the complicated nature of these conditions can be a barrier for the procedure to be applied in practice and the choice of allocation function in the example of the last section of the two papers (Eisele, 1994 and Eisele and Wooodroofe, 1995) violated their regularity conditions. Hu and Zhang (2004) developed a set of widely satisfied conditions for the allocation function for doubly-biased coin designs and proposed a specific allocation function. This specific allocation function was shown to generate asymptotically less variable allocation proportions (Hu et al., 2003) than other procedures, and the sequential maximum likelihood procedure was demonstrated to be a special case of the doubly-biased coin design under the specific allocation function. Geraldes et al. (2006) extended the idea of biased coin designs and proposed the doubly adaptive weighted difference design.

The third randomization procedure in response adaptive designs is the decision-

analytic procedure incorporating the Bayesian methodology. The learning on the treatment effect in the decision-analytic procedure is represented by the updating of the posterior distribution of parameters and the decision on the treatment allocation is made on the learning process. Early in the 1960s, Anscombe (1963) proposed the alternative formulation of sequential clinical trials from the point of view of Bayesian inference for ethical considerations. Later, Berry (1989) and Kass and Greenhouse (1989) argued that the alternative of randomized trial due to the ethical concern was to use the Bayesianly oriented methodology. The Bayesian method was used in Berry (2001, 2004), Berry and Eick (1995), Hardwick and Stout (1991), Muliere et al. (2006). Recently, Cheng and Berry (2007) proposed a r-optimal design which maximizes the expected utility in a Bayesian decision-analytic setting with an adaptive randomization allocation.

Other randomization procedures such as covariate-adjusted response adaptive randomization were considered by Rosenberger et al. (2001) and Biswas et al. (2006).

In brief, there are three major types of randomization procedures in response adaptive designs. Different procedures present different ways of information gathering in clinical trials. The common purpose of these procedures is to apply the collected information on treatment effects to assign more patients to the potentially better treatment, thus providing good medical practices for ethical considerations.

1.3 My Research

Response adaptive designs are very attractive to clinical and biostatistical researchers due to their efficiency and ethical advantages in clinical trials. However, the adaptation of the treatment allocation creates a dependence structure in the collected data and introduces more variability into the data, hence raises concerns about the validity of

statistical conclusions, power loss of statistical tests and experimental bias, etc. These issues lagged behind the application of response adaptive designs in practice. My dissertation focuses on the development of statistical inference methods for response adaptive designs and searching for a better treatment randomization procedure.

1.3.1 Issues and Literature Review

Response adaptive designs have been studied for decades. Statistical methodologies have been advanced and optimal adaptation procedures have been explored. However, the exploration of optimal treatment allocation procedures is very restricted and traditional statistical methodologies need to be justified for dependent data from response clinical trials, or new statistical methods must be developed.

Response adaptive designs tend to assign more patients to the better treatment based on the accumulated information of previously treated patients in the trial. Three classes of response adaptive designs developed in decades are discussed in the previous section. They are urn models, sequential estimation procedures and the decision-analytic models. In any of the response adaptive designs, interim information such as previous treatment allocation and the responses of previously treated patients are gathered for the modification of the probability of treatment allocation for the next patient. This adaptation of treatment allocation creates a dependency in the data collected from the trial and leads to difficulties and complications in statistical inference and in the development for the optimal allocation procedure.

Although statistical inference is very complicated for response adaptive designs, because of the dependency in the data from the trial, many statistical inference methods have been explored for adaptive designs in decades. Firstly, the exact distribution method was used for analyzing the data from the Michigan ECMO trial (Wei, 1988). Wei (1988) calculated the exact conditional p-value of the permutation test for the

Michigan ECMO trial in which the RPW design was used. Later, Wei et al. (1990) used the network method to derive the exact distribution of a sufficient statistic and derived the exact conditional and unconditional confidence intervals of the parameters for the ECMO data. Lin et al. (1991) employed the exact distribution method to the statistical inference of group sequential trials. However, the use of the exact distribution for the test statistic is computationally intensive for large scale clinical trials.

In addition, many researchers have investigated the limit theorems of allocation proportions and asymptotic properties of the maximum likelihood estimation of the parameters for different response adaptive models. Smythe (1996) and Bai et al. (2002) examined the urn composition process of GPU and demonstrated the asymptotic normality of the process. Inouse and Aki (2005) considered the multivariate distribution of the numbers of occurrences of different types of runs, and gave a recursive formula for the probability generating function of the GPU model. Under general assumptions on random generating matrices which determine how balls are added to the urn, Bai and Hu (2005) studied a very general urn model and established the strong consistency and asymptotic normality for both the urn composition and the proportion of treatment allocation. Eisele and Woodroofe (1995) proved the central limit theorems for the doubly adaptive biased coin design. Later, Hu and Zhang (2004) established the asymptotic properties of the proportions of treatment allocations for multi-treatment clinical trials with a doubly adaptive biased coin design. Recently, Zhang et al. (2007) explored the asymptotic properties of covariate-adjusted response-adaptive designs.

Several researchers also investigated the maximum likelihood estimation for general response adaptive models. Rosenberger et al. (1997) studied the maximum likelihood estimators for multi-parameter response adaptive designs. Regularity conditions were provided for the existence of the maximum likelihood estimator and its asymptotically

multivariate normality. Melfi and Page (2000) considered the estimation problem for general response adaptive designs and used the non-martingale approach to show the asymptotic normality of point estimators of the parameters. Melfi et al. (2001) then applied the martingale method to demonstrate the consistency and asymptotic normality of the point estimators for the adaptive design achieving a desired allocation proportion. Recently, Hu et al. (2006) also showed the asymptotic normality of the maximum likelihood estimator when studying the best response adaptive randomization procedures.

With the exception of maximum likelihood estimation, Coad and Woodroofe (1998) obtained some results on the bias of the maximum likelihood estimator for sequentially designed experiments. Coad and Ivanova (2001) derived the bias and variance of the maximum likelihood estimators of the probabilities of success and proposed bias-corrected estimators for adaptive urn designs. Cheng and Vidyashankar (2006) discussed the existence and asymptotic properties of the minimum Hellinger distance estimators for the randomized play-the-winner design. They established both consistency and asymptotic normality of the estimators.

Other methods such as nonparametric techniques are also used in the statistical analysis of response adaptive designs. Zhang and Rosenberger (2005) developed the log rank test for a wide class of randomization procedures including the adaptive randomization procedures. The Wilcoxon-Mann-Whitney score is used by Bandyopandhyay and Biswas (2004) to construct a test procedure for two univariate continuous populations.

Another concern on the use of response adaptive designs is the loss of power of statistical tests. The adaptation of treatment allocation introduces more variation to the estimators of parameters and to test statistics. Simulation results (Melfi and Page, 1998, Rosenberger et al., 2001, Ivanova, 2003) have demonstrated that a large

variance of the allocation proportion would reduce the power of the test. Hu and Rosenberger (2003) theoretically examined the relationship between the asymptotic power of Wald's test and the variance of allocation proportions. They found that less variability in the allocation proportion resulted in less loss in the asymptotic power of the test. Recently, Chen (2006) and Baldi Antognini (2007) provided theoretical analysis of power for the biased coin design.

The selection of treatment allocation procedures is a challenging problem in adaptive designs. On the one hand, a design is expected to assign as many patients as possible to the better treatment. On the other hand, too skewed treatment groups in a trial may result in the failure of drawing a valid statistical conclusion at the end of the trial. A response adaptive design has the advantage of balancing the individual ethics and collective ethics. However, different adaptive designs present very different tradeoffs between the individual and collective ethics.

A good response adaptive design is expected to assign more patients to the better treatment with a minimal loss in the power of the statistical test. Developing for an optimal design is complicated due to the adaptive process of treatment allocations. Rosenberger et al.(2001) obtained an optimal allocation proportion to minimize the expected number of treatment failures for a fixed power of the test. Although both the power and the expected number of failures were considered, this optimal allocation proportion doesn't depend on the desired power. Hu et al. (2006) established a lower bound on the asymptotic variances of the allocation proportions when the allocation proportions were asymptotically normally distributed. They concluded that Ivanova's DL design is the asymptotically best among the designs with the same allocation proportion as the DL's. The DBCD with allocation function $g(x, y)$ (proposed by Hu and Rosenberger (2003) and Hu and Zhang (2004)) was also claimed to be asymptotically best as the number of patients within the trial goes to infinity.

However, this comparison of designs is restricted to those with the same allocation proportion. Bandyopadhay and Bhattacharya (2006) developed a randomization rule which switches between the Neyman allocation and the myopically better treatment and conducted a simulation comparison of their randomization rule with the existing adaptive designs using the expected failure proportion criterion. Recently, Biswas et al. (2007) considered optimal response adaptive designs for continuous responses in Phase III trials. Cheng and Berry (2007) introduced a r-optimal design, a constrained adaptive randomized design in Bayesian decision-analytic setting, to maintain the randomness of treatment allocations. But the comparison of the r-optimal design with existing designs has not been conducted. In conclusion, different evaluation methods have been developed in the search for optimal response adaptive designs. However, the comparisons of adaptive designs are restricted to a particular class of designs. Furthermore, the optimal design may result in extremely unbalanced treatment groups and thus becomes unethical or difficult for statistical analysis.

In summary, both the tasks of searching for an optimal design and developing appropriate statistical inference for response adaptive designs are very challenging. Although statistical inference methodologies for response adaptive designs have been advanced, more traditional statistical inference methods need to be extended to dependent data from response adaptive designs, and efficient estimation and powerful tests need to be investigated, or new statistical methods must be developed. Besides, the exploration of optimal designs is very restricted and new methods to evaluate response adaptive designs need to be advanced.

1.3.2 Main Results

In my dissertation, I discuss several statistical inference methods such as the log likelihood ratio test, goodness-of-fit test and efficient estimation for a wide class of

response adaptive designs. The evaluation of response adaptive designs is also explored and a new randomization procedure is proposed, which is better than the existing procedures under common conditions.

Firstly, I investigate the randomization process of response adaptive clinical trials. The information gathering process in the adaptive trial has been formulated as a stochastic process, especially a Markov process for dichotomous responses. From the formulated stochastic process, the likelihood function for the observed data is derived. This explains why the format of the likelihood function for response adaptive designs is not affected by the use of adaptive treatment allocation and why there is more variation in the maximum likelihood estimators of the parameters for adaptive designs.

Then common test statistics including the log likelihood ratio statistic and goodness-of-fit test statistic are explored for response adaptive designs. I examine the limiting properties of proportions of treatment allocation using the theory of martingales and discuss the consistency and asymptotic normality of the maximum likelihood estimators for a wide class of adaptive designs. It is shown that the maximum likelihood estimators are strongly consistent and asymptotically normally distributed under some regularity conditions. These results hold for a wide class of response adaptive designs including the RPW design, the GPU model and the designs with a targeted allocation proportion. Under some regularity conditions, the logarithm of the likelihood ratio statistic for dependent data from a general class of response adaptive designs is proven to be asymptotically chi-square distributed. This provides a foundation for asymptotic analysis of adaptive clinical trials with multiple treatments. For response adaptive designs with dichotomous responses, under assumptions less restricted than those for general models, the estimated odds ratio and its logarithm are shown to follow asymptotically normal distributions. Moreover, the ordinary goodness-of-fit test statistic for two-by-two contingency tables with dependent data is proven to be

asymptotically chi-square distributed.

The third main result is on the efficient estimation for response adaptive designs. In the literature, the maximum likelihood estimators, bias-corrected estimators and minimum Hellinger distance estimators were discussed for response adaptive designs. An interesting question is which estimation is the best under what criterion. I discuss the issue of asymptotic efficiency of estimation for response adaptive designs of clinical trials. The asymptotic lower bound of exponential rates for consistent estimators is established and it is shown that under certain regularity conditions, the maximum likelihood estimator attains the asymptotic lower bound for response adaptive designs with binary observations. The estimation of the treatment effect is also investigated and the maximum likelihood estimator of the treatment effect is shown to be asymptotically efficient in the Bahadur sense under some regularity conditions for response adaptive designs with general responses.

At last, the evaluation of response adaptive designs is explored. I examine the optimality properties of response adaptive designs with a variance-penalized criterion. All response adaptive designs including those with different allocation proportions can be compared under this criterion. More importantly, the penalty criterion evaluates the performance of a design according to both the mean and the variability of the total responses. A good design under this criterion tends to allocate more patients to the better treatment and to increase the power of the test. I propose such a design and compare the design with some existing response adaptive designs. The asymptotic variance of the allocation proportion of our proposed design is shown to be smaller than that of the DL design if $p_A + p_B > \epsilon$, where ϵ is a pre-fixed number, a measure of tradeoff between individual and collective ethics. Simulation results indicate that our proposed design is better than other existing designs under the variance-penalized criterion, except for extreme values of the probabilities of success (such as very large

p_A and p_B or extremely small p_B). Potentially, our proposed design assigns a higher proportion of patients to the better treatment than the existing adaptive designs and the power of the statistical test remains good when the difference between p_A and p_B is not small (say, larger than or equal to 0.4). For trials with a large number of patients, the overall performance of our design is better than the existing designs.

In this dissertation, Chapter 2 introduces the formulation of response adaptive designs. The results on the log likelihood ratio statistic and the goodness-of-fit test statistic are presented in Chapter 3. The efficient estimation problem for response adaptive designs is discussed in Chapter 4. Chapter 5 examines the evaluation issue and the optimal properties of response adaptive designs. Chapter 6 concludes the dissertation and discusses further research directions.

Chapter 2

Formulation of Response Adaptive Designs

Response adaptive designs are characterized by randomized treatment allocation rules that are adaptive to previous responses. This chapter introduces the information gathering processes in an adaptive clinical trial and the Markov decision models for response adaptive designs with dichotomous responses.

2.1 Formulation of Adaptive Designs with General Responses

Suppose that trial subjects arrive sequentially and each receives one and only one of k treatments. Patients' responses $Y_{1j}, Y_{2j}, \cdots$ from treatment j are independent and identically distributed with the density function $f_j(y, \theta_j)$, where $\theta_j \in \Theta_j$, $1 \leq j \leq k$, is an unknown parameter. Denote $\boldsymbol{\theta} = (\theta_1, \theta_2, \cdots, \theta_k)^t$, where t stands for transpose. Besides making a statistical comparison of the alternative treatments at the conclusion of the trial, it is also desired to allocate trial subjects to the potentially best treatment as many as possible in order to balance collective and individual ethics. Response adaptive designs are aimed at this purpose by sequentially allocating treatments adaptive to responses so far accumulated in the trial.

Let $\boldsymbol{\delta}_i = (\delta_{i1}, \delta_{i2}, \cdots, \delta_{ik})$ be the i^{th} treatment assignment such that $\delta_{ij} = 1$ if the i^{th} patient receives treatment j and $\delta_{ij} = 0$ otherwise, and $\mathbf{y}_i = (Y_{i1}\delta_{i1}, Y_{i2}\delta_{i2}, \cdots, Y_{ik}\delta_{ik})$ be the corresponding response. Here we use the convention that if treatment j is not applied to patient i, then the response is $\mathbf{0}$. When the i^{th} patient ($i \geq 2$) is to be treated, the information available is given by the σ algebra $\mathcal{F}_{i-1}$ generated by $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_{i-1}, \mathbf{y}_{i-1})\}$.

A response adaptive design is defined by a sequence of possibly randomized allocation rules $\boldsymbol{\pi} = \{\boldsymbol{\pi}_i,\ i = 1, 2, \cdots\}$ such that each $\boldsymbol{\pi}_i = (\pi_{i1}, \pi_{i2}, \cdots, \pi_{ik})$ is given by the conditional treatment allocation probabilities $\pi_{ij} = P(\delta_{ij} = 1|\mathcal{F}_{i-1})$, $i \geq 2$, $\sum_{j=1}^{k} \pi_{ij} = 1$, and the initial possibly randomized treatment allocation probabilities $\pi_{1j} = P(\delta_{1j} = 1)$ are pre-fixed values (such as $1/k$), $1 \leq j \leq k$. Moreover, each randomized allocation rule $\boldsymbol{\pi}$ defines a probability measure $P_{\boldsymbol{\pi}}$ on the space of all possible sequences of treatment allocations and responses.

If $n, n = 1, 2, \cdots$, patients have been treated in the adaptive trial, let $N_j(n)$ be the number of patients allocated to treatment j and $\mathbf{X}_j(n) = (Y_{1j}, Y_{2j}, \cdots, Y_{N_j(n)j})$ be the corresponding observations on treatment j, where $j = 1, 2, \cdots, k$. Then $\sum_{j=1}^{k} N_j(n) = n$. Define

$$\mathbf{W}(n) = (N_1(n), N_2(n), \cdots, N_{k-1}(n), \mathbf{X}_1(n), \mathbf{X}_2(n), \cdots, \mathbf{X}_k(n)).$$

Clearly the information contained in $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_n, \mathbf{y}_n)\}$ is equivalent to the information contained in $\{\mathbf{W}(1), \mathbf{W}(2), \cdots, \mathbf{W}(n)\}$. So the treatment allocation for the $(n+1)^{th}$ patients depends on $\{\mathbf{W}(1), \mathbf{W}(2), \cdots, \mathbf{W}(n)\}$. Hence, $\{\mathbf{W}(n), n = 1, 2, \cdots\}$ becomes a stochastic process with a transition probability function specified by the allocation rule $\boldsymbol{\pi} = \{\boldsymbol{\pi}_n,\ n = 1, 2, \cdots\}$.

The transition probability function of the stochastic process $\mathbf{W}(n), n = 1, 2, \cdots$, under $P_{\boldsymbol{\pi}}$ is $\prod_{j=1}^{k} [\pi_{ij} f_j(y_{ij}, \theta_j)]^{\delta_{ij}}$.

For each observed sequence $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_n, \mathbf{y}_n)\}$, the likelihood function is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^{n}\prod_{j=1}^{k}[\pi_{ij} f_j(y_{ij}, \theta_j)]^{\delta_{ij}} = h(\boldsymbol{\pi}) \prod_{i=1}^{n}\prod_{j=1}^{k} f_j(y_{ij}, \theta_j)^{\delta_{ij}}$$

where $h(\boldsymbol{\pi}) = \prod_{i=1}^{n}\prod_{j=1}^{k} \pi_{ij}^{\delta_{ij}}$, and $\prod_{j=1}^{k} f_j(y_{ij}, \theta_j)^{\delta_{ij}}$ is the contribution of $(\boldsymbol{\delta}_i, \mathbf{y}_i)$ to the likelihood, $0^0 = 1$, and $\infty^0 = 1$.

Compared to the likelihood function for the data from an independent and identical distribution, the likelihood function for a response adaptive design is related to the randomized allocation rule $\boldsymbol{\pi}$. However, the second part in the likelihood function containing the parameters is not affected by $\boldsymbol{\pi}$ directly.

2.2 Markov Decision Models for Response Adaptive Designs with Dichotomous Responses

If the response in an adaptive trial is dichotomous (say, success or failure), the information gathering process can be formulated as a Markov decision process.

Let θ_j be the probability of success on treatment j. Then if treatment j is assigned to patient i, its response Y_{ij} follows the Bernoulli distribution $f_j(y_{ij}, \theta_j) = (\theta_j)^{y_{ij}}(1 - \theta_j)^{1-y_{ij}}$, where $y_{ij} = 1$ for a success and 0 for a failure.

After n patients have been treated in the adaptive trial, $N_j(n)$ is the number of patients allocated to treatment j as defined previously. Let $S_j(n)$ be the number of successes on treatment j, where $j = 1, 2, \cdots, k$. Then $N_j(n) = \sum_{i=1}^{n} \delta_{ij}, S_j(n) = \sum_{i=1}^{n} \delta_{ij} y_{ij}$. Clearly the information contained in $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_n, \mathbf{y}_n)\}$ is equivalent to the information contained in

$\{(N_1(1), \cdots, N_{k-1}(1), S_1(1), \cdots, S_k(1)), \cdots, (N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))\}$.

