University of Southampton
Faculty of Social Sciences
Social Statistics and Demography

Design of Sequential Experiments with Covariate Information

by

Mia Sato Tackney

A thesis submitted for the degree of
Doctor of Philosophy

October 2019

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL SCIENCES

SOCIAL STATISTICS AND DEMOGRAPHY

Thesis for the degree of Doctor of Philosophy

DESIGN OF SEQUENTIAL EXPERIMENTS WITH COVARIATE INFORMATION

by Mia Sato Tackney

Randomized experiments are considered the gold standard of study designs. When averaged over all possible randomizations, the treatment effect is free from bias and can also have a causal interpretation. In experiments involving human participants, it is common for participants to become available sequentially over time. Covariate information, such as age and sex, is taken and a treatment is assigned soon after. In this sequential setting, there are two main dangers of allocating treatments completely at random: replication of treatments may be unbalanced, and replication of treatments may be unbalanced for covariates. Both problems can lead to imprecision in the estimate of the treatment effect.

Covariate-adaptive schemes allocate treatments to patients to minimize some measure of imbalance in treatments and covariates. Minimization is a method commonly used in clinical trials which is appropriate for binary covariates. Alternatively, sequential optimal design based methodology allocates treatments to minimize the variance of the treatment effect under a specified model. We extend the optimal design based methodology to a nonmyopic setting, where treatment allocation for the current patient depends not only on the treatments and covariates of the patients in the study, but also the impact of possible future treatment assignments. The number of possible future decisions considered is called the horizon. Our simulation studies shows that there are very few examples with binary treatment where the non-myopic approach offers benefit over the myopic approach. One main limitation of the nonmyopic approach is that it involves computationally expensive recursive formulae which can only be implemented in limited contexts, for example for discrete treatments and for a horizon of no more than five. This motivated the development of a pseudo-nonmyopic approach which has a similar aim to the nonmyopic approach, but does not involve recursion. The horizon can be up until the end of the trial and the approach can also be used for continuous treatments.

We apply the sequential nonmyopic and pseudo-nonmyopic framework in the setting of personalized medicine. A trial for personalized medicine aims to identify effective combinations of treatments and biomarkers. In this context, our main result from simulations is that the pseudo-nonmyopic approach is more efficient than the myopic approach in the logistic model case with continuous treatment where there is a large interaction between the treatment and biomarker. In particular, this benefit is more pronounced when the biomarker relevant in the interaction is rare.

Table of Contents

List of Figures

List of Tables

List of Algorithms

Declaration of Authorship

I, Mia Sato Tackney, declare that this thesis entitled *Design of Sequential Experiments with Covariate Information* and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Either none of this work has been published before submission, or parts of this work have been published as:

Signed:

Date:

Acknowledgements

For their academic support and guidance, and for their inspiring dedication to statistics and education, I would like to thank my supervisors, Professor Dave Woods, Dr. Ilya Shpitser, and Professor Peter Smith. Right from the start, when I showed interest in studying Design of Experiments in Southampton, Dave responded with tremendous enthusiasm and support which continued throughout my Ph.D. I am also thankful to have had helpful discussions and advice from Dr. Kim May Lee and to have been part of a lively research group in Design of Experiments.

For funding support and for creating a nurturing community of Ph.D. students and staff, I would like to thank the South Coast Doctoral Training Partnership.

For being the coolest people I know, I would like to thank my parents, Toyoko Sato and Charlie Tackney. From the outskirts of Mt. Fuji in Japan, and Saddle Brook, New Jersey, they were the first in their families to pursue postgraduate education. Their choices meant that I was blessed with many opportunities to learn and a family that was completely supportive of my academic pursuits. Without them and all of my wonderful companions who are a joy in my life, this work would not have been possible. I would like to express special thanks to Francisco de Melo Virissimo, Astrid Ariazza and Jin Lee.

List of Symbols

Notation	Description
n	Number of units in an experiment
i	Current unit in a sequential experiment
k	Number of covariates in an experiment
t	Number of treatments in an experiment
t_i	Treatment given to ith unit
$\boldsymbol{t}_i$	$(t_1, t_2, ..., t_i)^\intercal$
$z_{i,s}$	Value of covariate s for unit i, $1 \leq s \leq k$
$\boldsymbol{z}_i$	$(z_{i,1}, z_{i,2}..., z_{i,k})^\intercal$
$\boldsymbol{Z}_i$	$\begin{pmatrix} \boldsymbol{z}_1^\intercal \\ \boldsymbol{z}_2^\intercal \\ \vdots \\ \boldsymbol{z}_i^\intercal \end{pmatrix}$
y_i	Response for ith unit
$\boldsymbol{y}_i$	$(y_1, y_2, ..., y_i)^\intercal$
$\boldsymbol{\beta}$	q-vector of unknown parameters for a model
$\overline{\boldsymbol{Z}}_i$	$(\mathbf{1} \ \boldsymbol{Z}_i)$
$\boldsymbol{X}_i$	Design matrix $\left(\overline{\boldsymbol{Z}}_i \ \boldsymbol{t}_i\right)$
$\boldsymbol{M}_i$	Information matrix $\boldsymbol{X}_i^\intercal \boldsymbol{X}_i$
Ψ	Objective function
$\Psi_{D_A}, \Psi_D, \Psi_G$	Objective functions for D_A-, D- and G-optimality, respectively
m	Number of linear combinations of interest for D_A optimality
$\boldsymbol{A}$	$q \times m$ matrix, each column gives coefficients of the linear combination for D_A optimality
N	horizon for the non-myopic approach
Φ_l	Expected loss after l units in the future
π_i	$P(y_i = 1)$ for a binary response
η_i	linear predictor for the ith unit for logistic regression

R	Number of subgroups of interest for the L-optimal objective function
$\boldsymbol{c}_r$	Indicator vector for the linear combination of parameters of interest in subgroup r,for $1 \leq r \leq R$
τ_r	Minimum uninteresting threshold for subgroup r
w_r	Weight for the rth subgroup of interest
M	Number of pseudo-designs in the pseudo-nonmyopic approach

Chapter 1

Introduction: Background, Aims and Outline

1.1 Background

1.1.1 Design of Experiments

Experiments are conducted to investigate scientific research questions which can inform decision-making. For example, clinical trials can determine whether a new treatment is better than the current best available treatment, and for which subpopulation it is most effective (Zhang et al., 2016). A chemicals experiment can find the optimal locations for taking measurements to identify the source of a contaminant (Huan and Marzouk, 2016). Tech giants such as Google and Amazon regularly conduct online experiments to determine the layout of a website or which widget is most attractive to customers to inform business decisions (Bakshy et al., 2014). In order to conduct such experiments, a number of operational decisions need to be made on how the data is collected: for example, how many subjects or units to include, at which locations or time points to take measurements, and how to decide which units receive which treatment. These decisions can have a direct impact on the quality of the data that is obtained and the conclusions that can be drawn from the experiment. Design of experiments provides a framework for defining what constitutes good decisions and developing methods for making them. This thesis focuses on how treatments should be allocated to the units of an experiment.

Assigning treatments at random or in an ad-hoc way could lead to inefficient estimates of treatment effects. One typically wishes to choose treatments in such a way that leads to statistical efficiency, meaning that the variance of estimators are as small as possible (Cox,

1958, p.7). Estimating the treatment effect as precisely as possible, or testing hypotheses about the treatment effect with high power are common ways in which designs are assessed (Atkinson, 1982; Hu et al., 2015). There may be other considerations in the experiment which are in conflict with the goal of making decisions for statistical efficiency (Rosenberger and Sverdlov, 2008). For example, in clinical experiments involving human participants, there is an ethical dimension that needs to be considered carefully. If it appears that one treatment is more effective than another, there is an ethical incentive to give proportionately more patients the effective treatment. If a particular treatment has harmful effects on patients, the treatment arm or the experiment itself may be dropped altogether. How to effectively balance multiple goals which potentially conflict with each other is a current area of research within the community.

1.1.2 Sequential Design with Covariates

Some units in an experiment share inherent characteristics which may be of importance to the results of an experiment. These characteristics which have the potential to explain the variation in response are referred to as covariates. For example, in an agricultural experiment, plots of land which are close to each other are more likely to experience similar weather conditions rather than plots that are further apart, so the location of the unit is a covariate (Bailey, 2008, p.55). In clinical trials, individual genetic or demographic characteristics of the patient are examples of covariates. A common experimental design when the covariates of all units are known is a complete block design, where units are grouped into blocks which have the same covariate value, and treatments are randomized so that each treatment is equally represented in each block (Atkinson et al., 2007, p.8); this leads to more precise estimates of the treatment effect.

In experiments in the social and medical sciences, potential participants often become available at different times throughout the recruitment period. Once a potential participant expresses interest, formal checks on eligibility and safety are made, and covariate information such as the participants' age, sex and medical history are taken before they enrol in the experiment (Pocock, 2013, p.67). In this setting, sequential design of experiments, where treatments are allocated to participants as soon as they enrol, is a practical approach and is also an active area of research.

The major challenge in sequential design of experiments is that, at each stage of the experiment, a treatment needs to be assigned to the current participant, based only on information about the participants already enrolled in the experiment, and the covariate information

about the current patient. It may not be known how many more people will enrol, and covariate information about future participants is unknown so a block design is not an option. A simple randomization strategy might lead to unequal replication of treatments (Efron, 1971). Further, it may lead to a situation where the replication of treatments is unequal for particular levels of a covariate. These issues could lead to bias and/or increased variability in estimates (Hu and Hu, 2012). The consequence is particularly pronounced in experiments with a small number of participants, which can be common in the social and medical sciences (Atkinson, 1982).

1.1.3 Overview of Existing Methods

Efron (1971) proposed the so-called “biased coin" that allocates treatments sequentially so that their replication is kept balanced while at the same time keeping a stochastic element to the allocation to mitigate selection bias. However, this does not consider the covariate information of the patients. Minimization is an approach aimed to keep the numbers of treatments roughly equal for each group of patients who have the same covariate combination and is now used extensively in clinical trials (Pocock and Simon, 1975; Taves, 1974). It has received some criticism for being based on measures of imbalance of covariates which are not theoretically grounded (Senn et al., 2010) and methods based on minimizing the variance of the parameter estimators in statistical models have been suggested instead by Atkinson (1982). Atkinson’s approach aims to minimize the variance of the estimator of the treatment effect within a linear model which describes the relationship between the treatments, covariates and response. Both minimization and Atkinson’s approach may be referred to as covariate-adaptive methods, as the criterion used to determine the treatment allocation takes into account the covariates of the patients, but does not consider their responses. We investigate these methods further in Chapter 2.

Response-adaptive approaches include methods such as the play-the-winner rule proposed by Wei and Durham (1978) which updates the probabilities for treatment assignment based on the response of the units only; it does not consider covariates. Some authors make the distinction between response-adaptive randomization (RAR) and covariate-adjusted response-adaptive (CARA) designs to distinguish between those that take covariates into account and those which do not (Rosenberger and Sverdlov, 2008; Zhu, 2015). For some statistical analyses which go beyond the linear model, such as generalized linear models, generalized linear mixed models and non-linear models, the criterion used for treatment allocation may depend on the model parameters. In this case, updating the estimates of the model parameters as more data becomes available can lead to more precise estimators of the treatment effect (Atkinson, 1999, p.257). CARA designs have been proposed by Atkinson

and Biswas (2005), Bandyopadhyay and Biswas (2001) and Rosenberger et al. (2001). Zhang et al. (2007) present a general framework for CARA designs for models including GLMs and provide some asymptotic properties of the estimators of the unknown parameters.

Adaptive designs that incorporate covariate information and also try to maximize the number of patients who receive the more effective treatment (once enough data has been accumulated in the trial to be able to identify a more effective treatment) are sometimes referred to as covariate-adjusted response-adaptive designs based on efficiency and ethics (CARAEE). Hu et al. (2015) propose a CARAEE design which assigns treatments based on the covariates, treatments and responses of the patients in the trial so far, and uses a tuning parameter to balance efficiency and ethics. Further, some work has been undertaken on relating bandit problems to clinical trials. Bandit problems are concerned with sampling sequentially from two populations in such a way to maximize the expected utility (Rigollet and Zeevi, 2010). In a clinical trials application, the two populations correspond to two possible treatments. A utility is defined which combines the merits of having precise estimators and also of maximizing the number of positive responses from patients, and the design aims to maximize this utility. The conflict between the goal of having high efficiency and the goal of keeping the experiment ethical by allowing a greater proportion of patients to receive the more effective treatment can be placed in a decision theoretic framework (Cheng and Berry, 2007; Mueller et al., 2007).

Table 1.1 provides a summary of the allocation methods used in sequential design of experiments and characterizes them based on the factors that the treatment assignment depends on.

Table 1.1: Summary of the types of allocation methods for sequential design of experiments and their features

Treatment assignment depends on:	**Allocation method**					
	Random	**Efron's biased coin**	**Covariate adaptive designs**	**Response adaptive designs**	**CARA**	**CARAEE**
Treatment	x	✓	✓	✓	✓	✓
Covariates	x	x	✓	x	✓	✓
Response	x	x	x	✓	✓	✓
Ethics	x	x	x	✓	x	✓

1.1.4 Nonmyopic Approaches

RAR, CARA and CARAEE designs are myopic approaches, meaning that we assume that the current patient is the last patient to enter the trial, and we choose a treatment which

minimizes the objective function using data up until the current patient. Non-myopic approaches, however, are able to consider the potential impact of the current treatment decision on future possible decisions (Huan and Marzouk, 2016). This relies on the method of dynamic programming to compute the expected value of the objective function, where the expectation is taken over unknown quantities of future patients. Most applications of nonmyopic approaches in clinical trials have been to achieve the goal of maximizing the total benefit of treatment to patients, such as methods using the Gittens index (Villar and Rosenberger, 2018). Nonmyopic approaches for a clinical trials based problem involving covariates where the objective is related to the estimation of parameters have not been explored explicitly in the literature. A major theme of this thesis is implementing the non-myopic approach in existing covariate-adaptive and CARA designs and to understand whether the non-myopic approach provides benefit in terms of efficiency. We explore this in Chapters 3, 4 and 5.

1.2 Clinical Trials

The most prominent application of sequential design with covariates is clinical trials. Clinical trials are experiments conducted in the biomedical or behavioural sciences which involve patients. They aim to test the safety and effectiveness of a new treatment, such as a drug, vaccination or dietary recommendation. The new treatment is compared to the current best treatment, or a placebo. A clinical trial is typically structured into the following four phases (MRC-CTU, 2014):

- **Phase I:** A small number of healthy individuals are recruited to test the safety of the treatment. Side effects are identified. Dose-escalation may be conducted, where the dose of the treatment is increased at intervals in order to measure the tolerance of the treatment and to ascertain some bounds on an appropriate dosage.
- **Phase II:** Once the safety of the drug is confirmed, a larger group of individuals (usually less than 100) who are affected by the disease or condition of interest are recruited. The efficacy of the treatment on the condition is tested. A number of different doses and timings are trialled to establish a treatment policy (Whitehead, 1992, p.3). The main purpose of a Phase II trial is to establish whether the new treatment works well enough to justify studying it further in a Phase III trial. In some Phase II trials, all patients are given the treatment of interest, while others involve more than one treatment and randomly allocate treatments so that it is possible to, for example, compare the treatment of interest to a placebo or the current best treatment (Cancer Research UK, 2015, Sambucini, 2015).

- **Phase III:** A trial in Phase III is much longer and typically involves hundreds or thousands of patients, possibly across different hospitals and countries. The experiment is randomized and the new treatment is compared to the current best treatment on the market or a placebo. A Phase III trial should provide enough evidence to allow drug authorities to make a decision on whether to make the treatment available on the market.
- **Phase IV:** Once authorities for drug regulations approve of the treatment, more information is gathered, such as any potential long-term side effects and effects on particular populations through monitoring the use of the drug in clinical practice (Whitehead, 1992, p.3).

Our interest primarily lies in Phase II and Phase III trials since they are randomized experiments and are typically sequential.

Personalized Medicine

The methods we are interested in may be particularly applicable for clinical trials for personalized medicine. In personalized medicine, the goal is to select treatments for patients which are tailored to individual characteristics (Kaplan, 2015). Treatments can be targeted at specific genetic or biological mechanisms called biomarkers, which we treat as a type of covariate. There is a need to be able to validate effective treatment-biomarker combinations through clinical trials (Lee and Wason, 2019). We show how an optimal design approach can be taken to design a sequential experiment which seeks to identify effective combinations of treatments and biomarkers. We compare myopic and non-myopic methods of finding the optimal treatment for such a problem in Chapter 6.

1.3 Aims of the Thesis

This thesis has five aims:

1. Assess two existing methods for designing sequential experiments with covariate information: minimization and Atkinson's method. We compare the performance of these two methods in a simulation study on performance measures which assess balance in treatment, balance in covariates, and optimality.
2. Extend Atkinson's approach so that it can be applied to a more general form of optimality criterion and for normal as well as binary responses.
3. Extend Atkinson's approach to a non-myopic framework for the binary response. Using simulations, we compare the performance of the myopic and non-myopic approaches

under a number of optimality criteria for both linear and non-linear models. We aim to understand if the non-myopic approach provides benefit over the myopic approach and why or why not.

4. Apply these methods to a problem in personalized medicine concerning the identification of effective treatment-biomarker combinations in a sequential setting.

5. Develop an `R` package with functions to construct designs using the above methods.

1.4 Outline of Thesis

In Chapter 2, we illustrate two naive approaches to the design problem to motivate the need for more sophisticated tools. We then describe minimization and the method by Atkinson (1982). We provide results from simulations comparing the two covariate-adaptive approaches to the naive approaches for a number of different scenarios. Chapter 3 gives a summary of the literature on the use of non-myopic approaches in clinical trials. We then describe how a non-myopic version of Atkinson's algorithm can be implemented in Chapter 4 and provide some simulation results comparing the myopic and non-myopic approaches. We also extend the non-myopic algorithm to allow for binary response and provide simulations to compare the myopic and non-myopic approaches when the response is binary. We then develop the idea of a pseudo-nonmyopic approach in Chapter 5 and again provide simulation results for the linear and logistic model cases. In Chapter 6, we introduce personalized medicine and the problem of identifying effective treatment-biomarker combinations. We show how the sequential design based on an optimal design approach can be used in this setting, and compare performances of myopic, non-myopic and pseudo-nonmyopic versions of the allocation scheme in some simulation studies in Chapter 7. Chapter 8 is a vignette of the `R` package. In Chapter 9, we outline the limitations of our current work and provide plans for future work.

Chapter 2

Myopic Approaches

In this thesis, we consider methods for treatment allocation in experimental settings where units become available sequentially and covariate information about units can be measured. Clinical trials are a prominent example of such a setting. We begin this chapter by describing some practical aspects of clinical trials to illustrate key ideas such as the need to balance treatment levels for covariates. We then describe some methods for sequential treatment allocation. Randomization and Efron's biased coin are two simple approaches. We describe their potential problems to motivate the need for covariate-adaptive approaches which include minimization and an optimal design approach. We illustrate the classic form of minimization used in clinical trials for binary covariates, as well as how it can be extended for continuous covariates. These approaches all make decisions about the current patient assuming that it is the last patient to be allocated in the trial; it does not take into account future patients. We refer to this approach as a *myopic* approach. We run simulations to compare the characteristics of designs produced by these myopic treatment allocation methods.

2.1 Sequential trials

For Phase II and Phase III trials, recruitment efforts are made in order to obtain volunteers over a specified period of time. When a volunteer expresses interest in participating, a number of covariates of the volunteer are noted, such as age, sex and hospital identification (Kendall, 2003). Covariates may also be referred to as prognostic information or baseline characteristics. Some of the covariates are used to check whether the volunteer meets the eligibility criteria for the trial, since there are exclusion criteria determined by the target population of the drug and any known information on subpopulations for which the treatment may have adverse effects. Subjects who do not meet the eligibility criteria, or cannot provide informed consent, are excluded from the trial (Kendall, 2003). Covariates are also included

so that analyses on subgroups can be conducted to see whether treatments are more effective for patients in combination with certain characteristics (Assmann et al., 2000).

Treatments are allocated to the patient soon after they are enrolled in the study. Since patients become available for the trial sequentially over time, their covariate information is also unknown until they enrol for the study. If there are only a few discrete covariates, permuted block designs can be used to allocate treatments, where patients who have the same combination of covariate values are grouped into a block. Suppose each block has $l \times t$ patients, where l is some positive integer and t is the number of treatments. A permutation of of size $l \times t$ is selected so that l replications of each treatment are allocated in each block. Since the patients in each block are more homogeneous than patients between blocks, the variance of the treatment effect should be reduced (Wu, 1985, p.51). However, block designs are not a feasible option if there are many covariates. Even with a few covariates, this method can be problematic since it is not guaranteed that enough patients for each covariate combination will arrive. Further, as the number of covariates increases, the number of required blocks increases exponentially.

A further consideration for sequential trials is that the total sample size may not be fixed; it may not be known how many available subjects can be recruited. Further, due to ethical considerations, there are stopping rules which determine when the trial should be terminated as a result of the responses of the patients enrolled so far. If the current treatment appears to be unsafe or ineffective, the trial may be stopped on the grounds that it is putting patients inadvertently at risk or wasting resources. If the treatment is deemed to be safe and effective, it may be unethical to withhold it from the control group, so the trial may also be justifiably terminated. The targeted value for treatment replication may be unknown since the sample size may be unknown. More information on stopping rules can be found in Stallard et al. (2001).

We consider treatment allocation approaches which are appropriate for the situation where patients, and their covariate information, become available sequentially and the sample size is unknown. Methods which take into account the covariate information of the patient are referred to as *covariate adaptive* methods, or *covariate adaptive* randomization (Rosenberger and Sverdlov, 2008). They can also be referred to as *dynamic allocation schemes* (Senn et al., 2010). Some sequential allocation strategies allow responses from the patients to be incorporated; these are usually called *response adaptive* or *covariate adjusted response adaptive* methods (Rosenberger and Sverdlov, 2008) and they will be considered in Chapters 4 and 6.

2.2 Notation

We first set up some notation that we will use throughout this thesis. We denote by n, the potentially unknown total number of patients in the trial. At a particular point in time, there are $i-1$ patients already enrolled in the trial, and we consider the best choice of treatment for patient i, $i \in \{1, 2, ..., n\}$. We assume that the treatment factor is binary with levels coded by 1 (e.g. new treatment) and -1 (e.g. control). It is typical in clinical trials to have a binary treatment structure, although treatment structures with three or more treatment levels occur occasionally. We denote by $t_i \in \{-1, 1\}$ the treatment given to the ith patient and define the following vectors:

$$\boldsymbol{t}_i = \Big(t_1, t_2, ..., t_i\Big)^{\mathsf{T}}, \tag{2.1}$$

associated with the treatments given to the i patients in the trial so far.

We denote by $z_{i,j}$ the value of covariate j for patient i, for $1 \leq j \leq k$, and define the vector $\boldsymbol{z}_i$:

$$\boldsymbol{z}_i = \Big(z_{i,1}, z_{i,2}, ..., z_{i,k}\Big)^{\mathsf{T}}, \tag{2.2}$$

associated with the covariates values of the ith patient. The $i \times k$ matrix $\boldsymbol{Z}_i$ consists of the all covariates values of patients 1 up to patient i:

$$\boldsymbol{Z}_i = \begin{pmatrix} \boldsymbol{z}_1^{\mathsf{T}} \\ \boldsymbol{z}_2^{\mathsf{T}} \\ \vdots \\ \boldsymbol{z}_i^{\mathsf{T}} \end{pmatrix}. \tag{2.3}$$

Finally, we denote the response of patient i by y_i and the following vectors are the responses for the i patients in the trial so far:

$$\boldsymbol{y}_i = \Big(y_1, y_2, ..., y_i\Big)^{\mathsf{T}}. \tag{2.4}$$

2.3 Randomization

The simplest approach to treatment allocation is complete randomization. Let us suppose that there are two treatments. For complete randomization, essentially, a coin is tossed

to decide the whether the new patient receives the treatment or the control, so we have that $\mathbb{P}(t_i = 1) = 0.5$ for all i. In an experimental setting, treatments are ideally assigned to experimental units randomly; the randomized controlled trial is deemed to be the gold standard of study designs (Rubin, 1978). Randomization achieves the following three aims (Cox and Reid, 2000, p.32):

1. It removes bias. Randomization ensures that treatment assignment is unconfounded with other variables. Confounding bias can occur when there is an extraneous variable that is highly correlated with the treatment assignment, as well as the response (Cox and Reid, 2000, p. 34). This would cause treatment effect to be under or over-estimated.

2. It makes causal inference possible. Randomization allows the ignorability assumption to be plausible, which means that the assignment mechanism is independent of the observed outcomes (Rubin, 1978). Any observed difference in response between the treatment and control groups can be attributed to the treatment effect, and not due to confounders.

3. It makes it possible to estimate variances and confidence intervals of treatment contrasts based on randomization-based inference under the assumption of unit-treatment additivity. There is no need to construct a probabilistic model for the response in order to obtain confidence intervals for the treatment effect.

In the context of clinical trials, there are adverse consequences of allocating treatments completely at random. For a sequence of treatment allocations obtained by randomization, it is possible that the replication of treatments is uneven across treatment groups. For small sample sizes, this can be particularly problematic. Further, the first aim of randomization listed above is only true in theory; over all possible randomizations, treatment assignment is completely uncorrelated with any other variable and there is no possibility of bias (Senn, 2004). For a given realization of randomized treatment assignment in an experiment, there is the possibility that treatment assignment is not ideal. It may be that, by chance, it highly correlated with an extraneous or prognostic variable. The permuted block method, by ensuring that treatments are equally replicated within each block, can alleviate both issues, and it is a method often used in the pharmaceutical industry (Senn, 2004). However, this approach is impractical when there is a large number of covariates.

2.4 Efron's biased coin

Efron's biased coin is an alternative approach to the permuted block design that is designed to promote balance in the replication of treatments (Efron, 1971). Efron's method states that, for the ith patient, the under-represented treatment so far in the trial should be

allocated with probability p, where $\frac{1}{2} < p \leq 1$. Complete randomization is equivalent to $p = \frac{1}{2}$, and $p = 1$ would be a deterministic approach, based purely on keeping replication balanced.

Efron recommended choosing $p = \frac{2}{3}$. The main reasons for this choice are as follows:

- With $p = \frac{2}{3}$, the experiment has asymptotic probability $\frac{1}{2}$ of having equal treatment replication for an even number of patients, and probability $\frac{3}{4}$ of being as close to possible to having equal replication for an odd number of patients.
- A comparison of treatment allocations using $p = \frac{2}{3}$ with a permuted block design with 5 patients per block show that the two designs have similar treatment replication, particularly for small sample sizes.
- Efron defines a measure, "excess selection bias," based on how well one can predict the treatment allocation of the new patient, given the treatment allocations of the previous patients. He shows that, asymptotically, this is low for $p = \frac{2}{3}$.
- One may wish to make randomization-based inference on the treatment effect (Imbens and Rubin, 2015, Chapter 5). Efron states that, for $p = \frac{2}{3}$, the standard deviation of the treatment effect cannot exceed 1.044 times its value under a completely randomized treatment allocation method.

Efron's biased coin is more practical than the permuted block approach since there is no need to wait for a certain number of patients to become available, and is also shown to have desirable properties when $p = \frac{2}{3}$. However, the risk of confounding between treatment allocation and covariates is still present. We now consider two main strands of covariate-adaptive approaches. We first describe minimization, which is used in clinical trials. Then, we provide background information on optimal design and introduce the method proposed by Atkinson (1982). Lastly, we briefly mention some other methods.

2.5 Minimization

Minimization is a collection of methods aimed to ensure that treatment levels are balanced with respect to the number of patients and also predefined covariates (see review article by Scott et al., 2002). Minimization was initially independently developed in the context of clinical trials by Taves (1974) and Pocock and Simon (1975). It is the most commonly used alternative to permuted block methods of randomization and is used by organizations such as the European Organization for Research and Treatment of Cancer (EORTC) and the Medical Research Council in the United Kingdom (Senn et al., 2010).

2.5.1 Binary covariates

The classic form of minimization is appropriate for binary covariates. Assuming that $i-1$ patients have already been assigned a treatment, we compare the effect of assigning patient i to treatment 1 and -1 on an imbalance measure. For covariate s, we calculate $n_1(z_{i,s})$, which is the number of patients already in the study who have the same value of covariate s as patient i, and have been given treatment 1:

$$n_1(z_{i,s}) = \sum_{j=1}^{i-1} \mathbb{I}(z_{j,s} = z_{i,s})\mathbb{I}(t_j = 1). \tag{2.5}$$

Then, the score for treatment 1, $S(1)$, is defined as the sum of $n_1(z_{i,s})$ across all covariates:

$$S(1) = \sum_{s=1}^{k} n_1(z_{i,s}). \tag{2.6}$$

We calculate the score for treatment -1, $S(-1)$, in an analogous way:

$$n_{-1}(z_{i+1,s}) = \sum_{j=1}^{i-1} \mathbb{I}(z_{j,s} = z_{i,s})\mathbb{I}(t_j = -1) \tag{2.7}$$

$$S(-1) = \sum_{s=1}^{k} n_{-1}(z_{i,s}) \tag{2.8}$$

We wish to allocate the treatment to patient i in way such that the score is minimized. Thus we use the following rule to assign patients to treatments:

- if $S(1) < S(-1)$, assign patient i to treatment 1 with probability p,
- if $S(1) > S(-1)$, assign patient i to treatment 1 with probability $1-p$,
- if $S(1) = S(-1)$, assign patient i to treatment 1 with probability $\frac{1}{2}$,

for $\frac{1}{2} < p < 1$. The value $p = \frac{2}{3}$ is often used in practice. Pseudocode for the minimization algorithm is given in Algorithm 1. Since there are several imbalance measures that may be used to calculate the score (such as alternatives for continuous covariates in Section 2.5.2), a function which calculates imbalance, *imb*, has been included as an argument of the algorithm.