So the treatment allocation for the $(n+1)^{th}$ patient depends on the observations of the stochastic process $\{(N_1(i), \cdots, N_{k-1}(i), S_1(i), \cdots, S_k(i)), i = 1, 2, \cdots, n\}$.

For each randomized allocation rule $\boldsymbol{\pi}$, the likelihood function for the observed sequence $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_n, \mathbf{y}_n)\}$ is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^{n}\prod_{j=1}^{k}[\pi_{ij}\theta_j^{y_{ij}}(1-\theta_j)^{1-y_{ij}}]^{\delta_{ij}} = \prod_{j=1}^{k}\left(\prod_{i=1}^{n}\pi_{ij}^{\delta_{ij}}\right)\theta_j^{S_j(n)}(1-\theta_j)^{N_j(n)-S_j(n)}.$$

Therefore, the statistic $(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))$ is sufficient and the maximum likelihood estimators for θ_j is $\hat{\theta}_j = S_j(n)/N_j(n),\ j = 1, 2, \cdots, k$. The statistic

$(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))$ summarizes the previous treatment allocations and accumulated information of the responses of previously treated patients in the trial. Thus, it is reasonable to set the randomized allocation $\boldsymbol{\pi}_{n+1}$ for the $(n+1)^{th}$ patient depending only on $(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))$. An allocation rule with this property is said to be a Markov randomization rule.

In brief, the information gathering process on the treatment effects in a binary response adaptive clinical trial can be formulated as the following stochastic process:

$$\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)), n = 1, 2, \cdots\}.$$

Due to the sufficiency of $(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))$, only Markov allocation rules $\boldsymbol{\pi} = \{\boldsymbol{\pi}_i,\ i = 1, 2, \cdots\}$, are to be considered, where $\boldsymbol{\pi}_{i+1}$ depends only on the current state

$(N_1(i), \cdots, N_{k-1}(i), S_1(i), \cdots, S_k(i))$. This class of allocation rules includes a wide range of adaptive designs, such as the RPW(α, β) design, the drop-the-loser design (Ivonava 2003), the optimal adaptive design (Rosenberger et al., 2000), the GPU design, and the doubly adaptive biased coin design with binary responses.

For example, in a trial with two treatments, say treatments A and B, the RPW(α, β) design assigns the i^{th} patient to treatment A with the probability

$$\pi_{iA} = P(\delta_{iA} = 1|\mathcal{F}_{i-1}) = \frac{\alpha + S_A(i-1)\beta + (N_B(i-1) - S_B(i-1))\beta}{2\alpha + (i-1)\beta}.$$

Instead, the probability in the optimal design proposed by Rosenberger (2000) is

$$\pi_{iA} = P(\delta_{iA} = 1|\mathcal{F}_{i-1}) = \frac{R^*(\hat{\theta}_A(i-1), \hat{\theta}_B)(i-1)}{1 + R^*(\hat{\theta}_A(i-1), \hat{\theta}_B(i-1))},$$

where $R^*(\theta_A, \theta_B)$ is the optimal proportion of patients assigned to treatment A to these assigned to treatment B and $\hat{\theta}_A = S_A/N_A, \hat{\theta}_B = S_B/N_B$. In the doubly biased coin designs with the allocation function developed by Hu and Zhang (2004), the probability of allocating the i^{th} patient to treatment A is

$$\pi_{iA} = \begin{cases} 1 & \text{if } \hat{x} = 0, \\ 0 & \text{if } \hat{x} = 1, \\ \frac{\hat{\rho}(\hat{\rho}/\hat{x})^\gamma}{\hat{\rho}(\hat{\rho}/\hat{x})^\gamma + (1-\hat{\rho})((1-\hat{\rho})/(1-\hat{x}))^\gamma} & \text{if } 0 < \hat{x} < 1, \end{cases}$$

where x is the proportion of patients to treatment A, $\hat{x} = N_A(i-1)/(i-1)$, ρ is the target proportion for treatment A and the estimated value $\hat{\rho}$ of ρ depends on $\hat{\theta}_A(i-1)$ and $\hat{\theta}_B(i-1)$. In all the above designs, the allocation probability π_{iA} depends only on the current state of the information process $\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)), n = 1, 2, \cdots, \}$. In other words, the allocation rules are Markovian in these adaptive designs.

Under a Markov allocation rule, the process

$$\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)), n = 1, 2, \cdots, \}$$

is a Markov process. The Markov property of $\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)), n = 1, 2, \cdots\}$ was firstly noticed by Wei et al (1990) in the randomized play-the-winner design for $k = 2$. Actually, it can be proved that $\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)), n = 1, 2, \cdots\}$ is a Markov process under the probability measure $P_{\boldsymbol{\pi}}$ specified by a Markov allocation rule $\boldsymbol{\pi}$ for any adaptive designs with dichotomous responses.

The decision model for response adaptive designs with dichotomous responses consists of

1. a state space $\{(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n)) \in \mathcal{N}^{2k-1}\}$,

2. an allocation rule $\boldsymbol{\pi} = \{\boldsymbol{\pi}_n,\ n = 1, 2, \cdots\}$,

3. a transition probability specified by $\boldsymbol{\pi}$,

4. an objective to minimize the total number of failures or to achieve other goals.

The decision model becomes a Markov decision model if the allocation rule $\boldsymbol{\pi}$ is Markovian. Based on the formulation, the objective function is used to compare two adaptive designs $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$.

In brief, the adaptation of treatment allocation in response adaptive designs is based on accumulated information such as previous treatment allocations and responses of previously treated patients. The information gathering process can be formulated as a stochastic process, especially a Markov process for designs with dichotomous responses. This formulation explains the format of the likelihood function of response adaptive designs.

Chapter 3

Likelihood Ratio Test and Goodness-of-fit Test

This chapter examines the extension of common statistical procedures such as the log-likelihood ratio test and the goodness-of-fit test to dependent data from a wide class of response adaptive designs.

3.1 Likelihood ratio test

Let $\mathcal{F}_n$ be the σ-algebra generated by $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, (\boldsymbol{\delta}_n, \mathbf{y}_n)\}$, representing the information available for allocating the treatment to the $(n+1)^{st}$ patient, $n = 1, 2, \cdots$. For each j, $j = 1, 2, \cdots, k$, $\sum_{i=1}^{n}(\delta_{ij} - \pi_{ij})$ is a martingale and $\sum_{i=1}^{\infty} i^{-2}E[(\delta_{ij} - \pi_{ij})^2|\mathcal{F}_{i-1}] < \infty$ since $|\delta_{ij} - \pi_{ij}| \leq 1$. Therefore, by the strong law of large numbers for martingales, $\frac{1}{n}\sum_{i=1}^{n}(\delta_{ij} - \pi_{ij}) \to 0$ almost surely. We have then

Lemma 3.1.1. $\frac{N_j(n)}{n} - \frac{\sum_{i=1}^{n}\pi_{ij}}{n} \to 0$ *almost surely,* $j = 1, 2, \cdots, k$ *under the allocation rule* π.

Lemma 3.1.1 is an extension of Proposition 1 in Melfi et al. (2001), whose allocation rule is limited with a target allocation proportion, but our allocation rule applies to a wide class of adaptive designs including Melfi's, the RPW design, the optimal adaptive

design in Rosenberger et al. (2000), the GPU model and the doubly adaptive biased coin design.

The log-likelihood function is

$$l(\boldsymbol{\theta}) = \ln L(\boldsymbol{\theta}) = \ln h(\boldsymbol{\pi}) + \sum_{i=1}^{n} g(i, \boldsymbol{\theta}),$$

where $g(i, \boldsymbol{\theta}) = \sum_{j=1}^{k} \delta_{ij} \ln f_j(y_{ij}, \theta_j)$. Obviously,

$$\frac{\partial l(\boldsymbol{\theta})}{\partial \theta_j} = \sum_{i=1}^{n} g'_{\theta_j}(i, \boldsymbol{\theta}) = \sum_{i=1}^{n} \delta_{ij} (\ln f_j(Y_{ij}, \theta_j))'_{\theta_j}.$$

Under the usual regularity conditions, the Fisher information $I_j(\theta_j) = -E\left((\ln f_j(X, \theta_j))''_{\theta_j}\right)$ is finite and positive and there exists a solution $\hat{\theta}_j$ of $\frac{\partial l(\boldsymbol{\theta})}{\partial \theta_j} = 0,\ j = 1, 2, \cdots, k$. Further we assume that the second moment of $(\ln f_j(X, \theta_j))''_{\theta_j}$ exists and is finite, and there exists $\alpha > 0, M < \infty$ such that $E(|(\ln f_j(Y_{ij}, \theta_j))'_{\theta_j}|^{2+\alpha}) < M$.

Let $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \hat{\theta}_2, \cdots, \hat{\theta}_k)^t$ and $\mathrm{l}'(\boldsymbol{\theta}) = \left(\frac{\partial l(\boldsymbol{\theta})}{\partial \theta_1}, \frac{\partial l(\boldsymbol{\theta})}{\partial \theta_2}, \cdots, \frac{\partial l(\boldsymbol{\theta})}{\partial \theta_k}\right)^t$.

Lemma 3.1.2. *If the allocation rule* $\boldsymbol{\pi} = \{\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \cdots\}$ *satisfies the condition that* $\frac{\sum_{i=1}^{n} \pi_{ij}}{n} \to v_j(\boldsymbol{\theta}) \in (0, 1)$, *as* $n \to \infty$, $j = 1, 2, \cdots, k$, *then as* $n \to \infty$,

(1) $\hat{\boldsymbol{\theta}} \to \boldsymbol{\theta}$ *almost surely,*

(2) $n^{-1/2} \boldsymbol{l}'(\boldsymbol{\theta}) \to N_k(\boldsymbol{0}, \Gamma(\boldsymbol{\theta}))$,

(3) $n^{1/2}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \to N_k(\boldsymbol{0}, \Gamma^{-1}(\boldsymbol{\theta}))$,

where $\Gamma(\boldsymbol{\theta}) = diag(v_1(\boldsymbol{\theta}) I_1(\theta_1), v_2(\boldsymbol{\theta}) I_2(\theta_2), \cdots, v_k(\boldsymbol{\theta}) I_k(\theta_k))$.

Proof. Notice that $g''_{\theta_j}(i, \boldsymbol{\theta}) = \delta_{ij} (\ln f_j(Y_{ij}, \theta_j))''_{\theta_j}$ and

$$E[g''_{\theta_j}(i, \boldsymbol{\theta}) | \mathcal{F}_{i-1}] = \pi_{ij} E\left[(\ln f_j(Y_{ij}, \theta_j))''_{\theta_j}\right] = -\pi_{ij} I_j(\theta_j).$$

Therefore, $\sum_{i=1}^{n} (g''_{\theta_j}(i, \boldsymbol{\theta}) + \pi_{ij} I_j(\theta_j))$ is a martingale. Since $E[(\ln f_j(Y_{ij}, \theta_j))''_{\theta_j}]^2$ is finite, $\sum_{i=1}^{\infty} i^{-2} E[(g''_{\theta_j}(i, \boldsymbol{\theta}) + \pi_{ij} I_j(\theta_j))^2 | \mathcal{F}_{i-1}] < \infty$. Using the strong law of large number for martingale,

$$\lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} g''_{\theta_j}(i, \boldsymbol{\theta}) = -\lim_{n \to \infty} \frac{\sum_{i=1}^{n} \pi_{ij} I_j(\theta_j)}{n} = -v_j(\boldsymbol{\theta}) I_j(\theta_j) \tag{3.1.1}$$

almost surely. Applying the strong law of large number for martingales to $\sum_{i=1}^{n} g'_{\theta_j}(i,\theta)$,

$$(\sum_{i=1}^{n} g'_{\theta_j}(i,\boldsymbol{\theta}))/n \to 0 \text{ almost surely .}$$

By Taylor expansion of $(\sum_{i=1}^{n} g'_{\theta_j}(i,\boldsymbol{\theta}))/n$ and equation (3.1.1),

$$(\sum_{i=1}^{n} g'_{\theta_j}(i,\boldsymbol{\theta}))/n = -v_j(\boldsymbol{\theta})I_j(\theta_j)(\hat{\theta}_j - \theta_j) + o_p(\hat{\theta}_j - \theta_j).$$

Hence part (1) is proved. Obviously,

$$\lim_{n\to\infty} \frac{\sum_{i=1}^{n} E[g''_{\theta_j}(i,\boldsymbol{\theta})|\mathcal{F}_{i-1}]}{n} = -\lim_{n\to\infty} \frac{\sum_{i=1}^{n} \pi_{ij} I_j(\theta_j)}{n} = -v_j(\boldsymbol{\theta})I_j(\theta_j)$$

almost surely and

$$n^{-1-\alpha/2} \sum_{i=1}^{n} E(|\delta_{ij}(\ln f_j(y_{ij},\theta_j))'_{\theta_j}|^{2+\alpha}|\mathcal{F}_{i-1}) \le \frac{M}{n^{\alpha/2}}.$$

Therefore parts (2) and (3) follow from Rosenberger et al. (2002). □

Lemma 3.1.2 implies strong consistency and asymptotic normality of the maximum likelihood estimator of $\boldsymbol{\theta}$ for any allocation rule $\boldsymbol{\pi} = \{\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \cdots\}$ satisfying the regularity condition $\sum_{i=1}^{n} \pi_{ij}/n \to v_j(\boldsymbol{\theta}) \neq 0, 1$. This condition is not restrictive and holds true for many adaptive designs, such as the RPW(0,1) (where $\sum_{i=1}^{n} \pi_i/n \to q_B/(q_A+q_B)$ almost surely), the optimal adaptive design of Rosenberger et al. (2001) (where for the sequential maximum procedure, $\sum_{i=1}^{n} \pi_i/n \to R^*/(1+R^*)$ almost surely), and the allocation rule in Melfi et al. (2001) if the desired allocation proportion is not 0 or 1.

Lemma 3.1.2 is more general than that in Hu et al. (2006) in which the response function is restricted to the exponential family. Our results are more general and include non-exponential families. Furthermore, Hu et al. (2006) and Rosenberger and Lachin (2002) did not consider strong consistency. Melfi and Page (2000) derived results similar to lemma 3.1.2 by a non-martingale theory, but their allocation rule is a special case of ours.

Let $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}^0$ be the null hypothesis, where $\boldsymbol{\theta}^0 \in \Theta$ and the parameter space $\Theta = \Theta_1 \times \Theta_2 \cdots \times \Theta_k$ is an open set. All results established in this section require the following regularity condition:

Regularity Condition 1: $\sum_{i=1}^{n} \frac{\pi_{ij}}{n} \to v_j(\boldsymbol{\theta}) \in (0,1)$, $j = 1, 2, \cdots, k$ almost surely.

Theorem 3.1.3. *Under regularity condition 1, the statistic*

$$-2[l(\boldsymbol{\theta}^0) - \max_{\Theta} l(\boldsymbol{\theta})] = 2[l(\hat{\boldsymbol{\theta}}) - l(\boldsymbol{\theta}^0)]$$

follows asymptotically the χ^2 distribution with k degrees of freedom when the null hypothesis is true.

Proof. The Taylor expansion of $l(\boldsymbol{\theta}) = \ln L(\boldsymbol{\theta}) = \ln h(\boldsymbol{\pi}) + \sum_{i=1}^{n} g(i, \boldsymbol{\theta})$ at $\boldsymbol{\theta}^0$ gives us

$$l(\hat{\boldsymbol{\theta}}) - l(\boldsymbol{\theta}^0) = \sum_{j=1}^{k} \left[(\hat{\theta}_j - \theta_j^0) \sum_{i=1}^{n} g'_{p_j}(i, \boldsymbol{\theta}^0) \right] + \frac{1}{2} \left[\sum_{j=1}^{k} (\hat{\theta}_j - \theta_j^0)^2 \sum_{i=1}^{n} g''_{\theta_j}(i, \boldsymbol{\theta}^0) \right] + o_p(1).$$

Hence,

$$2(l(\hat{\boldsymbol{\theta}}) - l(\boldsymbol{\theta}^0)) = 2(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)^t \mathrm{l}'(\boldsymbol{\theta}^0) + \sum_{j=1}^{k} (\hat{\theta}_j - \theta_j^0)^2 \sum_{i=1}^{n} g''_{\theta_j}(i, \boldsymbol{\theta}^0) + o_p(1). \quad (3.1.2)$$

It's obvious that

$$(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)^t \mathrm{l}'(\boldsymbol{\theta}^0) = n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]^t [n^{-\frac{1}{2}} \Gamma^{-\frac{1}{2}}(\boldsymbol{\theta}^0) \mathrm{l}'(\boldsymbol{\theta}^0)],$$

$$\sum_{j=1}^{k} (\hat{\theta}_j - \theta_j^0)^2 \sum_{i=1}^{n} g''_{\theta_j}(i, \boldsymbol{\theta}^0) = \sum_{j=1}^{k} [n^{\frac{1}{2}} (\hat{\theta}_j - \theta_j^0)]^2 \left[\frac{1}{n} \sum_{i=1}^{n} g''_{\theta_j}(i, \boldsymbol{\theta}^0) \right].$$

By Lemma (3.1.2), the first term in equation 3.1.2 is equivalent to $2[n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]^t [n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]$ and using equation (3.1.1), the second term in equation (3.1.2) is equivalent to $-[n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]^t [n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]$, where equivalence means that the two random variables have the same asymptotic distribution. Therefore, $2(l(\boldsymbol{\theta}) - l(\boldsymbol{\theta}^0))$ is equivalent to $[n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]^t [n^{\frac{1}{2}} \Gamma^{\frac{1}{2}}(\boldsymbol{\theta}^0)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0)]$ which follows

asymptotically the χ^2 distribution with k degrees of freedom when the null hypothesis is true. □

If the response functions $f_j(y, \theta_j)$, $j = 1, 2, \cdots, k$, have the same form $f(y, \theta_j)$, then the statistic in the following theorem can be used to test the null hypothesis $H_0' : \theta_1 = \theta_2 = \cdots = \theta_k$.

Theorem 3.1.4. *Under the regularity condition 1, the statistic*

$$-2\left[\max_{H_0'} l(\boldsymbol{\theta}) - \max_{\Theta} l(\boldsymbol{\theta})\right]$$

follows asymptotically the χ^2 distribution with $k-1$ degree of freedom when the null hypothesis $H_0' : \theta_1 = \theta_2 = \cdots = \theta_k$ is true.

Proof. Set $\boldsymbol{\theta}^0 = (\theta^0, \theta^0, \cdots, \theta^0)$. For the linear mapping $\theta_1 = \theta$, $\theta_2 = \theta, \cdots, \theta_k = \theta$, Theorem 11.2 in Billingsley (1961) shows that $2(\max_{H_0'} l(\boldsymbol{\theta}) - l(\boldsymbol{\theta}^0))$ is asymptotically χ^2 distributed with 1 degree of freedom, and $2(\max_{\Theta} l(\boldsymbol{\theta}) - \max_{H_0'} l(\boldsymbol{\theta}))$ is asymptotically χ^2 distributed with $k-1$ degrees of freedom. □

Although Billingsley (1961) derived results similar to theorems 3.1.3 and 3.1.4, his results apply only to stationary Markov processes. But the stochastic process $\{\mathbf{W}(m), m = 1, 2, \cdots\}$ is neither stationary nor Markov for general response adaptive designs. Hence our results are more general.

3.2 Goodness-of-fit test for contingency tables with dependent data

As an application of the results in Section 3.1, we consider hypothesis tests for response adaptive clinical trials with dichotomous response (say, success and failure). A traditional method for comparing dichotomous populations based on independent samples is the goodness-of-fit test. In this section, we extend this test to dependent data from

k dichotomous populations. Our focus is on the asymptotic sampling distribution of the test statistic.

Let θ_j be the probability of success on treatment j. Then if treatment j is assigned to patient i, its response $Y_{ij} \sim B(1, \theta_j)$, the Bernoulli distribution.

For each (randomized) allocation rule $\boldsymbol{\pi} = \{\boldsymbol{\pi}_n,\ n = 1, 2, \cdots\}$, the likelihood function for the observed sequence $\{(\boldsymbol{\delta}_1, \mathbf{y}_1), \cdots, \boldsymbol{\delta}_1, \mathbf{y}_1)\}$ is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^{n}\prod_{j=1}^{k}[\pi_{ij}\theta_j^{y_{ij}}(1-\theta_j)^{1-y_{ij}}]^{\delta_{ij}} = \prod_{j=1}^{k}\left(\prod_{i=1}^{n}\pi_{ij}^{\delta_{ij}}\right)\theta_j^{S_j(n)}(1-\theta_j)^{N_j(n)-S_j(n)},$$

where $N_j(n)$ is the number of patients allocated to treatment j and $S_j(n)$ is the number of successes on treatment j among the n patients, $j = 1, 2, \cdots, k$. Therefore, the statistic $(N_1(n), \cdots, N_{k-1}(n), S_1(n), \cdots, S_k(n))$ is sufficient and the unrestricted maximum likelihood estimators for θ_j is $\hat{\theta}_j = S_j(n)/N_j(n)$.

Let $H_0 : \theta_j = \theta_j^0,\ j = 1, 2, \cdots, k$, be the null hypothesis such that $0 < \theta_j^0 < 1,\ j = 1, 2, \cdots, k$. The parameter space is set to be $\Theta = (0, 1)^k$. It can be shown that for adaptive designs with dichotomous responses, $\sum_{i=1}^{n}\pi_{ij}/n \to N_j(n)/n$ almost surely. That is, under the condition that $\sum_{i=1}^{n}\pi_{ij}/n \to v_j(\boldsymbol{\theta}),\ 0 < v_j(\boldsymbol{\theta}) < 1$ almost surely, we have $N_j(n) \to \infty$ almost surely.

Writing $\lambda = L(\boldsymbol{\theta}^0)/L(\hat{\boldsymbol{\theta}})$ as the likelihood ratio, the following result is a consequence of Theorem 3.1.3.

Corollary 3.2.1. *Under the regularity condition 1, the statistic* $-2\ln\lambda$ *is asymptotically* χ^2 *distributed with* k *degrees of freedom when the null hypothesis* $H_0 : \theta_j = \theta_j^0,\ j = 1, 2, \cdots, k$ *is true.*

After n patients have been treated in an adaptive clinical trial with k treatments and dichotomous responses, the contingency table has the following data structure

Treatment	1	2	$\cdots$	k	Total
Success	$S_1(n)$	$S_2(n)$	$\cdots$	$S_k(n)$	$S(n)$
Failure	$N_1(n) - S_1(n)$	$N_2(n) - S_2(n)$	$\cdots$	$N_k(n) - S_k(n)$	$n - S(n)$
Total	$N_1(n)$	$N_2(n)$	$\cdots$	$N_k(n)$	n

where $S(n) = \sum_{j=1}^{k} S_j(n)$. The dependency among the entries is reflected in the observed treatment allocations.

The traditional form of the test statistic

$$\chi^2 = \sum_{all\ cells} \frac{(observed - expected)^2}{expected}$$

is applied for dependent data. For simplicity, denote $S_j(n) = S_j$ and $N_j(n) = N_j, j = 1, 2, \cdots, k$.

For the proof of the next result, let $X_n \sim Y_n$ denote that $\lim_{n\to\infty}(X_n - Y_n) = 0$ in probability, where X_n and Y_n are random variables.

Theorem 3.2.2. *Under the regularity condition 1, the test statistic*

$$\sum_{j=1}^{k} \left[\frac{(S_j - N_j\theta_j^0)^2}{N_j\theta_j^0} + \frac{(N_j - S_j - N_j(1-\theta_j^0))^2}{N_j(1-\theta_j^0)} \right]$$

for the null hypothesis $H_0 : \theta_j = \theta_j^0,\ j = 1, 2, \cdots, k$ is asymptotically χ^2 distributed with k degrees of freedom.

Proof. First of all, the likelihood ratio is

$$\lambda = \frac{L(\boldsymbol{\theta}^0)}{L(\hat{\boldsymbol{\theta}})} = \frac{\prod_{j=1}^{k} (\theta_j^0)^{S_j}(1-\theta_j^0)^{N_j - S_j}}{\prod_{j=1}^{k} \hat{\theta}_j^{S_j}(1-\hat{\theta}_j)^{N_j - S_j}}$$

and so

$$-2\ln\lambda = \sum_{j=1}^{k} 2\left[S_j \ln\frac{\hat{\theta}_j}{\theta_j^0} + (N_j - S_j)\ln\frac{1-\hat{\theta}_j}{1-\theta_j^0} \right].$$

For $j = 1, 2, \cdots, k$, we have $N_j^{1/3}(S_j/N_j - \theta_j^0) = [\sqrt{n}(S_j/N_j - \theta_j^0)](N_j^{1/3}/\sqrt{n})$. By Theorem 3.1.2, $\sqrt{n}(S_j/N_j - \theta_j^0) \to N(0, (v_j(\boldsymbol{\theta})I_j(\boldsymbol{\theta}))^{-1})$. Hence $N_j^{1/3}(S_j/N_j - \theta_j^0) \to 0$

in probability. Under regularity condition 1, $N_j \to \infty$ almost surely. It follows from Theorem 12.2 in Billingsley (1961) that

$$2\left[S_j \ln \frac{\hat{\theta}_j}{\theta_j^0} + (N_j - S_j) \ln \frac{1-\hat{\theta}_j}{1-\theta_j^0}\right] \sim \left[\frac{(S_j - N_j\theta_j^0)^2}{N_j\theta_j^0} + \frac{(N_j - S_j - N_j(1-\theta_j^0))^2}{N_j\theta_j^0}\right].$$

Therefore,

$$-2\ln\lambda \sim \sum_{j=1}^{k}\left[\frac{(S_j - N_j\theta_j^0)^2}{N_j\theta_j^0} + \frac{(N_j - S_j - N_j(1-\theta_j^0))^2}{N_j\theta_j^0}\right]$$

and the result follows from corollary 3.2.1. □

The above theorem investigates the asymptotic sampling distribution of the test statistic for given values of the parameters under the null hypothesis. If we wish to see if the k treatments are equally effective, we have the following.

Theorem 3.2.3. *Under the regularity condition 1, the test statistic*

$$\sum_{j=1}^{k}\left[\frac{(S_j - N_j\hat{\theta})^2}{N_j\hat{\theta}} + \frac{(N_j - S_j - N_j(1-\hat{\theta}))^2}{N_j(1-\hat{\theta})}\right]$$

for the null hypothesis $H_0 : \theta_1 = \theta_2 = \cdots = \theta_k = \theta$ is asymptotically χ^2 distributed with $k-1$ degree of freedom, where $\hat{\theta} = S(n)/n$ is the restricted (i.e. pooled) estimator for θ.

Proof. The likelihood ratio is $\lambda = \hat{\theta}^{S(n)}(1-\hat{\theta})^{n-S(n)} / \prod_{j=1}^{k}\hat{\theta}_j^{S_j}(1-\hat{\theta}_j)^{N_j-S_j}$. Then $-2\ln\lambda$ follows asymptotically the χ^2 distribution with $k-1$ degree of freedom by Theorem 3.1.4. However,

$$-2\ln\lambda = \sum_{j=1}^{k} 2\left[S_j \ln \frac{S_j}{N_j\hat{\theta}} + (N_j - S_j)\ln\frac{N_j - S_j}{N_j(1-\hat{\theta})}\right].$$

It follows from Theorem 3.1.2 that $\sqrt{n}(S_j/N_j - \theta) \to N(0, (v_j(\boldsymbol{\theta})I_j(\boldsymbol{\theta}))^{-1})$ for $j = 1, 2, \cdots, k$. Therefore $N_j^{1/3}(S_j/N_j - \theta) \to 0$ in probability. Similarly, $(N_j)^{1/3}(\hat{\theta} - \theta) \to 0$ in probability under the null hypothesis $\theta_1 = \theta_2 = \cdots = \theta_k$. The proof is completed by following steps similar to those taken in the proof of Theorem 3.2.1. □

In conclusion, this chapter discusses both the consistency and asymptotic normality of the maximum likelihood estimators for a wide class of response adaptive designs. Under regularity conditions, the logarithm of the likelihood ratio statistic $-2\ln\lambda$ for dependent data is shown to be asymptotically chi-square distributed. Moreover, the goodness-of-fit test is also extended to the data from adaptive designs with dichotomous responses.

Chapter 4

Efficient Estimation

Many researchers have considered the maximum likelihood estimation, the bias-corrected estimation and the minimum Hellinger distance estimation for response adaptive designs. To answer the question of which estimator is the best, this chapter is to investigate the efficient estimation for response adaptive designs.

The efficiency of response adaptive randomization procedures has been studied in depth by Hu and Rosenberger (2003). The variability of allocation proportions is affected by the treatment randomization procedure. Hu and Rosenberger (2003) explicitly established the relationship between the power of the test and the variability of allocation proportions and showed that the asymptotic power is a decreasing function of the asymptotic variance of the allocation proportions. There is extensive literature on the efficiency of estimation in statistics for independent data (see, for example, Bahadur (1971), Fu (1973), Bucklew (1990)). Anscombe (1949, 1952) discussed the large-sample problem for sequential estimation, however the assumption for his results to hold are not satisfied when considering response adaptive designs.

For simplicity, we write $P_{\boldsymbol{\pi}}$ as P. Denote $P_{\boldsymbol{\theta}}$ as the probability measure under $\boldsymbol{\pi}$ with parameter $\boldsymbol{\theta}$. The induced expectation is denoted as $E_{\boldsymbol{\theta}}$.

Throughout this section, we assume that the second moment of $\ln f_j(Y, \theta_j')$ exists

and is finite under $P_{\boldsymbol{\theta}}$ for any $\boldsymbol{\theta}'$, $j = 1, 2, \cdots, k$. Furthermore, we assume

Regularity condition 2: the fourth moments of both $(\ln f_j(Y, \theta_j))'_{\theta_j}$ and $[\ln(f_j(Y, \theta_j))']'_{\theta_j}$ exist and are finite.

Under regularity condition 1 and using the martingale strong law of large numbers, the limit of $\frac{1}{n} \ln \frac{L(\boldsymbol{\theta}')}{L(\boldsymbol{\theta})}$ exists and

$$\lim_{n\to\infty} \frac{1}{n} \ln \frac{L(\boldsymbol{\theta}')}{L(\boldsymbol{\theta})} = \sum_{j=1}^{k} v_j(\boldsymbol{\theta}) E_{\boldsymbol{\theta}} \left(\frac{f_j(Y, \theta'_j)}{f_j(Y, \theta_j)} \ln \frac{f_j(Y, \theta'_j)}{f_j(Y, \theta_j)} \right)$$

almost surely in the probability measure $P_{\boldsymbol{\theta}'}$. This limit is in fact the Kullback-Leibler number $k(\boldsymbol{\theta}', \boldsymbol{\theta}) = \lim_{n\to\infty} \left\{ \frac{1}{n} \ln \frac{L(\boldsymbol{\theta}')}{L(\boldsymbol{\theta})} \right\}$ almost surely in $P_{\boldsymbol{\theta}'}$. We now establish an asymptotic equivalence of $k(\boldsymbol{\theta}', \boldsymbol{\theta})$ by means of Fisher's information number $I_j(\theta_j) = E_{\boldsymbol{\theta}} \left((\ln f_j(X, \theta_j))'_{\theta_j} \right)^2$.

Lemma 4.0.4.

$$k(\boldsymbol{\theta}', \boldsymbol{\theta}) \sim \frac{1}{2} (\boldsymbol{\theta}' - \boldsymbol{\theta})^t \Gamma(\boldsymbol{\theta}) (\boldsymbol{\theta}' - \boldsymbol{\theta}) \quad as\ \boldsymbol{\theta}' \to \boldsymbol{\theta},$$

where $\Gamma(\boldsymbol{\theta}) = diag(v_1(\boldsymbol{\theta}) I_1(\theta_1), v_2(\boldsymbol{\theta}) I_2(\theta_2), \cdots, v_k(\boldsymbol{\theta}) I_k(\theta_k))$ *is the diagonal matrix.*

Proof: Write $r_j = \frac{f_j(Y, \theta'_j)}{f_j(Y, \theta_j)}$. By a Taylor expansion, $r_j \ln r_j \sim (r_j - 1) + \frac{1}{2}(r_j - 1)^2$ as $r_j \to 1$. Hence, $E_{\boldsymbol{\theta}}(r_j \ln r_j) \sim \frac{1}{2} E_{\boldsymbol{\theta}}(r_j - 1)^2$ as $r_j \to 1$. Expanding $f_j(x, \theta'_j)$ at θ_j, we have

$$(f_j(Y, \theta'_j) - f_j(Y, \theta_j)) \sim f'_j(Y, \theta_j)_{\theta_j} (\theta'_j - \theta_j) + \frac{1}{2} f''_j(Y, \theta_j)_{\theta_j} (\theta'_j - \theta_j)^2$$

as $\theta'_j \to \theta_j$. Under the regularity condition 2,

$$E_{\boldsymbol{\theta}}(r_j - 1)^2 \sim E_{\boldsymbol{\theta}} \left(\frac{f'_j(Y, \theta_j)_{\theta_j}}{f_j(Y, \theta_j)} \right)^2 (\theta'_j - \theta_j)^2 = I_j(\theta_j)(\theta'_j - \theta_j)^2$$

as $\theta'_j \to \theta_j$. Therefore, the result follows. □

4.1 Asymptotic efficiency of the maximum likelihood estimator

In this section, we discuss the asymptotic efficiency of maximum likelihood estimators. A measurable function $\boldsymbol{T}_n = (T_1^n, T_2^n, \cdots, T_k^n)^t$ is said to be a consistent estimator of $\boldsymbol{\theta} = (\theta_1, \theta_2, \cdots, \theta_k)^t$ if, for any $\varepsilon > 0$, we have $P_{\boldsymbol{\theta}}(d(\boldsymbol{T}_n, \boldsymbol{\theta}) < \varepsilon) \to 1$ as $n \to \infty$, where d is the Euclidean distance in $\mathbb{R}^k$.

Consider testing the null hypothesis $H_o : \boldsymbol{\theta} = \boldsymbol{\theta_0}$ against the alternative hypothesis $H_a : \boldsymbol{\theta} = \boldsymbol{\theta'}$. Let β be the probability of committing type II error, n be the sample size, and

$$\alpha_n(\beta) = \inf\{\alpha_n : \alpha_n \text{ is the size of a test for testing } H_o \text{ vs } H_a \text{ with power } (1-\beta)\}.$$

The following result is an extension of the Stein's Lemma for independent data. A similar result is derived in Bahadur (1971) for independent data.

Lemma 4.1.1. *Under regularity conditions 1 and 2,* $\lim_{n\to\infty}\left\{\frac{1}{n}\ln\alpha_n(\beta)\right\} = -k(\boldsymbol{\theta'}, \boldsymbol{\theta_0})$ *for any* β.

Proof. Without loss of generality, suppose $k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0) > 0$. Set $0 < \beta < 1$. By the Neyman-Pearson Lemma, there exists a test statistic

$$\phi_n^* = \begin{cases} 1 & \text{if } r_n > c_n \\ \xi_n & \text{if } r_n = c_n \\ 0 & \text{if } r_n < c_n \end{cases}$$

such that $E_{\boldsymbol{\theta'}}(\phi_n^*) = 1-\beta$ and $E_{\boldsymbol{\theta}_0}(\phi_n^*) = \alpha_n(\beta)$, where $0 \le \xi_n \le 1$, $r_n = L(\boldsymbol{\theta'})/L(\boldsymbol{\theta}_0)$ is the likelihood ratio. We show that $\lim_{n\to\infty}\left\{\frac{1}{n}\ln c_n\right\} = k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0)$ by the method of contradiction.

Assume that $\limsup_{n\to\infty}\left\{\frac{1}{n}\ln c_n\right\} = a > k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0)$. Since $\frac{1}{n}\ln r_n \to k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0)$ almost surely in $P_{\boldsymbol{\theta'}}$ by law of large numbers, for any $\varepsilon > 0$, $P_{\boldsymbol{\theta'}}\left(\frac{1}{n}\ln r_n < k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0) + \varepsilon\right) \to 1$ as $n \to \infty$. Set $\varepsilon = a - k(\boldsymbol{\theta'}, \boldsymbol{\theta}_0)$. Then $P_{\boldsymbol{\theta'}}\left(\frac{1}{n}\ln r_n \ge a\right) \to 0$ as $n \to \infty$. This is

contradictory to $P_{\boldsymbol{\theta}'}\left(\frac{1}{n}\ln r_n \geq \frac{1}{n}\ln c_n\right) \geq 1-\beta > 0$. Hence, $\limsup_{n\to\infty}\left\{\frac{1}{n}\ln c_n\right\} \leq k(\boldsymbol{\theta}', \boldsymbol{\theta}_0)$. Similarly, $\liminf_{n\to\infty}\left\{\frac{1}{n}\ln c_n\right\} \geq k(\boldsymbol{\theta}', \boldsymbol{\theta}_0)$.

The rest of the proof follows the similar idea as that in Bahadur (1971).

On the one hand, let d_n be any positive constant such that $d_n \geq r_n$. Since ϕ_n has power $1-\beta$ and achieves type I error $\alpha_n(\beta)$, then,

$$\begin{aligned}
\alpha_n(\beta) &= E_{\boldsymbol{\theta}_0}(\phi_n) \geq \int_{r_n \leq d_n} \phi_n dP_{\boldsymbol{\theta}_0} \\
&\geq \frac{1}{d_n}\int_{r_n \leq d_n} r_n \phi_n dP_{\boldsymbol{\theta}_0} \\
&= \frac{1}{d_n}\int_{r_n \leq d_n} \phi_n dP_{\boldsymbol{\theta}'} = \frac{1}{d_n}\left(\int \phi_n dP_{\boldsymbol{\theta}'} - \int_{r_n > d_n} \phi_n dP_{\boldsymbol{\theta}'}\right) \\
&\geq \frac{1}{d_n}\left((1-\beta) - \int_{r_n > d_n} \phi_n dP_{\boldsymbol{\theta}'}\right) \geq \frac{1}{d_n}((1-\beta) - P_{\boldsymbol{\theta}'}(r_n > d_n).
\end{aligned}$$

Take $d_n = e^{nk(\boldsymbol{\theta}', \boldsymbol{\theta}_0)+n\varepsilon}$, where $\varepsilon > 0$.

$$P_{\boldsymbol{\theta}'}(r_n > d_n) = P_{\boldsymbol{\theta}'}\left(\frac{1}{n}\ln r_n > k(\boldsymbol{\theta}', \boldsymbol{\theta}_0) + \varepsilon\right) \to 0 \text{ as } n \to \infty.$$

Thus,

$$\frac{1}{n}\ln P_{\boldsymbol{\theta}_0}(r_n \geq c_n) \geq \frac{1}{n}[-(nk(\boldsymbol{\theta}', \boldsymbol{\theta}_0) + n\varepsilon)] + \frac{1}{n}\ln[(1-\beta) - P_{\boldsymbol{\theta}'}(r_n > d_n)].$$

Therefore, for any $\varepsilon > 0$,

$$\liminf_{n\to\infty}\left\{\frac{1}{n}\ln P_{\boldsymbol{\theta}_0}(r_n \geq c_n)\right\} \geq -k(\boldsymbol{\theta}', \boldsymbol{\theta}_0) - \varepsilon.$$

Hence,

$$\liminf_{n\to\infty}\left\{\frac{1}{n}\ln \alpha_n(\beta)\right\} \geq -k(\boldsymbol{\theta}', \boldsymbol{\theta}_0).$$

On the other hand,

$$\begin{aligned}
\alpha_n(\beta) &= P_{\boldsymbol{\theta}_0}(r_n \geq c_n) = \int_{r_n \geq c_n} dP_{\boldsymbol{\theta}_0} \\
&\leq \frac{1}{c_n}\int_{r_n \geq c_n} r_n dP_{\boldsymbol{\theta}_0} = \frac{1}{c_n}\int_{r_n \geq c_n} dP_{\boldsymbol{\theta}_1} \leq \frac{1}{c_n}.
\end{aligned}$$

Therefore,

$$\limsup_{n\to\infty}\left\{\frac{1}{n}\ln\alpha_n(\beta)\right\} \le -\lim_{n\to\infty}\left\{\frac{1}{n}\ln c_n\right\} = -k(\boldsymbol{\theta}', \boldsymbol{\theta_0}). \qquad \square$$

Applying Lemma 4.1.1, we now establish an asymptotic lower bound of exponential rates for consistent estimators.