Algorithm 1 Minimization: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$, imbalance measure imb, and probability p, where $\frac{1}{2} < p < 1$. By default, $p = \frac{2}{3}$.

function MINIMIZE($\boldsymbol{Z}_n, imb, p$)
  $n \leftarrow$ length of vector $\boldsymbol{z}_1$ ▷ Number of patients
  $\boldsymbol{t} \leftarrow$ empty vector
  $\boldsymbol{Z} = [\boldsymbol{z}_1 ... \boldsymbol{z}_s ... \boldsymbol{z}_k]$
  $t_1 \leftarrow$ randomly select 1 or -1

  for i in 2:n **do**
    Calculate $S(1) = imb(\boldsymbol{Z}[1:i,], \boldsymbol{t})$
    Calculate $S(-1) = imb(\boldsymbol{Z}[1:i,], \boldsymbol{t})$

    if $S(1) > S(-1)$ **then**
      $t_i = 1$ with probability p
    else if $S(1) < S(-1)$ **then**
      $t_i = 1$ with probability $1 - p$
    else
      Assign $t_i = 1$ with probability $\frac{1}{2}$
    end if
    Append t_i to $\boldsymbol{t}$
  end for

  $\boldsymbol{X} = [\mathbf{1}\ \boldsymbol{Z}\ \boldsymbol{t}]$ ▷ Design matrix

  return $\boldsymbol{X}$

end function

2.5.2 Continuous covariates

The method of computing the scores defined in Equations (2.6) and (2.8) is appropriate only when covariates are binary. It can be easily generalized for discrete covariates with more than two levels. For continuous covariates, one can discretize each covariate by using a pre-defined threshold. This is usually what happens in practice in social and biomedical experiments (Hu and Hu, 2012). Another option is to discretize the covariate into two sets dynamically, where for each i in $2 \leq i \leq n$, one chooses the median of the covariate values observed so far as the threshold. However, one can argue that discretization results in the loss of information, and, further, choosing the threshold in an arbitrary way can introduce bias. For this reason, Hu and Hu (2012) suggested alternative approaches to calculating $Score(1)$ and $Score(-1)$ based on imbalance measures that are appropriate for the continuous case. Their two suggestions, K-S and Max.Imb, can be used when there is a single continuous covariate.

Kolmogorov-Smirnov (K-S) distance

The first suggestion by Hu and Hu (2012) was an imbalance measure based on the Kolmogorov-Smirnov (K-S) distance between two distribution functions. Suppose that patient i receives treatment t, for $t \in \{-1, 1\}$. We define by $\hat{F}_1^t$ the empirical distribution function of the covariate values for patients 1 up to i that received treatment 1, assuming that patient i received treatment t. Similarly, we define $\hat{F}_{-1}^t$ as the empirical distribution function of the covariate values for patients that received treatment -1. Let us denote the range of the covariate by $\mathcal{Z}$. The value of $Score(1)$ is given by the K-S distance between $\hat{F}_1^1$ and $\hat{F}_{-1}^1$:

$$Score(1) = \sup_{z \in \mathcal{Z}} \left| \hat{F}_1^1(z) - \hat{F}_{-1}^1(z) \right|. \tag{2.9}$$

Analogously, $Score(-1)$ is given by:

$$Score(-1) = \sup_{z \in \mathcal{Z}} \left| \hat{F}_1^{-1}(z) - \hat{F}_{-1}^{-1}(z) \right|. \tag{2.10}$$

This is the largest vertical distance between the two distribution functions. In Figure 2.1, we illustrate how the K-S distance is obtained. This imbalance measure focuses on the distributions of the covariates only.

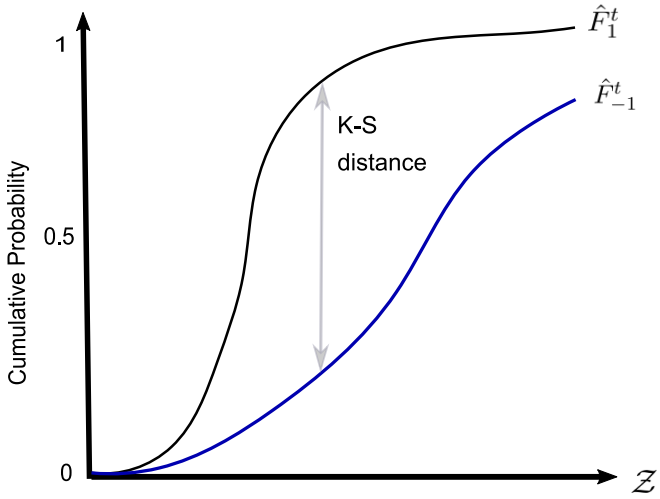

Figure 2.1: To obtain the Kolmogorov-Smirnov distance, the cumulative distribution functions $\hat{F}_1^t$ and $\hat{F}_{-1}^t$ are plotted, and the largest vertical distance between them is measured.

Maximum imbalance (Max.imb)

A second suggestion by Hu and Hu (2012) was called Maximum imbalance (Max.imb). Suppose we have an interval $\mathcal{I}$ in the range of the covariate values $\mathcal{Z}$. Let us denote by $\mathcal{U}$ the set for the indicies for the patients in $\mathcal{I}$. We denote by $\Delta(\mathcal{I}, \boldsymbol{t})$ the absolute difference in the number of patients in the two treatment groups with covariate values in $\mathcal{I}$:

$$\Delta(\mathcal{I}, \boldsymbol{t}) = \left| \sum_{j \in \mathcal{U}} \mathbb{I}(t_j = 1) - \mathbb{I}(t_j = -1) \right|. \tag{2.11}$$

We illustrate in Figure 2.2 how $\Delta(\mathcal{I}, \boldsymbol{t})$ is obtained for a small example. Now, given that we have allocated treatments to $i-1$ patients so $\boldsymbol{t}_{i-1}$ is known, suppose that we assign $t_i = 1$. Then, we define $Score(1)$ for patient i as:

$$Score(1) = \sup_{\mathcal{I} \in \mathcal{Z}} \{\Delta(\mathcal{I}, \boldsymbol{t}_i)\}. \tag{2.12}$$

We define $Score(-1)$ analogously, where we assume $t_i = -1$. The score is essentially the maximum absolute difference over all possible intervals after i units have been assigned a treatment. Hu and Hu (2012) demonstrated that this is a compromise between Efron's biased coin, which is concerned purely with treatment replication, and the K-S measure, which is concerned with the shape of the distribution. The Max.imb approach depends only on the ranks of the covariate values.

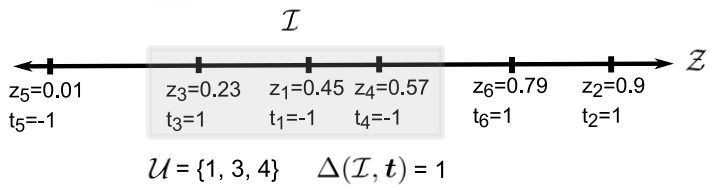

Figure 2.2: This small example illustrates how the value $\Delta(\mathcal{I}, \boldsymbol{t})$ is obtained. Here, there are six patients enrolled and their ordered covariate values along the range $\mathcal{Z}$ are indicated. The treatment assigned to each patient is also indicated. The grey shaded area is an interval $\mathcal{I}$, and associated with it is the set $\mathcal{U} = \{1, 3, 5\}$ of indices for the patients in $\mathcal{I}$. We obtain $\Delta(\mathcal{I}, \boldsymbol{t})$ by finding the absolute difference in the number of patients assigned to treatment 1 versus treatment -1, which in this case is 1.

Scott et al. (2002) stated that advantages of minimization are that it helps investigators to carefully think about prognostic factors before the study begins, and it can include more covariates than in a permuted block approach. There is also no need to split the sample into a large number of strata. One disadvantage is that the added complexity in the design

may harm recruitment. Further, two main criticisms of minimization, by Atkinson (1999) and Senn et al. (2010), are that, firstly, minimization considers only the marginal balance of covariates and any interdependence between them are ignored, and secondly, that the imbalance measures are somewhat ad-hoc and not theoretically grounded.

2.6 Optimal design

A different perspective on sequential treatment allocation using covariate information is given by Atkinson (1999), using optimal design theory. We introduce some of the main ideas in this section. Initially, in order to introduce the notion of D- and D_A-optimality, we assume that we wish to allocate treatments in a non-sequential setting. For simplicity, we omit subscripts and write $\boldsymbol{z}_s, \boldsymbol{t}$ and $\boldsymbol{y}$ to mean $\boldsymbol{z}_{s,n}, \boldsymbol{t}_n$ and $\boldsymbol{y}_n$, respectively. Thus the vectors $\boldsymbol{z}_s$ are known for all $s = 1, ..., k$, and we wish to determine the best choice of $\boldsymbol{t}$.

The relationship between the treatments, covariates and the responses can be framed as a statistical model. The goal of an experiment is often to estimate one or more unknown parameters of this model as accurately as possible. Optimal design methodology provides ways to select a design of an experiment that achieves this goal.

It is often plausible that, perhaps after some transformation, the responses, given some fixed factors, are approximately normally distributed. We specify the following linear model for the response, where we assume for simplicity that there are no interactions:

$$\begin{aligned} \boldsymbol{y} &= \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} && (2.13)\\ &= \overline{\boldsymbol{Z}}\boldsymbol{\gamma} + \boldsymbol{t}\alpha + \boldsymbol{\epsilon} && (2.14)\\ &= \mathbf{1}\gamma_0 + \boldsymbol{z}_1\gamma_1 + \boldsymbol{z}_2\gamma_2 + ... + \boldsymbol{z}_k\gamma_k + \boldsymbol{t}\alpha + \boldsymbol{\epsilon}, && (2.15) \end{aligned}$$

where the n-vector $\boldsymbol{\epsilon}$ contains mutually uncorrelated, normally distributed errors with constant variance σ^2. The $n \times q$ matrix $\boldsymbol{X}$ is the design matrix which can be partitioned as $\boldsymbol{X} = \left(\overline{\boldsymbol{Z}} \;\; \boldsymbol{t}\right)$ where $\overline{\boldsymbol{Z}} = (\mathbf{1} \;\; \boldsymbol{Z})$. The q-vector $\boldsymbol{\beta}$ contains the unknown parameters that we wish to study, and it can be partitioned into parameters that are not associated with treatments, $\boldsymbol{\gamma}$, and the treatment parameter α, such that $\boldsymbol{\beta} = (\boldsymbol{\gamma} \;\; \alpha)^\intercal$. We focus here on the case of the additive model where there are no interactions; however, in later chapters, we consider examples where there are treatment-covariate interactions.

The best linear unbiased estimator for our q unknown parameters $\boldsymbol{\beta}$ is given from the Gauss-Markov Theorem (Casella and Berger, 2002, p. 548),

$$\hat{\boldsymbol{\beta}} = (\boldsymbol{X}^\intercal \boldsymbol{X})^{-1} \boldsymbol{X}^\intercal \boldsymbol{y}.$$

Now, we denote the information matrix of the design by $\boldsymbol{M}$, where $\boldsymbol{M} = \boldsymbol{X}^\intercal \boldsymbol{X}$. The information matrix is linked to both the variance of the least squares estimator $\hat{\boldsymbol{\beta}}$ and the variance of the predicted response $\hat{y}(\boldsymbol{x})$ of a patient with covariate values $\boldsymbol{z}$ and assigned treatment t such that $\boldsymbol{x} = (1 \;\; \boldsymbol{z}^\intercal \;\; t)^\intercal$:

$$\mathbb{V}\text{ar } \hat{\boldsymbol{\beta}} = \sigma^2 \boldsymbol{M}^{-1}, \tag{2.16}$$

and,

$$\mathbb{V}\text{ar } \hat{y}(\boldsymbol{x}) = \sigma^2 \boldsymbol{x}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x}. \tag{2.17}$$

By choosing treatments in a way that optimizes some function of the design matrix $\Psi(\boldsymbol{X})$, we can minimize the variances given in either Equations (2.16) and (2.17).

2.6.1 Optimality criteria

D-optimum designs minimize the generalized variance of the estimator of $\hat{\boldsymbol{\beta}}$. That is, the volume of the confidence ellipsoid for $\boldsymbol{\beta}$ is minimized (Atkinson, 1999). Such a design is constructed by minimizing the objective function:

$$\Psi_D(\boldsymbol{X}) = \left|(\boldsymbol{X}^\intercal \boldsymbol{X})^{-1}\right| \tag{2.18}$$

$$= \left|\boldsymbol{M}^{-1}\right|. \tag{2.19}$$

Equivalently, one can maximize $|\boldsymbol{M}|$.

Using the partitioning of $\boldsymbol{X}$ into the treatment and covariate parts, the matrix $\boldsymbol{M}$ can be expressed as follows:

$$\boldsymbol{M} = (\boldsymbol{X}^\intercal \boldsymbol{X})^{-1} = \begin{pmatrix} \overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}} & \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t} \\ \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} & \boldsymbol{t}^\intercal \boldsymbol{t} \end{pmatrix}. \tag{2.20}$$

Using standard results on the inverse of partitioned matrices, we obtain:

$$\boldsymbol{M}^{-1} = (\boldsymbol{X}^\intercal \boldsymbol{X})^{-1} = \begin{pmatrix} \overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}} & \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t} \\ \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} & \boldsymbol{t}^\intercal \boldsymbol{t} \end{pmatrix}^{-1} \tag{2.21}$$

$$= (\boldsymbol{X}^\intercal \boldsymbol{X})^{-1} = \begin{pmatrix} \overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}} & \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t} \\ \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} & n \end{pmatrix}^{-1} \tag{2.22}$$

$$= \begin{pmatrix} \left(\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}} - \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}}\right)^{-1} & -(\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t} \boldsymbol{F}^{-1} \\ -\boldsymbol{F}^{-1} \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1}) & \boldsymbol{F}^{-1} \end{pmatrix}, \tag{2.23}$$

where $\boldsymbol{F}^{-1} = \left(n - \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t}\right)^{-1}$.

A D-optimal design minimizes the determinant of the above matrix. Such a design can be constructed using the exchange algorithm (Goos and Jones, 2011, p.36). In this algorithm, one wishes to find the optimal treatment allocation for column $\boldsymbol{t}$ of the design matrix $\boldsymbol{X}$. The algorithm starts by selecting treatment levels for $\boldsymbol{t}$ at random. Then, working through each entry i of this column for $i \in \{1, 2, ..., n\}$ the algorithm calculates $|\boldsymbol{M}^{-1}|$ when $t_i = 1$ and $t_i = -1$. The treatment level which leads to the minimum value of $|\boldsymbol{M}^{-1}|$ is selected. This produces a new design matrix with improved treatment allocation. Then, starting again at the top of the column $\boldsymbol{t}$, we calculate $|\boldsymbol{M}^{-1}|$ for treatment 1 and -1, and the optimal treatment is chosen. This procedure is repeated until no change is observed between passes of this loop. Finally, the whole procedure needs to be repeated a number of times with different random starting values, so it is more likely that the optimum found is global rather than local. Pseudo-code for this algorithm is given in Algorithm 2.

In assessing a particular design matrix $\boldsymbol{X}$, it can be helpful to compare it to another design $\boldsymbol{X}^*$. We define the relative D-efficiency of a design $\boldsymbol{X}$ relative to $\boldsymbol{X}^*$as

$$\text{Eff}_D = \left\{ \frac{\Psi_D(\boldsymbol{X}^*)}{\Psi_D(\boldsymbol{X})} \right\}^{1/p} = \left\{ \frac{\left|(\boldsymbol{X}^\intercal \boldsymbol{X})^{-1}\right|}{\left|(\boldsymbol{X}^{*\intercal} \boldsymbol{X}^*)^{-1}\right|} \right\}^{1/p}, \tag{2.24}$$

where we take the pth root so that the efficiency measure has the dimensions of a variance (Atkinson et al., 2007, p. 151).

Algorithm 2 Exchange Algorithm: function returns the design matrix where the treatments are allocated nonsequentially based on the D-optimality criterion. The arguments are the covariate values for the n patients, $\boldsymbol{Z}_n$ and m, the number of repetitions from different starting designs.

```
1: function COORD.EX(Z_n, m)
2:     n ← number of rows of Z_n                          ▷ number of units
3:     for j in 1 to m do
4:         t ← an n-vector of a random sequence of 1s and -1s
5:         Construct Xj according to the model
6:
7:         for i in 1 to n do
8:             X ← Xj
9:             Set t_i in X to 1
10:            m_plus ← |X^T X|
11:            Set t_i in X to -1
12:            m_minus ← |X^T X|
13:
14:            if m_plus > m_minus then
15:                Set t_i in X to 1
16:            else
17:                Set t_i in X to -1
18:            end if
19:        end for
20:
21:        if X == Xj then
22:            Break
23:        else
24:            Xj ← X
25:            store Xj in a list Xlist
26:        end if
27:    end for
28:    return the matrix in Xlist with largest value of |X^T X|
29: end function
```

The criterion of D-optimality assumes that one wishes to estimate all parameters as precisely as possible. However, one may have interest only in a subset of the parameters, or in some linear combination of them. In this case, the D_A-optimality criterion may be more appropriate, which is designed to be optimal for estimating m linear combinations in $\boldsymbol{\beta}$, which can be expressed as $\boldsymbol{A}^\intercal\boldsymbol{\beta}$, where $\boldsymbol{A}$ is a $q \times m$ matrix with $m < q$ (Atkinson et al., 2007, p.137). Each column of $\boldsymbol{A}$ gives the coefficients of a linear combination of interest. The criterion of D_A-optimality requires that $\Psi_{D_A}(\boldsymbol{X})$ is minimized (Atkinson, 1999):

$$\Psi_{D_A}(\boldsymbol{X}) = |\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A}|. \tag{2.25}$$

Relative D_A-efficiency of a design $\boldsymbol{X}$ relative to another design $\boldsymbol{X}^*$ is defined as

$$\text{Eff}_{D_A} = \left\{ \frac{\Psi_{D_A}(\boldsymbol{X}^*)}{\Psi_{D_A}(\boldsymbol{X})} \right\}^{1/b}, \tag{2.26}$$

where b is the number of non-zero rows in $\boldsymbol{A}$.

A special case of D_A-optimality is D_s-optimality, where interest lies in a subset of the parameters $\boldsymbol{\beta}$ (Cox and Reid, 2000, p.177). In this case, the matrix A is a $q \times q$ diagonal matrix were the diagonal entry is set to 1 if the parameter is of interest and zero otherwise.

In our setting, we wish to precisely estimate the treatment effect α. We focus only on the treatment effect and regard the parameters associated with the intercept and covariates $\gamma_1, ..., \gamma_k$ as nuisance parameters, so $\boldsymbol{A}^\intercal$ is a vector containing 1 in the last entry for the treatment effect, and $k+1$ preceding zeros: $\boldsymbol{A}^\intercal = (0 \;\; 0 \;\; 0 \;\; ... \;\; 1)$. For our choice of $\boldsymbol{A}$, the D_A-optimality criterion of minimizing $\left|\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A}\right|$ reduces to minimizing the bottom right entry of the matrix given in Equation (2.23):

$$\left|\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A}\right| = \frac{1}{n - \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t}}. \tag{2.27}$$

Minimizing the above quantity also minimizes the variance of the estimator of α, which is given by:

$$\mathbb{V}\text{ar }\hat{\alpha} = \frac{\sigma^2}{n - \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t}}. \tag{2.28}$$

We note that Atkinson (1999) referred to the following scalar as the loss after n patients, $\mathcal{L}_n$:

$$\mathcal{L}_n = \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}} (\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}})^{-1} \overline{\boldsymbol{Z}}^\intercal \boldsymbol{t}, \tag{2.29}$$

which is a measure of the confounding of the treatment with the covariates (through the measure of orthogonality given by the inner product $\boldsymbol{t}^\intercal \overline{\boldsymbol{Z}}$), and is also a measure of the interdependence between the covariates (through $\overline{\boldsymbol{Z}}^\intercal \overline{\boldsymbol{Z}}$). A D_A-optimal design minimizes the loss after n patients, and we will see in Section 2.6.2 that the loss can be used as an imbalance measure of a design. The co-ordinate exchange algorithm (Algorithm 2) can be modified to construct a D_A-optimal design (Goos and Jones, 2011, p.205).

Two other common criteria for optimality are A- and G-optimality. The A-optimal objective function aims to minimize the average variance of the parameter estimates:

$$\Psi_A(\boldsymbol{X}) = tr\left(\boldsymbol{M}^{-1}\right). \tag{2.30}$$

We define relative A-efficiency of a design $\boldsymbol{X}$ relative to another design $\boldsymbol{X}^*$ as

$$\text{Eff}_A = \frac{\Psi_A(\boldsymbol{X}^*)}{\Psi_A(\boldsymbol{X})}. \tag{2.31}$$

The G-optimal objective function is appropriate if the goal of the experiment is to find a good prediction model. It seeks to minimize the maximum prediction variance over a given set of design points $\boldsymbol{x}$ by minimizing the objective function:

$$\Psi_G(\boldsymbol{X}) = \max_{\boldsymbol{x}} \left\{\mathbb{V}\text{ar}\ \hat{y}(\boldsymbol{x})\right\}. \tag{2.32}$$

Relative G-efficiency is defined analogously to relative A-efficiency:

$$\text{Eff}_G = \frac{\Psi_G(\boldsymbol{X}^*)}{\Psi_G(\boldsymbol{X})}. \tag{2.33}$$

There is a close link between D- and G-optimal designs; the General Equivalence Theorem by Kiefer and Wolfowitz (1960) states that a design ξ which is D-optimal, is also G-optimal (Atkinson et al., 2007, p.122). We consider A- and G-optimal designs in more detail in, for example, Section 3.2.

2.6.2 Sequential optimal design approach

Atkinson (1982) proposed a method to construct D_A-optimal designs sequentially. Let us denote by $\boldsymbol{M}_i$ the information matrix for a design when patients $\{1, ..., i\}$ have been included in the experiment. We can express the determinant of the information matrix $|\boldsymbol{M}_i|$ using an updating formula (Rao, 1973, p.32), where $\boldsymbol{x}_i = (1 \ \ \boldsymbol{z}_i^\intercal \ \ t_i)^\intercal$ is the ith row of $\boldsymbol{X}_i$:

$$\begin{aligned}
|\boldsymbol{M}_i| &= |\boldsymbol{X}_i^\intercal \boldsymbol{X}_i| \\
&= \left|\boldsymbol{X}_{i-1}^\intercal \boldsymbol{X}_{i-1} + \boldsymbol{x}_i \boldsymbol{x}_i^\intercal\right| \\
&= \left|\boldsymbol{X}_{i-1}^\intercal \boldsymbol{X}_{i-1}\right| \left(1 + \boldsymbol{x}_i \left(\boldsymbol{X}_{i-1}^\intercal \boldsymbol{X}_{i-1}\right)^{-1} \boldsymbol{x}_i^\intercal\right)
\end{aligned}$$

$$= |\boldsymbol{M}_{i-1}| \left(1 + \boldsymbol{x}_i (\boldsymbol{M}_{i-1})^{-1} \boldsymbol{x}_i^\intercal\right)$$
$$= |\boldsymbol{M}_{i-1}| \left(1 + \frac{1}{\sigma^2} \mathbb{V}\text{ar } \hat{y}(\boldsymbol{x}_i)\right).$$

Given a design for the first $i-1$ patients, the ith treatment should be chosen such that the prediction variance, $\mathbb{V}\text{ar } \hat{y}(\boldsymbol{x}_i)$, is maximized (see Algorithm 2).

Now, in order to construct a D_A-optimal design sequentially, Atkinson (1982) proposed an objective function $d(\boldsymbol{z}_i, t_i)$ which is analogous to this prediction variance. We first rely on the result that, for the objective function given by $\Psi_{D_A}(\boldsymbol{M}_i) = \log \left|\boldsymbol{A}^\intercal \boldsymbol{M}_i^{-1} \boldsymbol{A}\right|$, the derivative with respect to the information matrix is given by:

$$\frac{\partial \Psi_{D_A}}{\partial \boldsymbol{M}_i} = \boldsymbol{M}_i^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_i^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}_i^{-1}. \tag{2.34}$$

Note that we have written it as a function of the information matrix $\boldsymbol{M}_i$, rather than the design matrix as we have done previously, in order to ease notation. We use the logarithm of the criterion for D_A-optimality as the logarithm is a monotone function which does not affect the location of the extrema and also eases calculation. Further, the log-determinant is a convex function and has only one minimum. The proof of this result is provided in Appendix A. Now, Atkinson (1982) denotes by $d(\boldsymbol{z}_i, t)$ the following quantity related to the derivative of Ψ_{D_A}:

$$d(\boldsymbol{z}_i, t) = \boldsymbol{x}_i^\intercal \frac{\partial \Psi_{D_A}}{\partial \boldsymbol{M}_{i-1}} \boldsymbol{x}_i$$
$$= \boldsymbol{x}_i^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}_i \tag{2.35}$$
$$= \begin{pmatrix} 1 & \boldsymbol{z_i}^\intercal & t \end{pmatrix} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \begin{pmatrix} 1 & \boldsymbol{z_i}^\intercal & t \end{pmatrix}^\intercal. \tag{2.36}$$

We wish to choose t for patient i to maximize the objective function $d(\boldsymbol{z}_i, t)$. It can be shown that, for the case where interest lies in estimating the treatment effect α precisely so that $\boldsymbol{A}^\intercal = (0\ 0\ 0\ \ldots\ 1)$, maximizing $d(\boldsymbol{z}_i, t)$ is equivalent to minimizing the loss after i trials, $\boldsymbol{t}^\intercal \overline{\boldsymbol{Z}}_i (\overline{\boldsymbol{Z}}_i^\intercal \overline{\boldsymbol{Z}}_i)^{-1} \overline{\boldsymbol{Z}}_i^\intercal \boldsymbol{t}$. The claim by Senn et al. (2010) that the sequential optimal design approach is equivalent to minimizing $\mathcal{L}_i$ (see Equation (2.29)) is proven in Appendix B.

Now, the treatment allocation method advocated by Atkinson (1982) states that, for patient i, the quantities $d(\boldsymbol{z}_i, 1)$ and $d(\boldsymbol{z}_i, -1)$ should be calculated, and treatment 1 is assigned with probability given by:

$$\frac{d(\boldsymbol{z}_i, 1)}{d(\boldsymbol{z}_i, 1) + d(\boldsymbol{z}_i, -1)}. \tag{2.37}$$

This allows us to add a component of randomness to the allocation scheme, which reduces the possibility of introducing selection bias. Pseudocode for the algorithm is provided in Algorithm 3. We note that, for the initial stages of this algorithm where $i < k + 2$, it is not possible to compute the inversions of $\boldsymbol{M}_{i-1}$ required in Equation (2.36). Regularization may be used to resolve this; we replace $\boldsymbol{M}_{i-1}$ by $(\boldsymbol{M}_{i-1} + \epsilon \boldsymbol{I}_{j\times j})$. By adding a negligible amount of noise to the diagonal, we can ensure that the matrix is invertible without noticeably affecting the results. In our simulations, we choose $\epsilon = 0.0001$.

Algorithm 3 Atkinson's sequential D-optimal design: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$, and a matrix $\boldsymbol{A}$ indicating the linear combinations of the parameters of interest.

function ATKINSON($\boldsymbol{Z}_n, \boldsymbol{A}$)
  $n \leftarrow$ number of row in $\boldsymbol{Z}_n$ ▷ Number of patients
  $\boldsymbol{t} \leftarrow$ empty vector
  $t_1 \leftarrow$ randomly select 1 or -1

  for i in 2:n **do**
    $\boldsymbol{X} \leftarrow [\mathbf{1}\ \ \boldsymbol{Z}_i\ \ \boldsymbol{t}]$
    $\boldsymbol{M} \leftarrow \boldsymbol{X}^\intercal \boldsymbol{X}$
    $\boldsymbol{x} \leftarrow [1\ \boldsymbol{z}_i\ 1]^\intercal$
    $d(\boldsymbol{z}_i, 1) \leftarrow \boldsymbol{x}_i^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x}_i$
    $\boldsymbol{x} \leftarrow [1\ \boldsymbol{z}_i\ -1]^\intercal$
    $d(\boldsymbol{z}_i, -1) \leftarrow \boldsymbol{x}_i^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x}_i$
    $p \leftarrow \frac{d(\boldsymbol{z}_i, 1)}{d(\boldsymbol{z}_i, 1) + d(\boldsymbol{z}_i, -1)}$

    Assign $t_i = 1$ with probability p and $t_i = -1$ with probability $1 - p$
    Append t_i to $\boldsymbol{t}$
  end for

  $\boldsymbol{X} = [\mathbf{1}\ \boldsymbol{Z}\ \boldsymbol{t}]$ ▷ Design matrix

  return $\boldsymbol{X}$

end function

2.7 Other methods

Re-randomization is an additional method that is used to try to achieve covariate imbalance in non-sequential settings. Treatments are randomly allocated to units and if it appears that covariate balance is not acceptable according to a predefined measure, it is necessary to re-do

the randomization. When covariate balance is at an acceptable level, the experiment is conducted (Morgan and Rubin, 2015). This method is, however, controversial; re-randomization can change the distribution of the test statistic and make common statistical procedures such as the t-test conservative or invalid (Morgan and Rubin, 2012).

Klotz (1978) outlined an approach to create a compromise between two conflicting goals of treatment assignment in clinical trials: firstly, the aim to make the assignment random and unpredictable, and secondly, the need for balance across strata. He defines an entropy measure, which quantifies the uncertainty pertaining to the probabilities for treatment assignment, and defines a measure of expected balance of treatments across strata. The construction of the design is framed as optimization problem, where one wishes to maximize the entropy, subject to a linear constraint on the expected balance.

Titterington (1983) further developed the work of Klotz (1978) and reformulated the problem as the minimization of the Kullback-Leibler distance between the vector of assignment probabilities and the equiprobable assignment vector.

2.8 Simulations

We wish to compare the performances of randomization, Efron's biased coin, minimization and the sequential optimal design approach. We run simulations to investigate how the methods compare in terms of balance in treatment, balance in covariates, and loss. We define the following three performance measures:

1. Treatment balance: we assess to what extent the methods are able to achieve equal replication of treatments. Since we assume that our treatment factor is binary, we define our performance measure as the proportion of patients who have been allocated treatment 1. More formally, we define by $\mathbf{1}_i^t$ an indicator vector for $\boldsymbol{t}_i$ which has jth entry equal to 1 when $t_j = 1$ and 0 otherwise. The treatment balance measure, denoted Ψ_t, is defined as:

$$\Psi_t = \frac{1}{i}\|\mathbf{1}_i^t\| \,, \tag{2.38}$$

   where $\|\cdot\|$ denotes the 1-norm.