Theorem 4.1.2. *Assume the regularity conditions 1 and 2 hold. For any consistent estimator* $\boldsymbol{T}_n$ *of* $\boldsymbol{\theta}_0$, *we have*

$$\liminf_{n\to\infty}\left\{\frac{1}{n}\ln P_{\boldsymbol{\theta_0}}(d(\boldsymbol{T}_n, \boldsymbol{\theta}_0) \ge \varepsilon)\right\} \ge -\inf_{\boldsymbol{\theta'}}\{k(\boldsymbol{\theta}', \boldsymbol{\theta}_0) : d(\boldsymbol{\theta}', \boldsymbol{\theta}_0) > \varepsilon\}.$$

Proof: For any given $\varepsilon > 0$, let $\boldsymbol{\theta}'$ be any point in $\boldsymbol{\Theta}$ such that $d(\boldsymbol{\theta}', \boldsymbol{\theta_0}) > \varepsilon$. Consider testing $H_o : \boldsymbol{\theta} = \boldsymbol{\theta_0}$ with the alternative hypothesis $H_a : \boldsymbol{\theta} = \boldsymbol{\theta}'$. Set $\alpha_n^* = \alpha_n(0.5)$. By Lemma 4.1.1, $\lim_{n\to\infty}\left\{\frac{1}{n}\ln\alpha_n^*\right\} = -k(\boldsymbol{\theta}', \boldsymbol{\theta_0})$. For any consistent estimator $\boldsymbol{T}_n$ of $\boldsymbol{\theta}_0$, define

$$\phi_n = \begin{cases} 1 & d(\boldsymbol{T}_n, \boldsymbol{\theta_0}) \ge \varepsilon \\ 0 & d(\boldsymbol{T}_n, \boldsymbol{\theta_0}) < \varepsilon \end{cases}$$

By the consistency of $\boldsymbol{T}_n$, $P_{\boldsymbol{\theta}'}\left(d(T_n, \boldsymbol{\theta}') < d(\boldsymbol{\theta}', \boldsymbol{\theta_0}) - \varepsilon\right) \to 1$ as $n \to \infty$. Hence, $E_{\boldsymbol{\theta}'}(\phi_n) \to 1$ as $n \to \infty$. Therefore for any $\boldsymbol{\theta}'$ such that $d(\boldsymbol{\theta}', \boldsymbol{\theta_0}) > \varepsilon$, $E_{\boldsymbol{\theta_0}}(\phi_n) \ge \alpha_n^*$. The result follows. $\square$

Theorem 4.1.2 is similar to Theorem 4.1 in Bahadur et al. (1980) and the result of Bucklew (1990, page 21). Bucklew obtained the lower bound of $\frac{1}{n}\ln P_{\boldsymbol{\theta}}(d(\boldsymbol{T}_n, \boldsymbol{\theta}) > \varepsilon)$ by the moment generating function method. However it is impossible to derive the closed form of the moment generating function for $\boldsymbol{T}_n$ in response adaptive designs because of data dependency. Bahadur et al. (1980) defined the Kullback-Leibler information number $k(\boldsymbol{\theta}', \boldsymbol{\theta})$ in a general but complicated way. We define $k(\boldsymbol{\theta}', \boldsymbol{\theta})$ in a simpler manner and derive the lower bound with a straightforward method.

Theorem 4.1.2 shows that the asymptotic lower bound of exponential rates for consistent estimators is of the type $-\inf_{\boldsymbol{\theta}'}\{k(\boldsymbol{\theta}', \boldsymbol{\theta}) : d(\boldsymbol{\theta}', \boldsymbol{\theta}) > \varepsilon\}$.

For a response adaptive design with dichotomous responses, let $S_j(n)$ be the number of successes from treatment j after n patients have been treated in the adaptive trial. Yi and Wang (2007b) showed that both $\left(\frac{1}{n}\sum_{i=1}^{n}\pi_{ij} - \frac{N_j(n)}{n}\right) \to 0$ and $\frac{S_j(n)}{N_j(n)} \to \theta_j$ almost surely as $n \to \infty, j = 1, 2, \cdots, k$. Therefore, the maximum likelihood estimator $\boldsymbol{T}_n^* = (S_1(n)/N_1(n), \cdots, S_k(n)/N_k(n))$ is a consistent estimator of $\boldsymbol{\theta}$. We show that this estimator achieves the asymptotic lower bound.

Theorem 4.1.3. *Consider a response adaptive design with dichotomous responses. If* $\frac{N_j(n)}{n} \to v_j(\boldsymbol{\theta}) \in (0,1)$ *almost surely as* $\to \infty$ *and* $v_j(\boldsymbol{\theta})$ *is continuous in* $\boldsymbol{\theta}$, $j = 1, 2, \cdots, k$, *then for very small* ε,

$$\liminf_{n\to\infty}\left\{\frac{1}{n}\ln P_{\boldsymbol{\theta}}(d(\boldsymbol{T}_n^*, \boldsymbol{\theta}) \geq \varepsilon)\right\} = -\inf_{\boldsymbol{\theta'}}\{k(\boldsymbol{\theta'}, \boldsymbol{\theta}) : d(\boldsymbol{\theta'}, \boldsymbol{\theta}) \geq \varepsilon\}$$

for the maximum likelihood estimator $\boldsymbol{T}_n^* = (S_1(n)/N_1(n), S_2(n)/N_2(n), \cdots, S_k(n)/N_k(n))$.

Proof: We only need to prove

$$\liminf_{n\to\infty}\left\{\frac{1}{n}\ln P_{\boldsymbol{\theta}}(d(\boldsymbol{T}_n^*, \boldsymbol{\theta}) \geq \varepsilon)\right\} \leq -\inf_{\boldsymbol{\theta'}}\{k(\boldsymbol{\theta'}, \boldsymbol{\theta}) : d(\boldsymbol{\theta'}, \boldsymbol{\theta}) \geq \varepsilon\}. \tag{4.1.1}$$

Note that for an adaptive trial with dichotomous responses,

$$k(\boldsymbol{\theta'}, \boldsymbol{\theta}) = \lim_{n\to\infty}\left\{\frac{1}{n}\ln\frac{L(\boldsymbol{\theta'})}{L(\boldsymbol{\theta})}\right\} = \sum_{j=1}^{k}\left[v_j(\boldsymbol{\theta'})\theta_j'\ln\frac{\theta_j'}{\theta_j} + v_j(\boldsymbol{\theta'})(1-\theta_j')\ln\frac{1-\theta_j'}{1-\theta_j}\right]$$

almost surely in $P_{\boldsymbol{\theta'}}$. Writing $T_j^n = S_j(n)/N_j(n)$, $j = 1, 2, \cdots, k$, then

$$nk(\boldsymbol{T}_n^*, \boldsymbol{\theta}) = \sum_{j=1}^{k} v_j(\boldsymbol{T}_n^*)\frac{n}{N_j(n)}\ln\left[\left(\frac{T_j^n}{\theta_j}\right)^{S_j(n)}\left(\frac{1-T_j^n}{1-\theta_j}\right)^{N_j(n)-S_j(n)}\right].$$

Since $\boldsymbol{T}_n^* \to \boldsymbol{\theta}$ almost surely as $n \to \infty$ and $v_j(\boldsymbol{\theta})$ is continuous in $\boldsymbol{\theta}$, $v_j(\boldsymbol{T}_n^*)\frac{n}{N_j(n)} \to 1$ almost surely as $n \to \infty$, $j = 1, 2, \cdots, k$. Therefore,

$$nk(\boldsymbol{T}_n^*, \boldsymbol{\theta}) = \ln\frac{\prod_{j=1}^{k}(T_j^n)^{S_j(n)}(1-T_j^n)^{N_j(n)-S_j(n)}}{\prod_{j=1}^{k}\theta_j^{S_j(n)}(1-\theta_j)^{N_j(n)-S_j(n)}} + h(\boldsymbol{T}_n^*),$$

where $h(\boldsymbol{T}_n^*) = \sum_{j=1}^k o_j(1) \ln \left[\left(\frac{T_j^n}{\theta_j} \right)^{S_j(n)} \left(\frac{1-T_j^n}{1-\theta_j} \right)^{N_j(n)-S_j(n)} \right]$ and $o_j(1) \to 0$ as $n \to \infty$, $j = 1, 2, \cdots, k$.

Let n_j and s_j be nonnegative integers, $j = 1, 2, \cdots, k$, and $\mathbf{z} = (n_1, \cdots, n_k, s_1, \cdots, s_k)$, and $\mathbf{t}_z = (s_1/n_1, \cdots, s_k/n_k)$, and $\Delta_n = \{\mathbf{z} : d(\mathbf{t}_z, \boldsymbol{\theta}) \geq \varepsilon, \sum_{j=1}^k n_j = n, 0 \leq s_j \leq n_j\}$, and Λ_n^M be the set of all sample paths leading to the observation $\mathbf{z}$ of $\mathbf{M}(n) = (N_1(n), \cdots, N_k(n), S_1(n), \cdots, S_k(n))$. Then

$$\begin{aligned} P_{\boldsymbol{\theta}}(\mathbf{M}(n) = \mathbf{z}) &= \sum_{\Lambda_n^M} \left(\prod_{i=1}^n \prod_{j=1}^k \pi_{ij}^{\delta_{ij}} (1 - \pi_{ij})^{1-\delta_{ij}} \right) \left(\prod_{j=1}^k \theta_j^{s_j} (1 - \theta_j)^{n_j - s_j} \right) \\ &= e^{-nk(\mathbf{t}_z, \boldsymbol{\theta}) + h(\mathbf{t}_z)} P_{\mathbf{t}_z}(\mathbf{M}(n) = \mathbf{z}). \end{aligned}$$

and therefore

$$P_{\boldsymbol{\theta}}(d(\boldsymbol{T}_n^*, \boldsymbol{\theta}) \geq \varepsilon) = \sum_{\mathbf{z} \in \Delta_n} P_{\boldsymbol{\theta}}(\mathbf{M}(n) = \mathbf{z}) \leq e^{-n \inf_{\Delta_n} \{k(\mathbf{t}_z, \boldsymbol{\theta})\}} \sum_{\mathbf{z} \in \Delta_n} e^{h(\mathbf{t}_z)} P_{\mathbf{t}_z}(\mathbf{M}(n) = \mathbf{z}).$$

For each fixed n, there are at most $(n+1)^{2k}$ points in Δ_n. So if we let h^* be the largest value of $h(\boldsymbol{T}_n^*)$, then

$$\frac{1}{n} \ln P(d(\boldsymbol{T}_n^*, \boldsymbol{\theta}) \geq \varepsilon) \leq -\inf_{\Delta_n} \{k(\mathbf{t}_z, \boldsymbol{\theta})\} + \frac{1}{n} h^* + \frac{(2k) \ln(n+1)}{n}.$$

Since $\frac{1}{n} h^* \to 0$ and $\frac{(2k) \ln(n+1)}{n} \to 0$ as $n \to \infty$, we have

$$\liminf_{n \to \infty} \left\{ \frac{1}{n} \ln P(d(\boldsymbol{T}_n^*, \boldsymbol{\theta}) \geq \varepsilon) \right\} \leq -\inf_{\Delta_n} \{k(\mathbf{t}_z, \boldsymbol{\theta})\} \leq -\inf_{\boldsymbol{\theta}'} \{k(\boldsymbol{\theta}', \boldsymbol{\theta}_0) : d(\boldsymbol{\theta}', \boldsymbol{\theta}_0) > \varepsilon\}.$$

So equation (4.1.1) follows.

The result follows from Theorem 4.1.2 and the format of $k(\boldsymbol{\theta}', \boldsymbol{\theta})$. □

A similar result appears as Example 5.4 in Bahadur (1971) for the multinomial distribution. Our result is an extension of Bahadur's result.

4.2 Asymptotically efficient estimation of treatment effects

Suppose that the treatment effect is described by a real-valued and differentiable function $g(\boldsymbol{\theta})$ in $\boldsymbol{\theta} \in \Theta$. For example, in a response adaptive design with dichotomous responses, $g(\boldsymbol{\theta})$ can be the difference of the success probabilities of two treatments or the odds ratio.

Let U_n be an estimator of $g(\boldsymbol{\theta})$. For any $\boldsymbol{\theta}$ and $\varepsilon > 0$, let $\tau_n(\varepsilon, \boldsymbol{\theta})$ be such that

$$P_{\boldsymbol{\theta}}(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon) = P_{\boldsymbol{\theta}}\left(|Z| \geq \varepsilon/\tau_n(\varepsilon, \boldsymbol{\theta})\right),$$

$0 \leq \tau_n(\varepsilon, \boldsymbol{\theta}) \leq \infty$, where Z follows the standard normal distribution. Obviously, U_n is consistent if and only if $\lim_{n\to\infty} \tau_n(\varepsilon, \boldsymbol{\theta}) = 0$. Such a $\tau_n(\varepsilon, \boldsymbol{\theta})$ is called the effective standard deviation of U_n given $\boldsymbol{\theta}$. Bahadur (1971) showed that if U_n is asymptotically normal with mean $g(\boldsymbol{\theta})$ and variance $u(\boldsymbol{\theta})/n$, then $\lim_{n\to\infty} n\tau_n^2(n^{-1/2}\varepsilon, \boldsymbol{\theta}) = u(\boldsymbol{\theta})$.

In this section, we derive the asymptotic lower bound of exponential rates for consistent estimators of the treatment effect, and show that this lower bound is achieved by the maximum likelihood estimator.

Theorem 4.2.1. *Assume regularity conditions 1 and 2 hold. If U_n is a consistent estimator of $g(\boldsymbol{\theta})$, then*

$$\liminf_{\varepsilon\to 0} \liminf_{n\to\infty} \left\{ \frac{1}{n\varepsilon^2} \ln P_{\boldsymbol{\theta}}(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon) \right\} \geq -\frac{1}{2w(\boldsymbol{\theta})},$$

where $w(\boldsymbol{\theta}) = \sum_{j=1}^{k} [v_j(\boldsymbol{\theta}) I_j(\theta_j)]^{-1} [(g(\boldsymbol{\theta}))'_{\theta_j}]^2$.

Proof. The proof follows an idea similar to that in Theorem 6.1 of Bahadur (1971). Without loss of generality, assume $w(\boldsymbol{\theta}) > 0$. For any $\boldsymbol{\theta}_0 \in \Theta$, set

$$\boldsymbol{\theta}^* = \boldsymbol{\theta}_0 + \varepsilon \Gamma^{-1}(\boldsymbol{\theta}_0)(g(\boldsymbol{\theta})'_{\theta_1}, (g(\boldsymbol{\theta})'_{\theta_2}), \cdots, (g(\boldsymbol{\theta})'_{\theta_k}))^t_{\boldsymbol{\theta}_0}.$$

For any consistent estimator U_n, we construct a test statistic ϕ_n for testing $H_o : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ with the alternative hypothesis $H_a : \boldsymbol{\theta} = \boldsymbol{\theta}^*$. Reject H_o under ϕ_n if $|U_n - g(\boldsymbol{\theta}_0)| \geq \lambda\varepsilon w(\boldsymbol{\theta}_0)$, where $0 < \lambda < 1$. Note that $g(\boldsymbol{\theta}^*) - g(\boldsymbol{\theta}_0) \sim \varepsilon w(\boldsymbol{\theta}_0)$ and $K(\boldsymbol{\theta}^*, \boldsymbol{\theta}_0) \sim \frac{1}{2}\varepsilon^2 w(\boldsymbol{\theta}_0)$ as $\varepsilon \to 0$. For fixed type II error β, $E_{\boldsymbol{\theta}^*}(\phi_n) \to 1 > 1 - \beta$ as $n \to \infty$ by the consistency of U_n. Therefore $E_{\boldsymbol{\theta}_0}(\phi_n) \geq \alpha_n(\beta)$. The proof is completed after applying Lemma 4.1.1 and setting $\lambda \to 1$. $\square$

An estimator U_n is said to be asymptotically efficient in the Bahadur sense if

$$\liminf_{\varepsilon \to 0} \liminf_{n \to \infty} \left\{ \frac{1}{n\varepsilon^2} \ln P_{\boldsymbol{\theta}}(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon) \right\} = -\frac{1}{2w(\boldsymbol{\theta})}.$$

The following result states that the maximum likelihood estimator of $g(\boldsymbol{\theta})$ is asymptotically efficient in the Bahadur sense.

Theorem 4.2.1 shows that $e^{-n\varepsilon^2/[2w(\boldsymbol{\theta})]}$ is the fastest rate of $P_{\boldsymbol{\theta}}(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon)$ converging to 0 exponentially. Bahadur (1971) derived a similar result for independent data, but our result applies to the case of dependent data from response adaptive designs.

For the standard normal distribution, $\left(\frac{1}{t} - \frac{1}{t^3}\right) \frac{1}{\sqrt{2\pi}} e^{-t^2/2} < P(Z \geq t) < \frac{1}{t} \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$. It can be shown that $\ln P(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon) \sim -\varepsilon^2/[2\tau_n^2(\varepsilon, \boldsymbol{\theta})]$ as $n \to \infty$. Applying Theorem 4.2.1, we have the following extension of Theorem 6.1 in Bahadur (1971).

Corollary 4.2.2. *Assume regularity conditions 1 and 2 hold. If U_n is a consistent estimator of $g(\boldsymbol{\theta})$, then*

$$\liminf_{\varepsilon \to 0} \liminf_{n \to \infty} \{n\tau_n^2(\varepsilon, \boldsymbol{\theta})\} \geq w(\boldsymbol{\theta}),$$

where $w(\boldsymbol{\theta}) = \sum_{j=1}^{k} [v_j(\boldsymbol{\theta}) I_j(\theta_j)]^{-1} [(g(\boldsymbol{\theta}))'_{\theta_j}]^2$.

Corollary 4.2.3. *Assume regularity conditions 1 and 2 hold. If the maximum likelihood estimator $\hat{\boldsymbol{\theta}}_n$ of $\boldsymbol{\theta}$ exists, then*

$$\liminf_{\varepsilon \to 0} \liminf_{n \to \infty} \left\{ \frac{1}{n\varepsilon^2} \ln P_{\boldsymbol{\pi}}(|g(\hat{\boldsymbol{\theta}}_n) - g(\boldsymbol{\theta})| \geq \varepsilon | \boldsymbol{\theta}) \right\} = -\frac{1}{2w(\boldsymbol{\theta})}.$$

The proof follows from the asymptotic variance $w(\boldsymbol{\theta}) = \sum_{j=1}^{k} [v_j(\boldsymbol{\theta}) I_j(\theta_j)]^{-1} [(g(\boldsymbol{\theta}))'_{\theta_j}]^2 / n$ of $g(\hat{\boldsymbol{\theta}}_n)$ and $\ln P(|U_n - g(\boldsymbol{\theta})| \geq \varepsilon) \sim -\varepsilon^2 / [2\tau_n^2(\varepsilon, \boldsymbol{\theta})]$ as $n \to \infty$ and $\varepsilon \to 0$.

In summary, this chapter examines the issue of asymptotic efficiency of estimation for response adaptive designs of clinical trials, from which the collected data set contains a dependency structure. The asymptotic lower bound of exponential rates for consistent estimators is established. Under certain regularity conditions, the maximum likelihood estimator is shown to achieve the asymptotic lower bound for response adaptive trials with dichotomous responses. Furthermore, it is proven that the maximum likelihood estimator of the treatment effect is asymptotically efficient in the Bahadur sense for response adaptive clinical trials.

Chapter 5

Response Adaptive Designs with a Variance-Penalized Criterion

A good response adaptive design is expected to assign more patients to the better treatment with minimal loss in the power of the statistical test. Evaluating the quality of a response adaptive design from multiple objectives is difficult because of the dependency in the data and the involvement of the unknown parameters in the design. This chapter discusses the evaluation of response adaptive designs with a variance-penalized criterion.

5.1 The variance-penalized criterion

Suppose that patients are recruited sequentially into a clinical trial and are treated with one and only one of two treatments A and B. Suppose that the responses from treatment k, $k = A,\ B$, are independent and follow a distribution $f_k(x, \boldsymbol{\theta}_k)$, where the unknown parameter vector $\boldsymbol{\theta}_k$ may consist of parameters such as mean μ and standard deviation σ of the distribution. We assume that the larger the mean response, the better the treatment. A design is good if it assigns as many patients as possible to the better treatment and achieves this goal with less variability.

Let δ_n be the treatment allocated to the n^{th} patient such that $\delta_n = 1$ for treat-

ment A, $\delta_n = 0$ for treatment B, and Y_n be the response following the distribution $f_k(x, \boldsymbol{\theta}_k)$ if the patient is assigned to treatment $k, k = A, B$. In an adaptive design, the treatment allocation δ_n for the n^{th} patient depends on the accumulated information $\{(\delta_1, y_1), \cdots, (\delta_{n-1}, y_{n-1})\}, n \geq 2$. The response adaptive design is specified by its randomized allocation rule.

Under the variance-penalized criterion, the objective function is the expected total responses minus a positive multiple of its variance. The objective is to maximize

$$\left\{ E\left(\sum_{i=1}^{n} Y_i\right) - \lambda Var\left(\sum_{i=1}^{n} Y_i\right)\right\},$$

which is called the variance-penalized mean, where $\lambda > 0$ is the penalty parameter.

Deriving the optimal design is very difficult in both template and computation of the optimal solution for general responses. However, if the responses are binary, the information gathering process can be formulated as a Markov process (Yi and Wang, 2007b), and the stochastic optimality problem of response adaptive designs becomes a variance-penalized Markov decision process. In this situation, there are algorithms to compute the optimal allocation rule π (see White, 1992 and Collins, 1997). However, the optimal allocation rule is deterministic and vulnerable to selection bias in a clinical trial.

Let $N_k = N_k(n)$ be the number of patients allocated to treatment k after n patients have been treated in the trial, $k = A, B$. Denote the responses of the N_k patients as $X_{1k}, \cdots, X_{N_k k}$. Then $\sum_{i=1}^{N_A} X_{iA} = \sum_{i=1}^{n} Y_i \delta_i$ and $\sum_{i=1}^{N_B} X_{iB} = \sum_{i=1}^{n} Y_i(1 - \delta_i)$. So $\sum_{i=1}^{n} Y_i = \sum_{i=1}^{N_A} X_{iA} + \sum_{i=1}^{N_B} X_{iB}$.

It can be shown that

$$E\left(\sum_{i=1}^{n} Y_i\right) = (\mu_A - \mu_B)E(N_A) + n\mu_B, \tag{5.1.1}$$

and

$$Var\left(\sum_{i=1}^{n} Y_i\right) = (\sigma_A^2 - \sigma_B^2)E(N_A) + n\sigma_B^2 + (\mu_A - \mu_B)^2 Var(N_A). \qquad (5.1.2)$$

where $\mu_A = E(X_{iA})$, $\mu_B = E(X_{iB})$, $\sigma_A = Var(X_{iA})$ and $\sigma_B = Var(X_{iB})$. Therefore, the optimality problem of response adaptive designs is equivalent to achieving

$$\max_{\pi} \left\{ [\mu_A - \mu_B - \lambda(\sigma_A^2 - \sigma_B^2)]E(N_A) - \lambda(\mu_A - \mu_B)^2 Var(N_A) \right\}.$$

A treatment is said to be better than another if its variance-penalized mean of the responses is larger. The first term of the criterion demonstrates that the design with a larger value of the objective function is expected to assign more patients to the better treatment. The second term indicates that for a fixed total number of patients n, the design with a larger value of the objective function has a smaller variance of the treatment allocation proportion, hence yields more power when conducting Wald's test according to the results of Hu and Rosenberger (2003). Therefore, the variance-penalized criterion prefers a design that allocates as many patients as possible to the better treatment while keeping Wald's test more powerful.

If $\sigma_A = \sigma_B$, then for a small value of λ, more weight is put on the mean number of patients assigned to the better treatment. A large value of λ emphasizes the power to draw a statistical conclusion. Therefore, the parameter λ is a measure of tradeoff between the individual and collective ethics.

In an adaptive design with binary responses, the response is 1 for a success and 0 for a failure. Our objective is to maximize the expected total number of successes with less variability. This is equivalent to achieving

$$\max_{\pi} \left\{ (p_A - p_B)[1 - \lambda(1 - p_A - p_B)]E(N_A) - \lambda(p_A - p_B)^2 Var(N_A) \right\},$$

where p_A and p_B are success probabilities of treatment A and B respectively. Without loss of generality, assume $p_A > p_B$. If $p_A + p_B > 1 - \frac{1}{\lambda}$, then $(1 - \lambda(1 - p_A - p_B)) > 0$.

When $0 < \lambda \leq 1$, $p_A + p_B > 1 - \frac{1}{\lambda}$ is always true. In this case, the treatment with a higher success probability has a larger variance-penalized mean. Under the variance-penalized criterion, we prefer the design that assigns more patients to the highly successful treatment with less variability of the allocation proportion. In the remaining part of this chapter, we choose $0 < \lambda \leq 1$.

5.2 Randomization procedures

The randomization procedures in typical urn models such as the RPW (Wei and Durham, 1978) and the DL (Ivanova, 2003) update the urn compositions sequentially based on the responses of previously treated patients. The urn will eventually contain a high proportion of balls representing the more successful treatment. The sequential randomization procedures are flexible to target any proportion of treatment allocation. This kind of procedure includes the DBCD (Eisele, 1994, Hu and Zhang, 2004) and the randomization procedure considered by Melfi et al.(2001) and Rosenberger et al. (2001).

We prefer a design that assigns more patients to the better treatment with less variability under the variance-penalized criterion. Our design is to target the treatment allocation proportion

$$\rho = \frac{q_B + \epsilon \min\{q_A, q_B\} \text{sign}(q_B - q_A)}{q_A + q_B},$$

where $q_A = 1 - p_A$, $q_B = 1 - p_A$ and $0 \leq \epsilon \leq 1$. If $\epsilon = 0$, this proportion reduces to the one in Ivanova's DL (2003). If $\epsilon > 0$, a higher proportion of patients is expected to be assigned to the better treatment. From the individual ethics point of view, we should choose ϵ as large as possible. However, when ϵ approaches 1, it is possible that only a few number of patients are allocated to the inferior treatment. This will make it inadequate to draw a statistical conclusion about the treatment effect. Intuitively,

ϵ is a kind of measure of tradeoff between individual and collective ethics.

We wish to achieve the target proportion ρ using DBCD with the allocation function

$$\begin{aligned} g(x,\rho) &= \frac{\rho(\rho/x)^\gamma}{\rho(\rho/x)^\gamma + (1-\rho)((1-\rho)/(1-x))^\gamma}, \\ g(0,\rho) &= 1, \\ g(1,\rho) &= 0. \end{aligned}$$

This allocation function was proposed by Hu and Rosenberger (2003), and Hu and Zhang (2004). Using Theorem 2.1 of Hu and Zhang (2004), the allocation proportion in our randomization procedure is asymptotically normal. That is,

$$\sqrt{n}\left(\frac{N_A}{n} - \rho\right) \to N(0, \tau^2),$$

where for $p_A \geq p_B$,

$$\begin{aligned} \tau^2 &= \left\{ \frac{(1-\epsilon)q_A q_B[(1-\epsilon)p_A q_B + p_B(q_B+\epsilon q_A)]}{(q_B+q_A)^3(q_B+\epsilon q_A)} \right. \\ &\quad \left. + \frac{1}{1+2\gamma}\left[\frac{(1-\epsilon)q_A(q_B+\epsilon q_A)}{(q_B+q_A)^2} + \frac{(1-\epsilon)q_A q_B[(1-\epsilon)p_A q_B + p_B(q_B+\epsilon q_A)]}{(q_B+q_A)^3(q_B+\epsilon q_A)}\right]\right\}, \end{aligned}$$

and for $p_A < p_B$,

$$\begin{aligned} \tau^2 &= \left\{ \frac{(1-\epsilon)q_A q_B[(1-\epsilon)p_B q_A + p_A(q_A+\epsilon q_B)]}{(q_B+q_A)^3(q_A+\epsilon q_B)} \right. \\ &\quad \left. + \frac{1}{1+2\gamma}\left[\frac{(1-\epsilon)q_B(q_A+\epsilon q_B)}{(q_B+q_A)^2} + \frac{(1-\epsilon)q_A q_B[(1-\epsilon)p_B q_A + p_A(q_A+\epsilon q_B)]}{(q_B+q_A)^3(q_A+\epsilon q_B)}\right]\right\}. \end{aligned}$$

It can be shown that if $p_A + p_B > 1-\epsilon$ and $\gamma \to \infty$, then

$$\tau^2 \leq q_A q_B(2-(q_A+q_B))/(q_A+q_B)^3, \tag{5.2.1}$$

where $q_A q_B(2-(q_A+q_B))/(q_A+q_B)^3$ is the asymptotic variance of the DL design. Hu et al. (2006) claimed that Ivanova's DL is the asymptotically best design in that

it attains the lower bound on the asymptotic variances of the allocation proportions among all the designs with the target proportion $q_B/(q_B + q_A)$ when the allocation proportions were asymptotically normal. The target proportion of our proposed design is $\rho = (q_B + \epsilon \min\{q_A, q_B\}\text{sign}(q_B - q_A))/(q_A + q_B)$, which is larger than $q_B/(q_B + q_A)$ for $p_A > p_B$ and smaller than $q_B/(q_B + q_A)$ for $p_A < p_B$. Therefore, the proposed design asymptotically assigns a higher proportion of patients to the better treatment with less variability than the DL design.

5.3 A comparison of designs

The variance-penalized criterion prefers the design that assigns more patients to the better treatment with less variability. Adaptive designs can be compared and evaluated under this criterion, even if the designs target different proportions. We compare our proposed design with some existing designs such as the RPW, DL, the design targeting the optimal allocation proportion (denoted as RSIHR, Rosenberger et al. (2001)) and the design proposed by Bandyopadhay and Bhattacharya (2006) (denoted as BB).

We use the RPW with an initial structure of 5 balls in the urn representing treatment A and 5 balls for treatment B, as recommended by Rosenberger (1999). The DL (Ivanova (2003)) is the design with 3 balls representing treatment A and B respectively and 1 immigration ball. The RSIHR design considers both the expected number of failures and the power of the test. The target proportion of subjects receiving treatment A in the RSIHR design is $\rho = \sqrt{p_A}/(\sqrt{p_A} + \sqrt{p_B})$. The allocation function $g(x, \rho)$ with $\gamma = 100$ is used to target the RSIHR proportion and our proposed proportion. The BB design is the combination of the Neyman allocation and the myopically better treatment. The sequential maximum likelihood procedure is employed to target the

BB proportion

$$\rho = \max\left(\frac{\sqrt{p_A q_A}}{\sqrt{p_A q_A} + \sqrt{p_B q_B}}, \zeta_A\right) I_{[p_A \geq p_B]} + \min\left(\frac{\sqrt{p_A q_A}}{\sqrt{p_A q_A} + \sqrt{p_B q_B}}, \zeta_A\right) I_{[p_A > p_B]},$$

where $\zeta_A = \Phi\left(\frac{p_A - p_B}{\sqrt{p_A q_A + p_B q_B}}\right)$ and Φ is the cumulative standard normal distribution function.

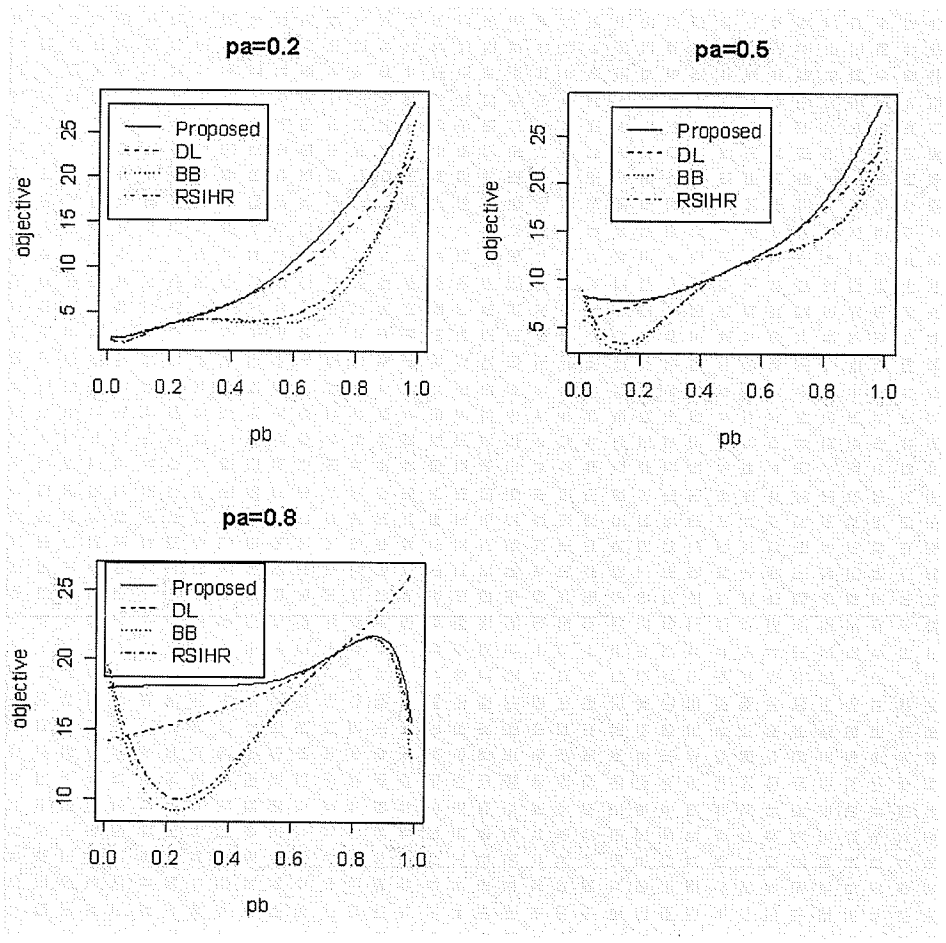

Figure 5.1: Comparison of the objective functions with $\lambda = 1/2$

DL: the drop the loss design; RSIHR:the optimal design proposed by Rosenberger et al.;BB: the design proposed by Bandyopadhyay and Bhattacharya; Proposed: $\epsilon = 1/2$.

Figure 5.1 gives a comparison of the objective functions of different designs for a total number of patients $n = 30$. The RPW is omitted because of its similar behavior

as the DL. From the graphs we see that if the success probability p_A of the standard treatment is less than 0.8, then our proposed design ($\epsilon = 1/2$) generates a higher expected number of successes with less variability than all other designs, except for extremely small values of p_B. But when both the success probabilities p_A and p_B are large, say both are larger than 0.8, then the DL is better than other designs for clinical trials with a small number of patients.

Tables 5.1, 5.2 and 5.3 describe the simulation results for a total number of patients of $n = 100$ with 10000 replications. The proposed designs I and II are defined with $\epsilon = 1/2$ and $\epsilon = 1/4$ respectively. Table 5.1 gives the simulated values of the objective function ($\lambda = 1/2$) of the DL, RSIHR, BB and our proposed designs. It seems that the proposed designs have objective values larger than all the others when p_A and p_B differ.

p_A	p_B	DL	RSIHR	BB	Proposed I	Proposed II
0.9	0.9	85.565	85.429	85.557	85.353	85.405
	0.7	76.529	61.987	63.823	79.729	78.863
	0.5	72.012	14.503	24.388	79.065	76.284
	0.3	69.260	-21.395	-4.249	78.250	74.334
0.7	0.7	59.114	59.543	59.560	59.361	59.730
	0.5	49.679	24.584	24.262	51.034	51.115
	0.3	42.928	-41.346	-42.617	50.716	47.233
	0.1	37.163	-33.771	-48.818	48.698	43.104
0.5	0.5	37.419	37.542	37.542	37.848	37.558
	0.3	28.600	-3.089	-4.170	29.507	29.942
	0.1	21.754	-43.147	-45.749	29.067	25.415

Table 5.1: Simulation of the variance-penalized mean

DL: the drop the loss design; RSIHR:the optimal design proposed by Rosenberger et al.;BB: the design proposed by Bandyopadhyay and Bhattacharya; Proposed I: $\epsilon = 1/2$; Proposed II: $\epsilon = 1/4$.

Without loss of generality, we assume that the treatment A is better than B. Table 5.2 gives the simulation results for the expected proportion of patients assigned to treatment A, together with the standard deviation of N_A/n given in parentheses.

Our proposed designs allocate higher expected proportions of patients to the superior treatment than the DL, although the standard deviations are a little bit large when p_A and p_B are close to each other. As the difference between p_A and p_B increases, the standard deviations of our proposed designs become smaller. Except for very small values of p_B such as 0.1, our proposed designs assign higher proportions of patients to the better treatment with much smaller standard deviations than both the RSIHR and BB designs.

p_A	p_B	DL	RSIHR	BB	Proposed I	Proposed II
0.9	0.9	0.499(0.065)	0.498(0.217)	0.502(0.220)	0.497(0.285)	0.500(0.210)
	0.7	0.639(0.062)	0.620(0.275)	0.663(0.623)	0.849(0.099)	0.803(0.08)
	0.5	0.730(0.045)	0.737(0.279)	0.786(0.257)	0.908(0.039)	0.869(0.045)
	0.3	0.788(0.036)	0.851(0.239)	0.881(0.216)	0.932(0.023)	0.902(0.031)
0.7	0.7	0.500(0.063)	0.496(0.349)	0.498(0.347)	0.496(0.256)	0.497(0.164)
	0.5	0.605(0.053)	0.630(0.374)	0.631(0.369)	0.772(0.126)	0.707(0.076)
	0.3	0.677(0.042)	0.771(0.342)	0.773(0.342)	0.844(0.035)	0.772(0.036)
	0.1	0.728(0.033)	0.918(0.223)	0.906(0.241)	0.872(0.018)	0.810(0.026)
0.5	0.5	0.500(0.048)	0.495(0.424)	0.502(0.422)	0.498(0.244)	0.502(0.146)
	0.3	0.577(0.040)	0.660(0.418)	0.653(0.419)	0.751(0.122)	0.676(0.062)
	0.1	0.635(0.031)	0.858(0.310)	0.849(0.315)	0.818(0.023)	0.731(0.023)

Table 5.2: Simulation of the expected proportion (standard deviation)
DL: the drop the loss design; RSIHR:the optimal design proposed by Rosenberger et al.;BB: the design proposed by Bandyopadhyay and Bhattacharya; Proposed I: $\epsilon = 1/2$; Proposed II: $\epsilon = 1/4$.