2. Treatment-covariate balance: we consider the extent to which the allocation methods lead to a design where the treatment is orthogonal to each of the covariates. We define

the covariate balance measure Ψ_z as:

$$\Psi_z = \frac{1}{i}\|\boldsymbol{t}_i^\intercal \boldsymbol{Z}_i\|, \tag{2.39}$$

where we divide by i so that we can make comparisons across different samples. Note that for a single covariate, we simply have:

$$\begin{aligned} \Psi_z &= \frac{1}{i}\|\boldsymbol{t}_i^\intercal \boldsymbol{z}_i\| \\ &= \frac{1}{i} \mid \boldsymbol{t}_i^\intercal \boldsymbol{z}_i \mid . \end{aligned} \tag{2.40}$$

3. Loss: We consider loss as a measure of optimality of the design, as established by Atkinson (1999). We define this performance measure Ψ_{loss} as:

$$\begin{aligned} \Psi_{loss} &= \frac{1}{i}\mathcal{L}_i \\ &= \frac{1}{i}\boldsymbol{t}_i^\intercal \overline{\boldsymbol{Z}}_i(\overline{\boldsymbol{Z}}_i^\intercal \overline{\boldsymbol{Z}}_i)^{-1}\overline{\boldsymbol{Z}}_i^\intercal \boldsymbol{t}_i, \end{aligned} \tag{2.41}$$

where we again divide by i so that we can make comparisons across different samples.

Our simulations have the following basic structure:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) Treatment allocation methods are used to generate designs sequentially.

(c) Designs are evaluated using the performance measures $\Psi_t, \Psi_z, \Psi_{loss}$ at each sample size.

II (a)-(c) is repeated above 100 times to obtain a distribution of the performance measure for each sample size.

2.8.1 Binary covariates

Firstly, we consider five covariates which are independent realizations from Bernoulli($p = 0.5$) distributions. We use the following allocation methods which are appropriate for multiple binary covariates:

1. Randomization

2. Efron's biased coin with $p = \frac{2}{3}$
3. Classic minimization with $p = \frac{2}{3}$
4. Sequential optimal design approach

Figure 2.3 displays show the performance of the four allocation methods. We have the distribution of the performance measures Ψ_t, Ψ_z and Ψ_{loss} for sample sizes between 1 and 100. The black line indicates the median, the dark grey is for the $40\% - 60\%$ region of the distribution and the light grey is the $10\% - 90\%$ region.

For the balance in treatment Ψ_t, we see in the first row of Figure 2.3 that the variability of the performance measure is reduced dramatically when Efron's biased coin is used instead of randomization. Since Efron's method is designed specifically to keep treatment replication balanced, this result is expected. The classic minimization approach appears to result in a greater variability than Efron's biased coin; this is also expected, since minimization aims to balance covariates as well as treatment replication. The sequential optimal design method results in greater variability than minimization or Efron's biased coin.

For treatment-covariate balance, as measured by Ψ_z, we observe in the second row of Figure 2.3 that randomization, Efron's biased coin and minimization perform similarly. For larger sample sizes, minimization appears to have a slight advantage over randomization and Efron's method. As minimization is designed to achieve treatment-covariate balance and not just balance in treatments, this is expected. Interestingly, we see that the sequential optimal design approach provides an even better performance, especially for small samples.

Senn et al. (2010) presented results from simulation studies for the loss. They produced plots to display the loss versus sample size for five binary covariates that are generated from the geometric distribution, with four different values of the success probability: 0.0001, 0.01, 0.5, and 0.8. They compared the randomized, minimization and sequential optimal design methods, and showed that the sequential optimal design method results in minimal loss. In our simulations, we additionally consider Efron's biased coin as an allocation method and show results of Ψ_t and Ψ_z as performance measures.

For our simulation study, we first consider five binary independent covariates. The results in the third row of Figure 2.3 indicate that the performance of the four methods are similar to that of covariate balance. It is clear that the sequential optimal design method outperforms the other methods. One can consider loss as measuring two components of treatment allocation: firstly, it is a measure of orthogonality between treatments and covariates, and

secondly, it measures the interdependence between the covariates. Since our covariates in this case have been generated independently, we expect the loss to behave similarly to the orthogonality measure for covariate balance: $\Psi_z = \frac{1}{i}\|\boldsymbol{t}^\intercal \boldsymbol{Z}_i\|$. Figure 2.4 displays the medians of the distributions of Ψ_z for all four methods on the left panel and the medians of the distributions of Ψ_{loss} on the right panel to ease comparison.

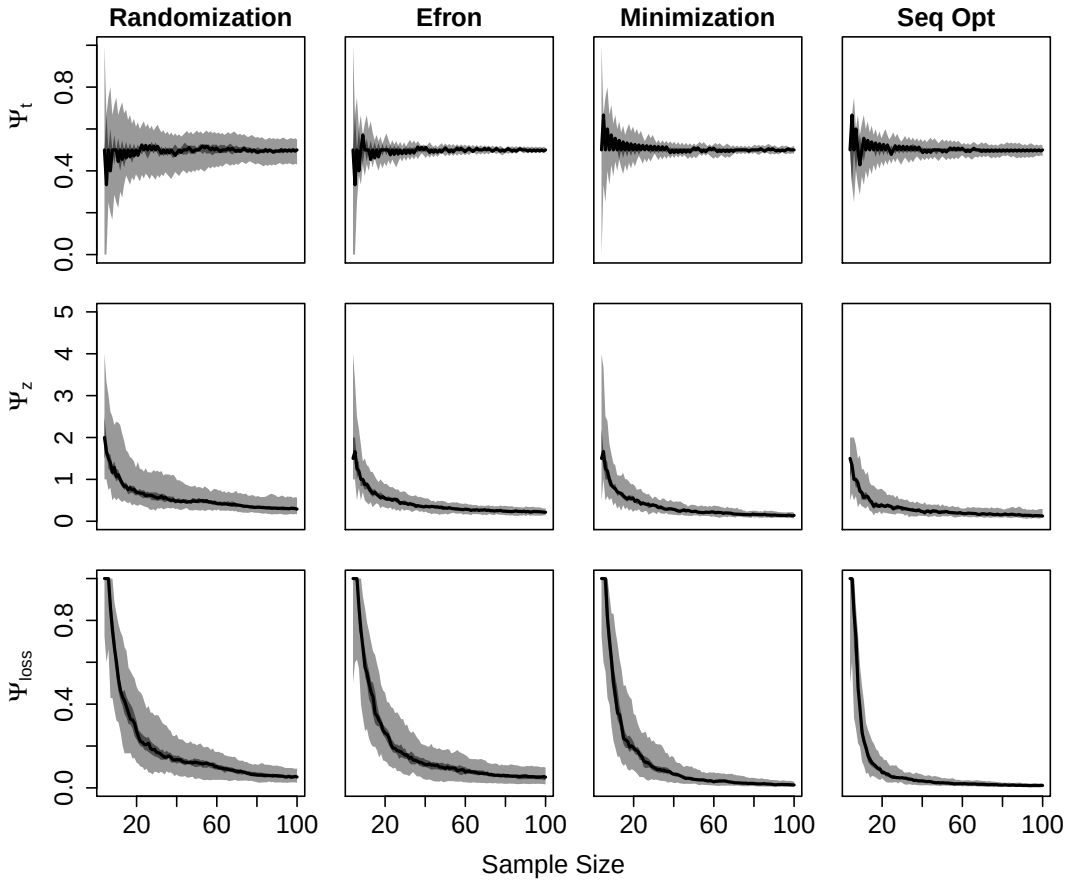

Figure 2.3: Distributions of Ψ_t, Ψ_z and Ψ_{loss} as sample size increases for five independent Bernoulli($p = 0.5$) covariates, based on 100 simulations. Four allocation methods are considered: Randomization, Efron's biased coin, the classic form of minimization, and the sequential optimal design method. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

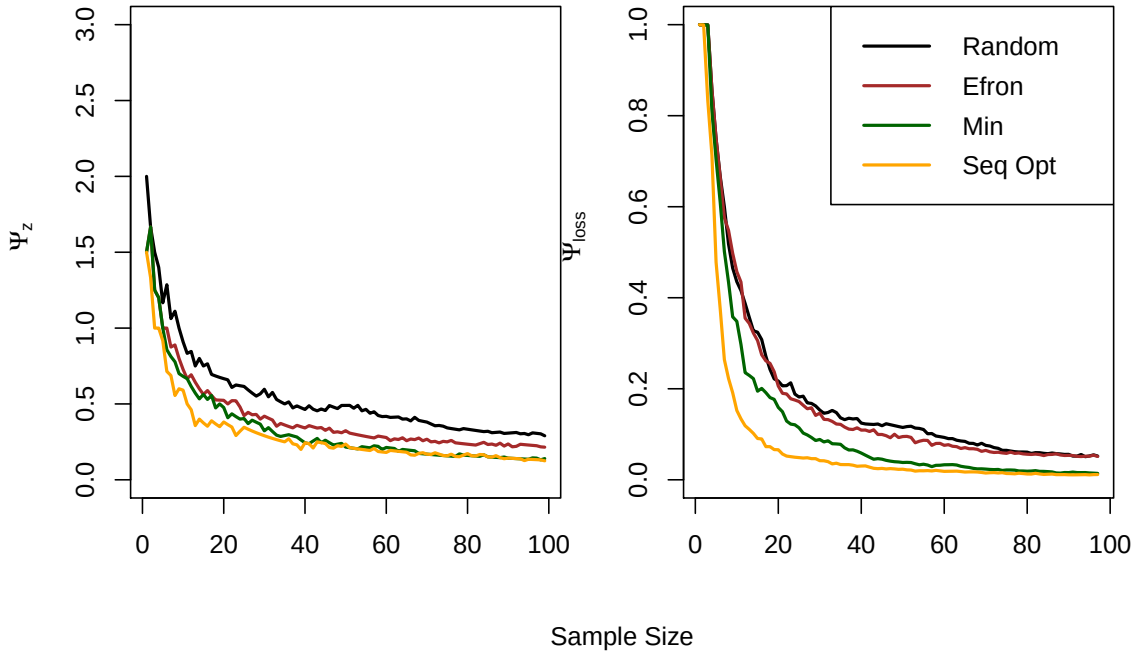

Figure 2.4: Medians of the distributions of Ψ_z and Ψ_{loss} as sample size increases for five independent Bernoulli($p = 0.5$) covariates, based on 100 simulations. Four allocation methods are considered: Randomization (black), Efron's biased coin (brown), the classic form of minimization (green), and the sequential optimal design method (yellow).

Additionally, we consider how the the performance of minimization and the sequential optimal design approach compare for covariate balance when the five covariates are highly correlated. We consider the same simulation procedure which produced Figure 2.3, but we generate the five covariates from Bernoulli($p = 0.5$) distributions with pairwise Pearson correlation 0.8, where the marginal probability that covariate s takes value one is equal to 0.5 for $s \in \{1, 2, ..., 5\}$. This is achieved by using the $\mathbb{R}$ package `bindata`. The results are shown in Figure 2.5. We find that minimization and the sequential optimal design approaches have very similar performance, both in terms of the variability and medians of the performance measure. This demonstrates the fact that the sequential optimal design approach is able to take into account the inter-dependence between the covariates; when covariates are highly correlated, the sequential optimal design and minimization methods perform similarly. However, when the covariates are independent, the sequential optimal design approach has an improved performance. Figure 2.6 displays the medians of the distributions of Ψ_z for all four methods on the left panel and the medians of the distributions of Ψ_{loss} on the right panel.

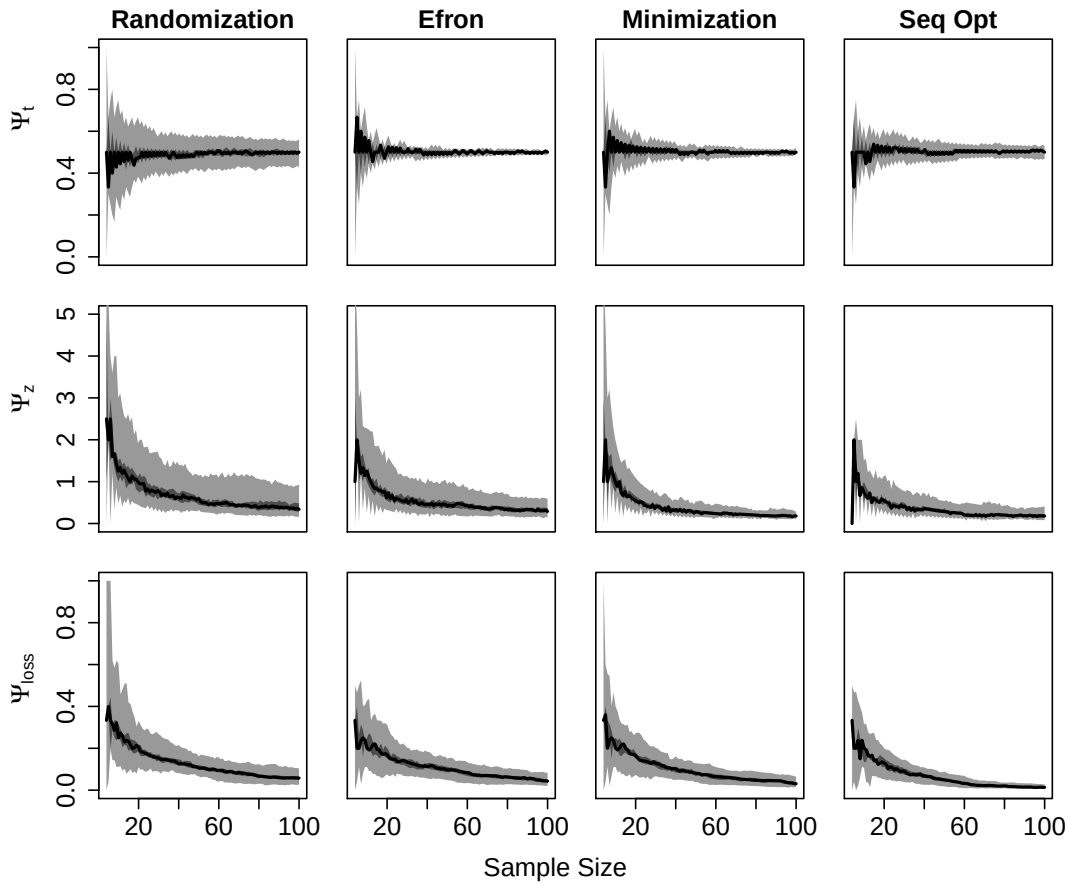

Figure 2.5: Distributions of Ψ_t, Ψ_z and Ψ_{loss} as sample size increases for five Bernoulli($p = 0.5$) covariates which have correlation 0.8, based on 100 simulations. Four allocation methods are considered: Randomization, Efron's biased coin, the classic form of minimization, and the sequential optimal design method. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

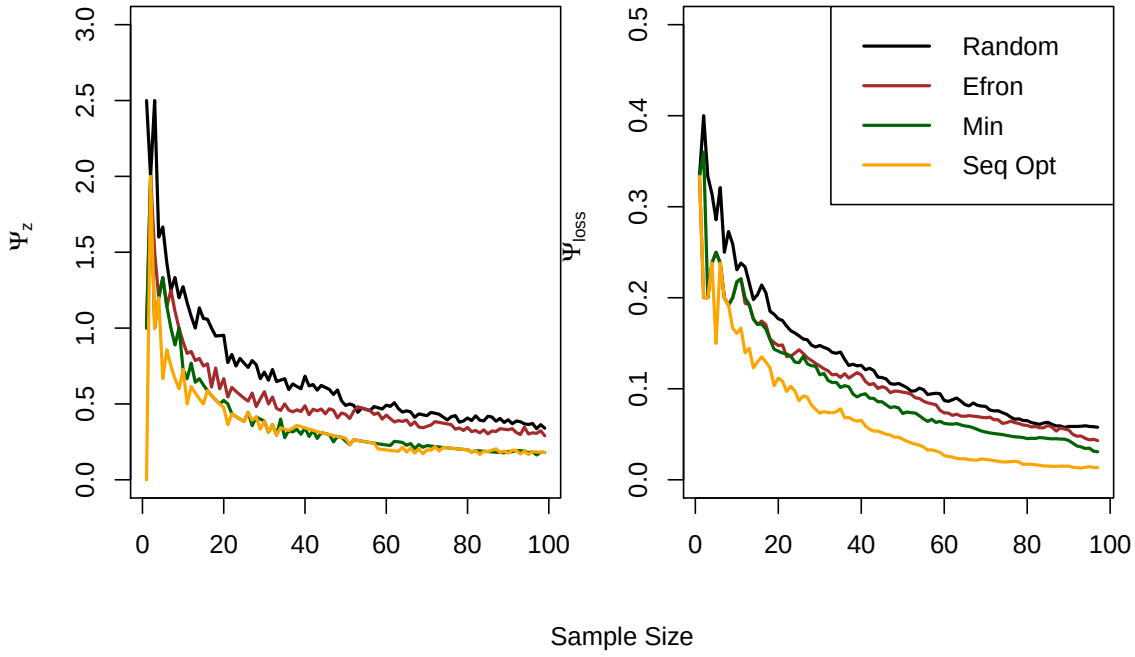

Figure 2.6: Medians of the distributions of Ψ_z and Ψ_{loss} as sample size increases for five Bernoulli($p = 0.5$) covariates which have correlation 0.8, based on 100 simulations. Four allocation methods are considered: Randomization (black), Efron's biased coin (brown), the classic form of minimization (green), and the sequential optimal design method (yellow).

The the two previous simulations, we have assumed that all covariates are generated with $p = 0.5$. In Figure C.1 in Appendix C, we consider the case where five covariates from Bernoulli($p = 0.8$) distributions are generated with correlation 0.8. The results are similar to that of Figure 2.5.

2.8.2 A single continuous covariate

We now consider the following six allocation methods which are appropriate for a single continuous covariate. We generate a single continuous covariate with distribution given by Unif($a = -\sqrt{3}, b = \sqrt{3}$):

1. Randomization
2. Efron's biased coin with $p = \frac{2}{3}$
3. Classic minimization with dynamic discretization at the median and $p = \frac{2}{3}$
4. Minimization with K-S measure of imbalance and $p = \frac{2}{3}$

5. Minimization with Max.imb measure of imbalance and $p = \frac{2}{3}$
6. Sequential optimal design approach

Figure 2.7 shows the results. For treatment balance, we see in the first row that the performance of the methods applicable to binary covariates appear similar to that in Figure 2.3. We note that the K-S minimization method appears to perform similarly to randomization; the variance is quite large and it takes longer than other approaches for the median to appear to converge to a 50-50 split. The Max.imb procedure appears to perform similarly to the classic minimization method, and finally, for a continuous covariate, the advantage of using the sequential optimal design method for small sample sizes is more obvious. The variability when sample size is less than 20 is reduced for the sequential optimal design method, compared to other methods.

For treatment-covariate balance, we observe in the second row of Figure 2.7 that the minimization approach of discretizing the covariate dynamically at the median appears not to improve Ψ_z compared to Efron's approach. This is likely to be because the performance measure is calculated using the continuous values of the covariate, and not in the discretized form. The Max.imb method performs similarly to the classic form of minimization. The K-S approach has poor performance and appears not to show much improvement over simple randomization. This is likely due to the fact that the Kolmogorov-Smirnov statistic is based only on the vertical distance between two distributions, which is perhaps too coarse of a summary statistic for our purposes. However, simulation studies by Hu and Hu (2012) show that if the K-S distance between the covariate distributions of the two treatment groups is used as the performance measure (which we have not done), the K-S form of minimization outperforms the classic and Max.imb forms of minimization, as well as Efron's biased coin. We note that the sequential optimal design method again shows effectiveness in ensuring treatment-covariate balance, with a particular improvement in small samples.

The third row of Figure 2.7 shows the distributions of Ψ_{loss} for a single continuous covariate. Again, the results are similar to that of covariate balance. We display the medians of Ψ_z and Ψ_{loss} of all six allocation methods in Figure 2.8.

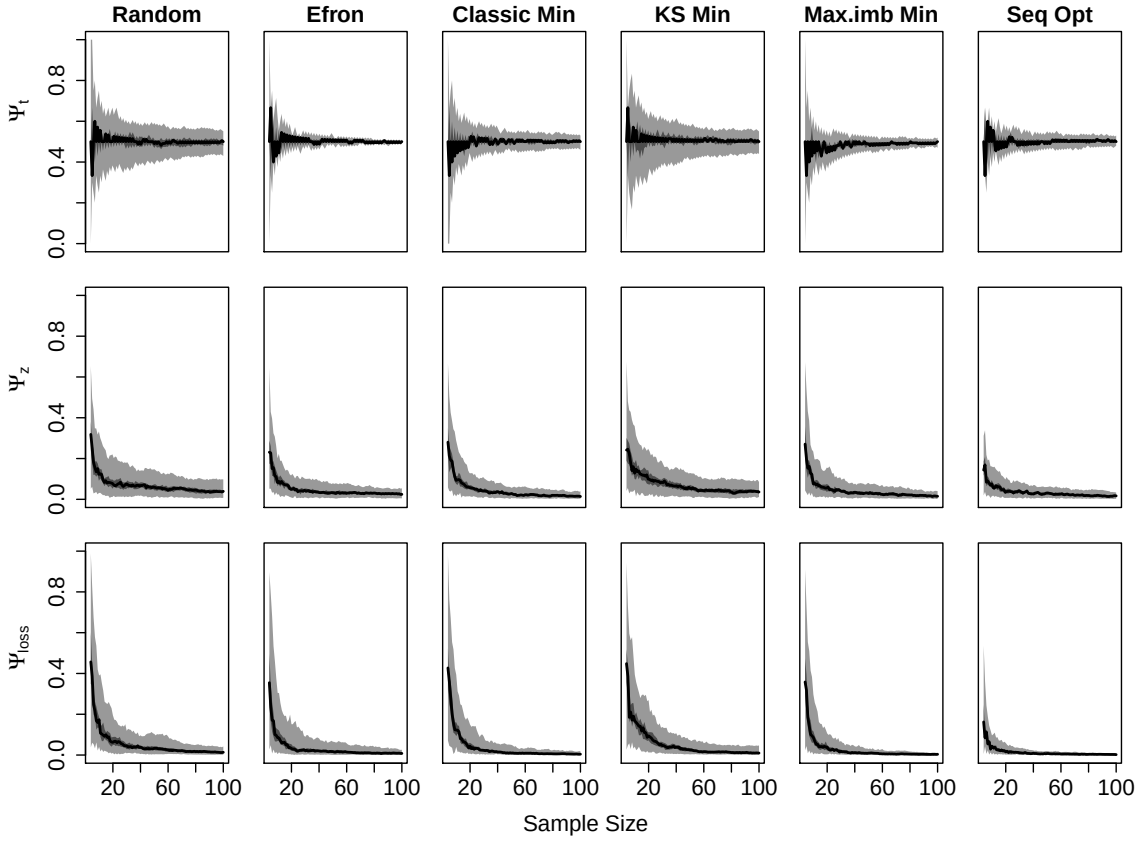

Figure 2.7: Distributions of Ψ_t, Ψ_z and Ψ_{loss} as sample size increases for 1 continuous Unif($a = -1.73, b = 1.73$) covariate, based on 100 simulations. Six allocation methods are considered: Randomization, Efron's biased coin, the classic form of minimization, minimization using the K-S measure, minimization using the Max.imb measure, and the sequential optimal design method. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

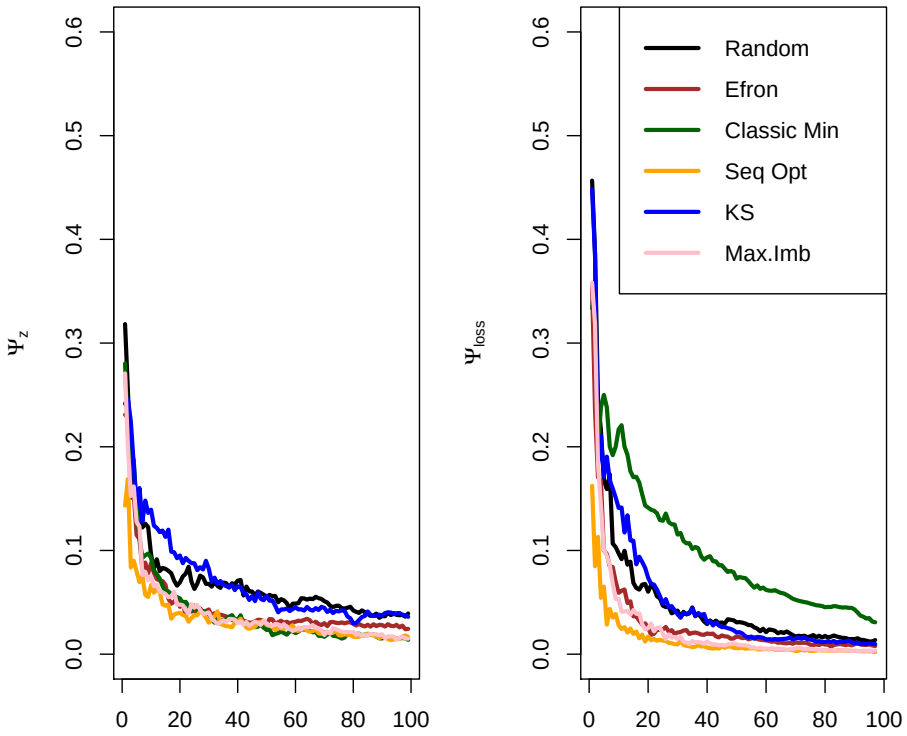

Figure 2.8: Medians of the distributions of Ψ_z and Ψ_{loss} as sample size increases for one Unif($a = -1.73, b = 1.73$) covariate, based on 100 simulations. Six allocation methods are considered: Randomization (black), Efron's biased coin (brown), the classic form of minimization (green), and the sequential optimal design method (yellow), the K-S form of minimization (blue) and the Max.imb form of minimization (pink).

In Figure C.3 in Appendix C, we consider the case of a skewed continuous covariate. We generate 100 observations from a Beta($\alpha = 5, \beta = 1$) distribution. The results are similar to that of Figure 2.7.

2.9 Conclusion

In this chapter, we described the setting of clinical trials then reviewed existing literature on sequential design of experiments with covariate information. We described simple approaches such as randomization and Efron's biased coin, and motivated the need for covariate adaptive approaches such as minimization and optimal design based methods. The methods we considered are myopic in the sense that decisions are made in order to optimize some kind of imbalance function or objective function at the current stage of the experiment.

Our contributions from this chapter are the proof of the claim by Senn et al. (2010) that

the optimal design approach is equivalent to minimizing $\mathcal{L}_i$, given in Appendix B. While Atkinson (1999) provided simulations comparing the performances in terms of the loss for randomization, Efron's biased coin, minimization and the sequential optimal design approach, we consider two other performance measures in our simulations: treatment balance and treatment-covariate balance. We also run simulations comparing these four approaches in the case of a continuous covariate, which has not been undertaken before.

We have seen from the simulation studies that, for both the binary and continuous cases, Efron's biased coin significantly improves balance in treatment over the randomization approach. Minimization does not appear to offer gain over Efron's approach. However, the sequential optimal design method appears to be more effective than Efron's biased coin for balancing treatment levels for sample sizes that are less than 20.

For the binary case, we observe that minimization is better than Efron's biased coin and randomization for treatment-covariate balance and loss. The sequential optimal design method provides a further improvement in both performance measures, especially for small samples. Further, we demonstrated that the sequential optimal design method shows improvement over the classic form of minimization when binary covariates are independently generated; when there are multiple binary covariates that are highly correlated, the performance of these two methods for treatment-covariate balance are very similar.

For the continuous case, the classic and Max.imb forms of minimization perform similarly to Efron's biased coin for covariate balance and loss. The K-S form of minimization overall shows poor performance, with results that appear worse than the Efron method for both performance measures. Finally, we note that the optimal design method appears to be the most effective of the approaches; again, this is most prominent for small sample sizes. Further, the optimal design method has the advantage that it is flexible in its requirements for the covariates; it can handle any number of continuous and discrete covariates, as well as interactions.

We note that we did not consider any interactions in this chapter, as minimization is not an approach that is able to take into account interactions. Later in this thesis, particularly in Chapter 6 and Chapter 7, treatment-covariate interactions play a crucial role. We propose extensions of the sequential optimal design approach in Chapter 4.

Chapter 3

Introduction to nonmyopic approaches

In the previous chapter, we assessed the impact on the performance of a design of applying a particular treatment to the current patient in the trial. This assumes that the current patient is the last patient in the experiment. We do not consider how the selected treatment may impact the objective function at a later point in time, when more patients have entered the trial. When taking a myopic or greedy approach, we make decisions that are optimal for the objective function at the current time without taking future decisions into acccount. Now, when taking a nonmyopic approach, one no longer assumes that the current patient is the last patient, and takes into account decisions that may be made about future possible patients in the trial. The number of future patients in the trial considered is called the *horizon*. It is possible that a decision is taken at point i which may be suboptimal for time i, but later leads to a more efficient result. We provide an overview of nonmyopic approaches in the literature and provide simulations which illustrate the potential for nonmyopic approaches in sequential design which we discuss in Chapters 4 and 5.

3.1 Nonmyopic methods in the literature

Having a nonmyopic approach to the treatment allocation problem means that the optimization involves multiple stages. Not only is it important to consider the impact of the decision at the time of patient i, but we consider future patients, possibly up to patient n. The state at stage i comprises the information that is known at that stage, which in our example includes the covariates of patients 1 up to i, as well as the treatments and responses of patients 1 up to $i-1$. The decision about the treatment t_i at stage i is made based on the state $S_{i+1} = (\boldsymbol{Z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_i)$. Based on that decision, there is a transition function f_i that

outputs the state of the next stage, $S_i = f_i(S_{i-1}, t_i)$. In our case, this transition function involves a linear or logistic model linking the responses to the treatments and covariates. There is a need to balance two conflicting aims in the decision making:

1. The aim to explore and to make decisions which may not be optimal given the current state, but may lead to gain of information which can inform future decision making.
2. The aim to exploit in order to make the best decision given the current state.

In this thesis, we will look at problems in which our objective function Ψ depends on the parameters $\boldsymbol{\beta}$ of the model for the response. Selecting a treatment for patient i which minimizes Ψ given our current estimate of $\boldsymbol{\beta}$ would be an exploitative decision; selecting a treatment for patient i which leads to gaining more information about $\boldsymbol{\beta}$ would be an exploratory decision. In CARAEE for clinical trials, one wishes to balance the aim of estimating the parameters (exploration) with the aim of giving patients the best possible treatment (exploitation).

In this section, we begin by providing an overview of dynamic programming, which is an approach to tackling multistage optimization problems such as nonmyopic methods. Nonmyopic approaches have been studied in a number of diverse contexts; we provide an overview of their use in control theory, clinical trials and also longitudinal studies.

3.1.1 Dynamic programming

Dynamic programming is an approach for solving multistage optimization problems (see, for example, Powell, 2009). The overall problem is broken into different stages, which often correspond to points in time, and each stage of the problem can be optimized. The key idea is that the overall sequence of decisions will be optimal for the entire problem (Bradley et al., 1977, p. 320). The subproblem at different stages may have similar structures; these similarities are leveraged to make the overall problem more tractable. The optimal design can be obtained by forward or backward induction. We focus on backward induction since it is the approach that is usually most appropriate in problems involving uncertainties (Bradley et al., 1977, p. 328). In backward induction, we start by finding the optimal decision at the end of the sequence of decisions, taking into account all possible treatments and covariates that may have been observed up until that point. Then, one can work backwards and obtain the optimal policy taking expectations of unknown quantities (Bradley et al., 1977, p.330).

There is no single formulation of a backward dynamic programming solution; we illustrate the overall approach with a simple generic example. Let us suppose that, at a particular stage i, we have the state $S_i = (\boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ and we consider the decision t_i for $t_i \in \{-1, 1\}$.

We denote here the objective function for selecting t_i as

$$\Psi(S_i, t_i), \tag{3.1}$$

which we want to minimize. In previous chapters, Ψ was written as a function of the design matrix. Given that a treatment t_i has been selected, there is a transition function f_i which outputs the state of the following stage:

$$S_{i+1} = f_i(S_i, t_i). \tag{3.2}$$

Now, suppose we wish to find the optimal treatment for a patient at some future stage N, for $N > i$. The optimal treatment $t_N^*(S_N)$, is given by

$$t_N^*(S_N) = \underset{t_N \in \{-1,1\}}{\operatorname{argmin}} \left\{ \mathbb{E}\left(\Psi\left(S_N, t_N\right)\right) \right\}, \tag{3.3}$$

subject to

$$S_m = f_m(S_{m-1}, t_m^*\left(S_m\right)), \tag{3.4}$$

$$t_m^*\left(S_m\right) \in \{-1, 1\}, \tag{3.5}$$

for $m \in \{i, ..., N\}$. In Equation (3.3), we take expectations over the unknown quantities in the state. Thus, considering all the possible configurations of the state S_N, one can work backwards using the recursion, which eventually provides us with an optimal decision for patient i, taking into consideration the future possible decisions for up to patient N.

Dynamic programming is computationally expensive when the number of stages is large, or the states are multidimensional due to the nested optimizations and expectations given in Equation (3.3). Further, when the response is continuous, Equation (3.3) requires nested evaluation of typically analytically intractable integrals which further adds computational complexity. Approximate dynamic programming refers to strategies aimed to reduce dimensionality and to make multistage optimization problems feasible in the face of these challenges (Powell, 2009). Common strategies include reducing the number of states via aggregation, using a sufficient statistic to summarize multidimensional states, and approximating the expectation via Monte-Carlo simulation.

Huan and Marzouk (2016) provided a recent overview of approximate dynamic programming in the context of Bayesian experimental design. A probabilistic model is given to describe the

relationship between the treatments and response. The state is comprised of the parameters of the model, and the posterior distributions of the parameters are updated at each stage using Bayes' rule. The model may be non-linear, so a response-adaptive approach is necessary to update current knowledge of the parameters. The state is multidimensional and is approximated by a grid representation to make dynamic programming computationally feasible. In this case, the utility function is the Kullback-Leiber divergence between the final posterior distribution to the prior. They provided an illustrative numerical example, as well as an example of the use of this methodology in an environmental experiment. In this second example, a chemical contaminant is exposed to the air unexpectedly. The chemical diffuses in the air and is also carried by the wind. A car or an aerial vehicle is sent to search for the source of the contaminant. The optimization problem at hand is choosing a series of locations for each time point that the vehicle should go to, in order to maximize expected information gain. We can see the analogy to our set-up: the selection of locations is the treatment that needs to be assigned for each time point. When the treatment is a location, it is continuous and potentially multi-dimensional, which makes calculating the expected values of the objective function more computationally involved. Knowledge about the wind current can be seen as a covariate; knowing the direction of the wind, or having a distribution for it for a few time points in advance can improve efficiency of the design. This is analogous to our problem where we have information about the covariates of a few future patients. Huan and Marzouk (2016) showed the advantage of taking a nonmyopic approach in this example; the nonmyopic approach is able to take into account the fact that wind will move the plume over time, whereas the greedy approach does not. We return to this example later in the thesis in light of our findings.

3.1.2 Clinical trials

In clinical trials applications, the need to consider both ethics and efficiency have motivated authors to consider a nonmyopic approach. For example, Cheng and Berry (2007) constructed an optimal adaptive design for a sequential randomized clinical trial, where one wishes to maximize the number of patients who have been treated effectively. This design initially starts with equal allocation for each of k possible treatment arms. As the trial continues and responses of treated patients become available, a larger proportion of patients are allocated to the better treatment. Dynamic programming with backward induction is utilized to make a decision on which treatment to give to the current patient, by defining a loss function that incorporates the responses of the trial so far, and the expected response of the current patient. Dynamic programming was also implemented recently in a clinical trials setting by Ondra et al. (2019) who consider Bayesian adaptive two-stage designs which consider utility functions for the total revenue earned and also the total societal health benefit.

A nonmyopic approach can also be beneficial for dose-finding in clinical trials. A problem proposed by Mueller et al. (2007) involves an experiment where, after observing the response of the current patient, a decision is made to terminate the trial, continue with dose-finding, or switch to a trial intended to provide evidence of efficacy for drug authorities. A utility function is defined in terms of the responses of the trial so far, the decisions made in the trial so far, the parameters of the probability model of the responses, and costs for each decision. Backward induction is used to find the optimal decision at every step by alternating between taking the expectation to find the posterior expected utility and maximizing to find the optimal decision. They use a low-dimensional summary statistic of the previous decisions and data to reduce computational complexity. Bartroff and Lai (2010) look at the use of approximate dynamic programming methods for dose-response problems. These applications of dynamic programming for clinical trials, however, do not consider covariates.

3.1.3 The bandit problem

A considerable amount of work has been undertaken on the bandit problem, which is relevant to response-adaptive treatment allocation in clinical trials. In the original bandit problem, there are several slot machines. One wishes to obtain an optimal strategy for which levers to pull, and in which order, to maximize gains. Suppose that there are t slot machines, which correspond to t treatments. At time i, for $i \in \{1, ..., n\}$, the slot machine $m(i)$ is chosen and a reward $y_{m(i)}$ is observed. The goal is to maximize the sum of the rewards, $\sum_{i=1}^{n} y_{m(i)}(i)$. On the one hand, it is necessary to make decisions that aid in understanding how the machines behave, but on the other hand, it is important to make decisions that lead to higher gains. Similarly, in clinical trials, one wishes to choose treatments to gain understanding of how patients respond, but one also wishes to assign patients to a treatment that will benefit them most. The multi-armed bandits as a sequential optimization problem was first introduced by Robbins (1952). Let us denote by $y_m(i)$ the response from the truly optimal treatment for patient i. One wishes to find a strategy for choosing the machine which minimizes the regret $R(i)$:

$$R(i) = \max_{m \in \{1,...,t\}} \mathbb{E}\left(\sum_{i=1}^{n} y_m(i)\right) - \mathbb{E}\left(\sum_{i=1}^{n} y_{m(i)}(i)\right). \tag{3.6}$$

Lai and Robbins (1985) proved a lower bound for $R(i)$. From a frequentist perspective, the bound is achieved by the Upper Confidence Bound (UCB) algorithm, which selects at time i the machine with the maximal upper confidence bound for the expected reward (Agrawal, 1995). Depending on the choice of the confidence level, this approach can favor exploration (for low confidence level) or exploitation (for high confidence level). Some work has also been

done on Upper Confidence Bounds from a Bayesian perspective; see, for example, Kaufmann et al. (2012).

Thompson (1933) proposed a Bayesian approach for balancing exploration and exploitation. Applicable for the case where the rewards $y_{m(i)}(i)$ are binary, we denote by θ_m the probability of success of each machine m. Independent prior distributions over each θ_m are specified. As the arms are pulled and rewards are observed, the distribution is updated by Bayes' rule. Often, a beta prior is used due to the convenience of the conjugacy property. The algorithm then randomly samples from the posterior distribution to obtain an estimate of the success probability $\hat{\theta}_m$ for each m, and pulls the arm with the largest estimate. Thus, compared to a greedy algorithm which would simply select the arm with the highest expected posterior mean, by randomly sampling from the posterior to obtain an estimate, the algorithm also incorporates exploration.

While the bandit problem probes into the conflict between exploration and exploitation, and the potential for exploration allows for decisions to be made that can be more beneficial for future patients, the solutions to the bandit problems that we have explained so far do not have an explicitly nonmyopic framework. However, there is a specific class of bandit problems in which using dynamic programming with forward induction is optimal. The following conditions are required (Mahajan, 2008):

1. Only one bandit is being operated on at one time.
2. Bandits other than the one being operated on are "frozen": they are not affected and do not contribute reward (this is not true for contextual bandits where there are covariates, as we describe later in this section).
3. Bandits are independent.
4. There is no time-dependence and the horizon is infinite.
5. Rewards are discounted geometrically, so for some $0 < \alpha < 1$, the cumulative sum of rewards is given by $\sum_{i=1}^{n} \alpha^i y_{m(i)}(i)$. This discount is needed for reasons of tractability (Villar et al., 2015).

It was shown by Gittins and Jones (1979) that the optimal design could be found by a forward induction algorithm. If the response y is binary, allocation can be performed based on an index which takes into account the posterior distribution, given the current state. The Gittens index is difficult to compute exactly, so Brezzi and Lai (2000) developed an approximate Gittens index based on a Gaussian approximation of the posterior distribution. Further work on the Gittens index has been conducted by Williamson et al. (2017), Coad (1991) and Wang (1991) for binary responses and Smith and Villar (2017) for continuous

responses. Villar and Rosenberger (2018) considered such approaches in contexts involving covariates for binary outcomes. The major limitations of the Gittens index are that it only works with a geometric discount reward, it requires the horizon to be infinite, and further, Scott (2010) criticizes the Gittens index for being an inconsistent estimator of the optimal arm. It is also not optimal when there are covariates involved.

Bandits with Covariates

Bandits which include covariate information are often called *contextual* bandits; some authors refer to them as *restless* bandits since their state, $S_i = (\boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$, can change from one stage to the next due to the covariates, regardless of the choice of treatment (Mahajan, 2008). The first look into the one-armed bandit problem with one covariate was by Woodroofe (1979). When the horizon is infinite and there is a geometric discount, they showed that the myopic algorithm is asymptotically optimal. Non-parametric approaches to bandits with covariates have been considered by Yang and Zhu (2002), Rigollet and Zeevi (2010) and Perchet and Rigollet (2013). The applicability of contextual bandits for online decision-making has made them a popular research topic. Recent work includes considering some nonlinearity in the model assumptions for the contextual bandits (Greenewald et al., 2017), and considering adversial bandits, where no statistical assumptions are made on how rewards are generated (Agrawal and Goyal, 2013).

Villar and Rosenberger (2018) proposed a reformulation of the bandit problem with covariates where each treatment-covariate combination is treated as an arm. However, the allocation rule is deterministic with no element of randomization, and there is the need to assume an infinite number of patients in the trial.

Longitudinal studies

We consider in this project clinical trials in which each patient is treated only once and their response is measured only once. In some studies involving chronic conditions, longitudinal data may be obtained where patients undergo multiple treatments over time (Clairon et al., 2017); an experimental unit in this context is a patient-timepoint combination. Dynamic treatment regimes consider, at each stage, the optimal treatment based on information about the covariates, treatments and responses up until that state for a particular patient. It is necessary to consider the effect in the long-run that the current treatment may have on the choice of future treatments and the subsequent responses (Murphy, 2003). Dynamic programming and backwards induction are useful in these settings in order to find the

optimal treatment strategy over a time period in the sense that it results in the highest average response at the end of the time period (Murphy, 2003). Chakraborty and Moodie (2013) provided an overview of dynamic treatment regimes and the two main strategies for obtaining optimal designs. The first is Q-learning, or quality learning, in which the conditional expectation of the response given the previous states and treatments (Q-functions) is modeled, potentially in a non-parametric way (Watkins, 1992). The second method is A-learning, or advantage learning, which models just part of the Q-function, the regret function, given by Equation (3.6), which is the difference in outcomes between the given and optimal treatments (Chakraborty and Moodie, 2013, p. 39). Robins (2004) demonstrated two A-learning methods for high dimensional data. An example of an A-learning approach to a problem in control theory is provided by Clairon et al. (2017). There are many other methods for dynamic treatment regimes, such as g-estimation and outcome weighted learning (Chakraborty and Moodie, 2013).

3.2 Simulations to motivate the potential of the nonmyopic approach

Before we explicitly introduce nonmyopic approaches to design of sequential experiments with covariates in Chapters 4 and 5, we first run some simulations to compare the design performance, at each sample size, of two methods of constructing a design:

1. A sequential approach where, for patient i, the optimal treatment for patient i is assigned. The covariates for patients 1 up to i are assumed when the decision for patient i is made.

2. A non-sequential approach where treatments are allocated for all patients using the exchange algorithm, so treatments are assigned in order to achieve an optimal design for all n patients. The covariate values for all patients from 1 up to n are assumed.

By comparing the performance at each value of sample size between 1 and n, we can obtain an indication of the settings in which a nonmyopic approach may be more efficient than a myopic approach. When evaluating the optimality for patient i, the decision made by the first approach is based on a myopic outlook and the decision made by the second approach is nonmyopic since it considers all patients up until the end of the experiment. If we find that the design constructed with the exchange algorithm, where we assume full knowledge about the covariates, is more efficient than the sequential design, this would indicate that there is potential for nonmyopic and pseudo-nonmyopic approaches, where we assume some knowledge about the covariates, to be beneficial over the myopic approach. The features of the simulations in this chapter are summarized in Table 3.1, and the simulations have the

following structure:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution. The run order of the design is fixed by these covariates.

(b) Two designs for the patients generated in part (a) are constructed using:

- A point sequential design using the sequential optimal design approach for the linear model with the D-optimal objective function (see Algorithm 3).
- A batch design design using the exchange algorithm the linear model with the D-optimal objective function (see Algorithm 2) with 10 repetitions from different starting designs.

(c) Designs are evaluated using the performance measure Ψ_D, given by Equation (2.18), at each sample size.

II (a)-(c) above is repeated 1000 times to obtain a distribution of the performance measure for each sample size, and to assess the impact of different orderings of the covariates (as the batch design does not account for run order).

We run the above simulation also using the performance measures Ψ_A and Ψ_G, given by Equations (2.30) and (2.32), for design selection. We consider only the case of a continuous response.

Table 3.1: Settings for simulations in Chapter 3

Example	**Model**	**Covariates**
1	$\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z} + \beta_2 \boldsymbol{t}$	Binary: $\mathbb{P}(z_i = 1) = 0.5$
2	$\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z} + \beta_2 \boldsymbol{t}$	Continuous: $z_i \sim \text{Unif}(0, 1)$
3	$\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z}_1 + \beta_2 \boldsymbol{z}_2 + \beta_3 \boldsymbol{t}$	Two binary covariates: $\mathbb{P}(z_{1,i} = 1) = 0.01i$ $\mathbb{P}(z_{2,i} = 1) = 0.3$
4	$\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z}_1 + \beta_2 \boldsymbol{z}_2 + \beta_3 \boldsymbol{t} + \beta_4 \boldsymbol{z}_1 \boldsymbol{t} + \beta_5 \boldsymbol{z}_2 \boldsymbol{t}$	Two binary covariates: $\mathbb{P}(z_{1,i} = 1) = 0.01i$ $\mathbb{P}(z_{2,i} = 1) = 0.3$

Example 1

We assume a linear model with one covariate and one binary treatment. The covariate z is binary and we have that $\mathbb{P}(z_i = 1) = \mathbb{P}(z_i = -1) = 0.5$. We say that this covariate is static, since the probability $\mathbb{P}(z_i = 1)$ is constant for all i.

In Figure 3.1, we show, at each sample size, the relative efficiencies of the sequential design versus the batch design (see Equations (2.24), (2.31), (2.33) for D-, A-, and G-efficiency,

respectively). For D-efficiency, we plot the ratio of the objective function Ψ_D taken to the power of $\frac{1}{3}$ as there are three parameters in the model, in order to have a quantity on the same scale as the variance. Values above the red line indicate that the design obtained using the exchange algorithm is more efficient; values below the red line indicate that the sequential design is more efficient. We observe that, initially, when the sample size is small, the sequential approach has better performance. However, as the number of patients becomes larger, the two approaches become similar in performance. With sample size greater than 60, we can see that, on average, the batch design and sequential designs perform very similarly with the batch design having a very slight advantage in D- and A-efficiencies. In the case of G-efficiency, the batch design becomes more efficient as sample size increases; at sample size of 100, on average, the batch design has an efficiency of around 1.08 times that of the sequential design.

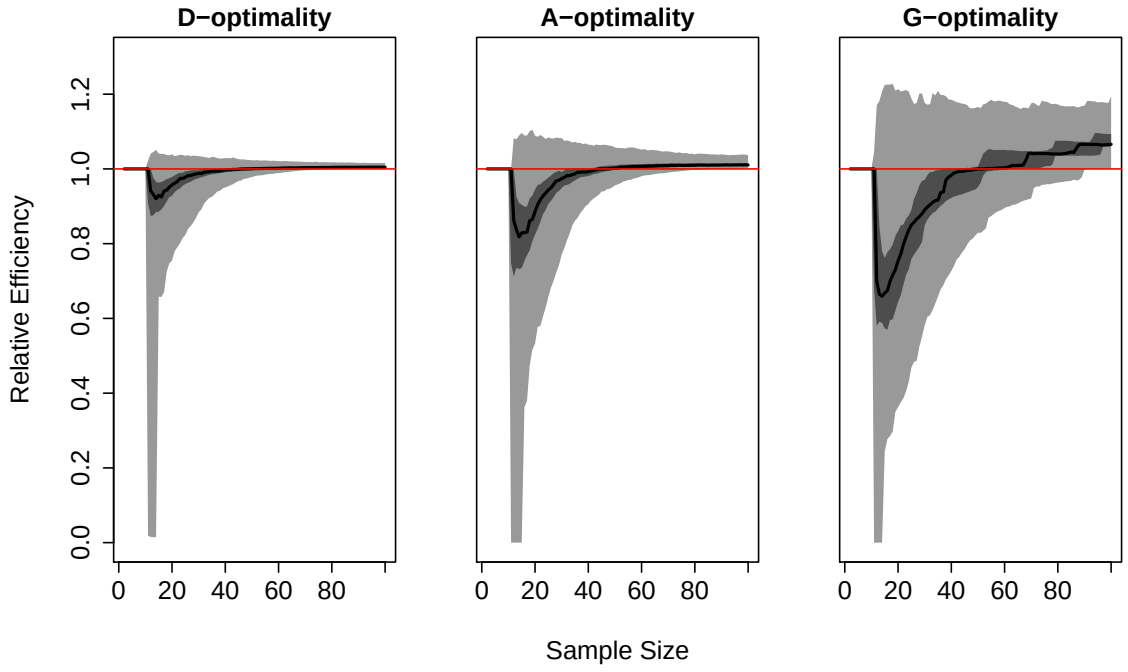

Figure 3.1: Example 1: Distributions of relative efficiencies of the sequential design vs the exchange algorithm design. We show D-efficiency (left), A-efficiency (middle) and G-efficiency (right). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

Example 2

In the second simulation, we use the same model but generate the covariate from a continuous uniform distribution such that $z_i \sim \text{Unif}(0, 1)$. Figure 3.2 shows that the conclusions are similar to that of the first simulation. For sample sizes that are small (less than around 30), the sequential design is more efficient, but for sample sizes greater than 60, the two approaches perform relatively similarly. There appears to be a very slight advantage for the batch designs at this larger sample size.

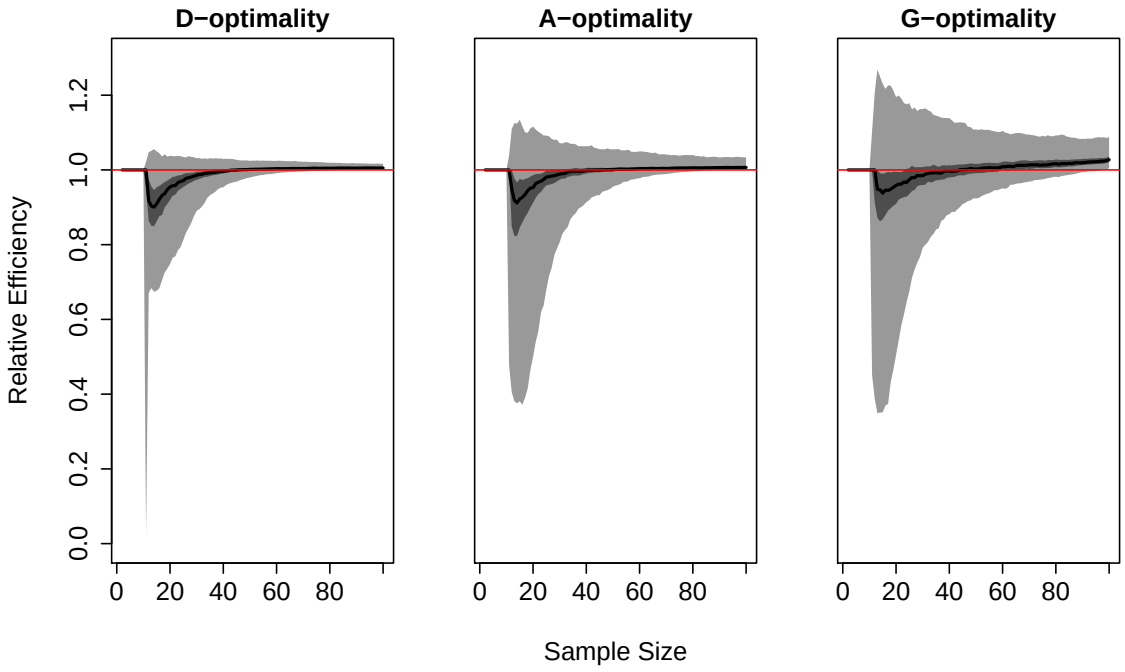

Figure 3.2: Example 2: Distributions of relative efficiencies of the sequential design vs the batch design. We show D-efficiency (left), A-efficiency (middle) and G-efficiency (right).

Example 3

Next, we consider examples with two covariates. We assume that our model includes terms for the two covariates, the treatment, and no interactions. In the fourth simulation, we include the covariate-treatment interaction terms. In both simulations, we have the same settings for the covariates. The first covariate is dynamic in the sense that its distribution changes over time; we have that $\mathbb{P}(z_{i,1} = 1) = 0.01i$. The second covariate is static as the probability of its value being equal to 1 is constant for all i: we have that $\mathbb{P}(z_{i,2} = 1) = 0.3$.

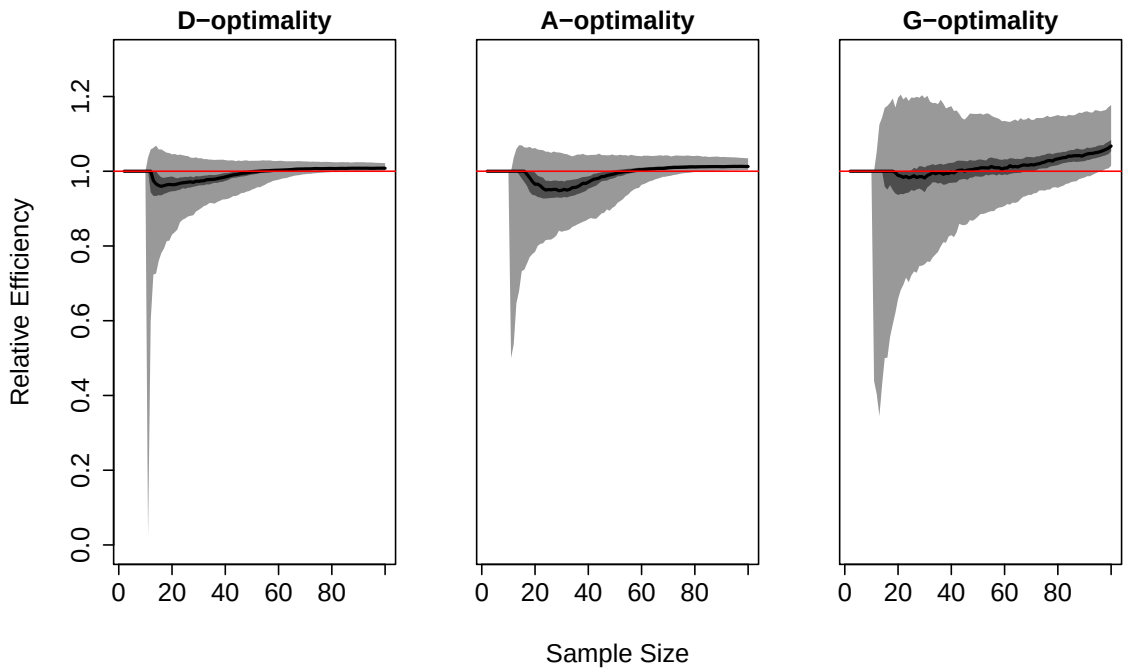

Figure 3.3: Example 3: Distributions of relative efficiencies of the sequential design vs the batch design. We show the results for D-efficiency (left), A-efficiency (middle) and G-efficiency (right).

Similarly to the cases with one covariate, we find in the example with two covariates where one is static and one is dynamic, that there is potential for the nonmyopic approach to be beneficial over the myopic when the sample size is large. In Figure 3.3, we find that, even with sample size of 40, it appears that the G-optimal batch designs are more efficient than the G-optimal sequential designs.

Example 4

Our fourth example has the same set-up as the previous example, except treatment-covariate interactions are included in the model. In Figure 3.4, we find that initially, the sequential design is much more efficient than the design obtained with the exchange algorithm. However, with sample size above 80, the exchange algorithm designs have a clear advantage, especially in the case of D- and G-optimality.

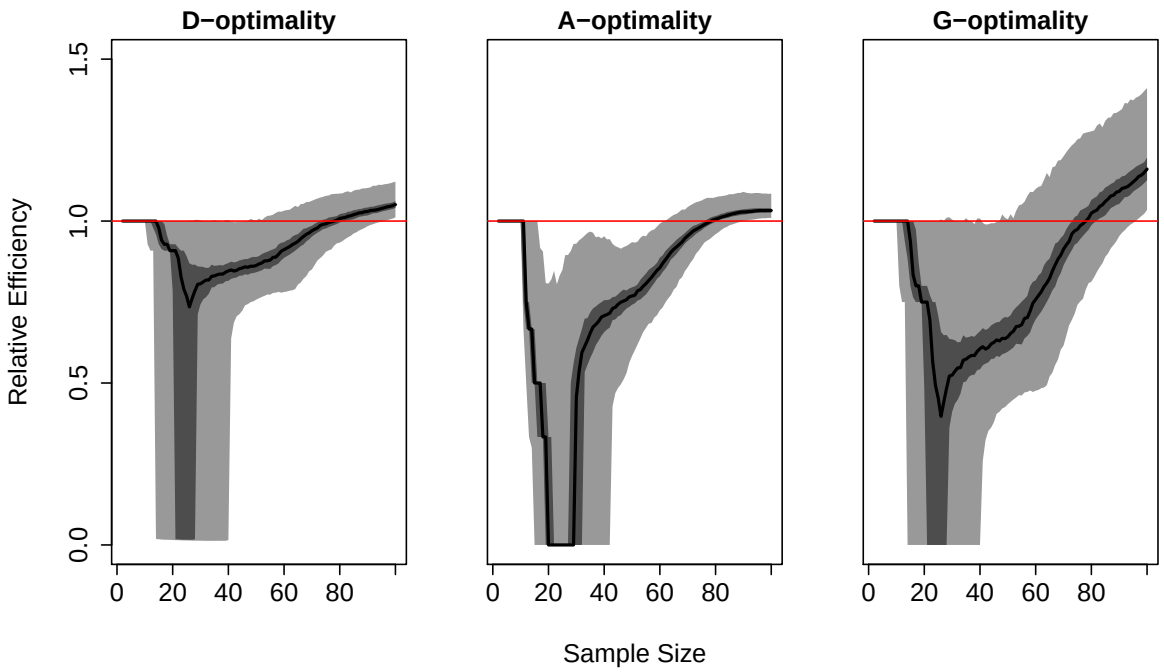

Figure 3.4: Example 4: Distributions of relative efficiencies of the sequential design vs batch design. We show the results for D-efficiency (left), A-efficiency (middle) and G-efficiency (right).

In examples with the linear model, the sequential approach initially leads to more efficient designs. As sample size increases, the two methods appear to have similar values of the objective function. With sample size greater than 60, it appears that the exchange algorithm designs are slightly more efficient in the simulations considered above. This indicates that there is a potential for the nonmyopic approach to be more beneficial than the myopic approach after a certain sample size, at least for the linear model. In further chapters, we also consider the case of a binary response, as well as other optimality criteria.

3.3 Conclusion

This chapter introduced nonmyopic approaches from the literature, focusing on dynamic programming and the bandit problem as two common approaches to problem solving where one wishes to make optimal decisions considering future possible decisions. We showed how they are used in the setting of clinical trials, for example in the context of the bandit problem and also for dynamic treatment regimes. We note that nonmyopic approaches are also used in areas such as control theory, though with a slightly different nomenclature; see,

for example, Gautier and Pronzato (1998) and Bertsekas (1976).

In Section 3.2, we demonstrated, through simulations where we compare the efficiency of designs constructed using the exchange algorithm and the sequential approach, the potential benefit of using nonmyopic approaches. In the context of the linear model, we find that the efficiency of designs constructed using the exchange algorithm is improved for sample sizes above 60. It appears that this improvement is particularly pronounced for G-optimal designs. In Chapter 4, we introduce the nonmyopic approach applied to our setting for both linear and logistic models.

Chapter 4

Nonmyopic approach

Chapter 3 provided an overview and examples from the literature of the nonmyopic approach to treatment allocation. We also provided a simulation study to motivate this approach for the linear model in the context of sequential design with covariates. In this chapter, we set up notation for the nonmyopic approach which unifies the notation for the sequential optimal design methods in Chapter 2 and also the notation for the pseudo-nonmyopic approach which we introduce in Chapter 5. Essentially, we extend the work of Atkinson (1982) so that it can be applied to any optimality criterion and also for a nonmyopic outlook.

Firstly, we describe an approach for selecting the treatment for unit i in a way that takes into account future possible decisions which may be made for the treatment of unit $i+1$. Secondly, we generalize this method so that we can consider the effect of the current decision on the next N possible future decisions. Thirdly, we generalize the method so it can be applied to a wider class of covariates such as having multiple discrete covariates with more than two levels and multiple continuous covariates. We show some simulation results comparing the myopic algorithm to the nonmyopic algorithm for a number of settings. We then introduce the logistic model and experiments for binary data, and extend the nonmyopic algorithm for this setting. We provide some simulation results comparing the myopic algorithm to the nonmyopic algorithm for the binary response case.