Table 5.3 gives the simulated power of the Wald's test for the data collected from different designs. The statistical powers of our proposed designs are higher than those of both RSIHR and BB designs. If the difference between p_A and p_B is small, say 0.2, the statistical powers of our proposed designs are smaller than that of the DL. But as the difference increases, the powers of our proposed designs increase and the power of the proposed design II is close to that of the DL design. The table entries of 1 are rounded values.

From Tables 5.2 and 5.3, it is clear that the parameter ϵ in our proposed design

measures a tradeoff between individual and collective ethics. The design with a large value of ϵ assigns a high proportion of patients to the better treatment and also results in more loss of the statistical power than a design with a small value of ϵ. Because of this compromise, no significant difference exists in our two proposed designs in terms of the variance-penalized objective function (Table 5.1).

p_A	p_B	DL	RSIHR	BB	Proposed I	Proposed II
0.9	0.9	0.049	0.060	0.056	0.050	0.050
	0.7	0.737	0.458	0.448	0.468	0.606
	0.5	0.997	0.560	0.578	0.900	0.976
	0.3	1	0.438	0.464	0.995	1
0.7	0.7	0.050	0.051	0.050	0.050	0.050
	0.5	0.524	0.224	0.228	0.361	0.471
	0.3	0.983	0.290	0.305	0.899	0.964
	0.1	1	0.220	0.249	1	1
0.5	0.5	0.050	0.049	0.051	0.049	0.050
	0.3	0.526	0.144	0.144	0.385	0.480
	0.1	0.996	0.175	0.206	0.972	0.992

Table 5.3: Simulation of statistical power

DL: the drop the loss design; RSIHR:the optimal design proposed by Rosenberger et al.;BB: the design proposed by Bandyopadhyay and Bhattacharya; Proposed I: $\epsilon = 1/2$; Proposed II: $\epsilon = 1/4$.

In the zidovudine trial (Connor et al (1994)), the HIV infection rates for the AZT group and the placebo group were observed to be $p_A = 0.916$ and $p_B = 0.748$ respectively for a total number of patients $n = 477$. Yao and Wei (1996) and Ivanova (2003) redesigned the trial using the RPW and DL designs. In Ivanova's study, the DL design was compared with the RPW and the other two designs, one targeting the optimal allocation that maximizes the power while testing treatment difference against 0, another targeting the optimal allocation that maximizes the power while testing the log odds ratio against 0. Ivanova showed that the DL design potentially assigns more patients to the AZT treatment than the two designs with optimal proportions, and is less variable than the RPW. In this chapter, we compare our proposed design with

the DL, RPW, RSIHR and BB designs assuming instantaneous responses. Table 5.4 compares various designs based on the simulation results for the expected proportion allocated to the AZT treatment (treatment A), the expected total number of HIV free infants and tail probabilities. In Table 5.4, S denotes the total number of HIV free infants in the trial and $\epsilon = 1/2$ is given in our proposed design.

The overall performance of our proposed design is very good. Table 5.4 shows that the proportion of patients assigned to the AZT treatment using our proposed design is at least 16% larger than the proportions from the other designs, and results in a higher proportion of success with less variability. The probability that our proposed design assigns more than 80% patients to the better treatment is 0.99, which is much higher than the other designs. The upper and lower tail probabilities (with cutoff points of 0.95 and 0.05) of our design are almost the same as those of the RPW and DL designs. That means that it does not lead to extremely unbalanced treatment groups. But both the RSIHR and BB designs do so with more than 25% and 21% probabilities respectively. Furthermore, both the RSIHR and BB designs allocate higher proportions of patients (more than 95%) to the inferior treatment with about 6% probability. Our proposed design as well as the RPW and DL designs do not have this problem.

Compared with other adaptive designs, our proposed design has the potential to assign more patients to the AZT treatment with less variability, and does not result in extremely unbalanced groups. According to Hu and Rosenberger's (2003) result, a higher proportion of AZT treatment allocation does not necessarily result in a loss of power when the variance of the treatment allocation proportion is small. Moreover the sample size of the placebo group is still reasonably large. Furthermore if needed, we can decrease the value of ϵ to achieve a desired power.

In summary, this chapter discusses response adaptive designs with a variance-

	DL	RPW	Proposed	RSIHR	BB
$E(N_A/n)$	0.701	0.653	0.873	0.609	0.595
S.D of N_A/n	0.038	0.081	0.028	0.262	0.250
$E(S/n)$	0.864	0.859	0.894	0.850	0.848
S.D of S/n	0.016	0.020	0.015	0.046	0.045
$p(N_A/n \geq 0.80)$	0.003	0.027	0.991	0.253	0.221
$p(N_A/n \geq 0.95)$	0	0	0	0.253	0.213
$p(N_A/n \leq 0.05)$	0	0	0	0.069	0.064

Table 5.4: Comparison of alternative designs for the zidovudine trial
DL: the drop the loss design; RPW: the played-the-winner design; RSIHR:the optimal design proposed by Rosenberger et al.;BB: the design proposed by Bandyopadhyay and Bhattacharya; Proposed: $\epsilon = 1/2$.

penalized criterion. A new design is proposed and compared with other existing designs according to the criterion.

We have investigated the properties of the variance-penalized criterion. This criterion evaluates the performance of a design according to both the mean number of patients assigned to the better treatment and the power of the statistical test for the data collected from the design. The variance-penalized criterion prefers the design that allocates more patients to the better treatment and at the same time keeps the statistical power at a high level.

In this chapter, we have also proposed a new proportion of treatment allocation and used the DBCD to target the proportion. The asymptotic variance of the allocation proportion of our proposed design is shown to be smaller than that of the DL if $p_A + p_B > 1 - \epsilon$ and $\gamma \to \infty$. Simulation results suggest that our proposed design is better than other existing designs under the variance-penalized criterion, except for extreme values of the probabilities of success (such as very large p_A and p_B or extremely small p_B). Potentially, our proposed design assigns a higher proportion of patients to the better treatment and the power of the statistical test remains competitive when the difference between p_A and p_B is not small (say, larger than or equal to 0.4). For

a fixed large number of patients, the overall performance of our design is better than the existing designs.

In brief, the variance-penalized criterion considers both the number of patients assigned to the better treatment and the power of the statistical test. Our proposed design has good performance in common clinical situations.

Chapter 6

Conclusion and Further Research

A traditional randomized clinical trial provides a powerful tool for the comparison of treatment effects. The balanced randomization is often regarded as a gold standard for clinical research. However, such a randomized design becomes ethically infeasible in desperate medical situations because individual ethics is often sacrificed. For a more ethical balance between individual and collective ethics, alternative designs such as response adaptive designs have been proposed and employed in some clinical trials. A response adaptive design adapts the treatment allocation based on accumulated information of the treatment effect to assign more patients to the potentially better treatment. However, the adaptation of the treatment allocation creates a dependence structure in the collected data and raises concerns about the validity of conventional statistical inference, the loss of power of testing hypotheses, experimental bias. My dissertation focused on the statistical inference and the optimality of response adaptive designs.

The issue of statistical inference for response adaptive clinical trials has been both important and challenging. Due to the adaptation of the treatment allocation, data collected from response adaptive designs are dependent, and hence traditional statistical inference assuming independent observations is not applicable without modifica-

tion. I studied the treatment randomization processes of response adaptive clinical trials. The information gathering process in the trial has been formulated as a stochastic process, especially a Markov process for dichotomous responses. Then the logarithm of the likelihood ratio test and goodness-of-fit test were extended to dependent data from adaptive trials. I also examined the issue of asymptotic efficiency of estimation in response adaptive designs of clinical trials. An asymptotic lower bound of exponential rates for consistent estimators was established and the maximum likelihood estimator of the treatment effect was shown to be asymptotically efficient in the Bahadur sense for response adaptive clinical trials.

In addition, I investigated the optimality properties of the designs. How to select treatment allocation procedures is a commonly difficult issue in adaptive designs. On the one hand, a design is expected to assign patients to the better treatment as many as possible. On the other hand, unbalanced treatment groups can reduce the power of statistical test and hence result in failure of drawing a valid conclusion at the end of the trial. Response adaptive designs have the advantage of balancing the individual ethics and collective ethics. But different adaptive designs present very different tradeoffs between individual ethics and collective ethics. This thesis proposed to use a variance-penalized criterion for the evaluation of response adaptive designs. It was shown that this criterion accesses the performance of a response adaptive design based on both the expected number of patients assigned to the better treatment and the power of the statistical test. A new proportion of treatment allocation was proposed and simulation studies were conducted to compare the proposed design with some existing designs. The proposed design has the advantage of assigning more patients to the potentially better treatment with lower loss of power for testing hypotheses in common clinical trial conditions.

However, the optimal treatment allocation under the variance-penalized criterion

is deterministic, which is vulnerable to selection bias in clinical trials. Searching for an optimal randomization allocation is still under study. Constrained dynamic programming techniques will be employed and algorithms will be developed to search for an optimal solution in my further research. Moreover, although the log-likelihood ratio test and goodness-of-fit test are extended to data from response adaptive clinical trials, the power performances of the two statistical tests are not very good. The most powerful test statistic for response adaptive designs is waiting for further exploration.

Appendix

Proofs of equations and inequality

Proof of equation (5.1.1):

Since $\sum_{i=1}^{n} Y_i = \sum_{i=1}^{N_A} X_{iA} + \sum_{i=1}^{N_B} X_{iB}$ and $N_B = n - N_A$,

$$\begin{aligned} E\left(\sum_{i=1}^{n} Y_i\right) &= E\left\{E\left(\sum_{i=1}^{N_A} X_{iA} + \sum_{i=1}^{N_B} X_{iB}\right) \middle| N_A\right\} \\ &= E\{N_A E(X_{iA}\} + E\{N_B E(X_{iB}\} = (\mu_A - \mu_B)E(N_A) + n\mu_B. \end{aligned}$$

Proof of equation (5.1.2):

$$\begin{aligned} Var\left(\sum_{i=1}^{n} Y_i\right) &= E\left[\left(\sum_{i=1}^{N_A} X_{iA} + \sum_{i=1}^{N_B} X_{iB}\right) - E\left(\sum_{i=1}^{N_A} X_{iA} + \sum_{i=1}^{N_B} X_{iB}\right)\right]^2 \\ &= E\left[\left(\sum_{i=1}^{N_A} X_{iA} - E\left(\sum_{i=1}^{N_A} X_{iA}\right)\right) + \left(\sum_{i=1}^{N_B} X_{iB} - E\left(\sum_{i=1}^{N_B} X_{iB}\right)\right)\right]^2 \\ &= E\left[\left(\sum_{i=1}^{N_A} X_{iA} - E\left(\sum_{i=1}^{N_A} X_{iA}\right)\right)^2 + \left(\sum_{i=1}^{N_B} X_{iB} - E\left(\sum_{i=1}^{N_B} X_{iB}\right)\right)^2\right. \\ &\quad \left. +2\left(\sum_{i=1}^{N_A} X_{iA} - E\left(\sum_{i=1}^{N_A} X_{iA}\right)\right)\left(\sum_{i=1}^{N_B} X_{iB} - E\left(\sum_{i=1}^{N_B} X_{iB}\right)\right)\right] \\ &= Var\left(\sum_{i=1}^{N_A} X_{iA}\right) + Var\left(\sum_{i=1}^{N_B} X_{iB}\right) + 2Cov\left(\sum_{i=1}^{N_A} X_{iA}, \sum_{i=1}^{N_B} X_{iB}\right) \end{aligned}$$

Since

$$Var\left(\sum_{i=1}^{N_A} X_{iA}\right) = E(N_A)\sigma_A^2 + \mu_A^2 Var(N_A),$$

$$Var\left(\sum_{i=1}^{N_B} X_{iB}\right) = E(N_B)\sigma_B^2 + \mu_B^2 Var(N_B),$$

and

$$
\begin{aligned}
& Cov\left(\sum_{i=1}^{N_A} X_{iA}, \sum_{i=1}^{N_B} X_{iB}\right) \\
= \; & E\left\{Cov\left(\sum_{i=1}^{N_A} X_{iA}, \sum_{i=1}^{N_B} X_{iB}\middle| N_A\right)\right\} + Cov\left(E\left(\sum_{i=1}^{N_A} X_{iA}\middle| N_A\right), E\left(\sum_{i=1}^{N_B} X_{iB}\middle| N_A\right)\right) \\
= \; & Cov(N_A E(X_{iA}), N_B E(X_{iB})) \\
= \; & \mu_A \mu_B Cov(N_A, N_B) \\
= \; & \mu_A \mu_B Cov(N_A, n - N_B) \\
= \; & -\mu_A \mu_B Var(N_A)
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
& Var\left(\sum_{i=1}^{n} Y_i\right) \\
& = E(N_A)\sigma_A^2 + \mu_A^2 Var(N_A) + E(N_B)\sigma_B^2 + \mu_B^2 Var(N_B) - 2\mu_A\mu_B Var(N_A) \\
& = E(N_A)\sigma_A^2 + \mu_A^2 Var(N_A) + E(n - N_A)\sigma_B^2 + \mu_B^2 Var(n - N_A) - 2\mu_A\mu_B Var(N_A) \\
& = (\sigma_A^2 - \sigma_B^2)E(N_A) + n\sigma_B^2 + (\mu_A - \mu_B)^2 Var(N_A).
\end{aligned}
$$

Proof of asymptotic variance τ^2 on page 52:

We show that for $p_A \geq p_B$,

$$\begin{aligned}\tau^2 &= \left\{\frac{(1-\epsilon)q_Aq_B[(1-\epsilon)p_Aq_B+p_B(q_B+\epsilon q_A)]}{(q_B+q_A)^3(q_B+\epsilon q_A)}\right. \\ &\quad \left.+\frac{1}{1+2\gamma}\left[\frac{(1-\epsilon)q_A(q_B+\epsilon q_A)}{(q_B+q_A)^2}+\frac{(1-\epsilon)q_Aq_B[(1-\epsilon)p_Aq_B+p_B(q_B+\epsilon q_A)]}{(q_B+q_A)^3(q_B+\epsilon q_A)}\right]\right\},\end{aligned}$$

and for $p_A < p_B$,

$$\begin{aligned}\tau^2 &= \left\{\frac{(1-\epsilon)q_Aq_B[(1-\epsilon)p_Bq_A+p_A(q_A+\epsilon q_B)]}{(q_B+q_A)^3(q_A+\epsilon q_B)}\right. \\ &\quad \left.+\frac{1}{1+2\gamma}\left[\frac{(1-\epsilon)q_B(q_A+\epsilon q_B)}{(q_B+q_A)^2}+\frac{(1-\epsilon)q_Aq_B[(1-\epsilon)p_Bq_A+p_A(q_A+\epsilon q_B)]}{(q_B+q_A)^3(q_A+\epsilon q_B)}\right]\right\}.\end{aligned}$$

Proof: According to Theorem 2.1 in Hu and Zhang (2004), the asymptotic variance τ^2 is

$$\tau^2 = \frac{\rho(1-\rho)}{1-2\frac{\partial g}{\partial x}\big|_{(\rho,\rho)}}+\frac{2\left(\frac{\partial g}{\partial y}\big|_{(\rho,\rho)}\right)^2\sigma_3^2}{\left(1-\frac{\partial g}{\partial x}\big|_{(\rho,\rho)}\right)\left(1-2\frac{\partial g}{\partial x}\big|_{(\rho,\rho)}\right)},$$

where

$$\sigma_3 = \begin{pmatrix}\frac{\partial\rho}{\partial p_A} & \frac{\partial\rho}{\partial p_B}\end{pmatrix}\begin{pmatrix}\frac{p_Aq_A}{\rho} & \\ & \frac{p_Bq_B}{1-\rho}\end{pmatrix}\begin{pmatrix}\frac{\partial\rho}{\partial p_A}\\ \frac{\partial\rho}{\partial p_B}\end{pmatrix}.$$

For the allocation function $g(x,y)$,

$$\frac{\partial g}{\partial x} = \frac{-\gamma y^{1+\gamma}x^{-\gamma-1}(1-y)^{1+\gamma}(1-x)^{-\gamma}-\gamma y^{1+\gamma}x^{-\gamma}(1-y)^{1+\gamma}(1-x)^{-\gamma-1}}{[y^{1+\gamma}x^{-\gamma}+(1-y)^{1+\gamma}(1-x)^{-\gamma}]^2},$$

and

$$\frac{\partial g}{\partial y} = \frac{(1+\gamma)y^{\gamma}x^{-\gamma}(1-y)^{1+\gamma}(1-x)^{-\gamma}+(1+\gamma)y^{1+\gamma}x^{-\gamma}(1-y)^{\gamma}(1-x)^{-\gamma}}{[y^{1+\gamma}x^{-\gamma}+(1-y)^{1+\gamma}(1-x)^{-\gamma}]^2}.$$

So,

$$\begin{aligned}\frac{\partial g}{\partial x}\Big|_{(\rho,\rho)} &= -\gamma \\ \frac{\partial g}{\partial y}\Big|_{(\rho,\rho)} &= 1+\gamma.\end{aligned}$$

When $p_A > p_B$, for the target proportion $\rho(p_A, p_B)$,

$$\rho(p_A, p_B) = \frac{q_B + \epsilon q_A}{q_B + q_A},$$
$$1 - \rho(p_A, p_B) = \frac{(1-\epsilon)q_A}{q_B + q_A}.$$

Then,

$$\frac{\rho(p_A, p_B)}{\partial p_A} = \frac{\epsilon(q_B + q_A) - (q_B + \epsilon q_A)}{(q_B + q_A)^2} = \frac{-(1-\epsilon)q_B}{(q_B + q_A)^2},$$
$$\frac{\rho(p_A, p_B)}{\partial p_B} = \frac{(q_B + q_A) - (q_B + \epsilon q_A)}{q_B + q_A} = \frac{(1-\epsilon)q_A}{(q_B + q_A)^2}.$$

Therefore,

$$\begin{aligned}
\sigma_3^2 &= p_A q_A \frac{q_B + q_A}{q_B + \epsilon q_A} \frac{(1-\epsilon)^2 q_B^2}{(q_B + q_A)^4} + p_B q_B \frac{q_B + q_A}{(1-\epsilon)q_A} \frac{(1-\epsilon)^2 q_A^2}{(q_B + q_A)^4} \\
&= \frac{p_A q_A (1-\epsilon)^2 q_B^2}{(q_B + q_A)^3 (q_B + \epsilon q_A)} + \frac{p_B q_B (1-\epsilon) q_A}{(q_B + q_A)^3} \\
&= \frac{p_A q_A (1-\epsilon)^2 q_B^2 + p_B q_B (1-\epsilon) q_A (q_B + \epsilon q_A)}{(q_B + \epsilon q_A)(q_B + q_A)^3} \\
&= \frac{(1-\epsilon) q_A q_B [(1-\epsilon) p_A q_A + p_B (q_B + \epsilon q_A)]}{(q_B + q_A)^3 (q_B + \epsilon q_A)}.
\end{aligned}$$

Hence,

$$\begin{aligned}
\tau^2 &= \frac{\rho(1-\rho)}{1+2\gamma} + \frac{2(1+\gamma)^2}{(1+\gamma)(1+2\gamma)} \sigma_3^2 \\
&= \frac{1}{1+2\gamma} \frac{(q_B + \epsilon q_A)(1-\epsilon) q_A}{(q_B + q_A)^2} + \frac{2(1+\gamma)}{1+2\gamma} \frac{(1-\epsilon) q_A q_B [(1-\epsilon) p_A q_A + p_B (q_B + \epsilon q_A)]}{(q_B + \epsilon q_A)(q_B + q_A)^3} \\
&= \frac{(1-\epsilon) q_A q_B [(1-\epsilon) p_A q_A + p_B (q_B + \epsilon q_A)]}{(q_B + q_A)^3 (q_B + \epsilon q_A)} \\
&\quad + \frac{1}{1+2\gamma} \left[\frac{(1-\epsilon) q_A (q_B + \epsilon q_A)}{(q_B + q_A)^2} \frac{(1-\epsilon) q_A q_B [(1-\epsilon) p_A q_A + p_B (q_B + \epsilon q_A)]}{(q_B + q_A)^3 (q_B + \epsilon q_A)} \right].
\end{aligned}$$

The result on τ^2 for $p_A \leq p_B$ can be proven similarly.