4.1 Linear model

4.1.1 The nonmyopic algorithm looking one patient ahead

We wish to extend the myopic algorithm so that decisions on the allocation of treatments for unit i are made based on the expected loss after $i+1$ units, rather than the loss after i units. Firstly, for the sake of simplicity, we assume that we have only one binary covariate. We also assume that we know the true distribution of this covariate; that is, we know $\mathbb{P}(z_i = 1)$ and $\mathbb{P}(z_i = -1)$ and they in fact do not depend on i. When the true distribution is unknown, it can be estimated by the empirical distribution of the covariates of the first i patients. The covariate distribution remains the same regardless of how many patients have enrolled in the experiment so far. As before, our treatment structure is binary and we assume that a linear model for the response, such as Equation (2.13), is appropriate.

Let us denote the objective function after i units when treatment t_i is assigned to patient i as

$$\Psi(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}), \tag{4.1}$$

where we emphasize the dependence on the vector of previous treatments $\boldsymbol{t}_{i-1} = (t_1, t_2, ..., t_{i-1})^\intercal$ and previous and current covariates $\boldsymbol{z}_i = (z_1, z_2, ..., z_i)^\intercal$. The function Ψ could be, for example, the D-, D_A-, A- or G-optimal objective functions (see Equations (2.18), (2.25), (2.30) and (2.32)).

Now, we wish to consider the impact of assigning treatment t_i on the objective function after $i+1$ patients. We consider the final decision that we are going to make: the choice of treatment $i+1$ for the future patient. For each combination of values that t_i and z_{i+1} may take, we compute the optimal choice of treatment for patient $i+1$. We denote by $t^*_{i+1}(z_{i+1}, t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$ the optimal choice of treatment for patient $i+1$, for a particular choice of z_{i+1} and t_i, given all previous treatments and covariates:

$$t^*_{i+1}(z_{i+1}, t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}) = \underset{t_{i+1}}{\operatorname{argmin}}\, \Psi\left(t_{i+1} \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i\right).$$

The star in the notation indicates that the optimal decision concerns a future patient. We suppress the conditioning in the notation and simply write $t^*_{i+1}(z_{i+1}, t_i)$. We wish to choose a treatment t_i which is optimal given the information about the patients up until patient i, and also given possible decisions which could be made for patient $i+1$. The quantity $t^*_{i+1}(z_{i+1}, t_i)$ is known if both z_{i+1} and t_i are known. However, since z_{i+1} is unknown, we need to find the expectation of the objective function after patient $i+1$ over the two possible

covariate values, assuming that the optimal treatment has been chosen for the $(i+1)$th patient:

$$\begin{aligned}&\mathbb{E}_{z_{i+1}}\left[\Psi\left(t_{i+1}^{*}(z_{i+1},t_i)\mid \boldsymbol{z}_i,\boldsymbol{t}_i\right)\right]\\&=\mathbb{P}(z_{i+1}=1)\Psi\left(t_{i+1}^{*}(z_{i+1}=1,t_i)\mid \boldsymbol{z}_i,\boldsymbol{t}_i\right)\\&+\mathbb{P}(z_{i+1}=-1)\Psi\left(t_{i+1}^{*}(z_{i+1}=-1,t_i)\mid \boldsymbol{z}_i,\boldsymbol{t}_i\right).\end{aligned}\tag{4.2}$$

We denote the expected value of the objective function after one patient in the future as $\Phi_1(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$:

$$\Phi_1(t_i \mid \boldsymbol{z}_i,\boldsymbol{t}_{i-1})=\mathbb{E}_{z_{i+1}}\left[\Psi\left(t_{i+1}^{*}(z_{i+1},t_i)\mid \boldsymbol{z}_i,\boldsymbol{t}_i\right)\right].\tag{4.3}$$

The optimal decision is given by $\operatorname{argmin}_{t_i}\Phi_1(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$. For a sequential implementation of a nonmyopic algorithm, at stage i, $\Phi_1(t_i = 1 \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$ and $\Phi_1(t_i = -1 \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$ are computed, and treatment 1 is given to patient i with probability given by:

$$\frac{\Phi_1(t_i=1\mid \boldsymbol{z}_i,\boldsymbol{t}_{i-1})^{-1}}{\Phi_1(t_i=1\mid \boldsymbol{z}_i,\boldsymbol{t}_{i-1})^{-1}+\Phi_1(t_i=-1\mid \boldsymbol{z}_i,\boldsymbol{t}_{i-1})^{-1}}.\tag{4.4}$$

In Appendix H, we provide a simple example to illustrate the myopic approach and nonmyopic approach for the linear model case, as well as the pseudo-nonmyopic approach which we describe in Chapter 5.

4.1.2 Extension one: Longer Horizon

We refer to the number of future decisions considered as the horizon, denoted N. The assignment of the ith treatment then becomes a multistage optimization problem with N stages, as it becomes necessary to consider whether the decision is optimal for the objective function at each stage for patient $i+2, i+3, ..., i+N$. A decision-theoretic framework is used in order to provide a formal structure to the problem (Parmigiani, 2002).

We extend the notion of the expected loss in Equation (4.2) so that it can be generalized for a larger number of steps into the future. Suppose that we wish to assign a treatment to unit ν, considering the impact one step ahead into the future, for $\nu = \{1, 2, ..., n-1\}$. Given that we know $\boldsymbol{z}_\nu$ and $\boldsymbol{t}_{\nu-1}$, the expected objective function is given by:

$$\Phi_1(t_\nu \mid \boldsymbol{z}_\nu,\boldsymbol{t}_{\nu-1})=\mathbb{E}_{z_{\nu+1}}\left[\Psi\left(t_{\nu+1}^{*}(z_{\nu+1},t_\nu)\mid \boldsymbol{z}_{\nu+1},\boldsymbol{t}_\nu\right)\right],\tag{4.5}$$

where the subscript for Φ indicates the horizon. Now, we consider the impact of the decision

to assign t_ν on the loss after N additional unseen units. The following recursive relationship can be used to find the expected objective function after $\nu + N$ units, when treatment t_ν has been assigned:

For $N > 0$:

$$\begin{aligned}\Phi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}) =& \mathbb{E}_{z_{\nu+1}}\left[\Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1}, t_\nu) \mid \boldsymbol{z}_{\nu+1}, \boldsymbol{t}_\nu\right)\right] \\ =& \mathbb{P}(z_{v+1} = 1)\left[\Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1} = 1, t_\nu) \mid \boldsymbol{z}_\nu, z_{\nu+1} = 1, \boldsymbol{t}_\nu\right)\right] \\ &+ \mathbb{P}(z_{v+1} = -1)\left[\Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1} = -1, t_\nu) \mid \boldsymbol{z}_\nu, z_{\nu+1} = -1, \boldsymbol{t}_\nu\right)\right],\end{aligned} \tag{4.6}$$

and for $N = 0$, we have:

$$\Phi_0(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}) = \Psi(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}), \tag{4.7}$$

which is simply the myopic loss after ν units; the subscript 0 indicates that we are considering no steps into the future. Thus, treatment 1 is assigned to patient t_ν with the probability given by

$$\frac{\Phi_N(t_\nu = 1 \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})^{-1}}{\Phi_N(t_\nu = 1 \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})^{-1} + \Phi_N(t_\nu = -1 \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})^{-1}}. \tag{4.8}$$

We are using backward induction since the recursion in Equation (4.6) requires that one finds the optimal treatment which minimizes the expected objective function for the Nth unit in the horizon, then the optimal treatment for the $N-1$ unit, and so on until one reaches the current unit.

It is possible to see that our myopic algorithm is a special case of this nonmyopic set-up; it is an allocation scheme based on $\Phi_0(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1})$ as an objective function. Equation (4.2) is also a special case where we look one step ahead in the future and use $\Phi_1\left(t^*_{i+1}(z_{i+1}, t_i) \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i\right)$ as an objective function.

4.1.3 Extension two: more complex covariates

So far, we have assumed that we have one binary covariate. Let us now extend the problem statement so that the covariate values are discrete and take values from a sample space S. We denote by $\mathbb{P}(z_i = z)$ the probability that the ith patient takes covariate value z, for $z \in S$. We assume that these probabilities do not depend on the number of patients in the trial so far. In this setting, $\Phi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})$ for $N > 0$, originally defined in Equation (4.6)

for a single binary covariate, is restated as follows:

$$\begin{aligned}\Phi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}) =& \mathbb{E}_{z_{\nu+1}}\left[\Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1}, t_\nu) \mid \boldsymbol{z}_{\nu+1}, \boldsymbol{t}_\nu\right)\right] \\ =& \sum_{z \in S} \mathbb{P}(z_{\nu+1} = z)\Phi_{N-1}\left(t^*_{\nu+1}(z, t_\nu) \mid \boldsymbol{z}_\nu, z, \boldsymbol{t}_\nu\right),\end{aligned} \tag{4.9}$$

and $\Phi_0(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})$ is defined as in Equation (4.7). We note that this can be easily extended further to the case where there are k covariates observed for each unit. In this case, the probability mass function gives a probability for each possible combination of the k covariates.

For a further extension, suppose we have k continuous covariates, and the covariate values for unit i is given by the k-vector $\boldsymbol{z}_i$. Its joint probability distribution function is denoted $f(\boldsymbol{z}_i)$. For $N > 0$, we redefine $\Phi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})$ as follows:

$$\begin{aligned}\Phi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}) =& \mathbb{E}_{z_{\nu+1}}\left[\Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1}, t_\nu) \mid \boldsymbol{z}_{\nu+1}, \boldsymbol{t}_\nu\right)\right] \\ =& \int \Phi_{N-1}\left(t^*_{\nu+1}(z_{\nu+1}, t_\nu) \mid \boldsymbol{z}_{\nu+1}, \boldsymbol{t}_\nu\right) f(\boldsymbol{z}_{\nu+1}) d\boldsymbol{z}_{\nu+1},\end{aligned} \tag{4.10}$$

and $\Phi_0(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1})$ is defined as in Equation (4.7).

4.1.4 Simulations

We wish to compare D_A-, D- and G-optimal designs that are constructed sequentially using myopic and nonmyopic methods. Further, we wish to compare the nonmyopic approach where we assume the true distribution for the covariates, and the nonmyopic approach where we use the empirical distribution. In our binary covariate case, we obtain an approximation of the empirical distribution of the covariate by finding the proportion of observed patients with each covariate value. Our simulations have the following structure, which allows us to compare seven ways of constructing sequential designs:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) Seven designs for the patients in part (a) are constructed using:

- A myopic D_A-optimal design.
- A nonmyopic D_A-optimal design with horizon $N = 1$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 1$, with the empirical covariate distribution assumed.

- A nonmyopic D_A-optimal design with horizon $N = 2$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 2$, with the empirical covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 5$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 5$, with the empirical covariate distribution assumed.

(c) Designs are evaluated using the performance measure Ψ_{D_A}, given by Equation (2.25), at each sample size between 5 and 100, inclusive.

II (a)-(c) above is repeated 20 times to obtain a distribution of the performance measure for each sample size.

We run the above simulation also using the performance measures Ψ_D and Ψ_G from D- and G-optimality, given by Equations (2.18) and (2.32), for design selection .

Results for one covariate

Firstly, we run a simulation in which 100 covariates are generated from the Bernoulli($p = 0.5$) distribution. We assume the relationship between the response, covariate and treatment is given by the linear model $\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z} + \beta_2 \boldsymbol{t}$. We observe in Figure 4.1 that the plots showing the distribution of Ψ_{D_A} are similar across the seven methods. When the correct covariate distribution is known, the results for the myopic designs look almost identical to the nonmyopic designs. When the correct covariate distribution is unknown and the empirical distribution is assumed, it appears that the variability of Ψ_{D_A} is higher for the nonmyopic designs compared to the case where the correct covariate is distribution known. In Figure D.1 in the Appendix, we plot the relative efficiencies of the nonmyopic designs against the myopic design. It appears that, in the early stages of the experiment, the myopic approach is more efficient. However, the two approaches soon converge to the same efficiency. When the correct covariate distribution is assumed, the nonmyopic approach has the same efficiency as the myopic approach.

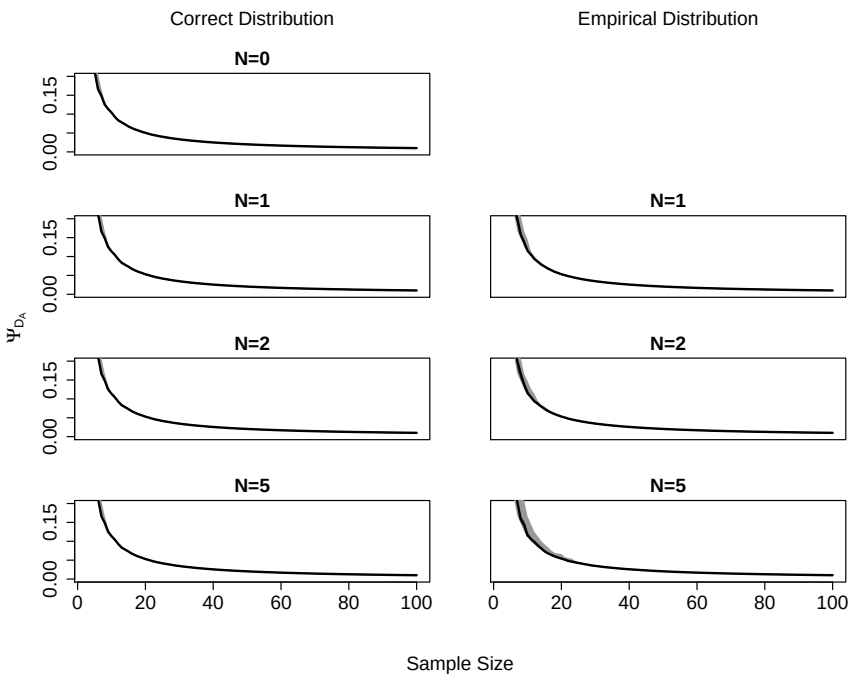

Figure 4.1: Distributions of Ψ_{D_A} for D_A-optimal designs for the linear model are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

Under D-optimality, in Figure 4.2, we observe that the distribution of Ψ_D appears fairly similar for all seven methods; Figure D.2 in the appendix shows that, similarly to the case of D_A-optimality, the myopic approach is initially more efficient, but then the myopic and nonmyopic approaches become equally efficient.

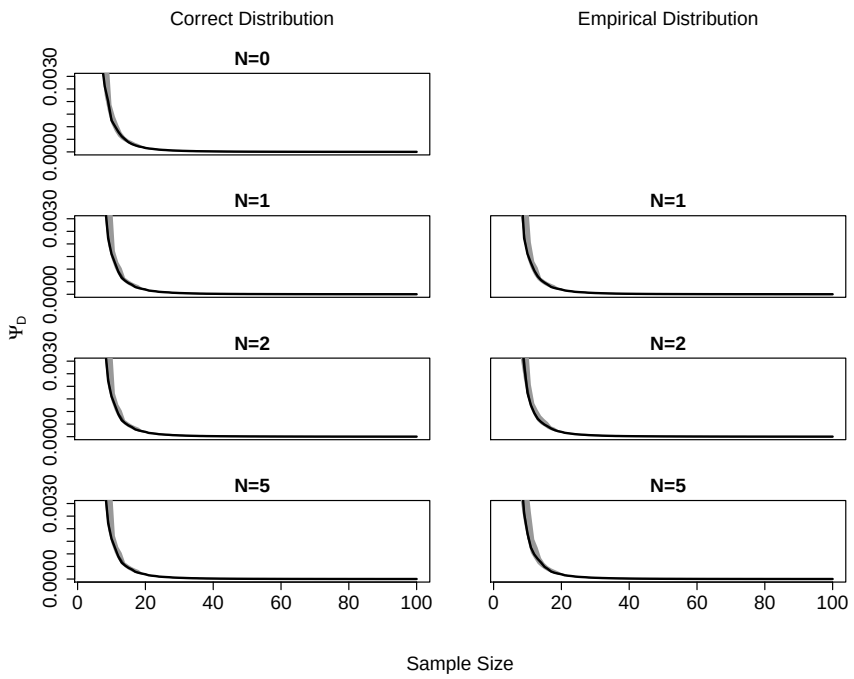

Figure 4.2: Distributions of Ψ_D for D-optimal designs for the linear model are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

Similarly, for designs constructed to be G-optimal, we observe in Figure 4.3 that the distribution of Ψ_G is extremely similar for all seven methods. We observe greater variability in the plots for G-optimality, compared to D- and D_A-optimality. Figure D.3 in the appendix shows a similar pattern in the efficiency plots to the previous two simulations.

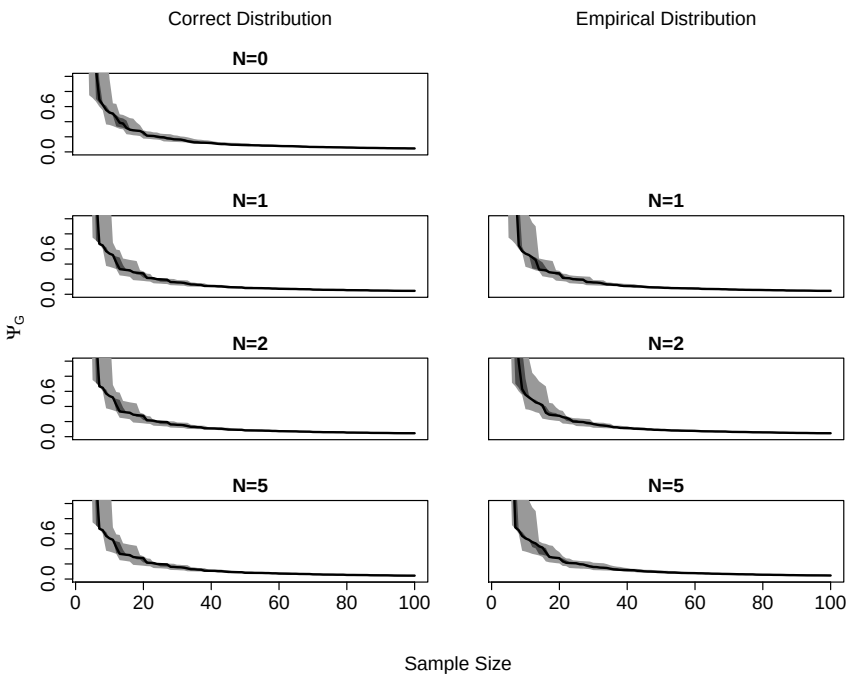

Figure 4.3: Distributions of Ψ_G for G-optimal designs for the linear model are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

We conclude that for the case where we have one covariate only, there is no benefit to using the nonmyopic approach over the myopic approach. We repeated the simulation where there are two covariates, both generated from a Bernoulli(0.5) distribution. Let us denote by $\boldsymbol{z}_{\cdot 1}$ and $\boldsymbol{z}_{\cdot 2}$ the vectors containing values of the first and second covariates, respectively:

$$\boldsymbol{z}_{\cdot 1} = \left(z_{1,1}, z_{2,1}, ..., z_{n,1}\right)^{\mathsf{T}}, \tag{4.11}$$

$$\boldsymbol{z}_{\cdot 2} = \left(z_{1,2}, z_{2,2}, ..., z_{n,2}\right)^{\mathsf{T}}, \tag{4.12}$$

The relationship between the covariates, treatment and response are given by the model $\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z}_{\cdot 1} + \beta_2 \boldsymbol{z}_{\cdot 2} + \beta_3 \boldsymbol{t}$. The results are shown in Appendix D.2 and we draw the same conclusion that there appears to be no benefit to using the nonmyopic approach over the myopic approach.

Dynamic covariates

We now investigate what happens when the covariates are generated in a non-uniform way. We consider a case where there is one binary covariate z generated from a Bernoulli(p_i) distribution, where p_i is the ith element of $\boldsymbol{p}$:

$$\boldsymbol{p} = \left[\underbrace{0.1, ..., 0.1}_{25 \text{ times}}, \underbrace{0.9, ..., 0.9}_{25 \text{ times}}, \underbrace{0.1, ..., 0.1}_{25 \text{ times}}, \underbrace{0.9, ..., 0.9}_{25 \text{ times}}\right]. \tag{4.13}$$

We run a simulation similar to that described in Section 4.1.4, but the covariate is generated dynamically. In the case where we assume that we know the correct distribution, we assume complete knowledge of the future p_i. Figures 4.4 and 4.5 show the distributions of Ψ_{D_A} and the relative efficiencies, respectively. We observe that, assuming the correct distribution, the myopic and nonmyopic approaches appear equally efficient. If the true distribution is unknown and the empirical distribution is assumed, the myopic approach is initially more efficient.

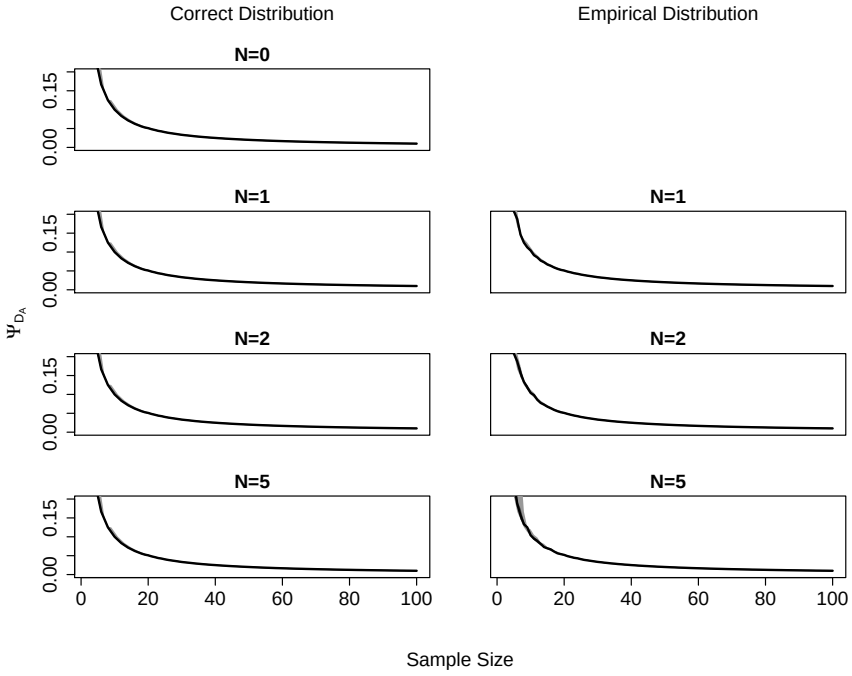

Figure 4.4: Distributions of Ψ_{D_A} for designs for the linear model for one covariate generated dynamically are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

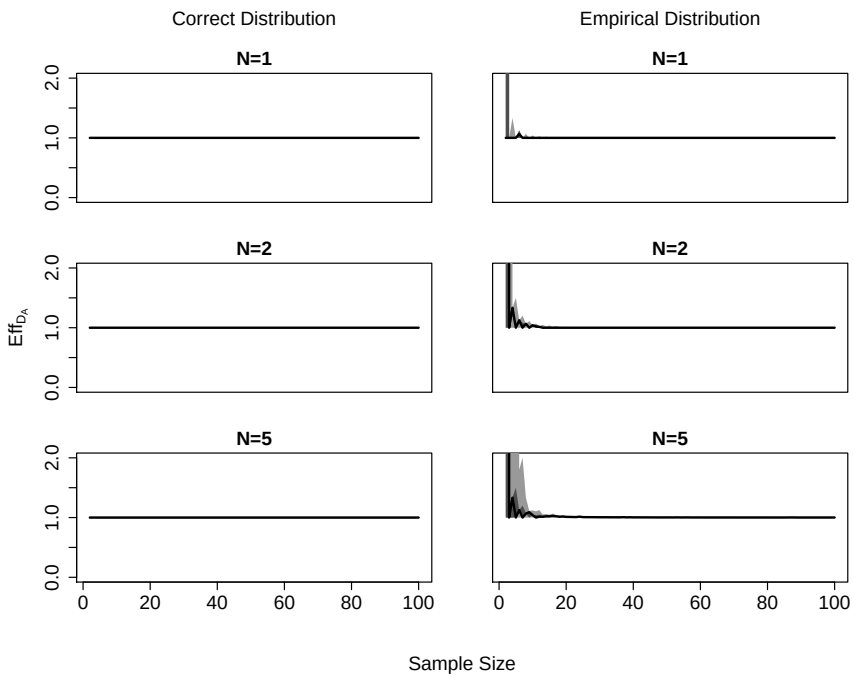

Figure 4.5: Distributions of D_A-efficiencies of nonmyopic designs for the linear model from Figure 4.4 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

In Figures 4.6 and 4.7, which show relative efficiencies, we see that for the case of D-optimality, when the correct distribution of the covariates is known, the nonmyopic approach appears to be very slightly more efficient than the myopic approach for sample sizes less than 30. When the empirical distribution is assumed, the myopic approach appears to be more efficient initially.

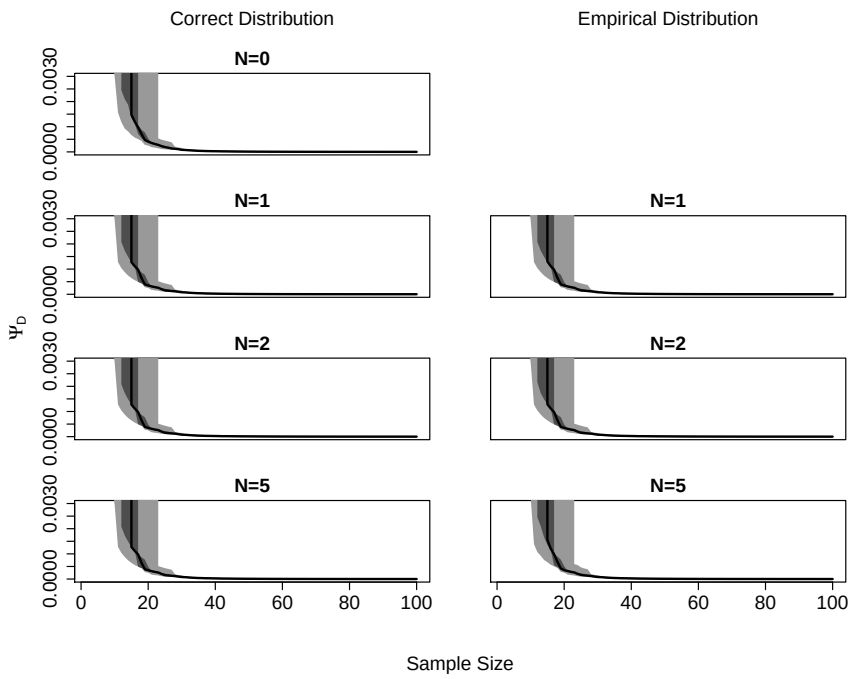

Figure 4.6: Distributions of Ψ_D for designs for the linear model for one covariate generated dynamically are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

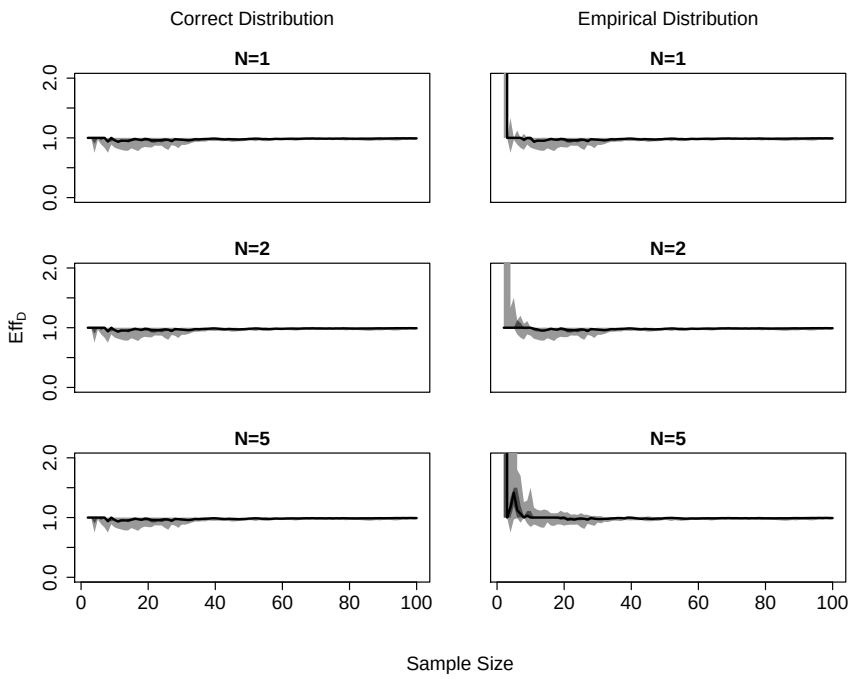

Figure 4.7: Distributions of *D*efficiencies of nonmyopic designs for the linear model from Figure 4.6 are shown, relative to the myopic design.

Finally, in the case of *G*-optimal designs, the plots in Figures 4.8 and 4.9 show that the nonmyopic approach appears to be slightly more efficient than the myopic approach, regardless of whether the correct or empirical distributions are used.