Proof of inequality (5.2.1):

If $p_A + p_B \geq (1-\epsilon)$, then

$$2 - (p_A + p_B) \leq (1+\epsilon).$$
$$q_A + q_B - \frac{1+\epsilon}{2} \leq \frac{1+\epsilon}{2}.$$

If $q_A + q_B \geq \frac{1+\epsilon}{2}$, we have

$$\left(q_A + q_B - \frac{1+\epsilon}{2}\right)^2 \leq \frac{(1+\epsilon)^2}{4}. \qquad (A1)$$

If $q_A + q_B \leq \frac{1+\epsilon}{2}$, then

$$\frac{1+\epsilon}{2} - (q_A + q_B) \leq \frac{1+\epsilon}{2},$$

since $q_A + q_B \geq 0$. Hence the inequality (A1) is still valid.

Firstly, we prove the inequality (5.2.1) when $p_A \geq p_B$.

From inequality (A1), it is obvious that

$$\left(q_A + q_B - \frac{1+\epsilon}{2}\right)^2 \leq 2(1-\epsilon)q_B + \frac{(1+\epsilon)^2}{4}.$$

To prove that as $\gamma \to \infty$,

$$\tau^2 \leq q_A q_B (2 - (q_A + q_B))/(q_A + q_B)^3,$$

we want

$$\frac{(1-\epsilon)[(1-\epsilon)p_A q_B + p_B(q_B + \epsilon q_A)]}{q_B + \epsilon q_A} \leq [2 - (q_A + q_B)]$$

i.e. $$\frac{(1-\epsilon)^2(1-q_A)q_B}{q_B + \epsilon q_A} + (1-\epsilon)(1-q_B) \leq [2 - (q_A + q_B)]$$

i.e. $$\frac{(1-\epsilon)^2(q_B - q_B q_A)}{q_B + \epsilon q_A} \leq 2 - (1-\epsilon) - q_A - q_B + q_B(1-\epsilon)$$

i.e. $$\frac{(1-\epsilon)^2(q_B - q_B q_A)}{q_B + \epsilon q_A} \leq 1 + \epsilon - (q_A + \epsilon q_B).$$

This is equivalent to

$$(1-\epsilon)^2 q_B - (1-\epsilon)^2 q_B q_A \leq (1+\epsilon)(q_B + \epsilon q_A) - (q_B + \epsilon q_A)(q_A + \epsilon q_B)$$

$$= (1+\epsilon)q_B + \epsilon(1+\epsilon)q_A - q_B q_A - \epsilon q_B^2 - \epsilon q_A^2 - \epsilon^2 q_A q_B,$$

i.e. $(1-\epsilon)^2 q_B \leq (1+\epsilon)q_B + \epsilon(1+\epsilon)q_A - \epsilon(q_A + q_B)^2,$

i.e. $\epsilon(q_A + q_B)^2 \leq 3\epsilon q_B - \epsilon^2 q_B + \epsilon(1+\epsilon)q_A,$

i.e. $(q_A + q_B)^2 \leq 3q_B - \epsilon q_B + (1+\epsilon)q_A$

$$= 2q_B - 2\epsilon q_B + (1+\epsilon)q_B + (1+\epsilon)q_A$$

$$= 2(1-\epsilon)q_B + (1+\epsilon)(q_A + q_B),$$

i.e. $$\left(q_A + q_B - \frac{1+\epsilon}{2}\right) \leq 2(1-\epsilon)q_B + \frac{(1+\epsilon)^2}{4}.$$

which is true. Hence, the result follows.

Secondly, for $p_A < p_B$, From inequality (A1) we have

$$\left(q_A + q_B - \frac{1+\epsilon}{2}\right)^2 \leq 2(1-\epsilon)q_A + \frac{(1+\epsilon)^2}{4}.$$

To prove that as $\gamma \to \infty$,

$$\tau^2 \leq q_A q_B(2 - (q_A + q_B))/(q_A + q_B)^3,$$

we want

$$\frac{(1-\epsilon)[(1-\epsilon)p_B q_A + p_A(q_A + \epsilon q_B)]}{q_A + \epsilon q_B} \leq [2 - (q_A + q_B)]$$

i.e. $$\frac{(1-\epsilon)^2(1-q_B)q_A}{q_A + \epsilon q_B} + (1-\epsilon)(1-q_A) \leq [2 - (q_A + q_B)]$$

i.e. $$\frac{(1-\epsilon)^2(q_B - q_B q_A)}{q_B + \epsilon q_A} \leq 2 - (1-\epsilon) - q_A - q_B + q_B(1-\epsilon)$$

i.e. $$\frac{(1-\epsilon)^2 q_A - (1-\epsilon)^2 q_B q_A}{q_A + \epsilon q_B} \leq 1 + \epsilon - (q_B + \epsilon q_A).$$

This is equivalent to

$$(1-\epsilon)^2 q_A - (1-\epsilon)^2 q_B q_A \leq (1+\epsilon)q_A + \epsilon(1+\epsilon)q_B - (q_B + \epsilon q_A)(q_A + \epsilon q_B)$$

$$= (1+\epsilon)q_A + \epsilon(1+\epsilon)q_B - q_B q_A - \epsilon q_B^2 - \epsilon q_A^2 - \epsilon^2 q_A q_B,$$

i.e. $(1-\epsilon)^2 q_A \leq (1+\epsilon)q_A + \epsilon(1+\epsilon)q_B - \epsilon(q_A+q_B)^2,$

i.e. $\epsilon(q_A+q_B)^2 \leq 3\epsilon q_A - \epsilon^2 q_A + \epsilon(1+\epsilon)q_B,$

i.e. $(q_A+q_B)^2 \leq 3q_A - \epsilon q_A + (1+\epsilon)q_B$

$$= 2q_A - 2\epsilon q_A + (1+\epsilon)q_A + (1+\epsilon)q_B$$

$$= 2(1-\epsilon)q_A + (1+\epsilon)(q_A+q_B),$$

i.e. $$\left(q_A + q_B - \frac{1+\epsilon}{2}\right) \leq 2(1-\epsilon)q_A + \frac{(1+\epsilon)^2}{4}.$$

which is true. Therefore, the result follows.

Therefore, as $\gamma \to \infty$,

$$\tau^2 \leq q_A q_B (2 - (q_A + q_B))/(q_A + q_B)^3.$$

Programs

```
##### RPW


d<-477

a<-5

b<-5

Add<-1

pa<-0.916

pb<-0.748


cutPoint<-1.96^2

r<-10000


distrSRPW<-function(a,b,Add,cutPoint,pa,pb,r,d){
    naCount<-rep(0,d+1)
    sCount<-rep(0,d+1)
    zCount<-0
for (n in 1:r){ na<-0
        nb<-0
        sa<-0
        sb<-0
    aBall<-a
    bBall<-b
    Add<-Add
    p<-aBall/(aBall+bBall)
```

```
for (i in 1:d){
    u<-runif(3,0,1)
    e<-c(0,0,0)

    if (u[1]<p){na<-na+1
        if (u[2]<pa){sa<-sa+Add
            aBall<-aBall+Add
            e[2]<-1}
        if (e[2]==0){bBall<-bBall+Add}
        e[1]<-1}
    if (e[1]==0){nb<-nb+1
        if (u[3]<pb){sb<-sb+Add
            bBall<-bBall+Add
            e[3]<-1}
         if (e[3]==0){aBall<-aBall+Add}
             }
    p<-aBall/(aBall+bBall)}

    s<-sa+sb
    naCount[na+1]<-naCount[na+1]+1
    sCount[s+1]<-sCount[s+1]+1

    paHat<-(sa+0.5)/(na+1)
    pbHat<-(sb+0.5)/(nb+1)
    z<-(paHat-pbHat)/sqrt(paHat*(1-paHat)/(na+1)+pbHat*(1-pbHat)/(nb+1))
    c1<-0
```

```
    if (z^2<cutPoint){c1<-1}
    if (c1==0){zCount<-zCount+1}
        }
    xNA<-0:d
    probNA<-naCount/r
    probS<-sCount/r
    power<-zCount/r
return(list(xNA=xNA,probNA=probNA,probS=probS,power=power))}

m<-100
dd<-(1+d)*m

ZrpwFinal.probNA<-matrix(1:dd,ncol=m) ZrpwFinal.power<-1:m

for (n in 1:m){Zrpw<-distrSRPW(a,b,Add,cutPoint,pa,pb,r,d)
        tt<-matrix(Zrpw$probNA,ncol=1)
        ZrpwFinal.probNA[,n]<-tt
        ZrpwFinal.power[n]<-Zrpw$power}
save(ZrpwFinal.probNA,file="ZrpwFinal.probNA")
save(ZrpwFinal.power,file="ZrpwFinal.power")



#####DL

Im<-1
```

```
aBall<-3

bBall<-3


distrSDL<-function(Im,aBall,bBall,pa,pb,d,r,cutPoint){
        naCount<-rep(0,d+1)
        sCount<-rep(0,d+1)
        c1<-0


for (n in 1:r){naAdd<-rep(0,d+1)
        sAdd<-rep(0,d+1)
        w<-0
        na<-0
        sa<-0
        sb<-0
        nb<-0


        urn<-c(Im,aBall,bBall)


while(w<d){u<-runif(4,0,1)
    delta<-0
    if (u[1]<urn[1]/sum(urn)){delta<-1
                urn<-urn+c(0,1,1)}
    if (delta==0){w<-w+1
              e<-c(0,0,0)
        if (u[2]<urn[2]/sum(urn[2]+urn[3])){na<-na+1
            if (u[3]<pa){sa<-sa+1
```

```
            e[2]<-1}
        if (e[2]==0){urn[2]<-urn[2]-1
            if (urn[2]<0){urn[2]<-0}}
    e[1]<-1}

    if (e[1]==0){nb<-nb+1
        if (u[4]<pb){sb<-sb+1
            e[3]<-1}
        if (e[3]==0){urn[3]<-urn[3]-1
            if (urn[3]<0){urn[3]<-0}}
            }
        }
    }

paHat<-(sa+0.5)/(na+1)
pbHat<-(sb+0.5)/(nb+1)
Def<-paHat-pbHat
varDef<-sqrt(paHat*(1-paHat)/(na+1)+pbHat*(1-pbHat)/(1+nb))
z<-Def/varDef
c2<-0
if (z^2<cutPoint){c2<-1}
if (c2==0){c1<-c1+1}

naAdd[na+1]<-1
s<-sa+sb
sAdd[s+1]<-1
```

```
    naCount<-naCount+naAdd
    sCount<-sCount+sAdd
        }
    power<-c1/r
    probNA<-naCount/r
    probS<-sCount/r
    xvalue<-0:d
return(list(xvalue=xvalue,probNA=probNA,probS=probS,power=power))}

dd<-(1+d)*m
ZDLFinal.probNA<-matrix(1:dd,ncol=m)
ZDLFinal.power<-1:m

for (n in 1:m){ZDL<-distrSDL(Im,aBall,bBall,pa,pb,d,r,cutPoint)
        tt<-matrix(ZDL$probNA,ncol=1)
        ZDLFinal.probNA[,n]<-tt
        ZDLFinal.power[n]<-ZDL$power}
save(ZDLFinal.probNA,file="ZDLFinal.probNA")
save(ZDLFinal.power,file="ZDLFinal.power")



##### DBC

gamma<-100
alpha<-10^(-7)
```

```
n0<-1

d<-d-2*n0
dd<-(1+d)*m

distrSDBC<-function(pa,pb,alpha,gamma,n0,d,r,cutPoint){
    naCount<-rep(0,d+n0)
    xNA<-n0:(d+n0)
    sCount<-rep(0,d+2*n0+1)
    xS<-0:(d+2*n0)
    zCount<-0

for (n in 1:r){
    sa<-0
    sb<-0
    na<-1
    nb<-1

    u1<-runif(2,0,1)
    if (u1[1]<pa) {sa<-sa+1}
    if (u1[2]<pb) {sb<-sb+1}

    paHat1<-sa/na
    paHat<-paHat1+alpha*(as.numeric(paHat1==0))-alpha*(as.numeric(paHat1==1))
    pbHat1<-sb/nb
```

```
pbHat<-pbHat1+alpha*(as.numeric(pbHat1==0))-alpha*(as.numeric(pbHat1==1))

EstiProp<-TargProp(paHat,pbHat)

x1<-1/2
y<-EstiProp

g1<-y*(y/x1)^gamma/(y*(y/x1)^gamma+(1-y)*((1-y)/(1-x1))^gamma)
p<-g1

for (i in 1:d){
u<-runif(3,0,1)
e<-0
if (u[1]<p){na<-na+1
    if (u[2]<pa){sa<-sa+1}
    e<-1}
if (e==0){nb<-nb+1
    if (u[3]<pb){sb<-sb+1}}

paHat1<-sa/na
paHat<-paHat1+alpha*(as.numeric(paHat1==0))-alpha*(as.numeric(paHat1==1))
pbHat1<-sb/nb
pbHat<-pbHat1+alpha*(as.numeric(pbHat1==0))-alpha*(as.numeric(pbHat1==1))

EstiProp<-TargProp(paHat,pbHat)
```

```
    x1<-na/(2*n0+i)
    y<-EstiProp

    g1<-y*(y/x1)^gamma/(y*(y/x1)^gamma+(1-y)*((1-y)/(1-x1))^gamma)
    p<-g1
        }
    s<-sa+sb
    naCount[na-n0+1]<-naCount[na-n0+1]+1
    sCount[s+1]<-sCount[s+1]+1

    paHat<-(sa+0.5)/(na+1)
    pbHat<-(sb+0.5)/(nb+1)
    Def<-paHat-pbHat
    varDef<-sqrt(paHat*(1-paHat)/(na+1)+pbHat*(1-pbHat)/(1+nb))
    z<-Def/varDef
    c2<-0
    if (z^2<cutPoint){c2<-1}
    if (c2==0){zCount<-zCount+1}
        }
    probNA<-naCount/r
    probS<-sCount/r
    power<-zCount/r
return(list(xNA=xNA,probNA=probNA,xS=xS,probS=probS,power=power))}
```

```
##### epsilin1<-0.5

TargProp<-function(pa,pb){qa<-1-pa
        qb<-1-pb
        rho<-(qb+epsilin1*min(qa,qb)*sign(pa-pb))/(qb+qa)
                return(rho)}
dd<-(d+1)*m

ZDBC.newFinal.probNA<-matrix(1:dd,ncol=m)
ZDBC.newFinal.power<-1:m

for (n in 1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.newFinal.probNA[,n]<-tt
        ZDBC.newFinal.power[n]<-ZDBC$power}
save(ZDBC.newFinal.probNA,file="ZDBC.newFinal.probNA")
save(ZDBC.newFinal.power,file="ZDBC.newFinal.power")

##### gamma<-0 epsilin1<-0.5

TargProp<-function(pa,pb){qa<-1-pa
        qb<-1-pb
        rho<-(qb+epsilin1*min(qa,qb)*sign(pa-pb))/(qb+qa)
                return(rho)}
dd<-(d+1)*m
```

```
ZDBC.newFinal0.probNA<-matrix(1:dd,ncol=m)
ZDBC.newFinal0.power<-1:m

for (n in 1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.newFinal0.probNA[,n]<-tt
        ZDBC.newFinal0.power[n]<-ZDBC$power}
save(ZDBC.newFinal0.probNA,file="ZDBC.newFinal0.probNA")
save(ZDBC.newFinal0.power,file="ZDBC.newFinal0.power")

##### TargProp<-function(pa,pb){rho<-sqrt(pa)/(sqrt(pa)+sqrt(pb))
                return(rho)}

ZDBC.RSIHRFinal.probNA<-matrix(1:dd,ncol=m)
ZDBC.RSIHRFinal.power<-1:m for (n in
1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.RSIHRFinal.probNA[,n]<-tt
        ZDBC.RSIHRFinal.power[n]<-ZDBC$power}
save(ZDBC.RSIHRFinal.probNA,file="ZDBC.RSIHRFinal.probNA")
save(ZDBC.RSIHRFinal.power,file="ZDBC.RSIHRFinal.power")



##### gamma<-0
```

```
TargProp<-function(pa,pb){rho<-sqrt(pa)/(sqrt(pa)+sqrt(pb))
                  return(rho)}
ZDBC.RSIHRFinal.probNA0<-matrix(1:dd,ncol=m)
ZDBC.RSIHRFinal.power0<-1:m for (n in
1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.RSIHRFinal.probNA0[,n]<-tt
        ZDBC.RSIHRFinal.power0[n]<-ZDBC$power}
save(ZDBC.RSIHRFinal.probNA0,file="ZDBC.RSIHRFinal.probNA0")
save(ZDBC.RSIHRFinal.power0,file="ZDBC.RSIHRFinal.power0")


##### gamma<-100
TargProp<-function(pa,pb){rho<-sqrt(pa*(1-pa))/(sqrt(pa*(1-pa))+sqrt(pb*(1-pb)))
                  return(rho)}
ZDBC.NeymanFinal.probNA<-matrix(1:dd,ncol=m)
ZDBC.NeymanFinal.power<-1:m for (n in
1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.NeymanFinal.probNA[,n]<-tt
        ZDBC.NeymanFinal.power[n]<-ZDBC$power}
save(ZDBC.NeymanFinal.probNA,file="ZDBC.NeymanFinal.probNA")
save(ZDBC.NeymanFinal.power,file="ZDBC.NeymanFinal.power")

##### gamma<-0
```

```
TargProp<-function(pa,pb){rho<-sqrt(pa*(1-pa))/(sqrt(pa*(1-pa))+sqrt(pb*(1-pb)))
                    return(rho)}

ZDBC.NeymanFinal0.probNA<-matrix(1:dd,ncol=m)
ZDBC.NeymanFinal0.power<-1:m for (n in
1:m){ZDBC<-distrSDBC(pa,pb,alpha,gamma,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.NeymanFinal0.probNA[,n]<-tt
        ZDBC.NeymanFinal0.power[n]<-ZDBC$power}
save(ZDBC.NeymanFinal0.probNA,file="ZDBC.NeymanFinal0.probNA")
save(ZDBC.NeymanFinal0.power,file="ZDBC.NeymanFinal0.power")


##### Band&Bhatt
TargProp<-function(pa,pb){Ney<-sqrt(pa*(1-pa))/(sqrt(pa*(1-pa))+sqrt(pb*(1-pb)))
            pi<-pnorm((pa-pb)/(sqrt(pa*(1-pa))+sqrt(pb*(1-pb))),0,1)
            if (pa>=pb){p<-max(Ney,pi)}
            if (pa<pb){p<-min(Ney,pi)}
        return(p=p)}
n0<-1 lambda<-0.5 alpha<-10^(-7) r<-10000 cutPoint<-1.96^2

d<-477
d<-d-2*n0

pa<-0.916
pb<-0.748
```

```
distrSBB<-function(pa,pb,alpha,n0,d,r,cutPoint){
    naCount<-rep(0,d+n0)
    xNA<-n0:(d+n0)
    sCount<-rep(0,d+2*n0+1)
    xS<-0:(d+2*n0)
    zCount<-0

    zDistr<-rep(0,r)

for (n in 1:r){
    sa<-0
    sb<-0
    na<-1
    nb<-1

    u1<-runif(2,0,1)
    if (u1[1]<pa) {sa<-sa+1}
    if (u1[2]<pb) {sb<-sb+1}

    paHat1<-sa/na
    paHat<-paHat1+alpha*(as.numeric(paHat1==0))-alpha*(as.numeric(paHat1==1))
    pbHat1<-sb/nb
    pbHat<-pbHat1+alpha*(as.numeric(pbHat1==0))-alpha*(as.numeric(pbHat1==1))
```

```
 p<-TargProp(paHat,pbHat)

for (i in 1:d){
 u<-runif(3,0,1)
 e<-0
 if (u[1]<p){na<-na+1
     if (u[2]<pa){sa<-sa+1}
     e<-1}
 if (e==0){nb<-nb+1
     if (u[3]<pb){sb<-sb+1}}

 paHat1<-sa/na
 paHat<-paHat1+alpha*(as.numeric(paHat1==0))-alpha*(as.numeric(paHat1==1))
 pbHat1<-sb/nb
 pbHat<-pbHat1+alpha*(as.numeric(pbHat1==0))-alpha*(as.numeric(pbHat1==1))

 p<-TargProp(paHat,pbHat)
     }
 s<-sa+sb
 naCount[na-n0+1]<-naCount[na-n0+1]+1
 sCount[s+1]<-sCount[s+1]+1

 paHat<-(sa+0.5)/(na+1)
 pbHat<-(sb+0.5)/(nb+1)
 Def<-paHat-pbHat
 varDef<-sqrt(paHat*(1-paHat)/(na+1)+pbHat*(1-pbHat)/(1+nb))
```

```
    z<-Def/varDef

    zDistr[n]<-z

    c2<-0
    if (z^2<cutPoint){c2<-1}
    if (c2==0){zCount<-zCount+1}
        }
    probNA<-naCount/r
    probS<-sCount/r
    power<-zCount/r
            sF<-xS
            probS<-probS
            MeanS<-sum(sF*probS)
            M2S<-sum(sF^2*probS)
            varS<-M2S-MeanS^2
            obj<-MeanS-lambda*varS

        MeanNA<-sum(xNA*probNA)
        M2NA<-sum(xNA^2*probNA)
        varNA<-M2NA-MeanNA^2
        varNAFinal<-varNA*(r/(r-1))
return(list(zDistr=zDistr,xNA=xNA,probNA=probNA,probS=probS,
power=power,objective=obj,MeanNA=MeanNA,VarNA=varNAFinal,
pa=pa,pb=pb,cutPoint=cutPoint))}
```

```
#################### m<-100 dd<-(d+1)*m


ZDBC.BBFinal.probNA<-matrix(1:dd,ncol=m)
ZDBC.BBFinal.power<-1:m


for (n in 1:m){ZDBC<-distrSBB(pa,pb,alpha,n0,d,r,cutPoint)
        tt<-matrix(ZDBC$probNA,ncol=1)
        ZDBC.BBFinal.probNA[,n]<-tt
        ZDBC.BBFinal.power[n]<-ZDBC$power}
save(ZDBC.BBFinal.probNA,file="ZDBC.BBFinal.probNA")
save(ZDBC.BBFinal.power,file="ZDBC.BBFinaldd.power")
```