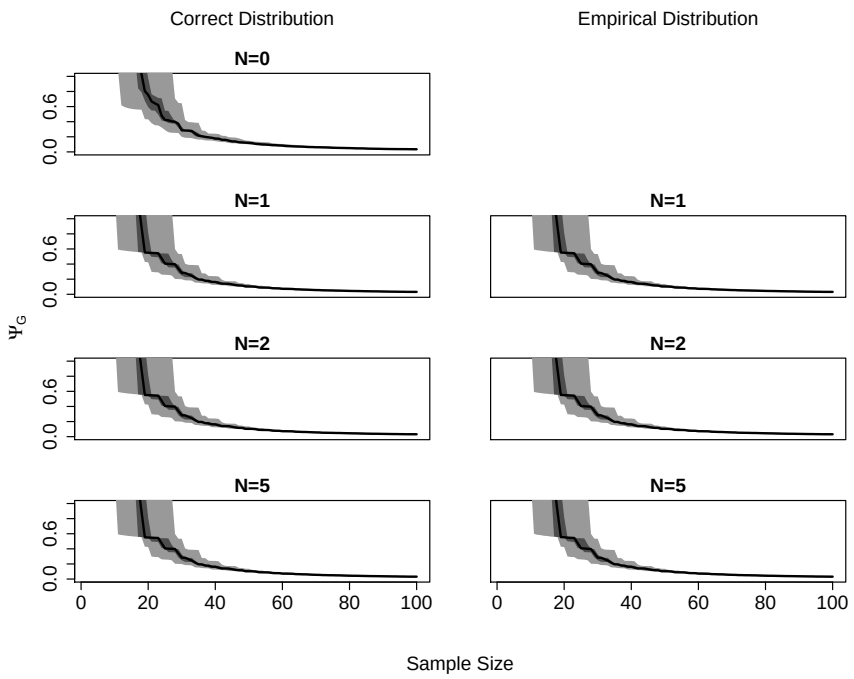

Figure 4.8: Distributions of Ψ_G for designs for the linear model for one covariate generated dynamically are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$, 2 and 5. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

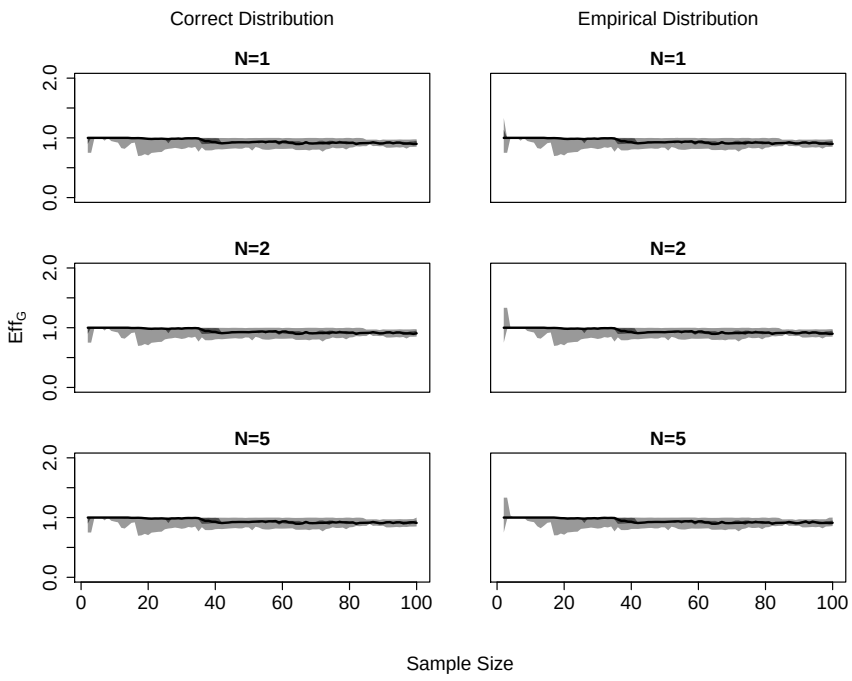

Figure 4.9: Distributions of G-efficiencies of nonmyopic designs for the linear model from Figure 4.8 are shown, relative to the myopic design.

Again, we observe some potential benefits to using a nonmyopic approach in these simulations involving a single dynamic covariate in the case of a linear model; we see that for the D-optimal case when the true covariate distribution is known and in the G-optimal case, there is a very slight improvement in efficiency for the nonmyopic approach. For all other settings, including the case where we have a single static covariate, it appears that the myopic approach is more efficient at the start of the experiment and the myopic and nonmyopic approaches achieve similar levels of efficiency.

4.2 Logistic model

In the previous section, we found that the nonmyopic approach provides a slight benefit over the myopic approach under a linear model when there is a dynamic covariate. We now consider the case of a binary response. A logistic GLM is used to model the response, and in our particular set-up, the information matrix is now contingent on the responses through estimated parameter values, as well as the covariates and treatments in the experiment so far. This added complexity may mean that the nonmyopic approach is better able to make decisions for the overall experiment. We provide an overview of the literature on design of experiments for binary data, then we describe the set-up for the design for the binary response case. Then, we demonstrate how the nonmyopic approach described in Section

4.1.1 can be modified to allow for a binary response. We then conduct simulations to test whether there is any benefit to using the nonmyopic approach over the myopic approach in the binary response setting.

4.2.1 Design of experiments for binary data

Early work on design of experiments for binary response focused on problems where a single continuous treatment is applied to a unit, and a binary response is observed. Such problems are common in experiments on bioassays, explosives, and drug dosage; the goal is to obtain the threshold value of the treatment that determines which of the two outcomes will occur (Abdelbasit and Plackett, 1983). Dixon and Mood (1948) developed what is known as the Bruceton analysis, or up-and-down method, to obtain a confidence interval for this threshold. It is a non-parametric method which sequentially determines the treatment value in the next run. Robbins and Monro (1951) developed a more general method involving a recursive updating rule which can estimate not only the threshold value, but any quantile of the response curve.

The non-parametric methods by Dixon and Mood (1948) and Robbins and Monro (1951) are easy to implement, but require large sample sizes and good starting values. Moreover, the Robbins-Monro method may be inaccurate in estimating extreme quantiles, even with large sample sizes. Wu (1985) criticized these procedures for not using the accruing data from the experiment in an efficient way and instead suggested that sequential designs with binary data should involve an updating rule that depends on a parametric model which captures all the data obtained in the experiment up to that point. Such a model could be a logistic model with two parameters: an intercept and coefficient for the treatment; the maximum likelihood estimates of the parameters are utilized to choose the next level of treatment. Wu (1985) proved properties of asymptotic efficiency of the MLEs and showed that they are optimal in the sense that they achieve the smallest asymptotic variance among a class of designs. A related method by Neyer (1994) is better able to handle the problem of biased maximum likelihood estimates when sample sizes are small. It involves, firstly, an initial search to approximate an interval of the treatment that leads to good estimates of quantiles of the response curve. Secondly, an updating rule based on a parametric model is used to obtain successive treatment levels using the D-optimality criterion. The estimates are refined in a third stage. Further, Joseph (2004) adapted the Robbins-Munro procedure to improve its performance on estimating extreme quantiles by allowing for a distributional form for the quantile curve whose parameters can have prior distributions.

From a broader modelling framework, designs for logistic regression allow for more than one

treatment factor and presence of covariates. Overviews of the theory of logistic regression, Generalized Linear Models and design issues related to them are provided by Khuri et al. (2006) and Woods and Atkinson (2015). Approaches to design of experiments for GLMs are not necessarily sequential. For GLMs, the information matrix, and through it, the objective function, depend on values of the model parameters. The paradox is that estimates of parameters are needed in order to design the experiment aimed to estimate these parameters in the first place. There are several strategies for overcoming this. Firstly, if an initial guess can be made from the results of a previous experiment, one can construct a locally optimal design, which assumes this initial guess as the parameter values when computing the optimality criterion. Work on locally optimal designs for binary data and one treatment was undertaken by Abdelbasit and Plackett (1983). Locally D-optimal designs for 2^k factorial experiments were developed by Yang and Zhu (2002). One disadvantage is that locally optimal designs may be inefficient if the initial guess is poor (Chaloner and Larntz, 1989). If a range for the parameter values is known, a method by King and Wong (2000) for minimax D-optimal designs for logistic regression may be appropriate. In situations where one does not have enough knowledge to make good initial estimates, Sitter and Wu (1999) suggested a two-stage optimal design where the second stage is intended to counteract the effect of inaccurate initial estimates in the first stage. Another strategy to deal with the problem of parameter dependence is to run the experiment sequentially, starting with an initial guess of the parameter values, and then updating the estimates after each run of the experiment (Atkinson, 1999, p.257). Thirdly, a Bayesian approach can also handle the parameter dependence issue. Chaloner and Larntz (1989) describe a non-sequential Bayesian approach to logistic regression, which requires prior distributions for all parameters, as well as their relative importance. Their work was extended by Woods et al. (2006) as well as Dror and Steinberg (2008) in a framework for GLMs which account for uncertainty in the link function and the linear predictor, and additionally allows sequential and batch-sequential implementations.

Work on sequential design of experiments on binary response has mostly focused on achieving optimality in a myopic sense. Nonmyopic approaches to design have been taken for bandit problems (as described in Section 3.1.2). The aim in the context of bandit problems is to obtain a strategy for making decisions sequentially for achieving a balance between exploration and exploitation; for example, in experiments involving human participants, one wishes to gain knowledge about the treatment effects, but at the same time, it is necessary to allocate patients to the treatment that is known to be most effective (Villar et al., 2015). The original two-arm bandit problem for Bernoulli responses was considered by Thompson (1933) from a Bayesian perspective, as well as Robbins (1952) from a frequentist point of view. Agrawal and Goyal (2012) produced theoretical results on the expected regret of Thompson's approach. Modifications of the bandit problem with Bernoulli response have been made by Bradt et al. (1956), Berry and Fristedt (1979) and Berry and Viscusi (1981).

There has been recent interest in the application of bandit-type problems for clinical trials. Villar et al. (2015) described bandits with Bernoulli outcomes and provided a simulation study comparing their performance to a purely randomized design. They considered both the finite-horizon case, where a finite total number of patients is assumed, as well as the infinite-horizon case (where $N \to \infty$). They showed that the nonmyopic approach allows a greater proportion of patients to receive the better treatment, but it suffers from having low statistical power to detect differences in treatment effect. They proposed a version of the bandit-type allocation scheme which aims to keep a stable number of patients for both treatments to keep power at a reasonable level. Work by Williamson et al. (2017) developed a bandit approach to treatment allocation for binary response for clinical trials for rare diseases. They proposed a Bayesian adaptive design that aims to maximize the total number of successes in the trial and induces a penalty if each treatment is not given to a minimum number of patients.

No work so far has explicitly mentioned how covariates should be dealt to design experiments with a binary response. In particular, the way in which covariates are incorporated into the construction of the design in a sequential setting can have a large impact on its efficiency.

4.2.2 Preliminaries

We now set up a notation for logistic regression which we then extend to the nonmyopic approach. Assuming a logistic regression for the response, we have

$$y_i \sim \text{Bernoulli}(\pi_i), \tag{4.14}$$

and the probability π_i is given by

$$\pi_i = \frac{\exp \eta_i}{1 + \exp \eta_i}, \tag{4.15}$$

where η_i is the linear predictor. For now, assume the following form for the linear predictor, where we have one binary covariate z_i and one binary treatment t_i:

$$\eta_i = \beta_0 + \beta_1 z_i + \beta_2 t_i, \tag{4.16}$$

where β_0 is the intercept, β_1 is the effect of the covariate and β_2 is the effect of the treatment. We also have

$$\eta_i = \log \frac{\pi_i}{1-\pi_i}. \tag{4.17}$$

The log-likelihood is given by

$$\begin{aligned} l(\beta_0, \beta_1, \beta_2 \mid y_1, ..., y_n) &= \sum_{i=1}^{n} \left(y_i \log \left(\frac{\pi_i}{1-\pi_i} \right) + \log(1-\pi_i) \right) \\ &= \sum_{i=1}^{n} \left(y_i \eta_i - \log(1+\exp \eta_i) \right) \\ &= \sum_{i=1}^{n} \left(y_i \left(\beta_0 + \beta_1 z_i + \beta_2 t_i \right) - \log(1+\exp\left(\beta_0 + \beta_1 z_i + \beta_2 t_i \right)) \right). \end{aligned}$$

The score $\boldsymbol{U}$ is given by the first derivatives of the log-likelihood, with respect to the parameters:

$$\boldsymbol{U} = \begin{pmatrix} \frac{dl}{d\beta_0} \\ \frac{dl}{d\beta_1} \\ \frac{dl}{d\beta_2} \end{pmatrix}, \tag{4.18}$$

where we have

$$\frac{dl}{d\beta_0} = \sum_{i=1}^{n} \left(y_i - \frac{\exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)}{1+\exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)} \right) \tag{4.19}$$

$$= \sum_{i=1}^{n} \left(y_i - \pi_i \right), \tag{4.20}$$

$$\frac{dl}{d\beta_1} = \sum_{i=1}^{n} \left(y_i z_i - \frac{z_i \exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)}{1+\exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)} \right) \tag{4.21}$$

$$= \sum_{i=1}^{n} z_i \left(y_i - \pi_i \right), \tag{4.22}$$

$$\frac{dl}{d\beta_2} = \sum_{i=1}^{n} \left(y_i t_i - \frac{t_i \exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)}{1+\exp(\beta_0 + \beta_1 z_i + \beta_2 t_i)} \right) \tag{4.23}$$

$$= \sum_{i=1}^{n} t_i \left(y_i - \pi_i\right). \tag{4.24}$$

The Hessian, denoted $\boldsymbol{H}$, is the matrix of second derivatives:

$$\boldsymbol{H} = \begin{pmatrix} \frac{d^2 l}{d^2 \beta_0^2} & \frac{d^2 l}{d\beta_0 d\beta_1} & \frac{d^2 l}{d\beta_0 d\beta_2} \\ \frac{d^2 l}{d\beta_1 d\beta_0} & \frac{d^2 l}{d\beta_1^2} & \frac{d^2 l}{d\beta_1 d\beta_2} \\ \frac{d^2 l}{d\beta_2 d\beta_0} & \frac{d^2 l}{d\beta_2 d\beta_1} & \frac{d^2 l}{d\beta_2^2} \end{pmatrix} \tag{4.25}$$

$$= \begin{pmatrix} \sum \pi_i(\pi_i - 1) & \sum z_i \pi_i(\pi_i - 1) & \sum t_i \pi_i(\pi_i - 1) \\ \sum z_i \pi_i(\pi_i - 1) & \sum z_i^2 \pi_i(\pi_i - 1) & \sum t_i z_i \pi_i(\pi_i - 1) \\ \sum t_i \pi_i(\pi_i - 1) & \sum t_i z_i \pi_i(\pi_i - 1) & \sum t_i^2 \pi_i(\pi_i - 1) \end{pmatrix}. \tag{4.26}$$

Now, the expected information matrix $\boldsymbol{I}$ is given by

$$\boldsymbol{I} = -\mathbb{E}\left(\boldsymbol{H}\right) \tag{4.27}$$

$$= \begin{pmatrix} \sum \pi_i(1 - \pi_i) & \sum z_i \pi_i(1 - \pi_i) & \sum t_i \pi_i(1 - \pi_i) \\ \sum z_i \pi_i(1 - \pi_i) & \sum z_i^2 \pi_i(1 - \pi_i) & \sum t_i z_i \pi_i(1 - \pi_i) \\ \sum t_i \pi_i(1 - \pi_i) & \sum t_i z_i \pi_i(1 - \pi_i) & \sum t_i^2 \pi_i(1 - \pi_i) \end{pmatrix}. \tag{4.28}$$

We can write the information matrix above as $\boldsymbol{I} = \boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X}$, where $\boldsymbol{W}$ is a diagonal matrix with entries $\pi_i(1 - \pi_i)$. The following iterative equation provides a scheme for obtaining the maximum likelihood estimates for the parameters $\boldsymbol{\beta} = \begin{pmatrix} \beta_0 & \beta_1 & \beta_2 \end{pmatrix}^\intercal$ above:

$$\boldsymbol{I}^{(m-1)} \boldsymbol{\beta}^m = \boldsymbol{I}^{(m-1)} \boldsymbol{\beta}^{(m-1)} + \boldsymbol{U}^{(m-1)}. \tag{4.29}$$

The information matrix now depends, through $\boldsymbol{W}$, on the values of the model parameters, which was not the case for the linear model in Section 2.6.1. When the true values of the model parameters are unknown, the MLEs can be substituted. The information matrix depends on the values of the model parameters and if estimated parameters are used, it then also depends on the responses. Given a design matrix $\boldsymbol{X}$ and values for the parameters $\boldsymbol{\beta}$, we can compute the matrix $\boldsymbol{W}$ and our D-, D_A-, A- and G-optimal objective functions then become

$$\Psi_D(\boldsymbol{X}, \boldsymbol{\beta}) = \left| (\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X})^{-1} \right|, \tag{4.30}$$

$$\Psi_{D_A}(\boldsymbol{X}, \boldsymbol{\beta}) = \left| \boldsymbol{A}^\intercal \left(\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X} \right)^{-1} \boldsymbol{A} \right|, \tag{4.31}$$

$$\Psi_A(\boldsymbol{X}, \boldsymbol{\beta}) = tr \left(\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X} \right)^{-1}, \tag{4.32}$$

$$\Psi_G(\boldsymbol{X}, \boldsymbol{\beta}) = \max_{\boldsymbol{x} \in \chi_G} \left\{ \boldsymbol{x}^\intercal \left(\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X} \right)^{-1} \boldsymbol{x} \right\}. \tag{4.33}$$

We can also define D-, D_A-, A- and G-efficiencies of a design $\boldsymbol{X}$ relative to another design $\boldsymbol{X}^*$ with parameter values $\boldsymbol{\beta}$ in the logistic model case as

$$\text{Eff}_D = \left\{ \frac{\Psi_D \left(\boldsymbol{X^*}, \boldsymbol{\beta} \right)}{\Psi_D \left(\boldsymbol{X}, \boldsymbol{\beta} \right)} \right\}^{1/p}, \tag{4.34}$$

$$\text{Eff}_{D_A} = \left\{ \frac{\Psi_{D_A} \left(\boldsymbol{X^*}, \boldsymbol{\beta} \right)}{\Psi_{D_A} \left(\boldsymbol{X}, \boldsymbol{\beta} \right)} \right\}^{1/b}, \tag{4.35}$$

$$\text{Eff}_A = \left\{ \frac{\Psi_A \left(\boldsymbol{X^*}, \boldsymbol{\beta} \right)}{\Psi_A \left(\boldsymbol{X}, \boldsymbol{\beta} \right)} \right\}, \tag{4.36}$$

$$\text{Eff}_G = \left\{ \frac{\Psi_G \left(\boldsymbol{X^*}, \boldsymbol{\beta} \right)}{\Psi_G \left(\boldsymbol{X}, \boldsymbol{\beta} \right)} \right\}, \tag{4.37}$$

where p is the number of parameters in the model and b is the number of non-zero rows in the matrix $\boldsymbol{A}$. We note that, in our simulations in this thesis, efficiencies are evaluated using the true parameter values to compute the objective function Ψ. In this way, we are able to compare the designs in terms of their performance with respect to the true parameter values, rather than their respective estimated parameter values.

In a sequential setting, it becomes necessary to observe the response y_i after the treatment for patient i is applied so that the objective function can be calculated. We denote a generic objective function when treatment i is assigned to patient i as

$$\Psi(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}), \tag{4.38}$$

to emphasize the dependence on the vector of previous responses $\boldsymbol{y}_{i-1} = (y_1, y_2, ..., y_{i-1})^\intercal$. We adapt the sequential approach provided in Section 2.6.2 so that the probability that patient i receives treatment 1, which was previously given in Equation (2.37), is changed to

$$\frac{\Psi_i(t_i = 1 \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})^{-1}}{\Psi_i(t_i = 1 \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})^{-1} + \Psi_i(t_i = -1 \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})^{-1}}. \tag{4.39}$$

In Algorithm 4, we provide pseudo-code for constructing a myopic sequential design for the binary response case.

Algorithm 4 Myopic sequential optimal design for binary response: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$, and the number of patients in initial design, n_0

function SeqOptL($\boldsymbol{Z}_n, n_0$)

  Initialization
  Construct initial design $\boldsymbol{X}_{n_0}$ using the exchange algorithm for the first n_0 patients assuming $\boldsymbol{\beta} = \mathbf{0}$.
  Observe responses $y_{n_0} = y_1, y_2, ..., y_{n_0}$
  Fit the model $\boldsymbol{y}_{n_0} \sim glm(\boldsymbol{X}_{n_0}, link = logit)$ to obtain the MLE $\hat{\boldsymbol{\beta}}_0$.

  for i in $n_0 + 1$ to n **do**
    Observe $z_{i,1}, ..., z_{i,k}$
    Calculate $\Psi(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ for each treatment.
    Sample treatment for patient i where probability of treatment 1 is given by Equation (4.39).
    Observe response y_i.
    Refit model $\boldsymbol{y}_i \sim glm(\boldsymbol{X}_i, link = logit)$ and update the parameter estimates $\hat{\boldsymbol{\beta}}_i$.
  end for

  return $\boldsymbol{X}$

end function

4.2.3 Separation

One issue that is important to be aware of with logistic regression is separation. Separation occurs when a linear combination of predictors perfectly predicts the response. It often occurs when predictors are binary (Gelman et al., 2008). Separation can result in the likelihood function becoming monotonic and maximum likelihood estimates of the regression coefficients tending to plus or minus infinity (Rainey, 2016).

There are a number of ways to deal with the problem of separation. Firth (1993) proposed a procedure where maximum likelihood estimates are penalized to reduce bias. Another common approach to stabilizing MLEs is to introduce a prior distribution for the regression coefficients which shrink parameter estimates, particularly large ones, to zero. Gelman et al. (2008) recommended a default prior model to be independent Cauchy distributions

where the intercept has location parameter set to zero and scale parameter set to ten, and for the slope coefficients, the location parameter is set to zero and scale parameter set to two. Other common choices are Jeffry's prior, but Gelman et al. (2008) argue that it is difficult to interpret and computation can be difficult for small datasets. They also show, through cross-validation in three different applications, that Cauchy's prior leads to better predictions than Gaussian or Laplace prior distributions.

We use the function `bayesglm` from the `R` package `arm` (Gelman and Su, 2016) to obtain posterior modes from designs that are built sequentially for Bernoulli GLMs with the use of prior distributions.

To check whether separation is likely to occur in simulations that we will conduct in Section 4.2.5, we look at the posterior modes when designs are fitted sequentially for a logistic GLM where $y_i \sim \text{Bernoulli}(\pi_i)$ with the linear predictor:

$$logit(\pi_i) = 1 + z_i + t_i. \tag{4.40}$$

We generate 100 covariate values z_i, for $1 \leq i \leq 100$, from a Bernoulli(0.5) distribution. Instead of observing y_i as described in Algorithm 4, we generate it from a Bernoulli distribution with parameter π_i given by:

$$\pi_i = \frac{\exp(1 + z_i + t_i)}{1 + \exp(1 + z_i + t_i)}. \tag{4.41}$$

We generate a sequential design using the MLE to estimate the coefficients of the GLM, as well as a sequential design using the Bayesian approach with Cauchy prior distributions to estimate the coefficients. We compare the parameter estimates from both methods. The initial random design has 10 units. We repeat this simulation 100 times. The boxplots in Figure 4.10 show the distribution of the estimates for the intercept $\beta_0 = 1$, coefficient for the covariate $\beta_1 = 1$ and and coefficient for the treatment $\beta = 1$ for the two methods at four different stages of the design: after 15 patients, after 30 patients, after 60 patients and finally, after all 100 patients have entered the trial. Note that the vertical axis changes for each row so that the boxplots corresponding to the MLE approach and the boxplots corresponding to the Bayesian appraoch can be can be compared easily.

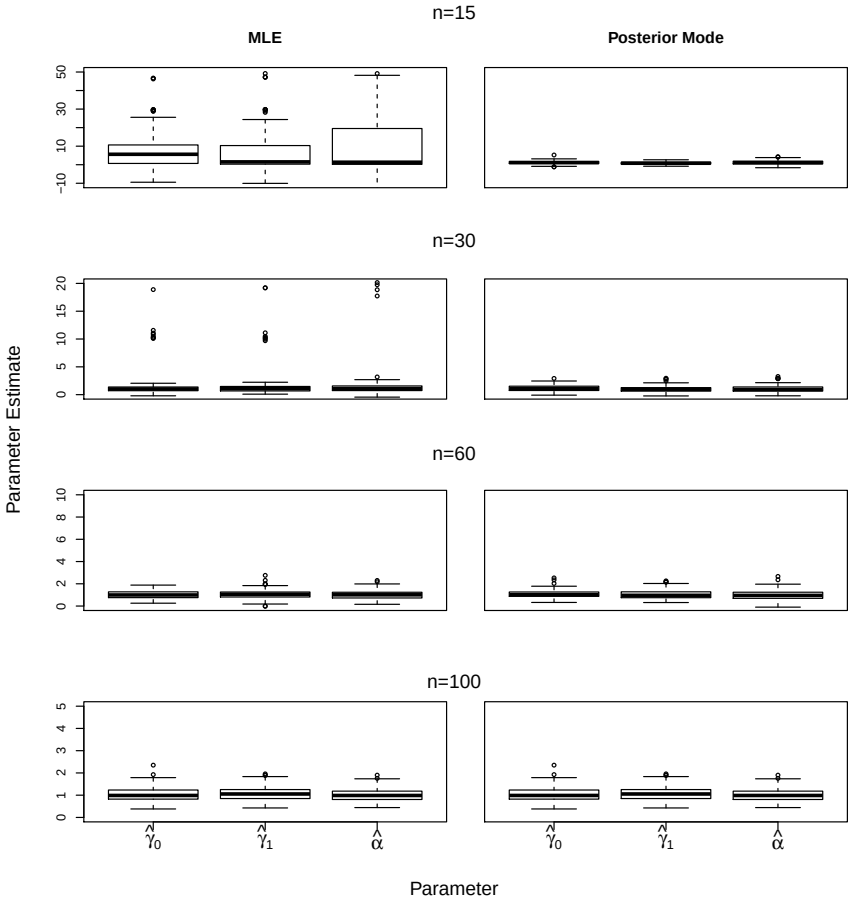

Figure 4.10: Boxplots to show the distribution of the parameter estimates when the MLE is used for model fitting (left panel) versus when a Bayesian approach with a Cauchy prior distribution is used (right panel). The first row displays boxplots of the estimates after 15 patients are treated in the simulation, the second row is after 30 patients, the third after 60 patients and the last row is after all 100 patients have been treated. Note that the vertical axis changes for each row so that the pairs of boxplots can be compared easily.

We observe that there is some evidence of separation occurring at the initial stages of the design: when $n = 15$, we observe that the Bayesian method leads to much more accurate and stable estimates than using the MLE. The advantage of using the Bayesian approach is also evident when $n = 30$. For the plots when $n = 60$ and $n = 100$, it appears that both methods lead to parameter estimates that are close to the true value of 1. The two sets of boxplots appear to be almost identical when $n = 100$. Thus we observe that separation is likely to be a problem for the early stages of the trial, but for the final parameter estimates at the end of the trials, the estimates provided by the two methods are extremely similar.

4.2.4 Nonmyopic approach for the binary response

We now describe the nonmyopic approach for the binary response. The objective function $\Psi(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ now depends on the previous covariates $\boldsymbol{z}_i$, the previous treatments $\boldsymbol{t}_{i-1}$ and also the previous responses $\boldsymbol{y}_{i-1}$ through the parameter estimates. For the sake of simplicity in notation, we assume without loss of generality that we have one covariate. In the nonmyopic approach, we wish to consider the impact of assigning treatment t_i on decisions on future possible patients. For example, for horizon $N = 1$, we consider the expected value of the objective function after $i+1$ patients. Suppose treatment t_i is assigned to patient i. There are two possible values that may be observed for the response y_i. Since Ψ depends on y_i, we need to consider the two possible responses that y_i may take, and then consider the possible values that z_{i+1} can take. For a given covariate value z_{i+1} for patient $i+1$, we denote by $t^*_{i+1}(z_{i+1}, t_i, y_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ the optimal choice of treatment for patient $i+1$ given z_{i+1} and t_i:

$$t^*_{i+1}(z_{i+1}, t_i, y_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}) = \underset{t_{i+1}}{\operatorname{argmin}}\, \Psi(t_{i+1} \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_i). \tag{4.42}$$

From here on, we suppress the conditioning and write $t^*_{i+1}(z_{i+1}, t_i, y_i)$ for simplicity. Now, we take the expectation of the objective function over two possible responses which may be obtained to find an expected value of the objective function over the unknown response:

$$\begin{aligned}\mathbb{E}_{y_i}\Psi(t_{i+1} \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_i) =& \mathbb{P}(y_i = 0 \mid \boldsymbol{z}_i, \boldsymbol{t}_i, \boldsymbol{y}_{i-1})\Psi(t_{i+1} \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_{i-1}, y) \\ &+ \mathbb{P}(y_i = 1 \mid \boldsymbol{z}_i, \boldsymbol{t}_i, \boldsymbol{y}_{i-1})\Psi(t_{i+1} \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_{i-1}, y),\end{aligned}$$

where $y_i \sim \text{Bernoulli}(\pi_i)$ with π_i given by:

$$\pi_i = \frac{\exp\left(\hat{\boldsymbol{\beta}}_{i-1}\,(1\ \boldsymbol{z}_i^\intercal\ t_i)^\intercal\right)}{1 + \exp\left(\hat{\boldsymbol{\beta}}_{i-1}\,(1\ \boldsymbol{z}_i^\intercal\ t_i)^\intercal\right)}. \tag{4.43}$$

Now, we denote by $\mathbb{P}(z_i = z)$ the probability that the ith patient has covariate value z. We assume that these probabilities do not depend on i. We denote by $\Psi_1(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ the expected value of the objective function when treatment t_i is assigned to patient i, taking into account the impact of the decision on one further decision in the future. We obtain an expectation over the possible covariate combinations of the optimality criterion:

$$\Psi_1(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}) = \mathbb{E}_{z_{i+1}}\mathbb{E}_{y_i}\Psi(t^*_{i+1}(\boldsymbol{z}_{i+1}, t_i, y_i) \mid \boldsymbol{z}_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_i) \tag{4.44}$$

$$= \sum_z \mathbb{P}(z_{i+1} = z)\mathbb{E}_{y_i}\Psi(t^*_{i+1}(z_{i+1}, t_i, y_i) \mid \boldsymbol{z}_i, z_{i+1}, \boldsymbol{t}_i, \boldsymbol{y}_i). \quad (4.45)$$

For a horizon greater than 1, we can use the following recursive relationship to find the optimal treatment for patient ν. The expected value of the objective function after $\nu + N$ patients, when treatment t_ν has been assigned, is given as follows:

For $N > 0$:

$$\begin{aligned}\Psi_N(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}, \boldsymbol{y}_{\nu-1}) =& \mathbb{E}_{z_{\nu+1}}\mathbb{E}_{y_\nu}\Psi_{N-1}(t^*_{\nu+1}(\boldsymbol{z}_{\nu+1}, t_\nu, y_\nu) \mid \boldsymbol{z}_{\nu+1}, \boldsymbol{t}_\nu, \boldsymbol{y}_\nu) \\ =& \sum_z \mathbb{P}(z_{\nu+1} = z)\mathbb{E}_{y_\nu}\Psi_{N-1}(t^*_{\nu+1}(\boldsymbol{z}_{\nu+1}, t_\nu, \boldsymbol{y}_\nu) \mid \boldsymbol{z}_\nu, z_{\nu+1}, \boldsymbol{t}_\nu, \boldsymbol{y}_\nu),\end{aligned} \quad (4.46)$$

and for $N = 0$, we have

$$\Psi_0(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}, \boldsymbol{y}_{\nu-1}) = \Psi(t_\nu \mid \boldsymbol{z}_\nu, \boldsymbol{t}_{\nu-1}, \boldsymbol{y}_{\nu-1}), \quad (4.47)$$

which is simply the myopic loss after ν patients.

We note that the nonmyopic approach for the logistic model case is considerably more computationally intensive than the myopic approach; Table 4.1 compares the running time in seconds for constructing a 100-patient design with one covariate for the linear and logistic case for a range of values for the horizon.

Table 4.1: Running time in seconds for constructing a design with one covariate and 100 patients. the initial design is one patient for the linear model case, and 10 patients for the logistic model case. We compare the myopic approach (horizon= 0) against the nonmyopic approach with horizon set to 1, 2, 3, 4, 5, and 10. Simulations were performed on a machine with a 2GHz processor and 64 GB of memory.

Horizon	**Linear Model**		**Logistic Model**	
	Static	**Dynamic**	**Static**	**Dynamic**
0	0.045	0.05	0.84	0.749
1	0.313	0.435	4.926	4.7
2	1.297	1.867	30.116	31.054
3	5.433	7.701	264.138	253.888
4	22.234	31.254	1916.986	2090.021
5	89.258	123.703	13929.6	16761.77
10	84671.05	120538.3	>60 hours	>60 hours

4.2.5 Simulations

Initial design

For the binary response case, there are a few details in the construction of the simulations that we did not have to consider in the linear model case. First, as we explained in Section 4.2.1, for the logistic model, the information matrix and the objective function depend on values of the model parameters. Therefore, estimates of parameters are needed in order to design the experiment aimed to estimate these parameters in the first place. We begin with an initial design where we use the exchange algorithm to allocate treatments to 10 units, under the assumption that $\boldsymbol{\beta}$ is a vector of zeros. We then generate the responses for the first 10 patients and fit the model to obtain the first estimate $\hat{\boldsymbol{\beta}}_0$. The `R` function `bayesglm` is used to fit the model for our simulations involving logistic regression to avoid problems in estimation due to separation. The prior used is a Cauchy distribution with mean zero and scale set to two for both the treatment and covariate parameters. When we compare sequential designs for the same set of covariates constructed with the same objective function, we make sure that the designs have the same initial design. Since our purpose is to compare myopic and nonmyopic approaches to decision making, we wish to control as many other factors as possible which can potentially add variability to the results.

Generating Responses

Another source of variability is the generation of the responses. Unlike the linear model case, for the logistic model, we need responses in order to generate the estimates of the model parameters and subsequently to evaluate the design under the objective function. When comparing sequential designs for the same set of covariates constructed with the same objective function, we generate the responses in the following way:

1. Generate a deviate u_i from the Unif$(0, 1)$ distribution.
2. Set

$$y_i = \begin{cases} 1 & \text{if } u_i \geq \pi_i \\ 0 & \text{if } u_i < \pi_i \end{cases}. \tag{4.48}$$

The deviates u_i are the same for simulations which compare the same objective function, so that we can ensure that the data generating mechanism is the same. In this way, we try to make sure that the only difference across two simulations is the approach to decision making: taking either a myopic approach or a nonmyopic approach.

Set up

We compare D_A-, D- and G-optimal designs for logistic regression that are constructed sequentially using myopic and nonmyopic methods. We run a simulation study in which 100 units of a covariate z are generated. The covariate has support $\{-1, 1\}$ and is generated such that $\mathbb{P}(z_i = 1) = 0.5$ and $\mathbb{P}(z_i = -1) = 0.5$ for all i. We assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = z_i + t_i$, and generate responses according to this model. We use this structure for the simulations:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) 100 deviates from a Unif$(0, 1)$ distribution are generated for the response.

(c) An initial design with 10 units is constructed using the exchange algorithm with D_A optimality as the objective function.

(d) Seven sequential designs using the covariates, random deviates for the responses, and initial design in part (a) are constructed using:

- A myopic D_A-optimal design.
- A nonmyopic D_A-optimal design with horizon $N = 1$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 1$, with the empirical covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 2$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 2$, with the empirical covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 3$, with the correct covariate distribution assumed.
- A nonmyopic D_A-optimal design with horizon $N = 3$, with the empirical covariate distribution assumed.

(e) Designs are evaluated using the performance measure Ψ_{D_A}, given by Equation (4.31), at each sample size between 10 and 100, inclusive. The true values of the parameters are used to calculate Ψ_{D_A}.

II (a)-(e) above is above 20 times to obtain a distribution of the performance measure for each sample size.

We run the simulation above using performance measures Ψ_D and Ψ_G, given by Equations (4.30) and (4.33), respectively.

Results for one covariate

We first consider the estimates of $\boldsymbol{\beta}$ produced by the seven different designs. Figure 4.11 shows the distributions of $\hat{\boldsymbol{\beta}}$ at each sample size between 11 and 100. The estimates appear to be centered around their true value, $\boldsymbol{\beta} = (0, 1, 1)^\intercal$, and the plots appear to be very similar across the seven methods.

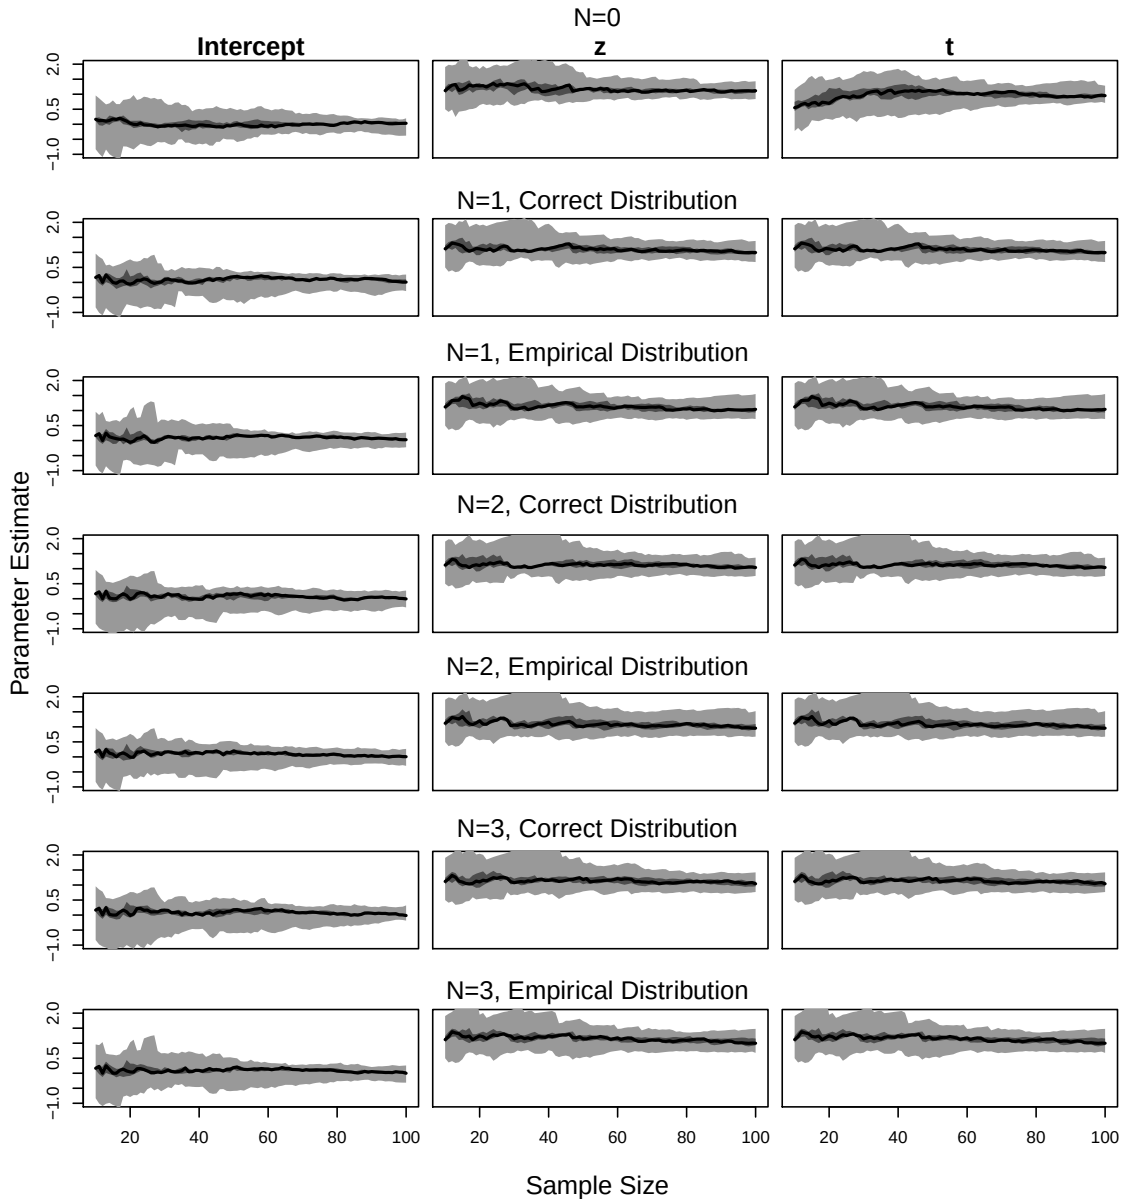

Figure 4.11: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one covariate are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

In Figure 4.12, we plot the distribution of Ψ_{D_A} for each sample size between 11 and 100. We observe that the value of this objective function decreases as sample size increases, as expected. We note that the plots look extremely similar across the seven methods. There

is is no noticeable difference between having horizon equal to one or three. In Figure 4.13, we plot the relative efficiencies of the nonmyopic designs against the myopic design, which again confirms that there is no observable difference across the methods in Ψ_{D_A}; the myopic approach is slightly more efficient for when sample size is below 30.

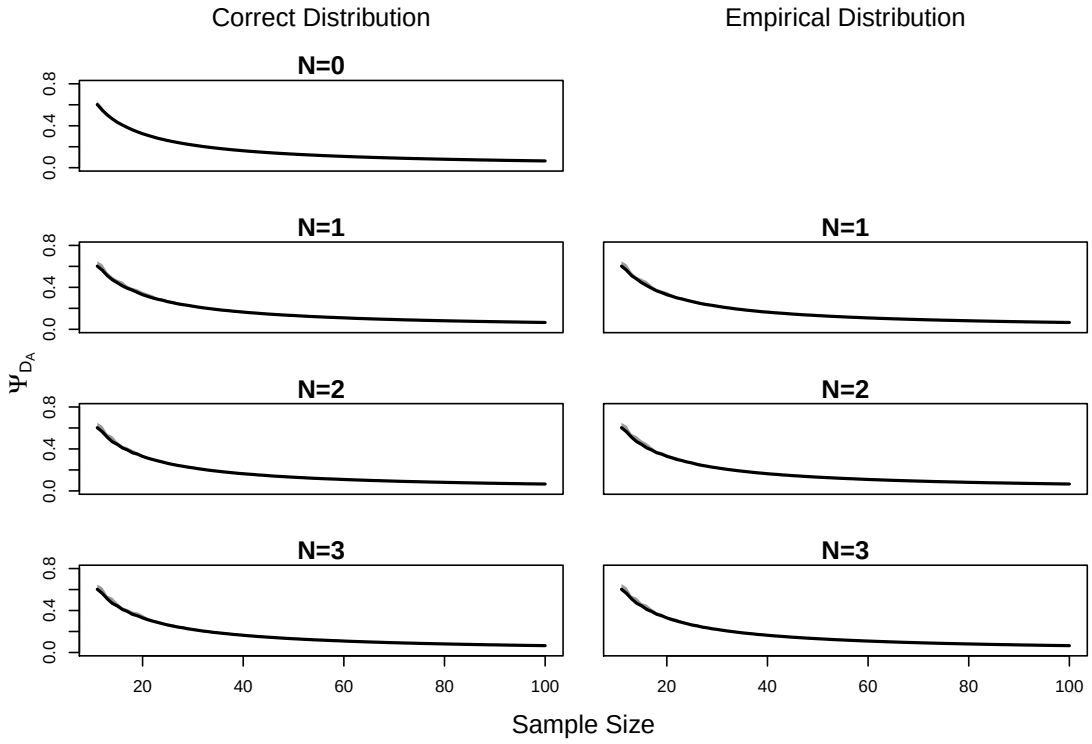

Figure 4.12: Distributions of Ψ_{D_A} for designs for the logistic model for one covariate are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

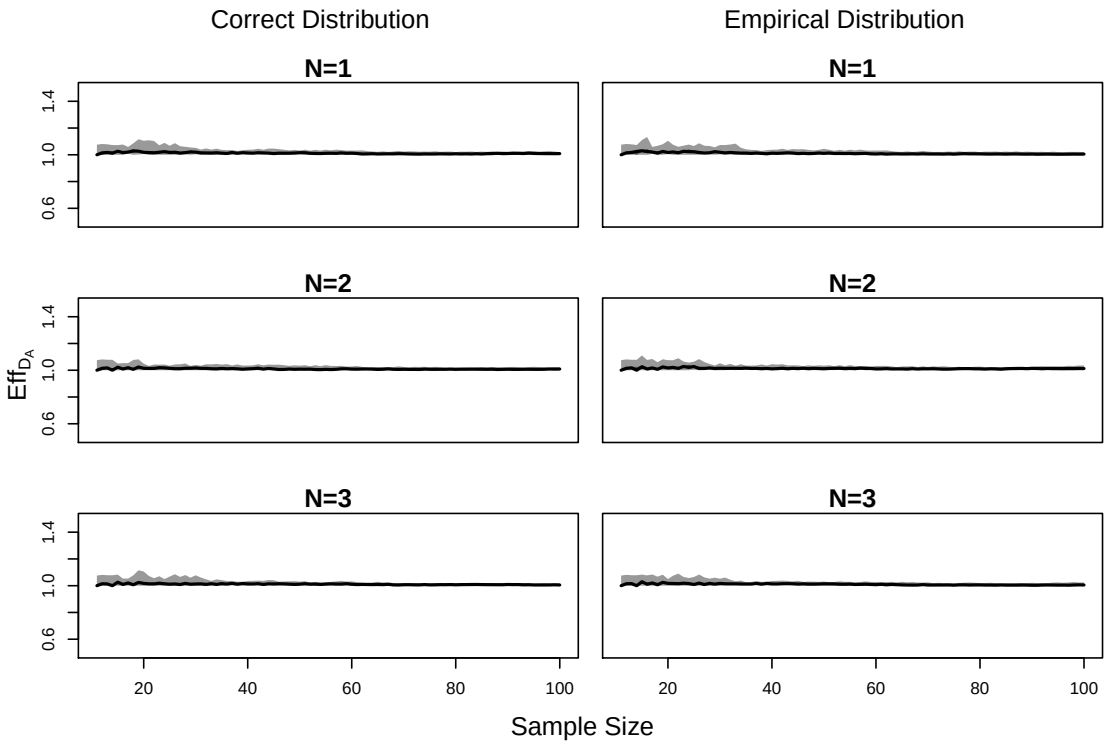

Figure 4.13: Distributions of the relative efficiencies of the nonmyopic D_A-optimal designs against the myopic D_A-optimal designs for the logistic model for one covariate are plotted against sample size. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

For the case with D-optimality, we see in Figure 4.14 that there appears to be more variability in the estimates for the nonmyopic approach, particularly in for the intercept and the treatment effect.

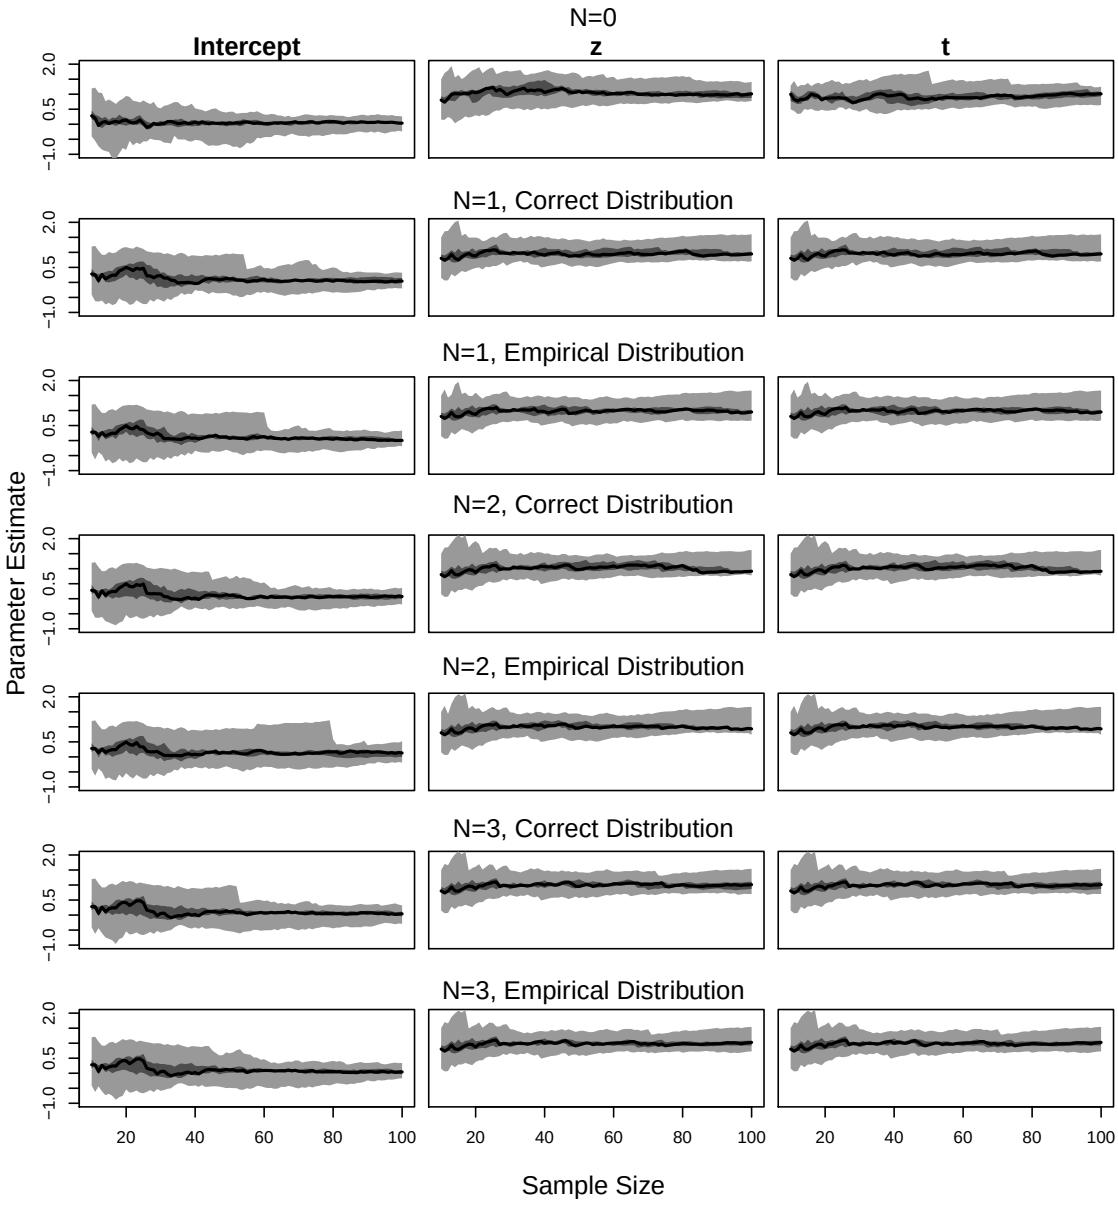

Figure 4.14: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one covariate are plotted against sample size. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

In Figure 4.15, it appears that the plots of the distributions of Ψ_D look extremely similar across the seven methods with very little variability across the 20 simulations. However, in Figure 4.16, we plot the relative efficiencies of the nonmyopic designs against the myopic design, which then shows that the myopic approach is more efficient than the nonmyopic

approaches for all sample sizes. For D-optimality, we are interested in not only the treatment effect but also the intercept and the effect of the covariate; since there is greater variability in the estimates of the intercept for the nonmyopic approach, this explains why the plot for D-efficiency is much more variable than that of D_A-efficiency.

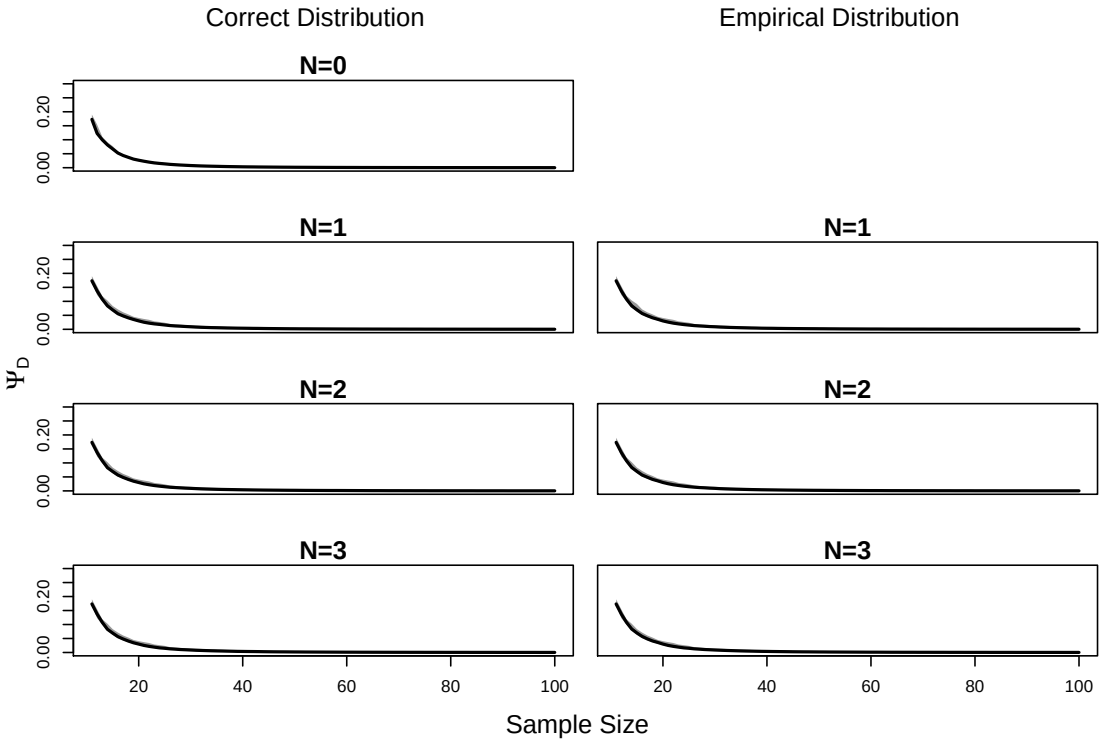

Figure 4.15: Distributions of Ψ_D for designs for the logistic model for one covariate are plotted against sample size. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

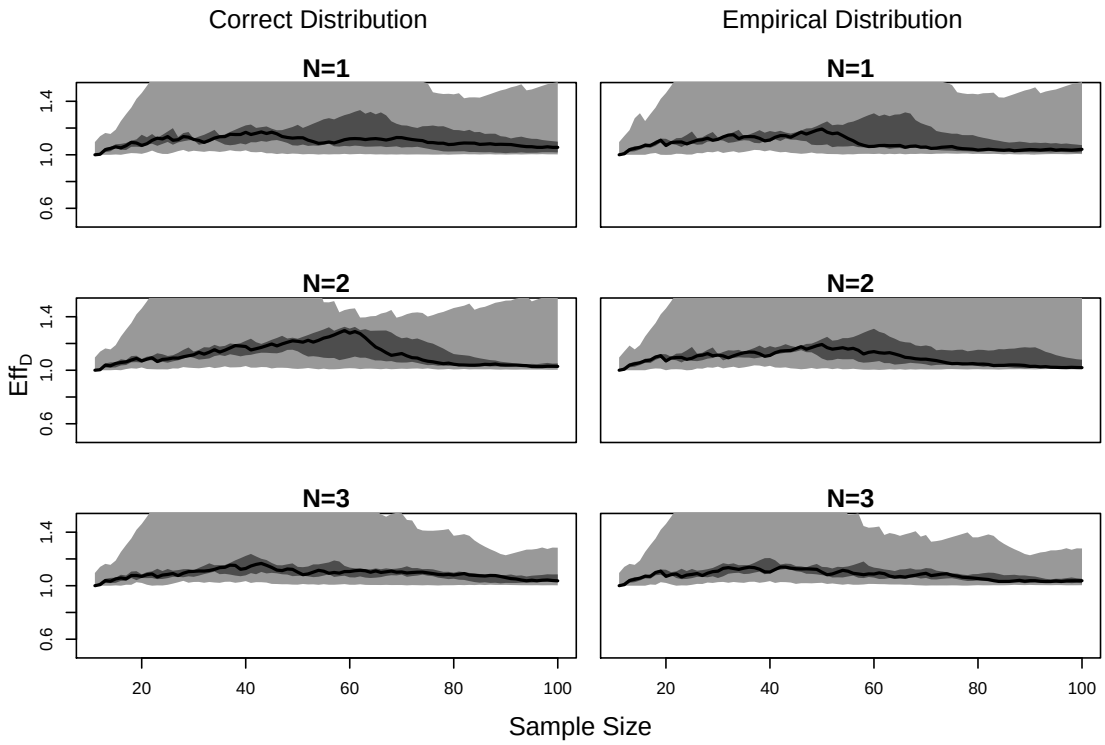

Figure 4.16: Distributions of the relative efficiencies of the nonmyopic D-optimal designs against the myopic D-optimal designs for the logistic model for one covariate are plotted against sample size. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

For the G-optimal designs, we observe in Figure 4.17 that the distributions of $\hat{\boldsymbol{\beta}}$ are similar across the seven methods.

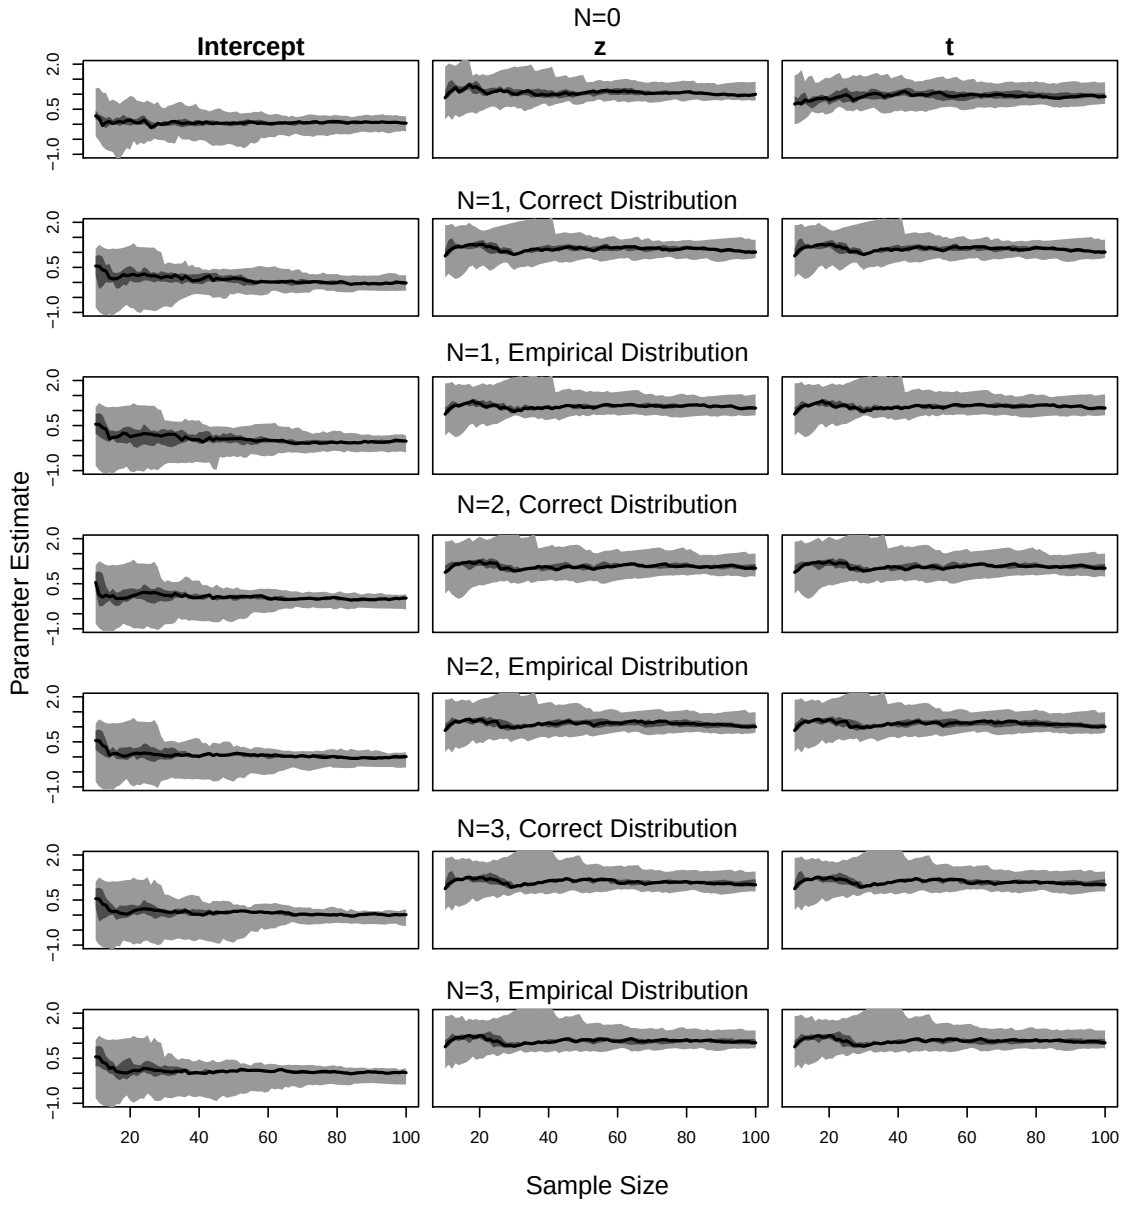

Figure 4.17: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one covariate are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

In Figure 4.18, we plot the distribution of Ψ_G for each sample size between 11 and 100. We observe that there is greater variability in the plots for the G-optimal designs compared to the D_A- and D-optimal designs. In Figure 4.19, we plot the relative efficiencies of the nonmyopic designs against the myopic design, which confirms again that the myopic approach is more efficient than the nonmyopic approaches.