Bibliography

[1] Anscombe, F.J., 1963. Sequential medical trials. Journal of the American Statistical Association 58, 365-383.

[2] Anscombe, F.J., 1949. Large-sample theory of sequential estimation. Biometrika 36, 455-458.

[3] Anscombe, F.J., 1952. Large-sample theory of sequential estimation. Proc. Cambridge Philos. Soc. 48, 600-607.

[4] Armitage, P., 1960. Sequential Medical Trials. Springfield, Illinois: Thomas.

[5] Athreya, K.B., Karlin, S., 1967. Limit theorems for the split times of branching processes. Journal of Mathematics and Mechanics 17, 257-277.

[6] Athreya, K.B., Karlin, S., 1968. Embedding of urn schemes into continuous time Markov branching processes and related limit theorems. Annals of Mathematical Statistics 39, 1801-1817.

[7] Bahadur, R.R., 1971. Some limit theorems in statistics. SIAM, Philadephia, PA.

[8] Bai. Z.D. and Hu. F., 2005. Asymptotics in randomized urn models. Annals of Applied Probability 15, 914-940.

[9] Bai, Z.D., Hu, F. and Shen, L., 2002. An adaptive design for multi-arm clinical trials. Journal of Multivariate Analysis 81, 1-18.

[10] Bai, Z.D., Hu, F., Zhang L.-X., 2002. Gaussian approximation theorems for urn models and their applications. Ann. Appl. Probab. 12, 1149-1173.

[11] Bahadur, R.R., Zabell, S.L., Gupta, J.C., 1980. Large deviations, tests, and estimates. In: Chakravarti, I.M. (Ed.), Asymptotic theory of statistical tests and estimation. Academic Press, New York-London-Toronto, pp. 33-64.

[12] Baldi Antognini, A., 2007. A theoretical analysis of the power of biased coin designs. (In Press, Available on line on August 15, 2007)

[13] Baldi Antognini, A. and Giovagnoli, A., 2005. On the large sample optimality of sequential designs for comparing two or more treatments. Sequential Analysis 24, 205-217.

[14] Bandyopadhyay, U., Bhattacharya, R., 2006. Adaptive allocation and failure saving in randomised clinical trials. J. Biopharma. Statist. 16:817-829.

[15] Bandyopadhyay, U., Biswas, A., 2004. An adaptive allocation for continuous response using Wilcoxon-Mann-Whitney score. J. Statist. Plann. Inf. 123, 207-224.

[16] Bartlett, R.H., Rolloff, D.W, Cornell, R.G., Andrews, A.F., Dillon, P.W., Zwischenberger, J.B., 1985. Extracorporeal circulation in neonatal repiratory failure: a prospective randomized trial. Pediatrics 76, 479-487.

[17] Bauer, P. and Brannath, W., 2004. The advantages and disadvantages of adaptive designs for clinical trials. Drug Discovery Today 9, 351-357.

[18] Bauer, P. and Einfalt, J., 2006. Application of adaptive designs - a review. Biometri. J. 48, 493-506.

[19] Bauer, P. and Kohne, K., 1994. Evaluation of experiments with adaptive interim analysis. Biometrics 50, 1029-1041.

[20] Berry, D.A, 1989. Comment: Ethics and ECMO. Statistical Science 4, 306-310.

[21] Berry, D.A., 2001. Adaptive trials and Bayesian statistics in drug development (with discussions). Biopharmaceutical Report 9, 1-11.

[22] Berry, D.A., 2004. Bayesian statistics and the efficiency and ethics of clinical trials. Statistical Science 19, 175-187.

[23] Berry, D.A. and Eick, S.G., 1995. Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. Statistics in Medicine 14, 231-246.

[24] Billingsley, P., 1961. Statistical inference for Markov processes. University of Chicago Press, Chicago.

[25] Biswas, A., Bhattacharya, R., and Zhang, L., 2007. Optimal Response-Adaptive Designs for Continuous Responses in Phase III Trials. Biometrical Journal 49, 928940.

[26] Biswas, A., Huang, H-H. and Huang, W-T., 2006. Covariate-adjusted adaptive designs for continuous responses in a phase III clinical trial: recommendation for practice. Journal of Biopharmaceutical Statistics 16, 227-239.

[27] Bucklew, J.A., 1990. Large deviation techniques in decision, simulation, and estimation. John Wiley & Sons, Inc., New York.

[28] Burman, C. F. and Sonesson, C., 2006. Are flexible designs sound? Biometrics 62, 664-683.

[29] Byar, D.P., Simon, R.M., Friendewald, W.T., Schlesselman, J. J., Demets, D.L., Ellenberg, J. H., Gail, M.H. and Ware, J.H., 1976. Randomized clinical trials: Perspective on some recent ideas. New England Journal of Medicine 295, 74-80.

[30] Chang, M., Chow, S.C. and Pong A., 2006. Adaptive design in clinical research: issues, opportunities, and recommendations. J. Biopharm. Stat. 16, 299-309.

[31] Chen, Y-P., 2006. The power of Efron's biased coin design. Journal of Statistical Planning and Inference 136, 1824-1835.

[32] Cheng, Y. and Berry, D.A., 2007. Optimal adaptive randomized designs for clinical trials. Biometrika 94, 673-689.

[33] Cheng, A.L., Vidyashankar, A.N., 2006. Minimum Hellinger distance estimation for randomized play the winner design. J. Statist. Plann. Inference 136, 1875-1910.

[34] Chow, S.C., Chang, M. and Pong A. 2005. Statistical consideration of adaptive methods in clinical development. J. Biopharm. Stat. 15, 575-591.

[35] Coad, D.S., Ivanova, A., 2001. Bias calculations for adaptive urn designs. Sequential Anal. 20, 91-116.

[36] Coad, D.S. and Woodroofe, M., 1998. Approximate bias calculations for sequentially designed experiments. Sequential Anal. 17, 1-31.

[37] Collins, E.J., 1997. Finite-horizon variance penalised Markov decision processes. OR Spektrum 19, 35-39.

[38] Connor, E.M., Sperling, R.S., Gelber, R. , Pavel, K., Gwendolyn, S., Jo O'Sullivan M., VanDyke R., Bey, M., Shearer, W., Jacobson, R.L., Jimenez, E., O'Neill, E., Brigitte, B., Delfraissy, J.-F. , Culnane, M., Coombs, R., Elkins, M., Moye, J., Stratton, P., Balsley, J., for The Pediatric AIDS Clinical Trials Group Protocol 076 Study Group, 1994. Reduction of Maternal-Infant Transmission of Human Immunodeficiency Virus Type 1 with Zidovudine Treatment. New England Journal of Medicine 331, 1173-1180.

[39] Dirienzo, A.G., 2000. Using urn models for the design of clinical trials. Sankhya 62, 43 - 69.

[40] Durham, S.D., Flournoy, N. and Li, W., 1998. Sequential designs for maximizing the probability of a favourable response. Canadian Journal of Statistics 3, 479-495.

[41] Durham, S.D., Yu, C.F., 1990. Randomized play-the-leader rules for sequential sampling from two populations. Probability in Engineering and Information Sciences 4, 355-367.

[42] Eggenberger, F. and Polya, G., 1923. Uber die statistik verketetter vorgange. Z. Ang. Math. Mech. 1, 279-289.

[43] Eisele, J.R., 1994. The doubly adaptive biased coin design for sequential clinical trials. J. Statist. Plann. Inference. 38, 249-261.

[44] Eisele, J.R., Woodroofe, M.B., 1995. Central limit theorems for doubly adaptive biased coin designs. Ann. Statist. 23, 234-254.

[45] Friedman, B., 1949. A simple urn model. Communica. Pure Appl. Math. 2, 59-70.

[46] Fu, J.C., 1975. The Rate of Convergence of Consistent Point Estimators. Ann. Statist. 3, 234 - 240.

[47] Fu, J.C., 1973. On a theorem of Bahadur on the rate of convergence of point estimators. Ann. Statist. 1, 745-749.

[48] Geraldes, M., Melfi, V., Page, C. and Zhang, H., 2006. The doubly adaptive weighted differences design. J. Statist. Plann. Inference. 136, 1923-1939.

[49] Gallo, P., Chunag-Stein, C., Dragalin, V., Gaydos, B., Krams, M. and Pinheiro, J., 2006. Adaptive designs in clinical drug development-an executive summary of the PhRMA working group. J. Biopharm. Stat. 16, 275-283.

[50] Hardwick, J.P. and Stout, Q.F., 1991. Bandit strategies for ethical sequential allocation. Computing Science and Statistics 23, 421-424.

[51] Hu, F., Rosenberger, W. F., 2003. Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. J. Amer. Statist. Assoc. 98, 671-678.

[52] Hu, F., Rosenberger, W.F., Zhang L.-X., 2006. Asymptotically best response-adaptive randomization procedures. J. Statist. Plann. Inference 136, 1911-1922.

[53] Hu, F., Zhang, L.-X., 2004. Asymptotic properties of doubly adaptive biased coin designs for multi-treatment clinical trials. Ann. Statistist. 32, 268-301.

[54] Inouse, K., Aki, S., 2005. A generalized Pólya urn model and related multivariate distributions. Ann. Inst. Statist. Math. 57, 49-59.

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

[56] Ivanova, A., Flournoy, N., 2001. A birth and death urn for ternary outcomes: stochastic processes applied to urn models. In probability and statistical models with applications (Charalambides, C.A., Koutras, M.V., Balakrishgnan, N., Eds.) Chapman and Hall/ CRC, Boca Raton, 583-600.

[57] Ivanova, A.V., Rosenberger, W.F., Durham, S.D., Flournoy, N., 2000. A birth and death urn for randomized clinical trials: Asymptotic methods. Sankhya B 62, 104-118.

[58] Jennison, C., Turnbull, B.W., 2006. Discussion of "executive summary of the PhRMA working group on adaptive disigns in clinical drug development." J. Biopharm. Stat. 16, 293-298.

[59] Kass, R.E., Greenhouse, J.B., 1989. Comment: A Bayesian Perspective, Statistical Science, 4, 310-317.

[60] Lin, D.Y., Wei, L.J., MeMets, D.L., 1991. Exact statistical inference for group sequential trials. Biometrics 47, 1399-1408.

[61] Lokhnygina, Y., 2006. Discussion of the "executive summary of the PhRMA working group on adaptive designs in clinical drug development". J. Biopharm. Stat. 16, 287-288.

[62] Lurie, P., Wolfe, S.M., 1997. Unethical trials of interventions to reduce perinatal transmission of the human immunodeficiency virus in developing countries. New England Journal of Medicine. 337, 853-856.

[63] Melfi, V.F., Page, C., 2000. Estimation after adaptive allocation. J. Statist. Plann. Inference 87, 353-363.

[64] Melfi, V.F., Page, C., 1998. Variability in adaptive designs for estimation of success probabilities. In: Folurnoy, N., Rosenberger, W.F., Wong, W.K. (Eds.), New Developments and Applications in Experimental Design. Institute of Mathematical Statistics, Hayward, CA, 106-114.

[65] Melfi, V.F., Page, C., Geraldes, M., 2001. An adaptive randomized design with application to estimation. Canad. J. Statist. 29, 107-116.

[66] McNally, H., Benett, C., Elbourne, D., Field, D., 2006. United Kingdom collaborative randomized trial of neonatal extrocorporeal membrane oxygenation: follow-up to age 7 years. Pediatrics 117, 845-854.

[67] Muliere, P., Paganoni, A.M. and Secchi, P. (2006) A randomly reinforced urn. Journal of Statistical Planning and Inference 136, 1853-1874.

[68] Palmer, C.R., Shahumyan, H., 2007. Implementing a decision-theoretic design in clinical trials: Why and how? Statistics in Medicine 26, 4939-4957.

[69] Palmer, C.R., Rosenberger, W.F., 1999. Ethics and practice: Alternative designs for Phase III randomized clinical trials. Control Clinical Trials. 20, 172-186.

[70] Posch, M. and Bauer, P., 2000. Interim analysis and sample size reassessment. Biometrics 56, 1170-1176.

[71] Proschan, M. A. and Hunsberger, S. A., 1995. Designed extension of studies based on conditional power. Biometrics 51, 13155-1324.

[72] Pullman, D., Wang, X., 2001. Adaptive designs, informed consent, and the ethics of research. Controlled Clinical Trials 22, 203-210.

[73] Rosenberger, W.F., 2002. Randomized urn models and sequential design. Sequential Analysis. 21, 1-28.

[74] Rosenberger, W.F., Flournoy, N., Durham, S.D., 1997. Asymptotic normality of maximum likelihood estimators from multiparameter reponse-driven designs. J. Statist. Plann. Inference 60, 69-76.

[75] Rosenberger, W.F. Lachin, J.M., 2002. Randomization in clinical trials: theory and practice. John Wiley and Sons, New York.

[76] Rosenberger, W.F., Stallard, N., Ivanova, A., Harper, C.N., Ricks, M.L., 2001. Optimal adaptive designs for binary response trials. Biometrics 57, 909-913.

[77] Rosenberger, W.F., Vidyashankar, A.N. and Agarwal, D.K., 2001. Covariate-adjusted response-adaptive designs for binary response. Journal of Biopharmaceutical Statistics 11, 227-236.

[78] Royall, R.M., 1991. Ethics and statistics in randomized clinical trials. Statistical Science 6, 52-88.

[79] Symthe, R.T., 1996. Central limit theorems for urn models. Stochastic Process. Appl. 65, 115-137.

[80] Sperling, R.S., Shapiro, D., Coombs, R.W., Todd, J.A., Herman, S.A.,McSherry, G., Jo O'Sullivan, M., Van Dyke, R., Jimenez, E., Rouzioux, C., Flynn, P.M., Sullivan, J.L., for The Pediatric AIDS Clinical Trials Group Protocol 076 Study Group, 1996. Maternal viral load, Zidovudine treatment, and the risk of transmission of human immunodeficiency virus type 1 from mother to infant. New England Journal of Medicine 338, 1621-1629.

[81] Sun, R., Cheung, S.H. and Zhang, L.-X., 2007. A generalized drop-the-loser rule for multi-treatment clinical trials. Journal of Statistical Planning and Inference 137, 2011-2023.

[82] Tamura, R.N., Faries, D.E., Anderson, J.S., Heiligenstein, J.H., 1994. A case study of an adaptive clinical trial in the treatment of out-patients with depressive disorder. J. Amer. Statist. Assoc. 89, 768-776.

[83] UK Collaborative ECMO Trial Group., 1998. The collaborative UK ECMO trial: Follow-up to 1 year of age. Pediatrics 101, 1.

[84] Wassmer, G., 2006. Discussion of "executive summary of the PhRMA working group on adaptive designs in clinical drug development". J. Biopharm. Stat. 16, 289-291.

[85] Wei, L.J., 1988. Exact two-sample permutation tests based on the randomized play-the-winner rule. Biometrika 75, 603-606.

[86] Wei, L.J., 1979. The generalized Pólyas urn for sequential medical trials. Ann. Statist. 7, 291-296.

[87] Wei, L.J., Durham, S., 1978. The randomized play-the-winner rule in medical trials. J. Amer. Statist. Assoc. 73, 840-843.

[88] White, D.J., 1992. Computational approaches to variance-penalised Markov decision processes. OR Spektrum 14, 79-83.

[89] Yao, Q., Wei L.J., 1996. Play the winner for phase II/III clinical trials. Statistics in Medicine. 15, 2413-2423.

[90] Yi, Y., Wang, X., 2007a. Goodness-of-fit test for response adaptive clinical trials. Statist. Prob. Letters. 77, 1014-1020.

[91] Yi, Y., Wang, X., 2007b. Statistical inference of adaptive designs with binary responses. Comm. Statist. Theory Methods 36, 1643-1653.

[92] Zhang, L.-X., Chan, W.S., Cheung, S.H., Hu, F., 2007. A generalized drop-the-loser urn for clinical trials with delayed responses. Statistica Sinica. 17, 387-409.

[93] Zelen, M., 1969. Play the winner rule and the controlled clinical trial. J. Amer. Statist. Assoc. 64, 131-146.

[94] Zhang, L.-X., Hu, F., Cheung, S.H., 2006. Asymptotic theorems of sequential estimation-adjusted urn models. Annals of Applied Probability. 16, 340-369.

[95] Zhang, L.-X., Hu, F., Cheung, S.H., Chan, W.S., 2007. Asymptotic properties of covariate-adjusted response-adaptive designs. Annals of Statistics 35, 1166-1182.

[96] Zhang, Y.Q., Rosenberger, W.F., 2005. On asymptotic normality of the randomization-based log rank test. Nonparametric Statistics 17, 833-839.