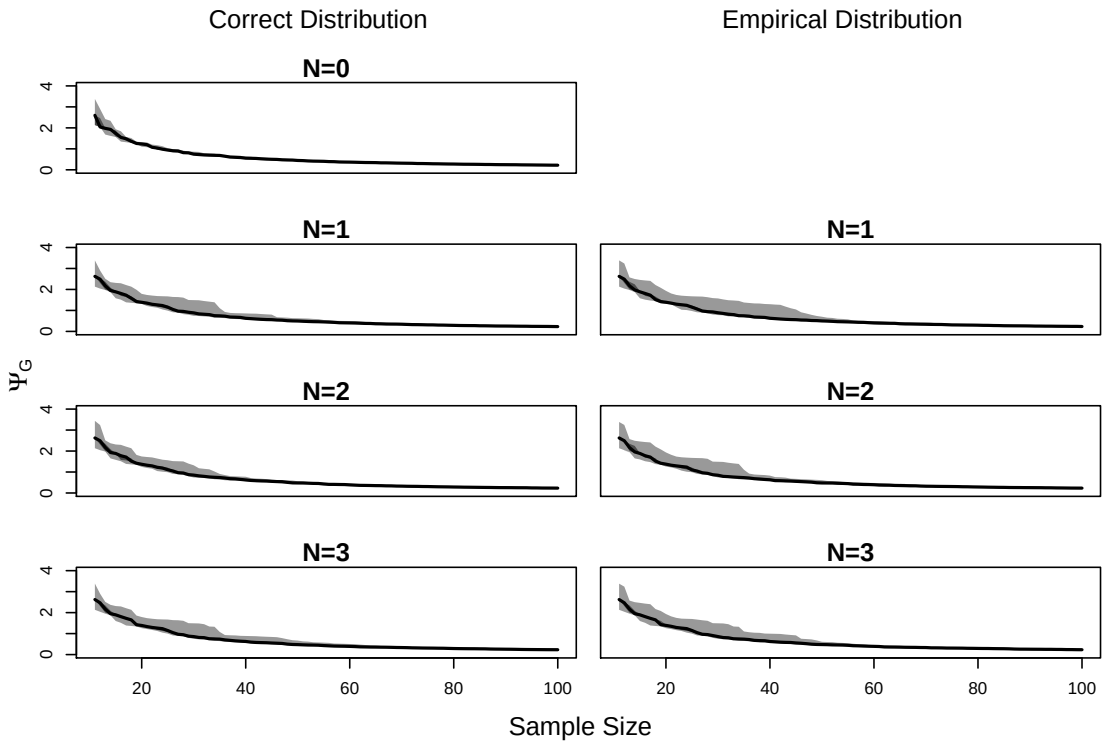

Figure 4.18: Distributions of Ψ_G for designs for the logistic model for one covariate are plotted against sample size. We show the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

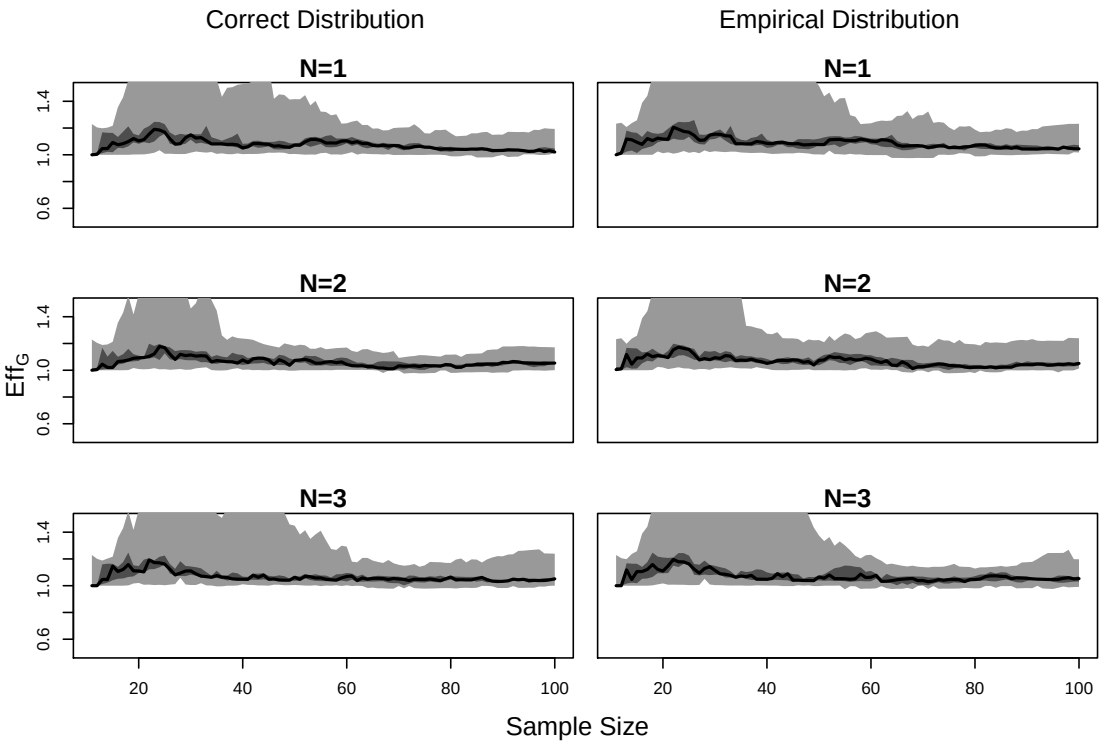

Figure 4.19: Distributions of the relative efficiencies of the nonmyopic G-optimal designs against the myopic G-optimal designs for the logistic model for one covariate are plotted against sample size. We show the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline. The black line indicates the median, the dark grey is for the $40\% - 60\%$ region of the distribution and the light grey is the $10\% - 90\%$ region.

We considered two further cases: one involves a single dynamic covariate, and the second involves two static covariates with treatment-covariate interactions assumed. The results are in Appendix E. In all cases, we found that the myopic approach is the most efficient.

In Chapter 9, we provide an overview of simulation results from all parts of the thesis and the overall conclusions that we draw from them.

4.3 Conclusion

In this chapter, we extended the work of Atkinson (1982) so that it can be applied to any optimality criterion under the linear and logistic model. We further developed it under a nonmyopic framework which can be applied to the linear model or the logistic model.

Based on simulations investigating the linear model case and the D-, D_A- and G-optimality criteria where we have one or two covariates and we consider a horizon of one up to five, there appears to be little to no benefit to using the nonmyopic approach over the myopic approach when there are static covariates, and slight benefit when there is a dynamic covariate.

For the logistic model case, we found that there is some evidence of separation occurring at the early stages of the trial and there is a benefit to putting prior distributions on the model parameters and taking a Bayesian approach to model fitting. Our set of simulation studies for the binary response shows that there is no evidence of an advantage for using the nonmyopic approach when the horizon is set to one, two, or three.

It is possible that, in the cases that we have considered, the horizon is too low for any benefit to be observed. It appears that the nonmyopic approach can lead to greater variability in the objective function because some information about future patients are considered and ends up being less efficient than the myopic approach. Further, the nested expectations and optimizations make it a computationally expensive approach and higher values of horizons would be difficult to implement. In Table 4.1, we see that the running times to construct a design with 100 patients grow quickly as the horizon increases. These results motivate us to develop an approximation to the nonmyopic approach which we introduce in Chapter 5.

Chapter 5

Pseudo-nonmyopic Approach

In Chapter 4, we considered the nonmyopic approach for allocating treatments to patients, where we compute the expected loss after N future patients are enrolled in the trial. One main limitation of this approach is that computing the nested expectations and minimizations over unknown quantities, such as in Equation (4.6) or Equation (4.46), requires recursive formulae which are computationally expensive. The number of calculations increases exponentially with each additional future patient in the horizon and, as a result, our simulations considered examples with horizon no more than five. We now explore a *pseudo-nonmyopic* approach which involves evaluating a related objective function with a similar aim without the use of recursion. The computational burden is reduced as nested expectations and minimizations are not necessary but we are still able to incorporate information about future possible decisions. We describe this novel approach for both the linear and logistic model cases, providing simulations to show how it compares to the myopic approach.

5.1 Linear Model

In this section, we use the notation $\boldsymbol{t}_{i:j}$ to indicate a vector of treatments allocated from the ith to the jth patient:

$$\boldsymbol{t}_{i:j} = (t_i, t_{i+1}, t_{i+2}, ..., t_j)^\intercal . \tag{5.1}$$

Similarly, for a matrix of covariate values from the ith to the jth patient, we write:

$$\boldsymbol{Z}_{i:j} = \begin{pmatrix} \boldsymbol{z}_i^\intercal \\ \boldsymbol{z}_{i+1}^\intercal \\ \vdots \\ \boldsymbol{z}_j^\intercal \end{pmatrix}, \tag{5.2}$$

where, as before, $\boldsymbol{z}_m = (z_{m,1}, ..., z_{m,k})^\intercal$ is a vector containing the k covariate values for patient m, for $m \in \{i, ..., j\}$.

In the pseudo-nonmyopic approach, in order to make a decision about the treatment of the ith patient, we generate M possible trajectories of covariate values for patient $i+1$ until patient n. We assume, as for the non-myopic approach, that we have a distribution $f_{\boldsymbol{z}}$ for the covariate $\boldsymbol{z}$. This may be the true distribution in the population (if it is known), or an empirical approximation based on the patients in the trial up until the ith patient. The covariate distribution may depend on time, in which case we refer to it as a dynamic covariate. For each of the M trajectories, we construct a *pseudo-design* in which we have the i patients and $(n-i-1)$ patients in the trajectory, and treatments allocated using an approach that we describe below. We look at the average losses of the M pseudo-designs where we assign $t_i = 1$, and compare it to the average loss of the M pseudo-designs when $t_i = -1$; we select t_i according to a probability that is weighted by these average losses.

This approach takes averages over simulated values of the covariates for patients $i+1$ up to n. Optimization based on Monte Carlo simulations of unknown quantities is typically conducted in a Bayesian setting for design of experiments (Woods et al., 2017), where values of the unknown parameters may be simulated from a prior distribution. See Gentle (2003) for an overview of Monte Carlo methods and Ryan (2003) for an application to Bayesian design of experiments.

We now describe the procedure in more detail. To allocate a treatment for patient i, for $i \in \{1, 2, ..., n-1\}$, we observe $\boldsymbol{z}_i$, and assume that we allocate treatment t_i, for $t_i \in \{-1, 1\}$. Based on the assumed covariate distribution $f_{\boldsymbol{z}}$, we generate M trajectories of covariate values for patient $i+1$ up to patient n. Let us denote by $\boldsymbol{z}_j^M$ the vector of covariate values for patient j, for $j \in \{i+1, ..., n\}$, in the mth trajectory, where $m \in \{1, ..., M\}$. We denote the matrix of covariate values in the mth trajectory as $\boldsymbol{Z}_{(i+1):n}^m$:

$$\boldsymbol{Z}_{(i+1):n}^m = \begin{pmatrix} \boldsymbol{z}_{i+1}^{m\intercal} \\ \vdots \\ \boldsymbol{z}_n^{m\intercal} \end{pmatrix}. \tag{5.3}$$

Then, we allocate treatments for the mth trajectory sequentially from the $i+1$th patient to the nth patient. Given the covariates of the first patient in the trajectory, $\boldsymbol{z}_{i+1}^m$, we choose the treatment $t_{i+1}^{*^m}$ which minimizes the objective function Ψ given t_i, and the treatments and covariates of previous patients:

$$t_{i+1}^{*^m}(\boldsymbol{z}_{i+1}^m, t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}) = \underset{t_{i+1}}{\operatorname{argmin}}\, \Psi\left(t_{i+1} \mid \boldsymbol{Z}_i, \boldsymbol{z}_{i+1}^m, \boldsymbol{t}_{i-1}, t_i\right). \tag{5.4}$$

To allocate a treatment for the jth patient in the trajectory, $j \in \{i+2, i+3, i+4, ..., n\}$, we assume that $\boldsymbol{t}_{(i+1):(j-1)}^{*^m}$ have been allocated to patients $i+1$ up to $j-1$ and choose the treatment $t_j^{*^m}$:

$$t_j^{*^m}\left(\boldsymbol{z}_j^m, t_{j-1}^* \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):(j-1)}^m, \boldsymbol{t}_{i-1}, t_i, t_{(i+1):(j-2)}^*\right) \tag{5.5}$$

$$= \underset{t_j}{\operatorname{argmin}}\, \Psi\left(t_j \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):(j)}^m, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}_{(i+1):(j-1)}^*\right). \tag{5.6}$$

For the mth trajectory, we obtain a pseudo-design which includes patient 1 up to i, as well as future patients $i+1$ up to n, assuming that the ith treatment is 1. Analogously, we can obtain the design where the ith patient receives treatment -1. We denote the objective function of the two designs as follows:

$$\Psi\left(t_i \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):n}^m, \boldsymbol{t}_{i-1}, t_i = 1, \boldsymbol{t}_{(i+1):(n-1)}^{*^m}\right), \tag{5.7}$$

$$\Psi\left(t_i \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):n}^m, \boldsymbol{t}_{i-1}, t_i = -1, \boldsymbol{t}_{(i+1):(n-1)}^{*^m}\right). \tag{5.8}$$

We define the average objective function across the M designs, assuming, firstly, that $t_i = 1$, and secondly, that $t_i = -1$, for $i = 1, ..., n-1$:

$$\overline{\Psi}(t_i = 1) = \frac{1}{M}\sum_{m=1}^{M} \Psi\left(t_i \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):n}^m, \boldsymbol{t}_{i-1}, t_i = 1, \boldsymbol{t}_{(i+1):(n-1)}^{*^m}\right), \tag{5.9}$$

$$\overline{\Psi}(t_i = -1) = \frac{1}{M}\sum_{m=1}^{M} \Psi\left(t_i \mid \boldsymbol{Z}_i, \boldsymbol{Z}_{(i+1):n}^m, \boldsymbol{t}_{i-1}, t_i = -1, \boldsymbol{t}_{(i+1):(n-1)}^{*^m}\right). \tag{5.10}$$

To select the treatment for patient n, we do not generate any future covariates. We simply compute the objective function:

$$\overline{\Psi}(t_n = t) = \Psi\left(t_n = t \mid \boldsymbol{Z}_n, \boldsymbol{t}_{n-1}\right). \tag{5.11}$$

We sample t_i from the set $\{-1, 1\}$ where the probability of selecting 1 is given by

$$\frac{\overline{\Psi}(t_i = 1)^{-1}}{\overline{\Psi}(t_i = 1)^{-1} + \overline{\Psi}(t_i = -1)^{-1}}. \tag{5.12}$$

This is the treatment that we allocate to patient i, for $i \in \{1, 2, ..., n-1\}$. Pseudocode is provided in Algorithm 5.

Algorithm 5 Sequential pseudo-nonmyopic design for continuous response: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$,and the distributions of the k covariates, $f_{\boldsymbol{z}_1}, ..., f_{\boldsymbol{z}_k}$

function PSEUDONONMY($\boldsymbol{Z}_n, f_{\boldsymbol{z}_1}, ..., f_{\boldsymbol{z}_k}$)

  for i in 1 to n **do**

    Observe $\boldsymbol{z}_i = (z_{i,1}, z_{i,2}, ..., z_{i,k})^\intercal$

    for t_i in $\{-1, 1\}$ **do**

      for m in 1 to M **do**

        Generate $\boldsymbol{Z}^m_{(i+1):n}$ using the assumed covariate distributions $f_{\boldsymbol{z}_1}, ..., f_{\boldsymbol{z}_k}$.

        Allocate treatments $t^{*^m}_{i+1}, t^{*^m}_{i+2}, ..., t^{*^m}_n$ along the trajectory:

$$t^{*^m}_{i+1}(\boldsymbol{z}^m_{i+1}, t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}) = \underset{t_{i+1}}{\operatorname{argmin}}\, \Psi\left(t_{i+1} \mid \boldsymbol{Z}_i, \boldsymbol{z}^m_{i+1}, \boldsymbol{t}_{i-1}, t_i\right). \tag{5.13}$$

$$\text{For } j \in \{i+2, i+4, ..., n\}: \quad t^{*^m}_j(\boldsymbol{z}^m_j, t^*_{j-1} \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j-1)}, \boldsymbol{t}_{i-1}, t_i, t^*_{(i+1):(j-2)}) = \underset{t_j}{\operatorname{argmin}}\, \Psi\left(t_j \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j)}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}^*_{(i+1):(j-1)}\right). \tag{5.14}$$

      end for

      Define $\overline{\Psi}(t_i) = \frac{1}{M} \sum_1 \Psi\left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}^{*^m}_{(i+1):(n-1)}\right)$

    end for

    Sample t_i from the set $\{-1, 1\}$ where the probability of selecting 1 is given by:

$$\frac{\overline{\Psi}(t_i = 1)^{-1}}{\overline{\Psi}(t_i = 1)^{-1} + \overline{\Psi}(t_i = -1)^{-1}}. \tag{5.15}$$

  end for

  $\boldsymbol{t} = (t_1, t_2, ..., t_n)^\intercal$

  $\boldsymbol{X} = [\boldsymbol{1}\ \boldsymbol{Z}\ \boldsymbol{t}]$ ▷ Design matrix

  return $\boldsymbol{X}$

end function

In Appendix H, we provide a simple example to illustrate the three model-based approaches for the linear model we consider in this thesis: the myopic approach, the nonmyopic approach and the pseudo-nonmyopic approach.

5.1.1 Simulations

We run two simulations to compare the myopic approach to the pseudo-nonmyopic approach. In this chapter, we focus on D-optimality as our objective function, but in Chapter 7, we look into other criteria that focus on specific combinations of covariates and treatments of interest. We would like to observe whether there is any benefit to using the pseudo-nonmyopic approach for the linear model case, and we would like to assess the role of M in the variability of the results. In the first simulation, we consider the case of a single binary covariate with no interaction assumed between treatment and covariate. The model is given by:

$$\mathbb{E}(y) = \beta_0 + \beta_1 \boldsymbol{z} + \beta_2 \boldsymbol{t}. \tag{5.16}$$

In the second simulation, we consider a more complex case where there are two covariates, one static and one dynamic, and treatment-covariate interactions.

Our simulations have the following structure:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) We consider the following five designs for the patients in part (a):

- A myopic D-optimal design.
- A pseudo-nonmyopic D-optimal design with $M = 10$.
- A pseudo-nonmyopic D-optimal design with $M = 50$.
- A pseudo-nonmyopic D-optimal design with $M = 100$.
- A pseudo-nonmyopic D-optimal design with $M = 200$.

(c) Designs are evaluated using the performance measure Ψ_D, given by Equation (2.18), at each sample size between 2 and 100, inclusive.

II (a)-(c) above is repeated 100 times to obtain a distribution of the performance measure for each sample size.

Example 1

In the first example, we generate the covariates from a Bernoulli(0.5) distribution. In the top row of Figure 5.1, we plot the distributions of Ψ_D evaluated at each sample size between 2 and 100 for each of the five different approaches. We observe that they appear to be similar across the five approaches. In the top row of Figure 5.1, we plot the distributions of

Ψ_D against sample size. The values of the objective function drop dramatically within 20 patients and appears to reach a stable value after that. We see very little variation across the five plots here, largely due to the scaling of the plot. In the bottom row of Figure 5.1, we plot the distributions of relative D-efficiencies (see Equation (2.24)) for the pseudo-nonmyopic approaches against the myopic approach. Initially, the myopic approach is more efficient than the pseudo-nonmyopic approach, but this difference becomes smaller as the sample size increases. It seems that, in this example, M does not appear to influence the variability of the plots in any way; there seems to be no reduction in performance when $M = 10$ as opposed to having $M = 200$.

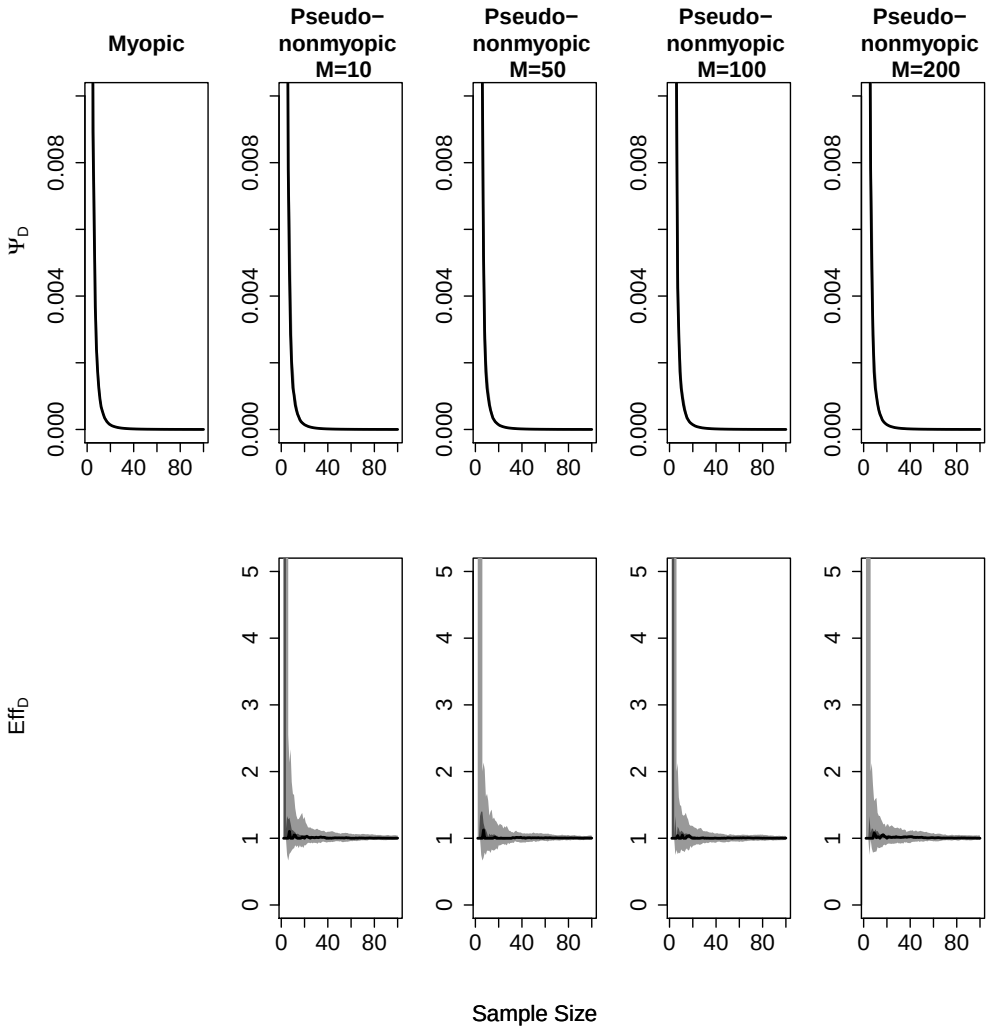

Figure 5.1: Top row: D-optimality against sample size for designs for a linear model with one static covariate. Bottom row: relative D-efficiency against sample size for designs for a linear model with one static covariate. Values below 1 indicate that the pseudononmyopic approach is more beneficial than the myopic approach. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

Example 2

Next, we consider an example with two covariates and treatment-covariate interactions to see whether we draw the same conclusion. We generate covariate values for 100 patients. There are two covariates; the first covariate is dynamic and is generated such that $\mathbb{P}(z_{1,i} = 1) = 0.01i$. The second covariate is static as the probability of its value being equal to 1 is constant for all i; we have that $\mathbb{P}(z_{2,i} = 1) = 0.3$.

The model is given by:

$$\mathbb{E}(y) = \beta_0 + \beta_1 \boldsymbol{z}_{\cdot 1} + \beta_2 \boldsymbol{z}_{\cdot 2} + \beta_3 \boldsymbol{t} + \boldsymbol{\beta}_4 \boldsymbol{z}_{\cdot 1} \boldsymbol{t} + \beta_5 \boldsymbol{z}_{\cdot 2} \boldsymbol{t}. \tag{5.17}$$

We see in the top row of Figure 5.2 that, again, the values of Ψ_D evaluated at each sample size appears to be similar across all methods. The value of the objective function is initially very large but it drops dramatically by sample size of around 30 and reaches a stable value after that. The bottom row of Figure 5.2 are plots of the relative efficiencies and they reveals that, on average, the myopic approach is still more efficient than the pseudo-nonmyopic approach. In particular, there is a very noticeable improvement in the myopic approach compared to the pseudo-nonmyopic approach when sample size is between 1 and 50.

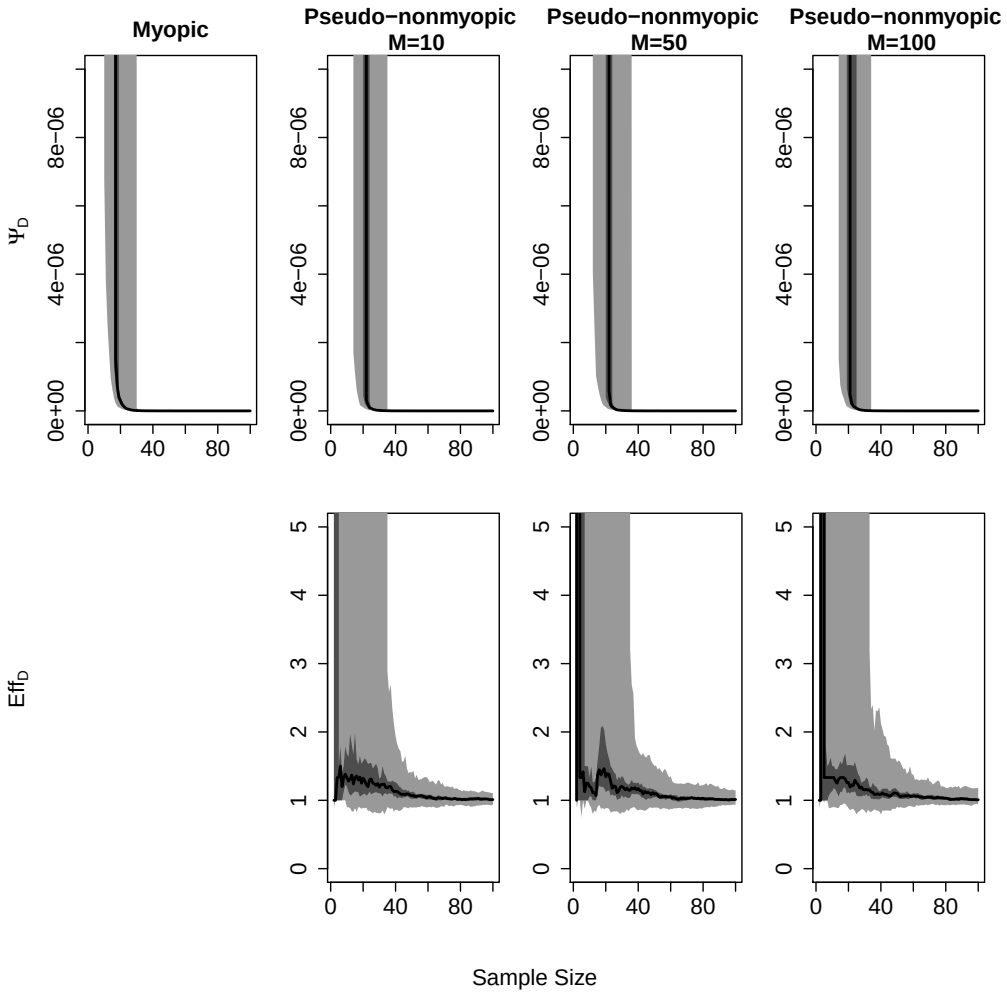

Figure 5.2: Top row: D-optimality against sample size for designs for a linear model with one static covariate. Bottom row: relative D-efficiency against sample size for designs for a linear model with one static covariate. Values below 1 indicate that the pseudo-nonmyopic approach is more beneficial than the myopic approach.

Similarly to what we found in Section 4.1.4 for the nonmyopic approach, we see that the myopic approach is generally more efficient than the pseudo-nonmyopic approach for the case of the linear model. We also observe that the choice of M does not make a noticeable impact on the values of Ψ_D.

5.2 Logistic Model

In the logistic model case, the objective function $\Psi(t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ is dependent on the previous covariates $\boldsymbol{Z}_i$, the previous treatments $\boldsymbol{t}_{i-1}$ and also the previous responses

$\boldsymbol{y}_{i-1}$, since we are using estimated parameter values. In order to create a design using the pseudo-nonmyopic approach for the logistic model, we begin by constructing an initial design $\boldsymbol{X}_{n_0}$ with n_0 patients using the exchange algorithm. We assume $\boldsymbol{\beta} = \mathbf{0}$ as an initial guess for evaluating the objective function in the construction of $\boldsymbol{X}_{n_0}$. We then generate responses for the first n_0 patients, $\boldsymbol{y}_{n_0}$, and fit the model to obtain the initial maximum likelihood estimates of the model parameters, $\hat{\boldsymbol{\beta}}_0$.

Then, to select a treatment for patient i, for $i \in \{1, 2, ..., n\}$, we observe $\boldsymbol{z}_i$. Similarly to the pseudo-nonmyopic approach for the linear model, we generate M possible trajectories for the covariates, $\boldsymbol{Z}^1_{(i+1):n}, \boldsymbol{Z}^2_{(i+1):n}, ..., \boldsymbol{Z}^m_{(i+1):n}$ and we allocate treatments sequentially along each trajectory. Given the first patient in the trajectory, $\boldsymbol{z}^m_{i+1}$, we choose the treatment $t^{*^m}_{i+1}$ which minimizes the objective function Ψ given t_i, and the treatments and covariates of previous patients and the estimate of $\boldsymbol{\beta}$ based on the responses of the previous patients, $\boldsymbol{y}_{i-1}$:

$$t^{*^m}_{i+1}\left(\boldsymbol{z}^m_{i+1}, t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}\right) = \underset{t_{i+1}}{\operatorname{argmin}}\, \Psi\left(t_{i+1} \mid \boldsymbol{Z}_i, \boldsymbol{z}^m_{i+1}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{y}_{i-1}\right). \tag{5.18}$$

To allocate a treatment for the next patient in the trajectory with covariate values $\boldsymbol{z}^m_{i+2}$, we then assume that $t^{*^m}_{i+1}$ has been allocated to patient $\boldsymbol{z}^m_{i+1}$ and choose the treatment $t^{*^m}_{i+2}$ which minimizes the objective function. We make the assumption that the future decisions are independent of the future responses. This means that we assume the same estimate for $\boldsymbol{\beta}$ as in the Equation (5.18) and do not update it. We continue in this way until all patients in the trajectory have been allocated a treatment:

For each j in $\{i+2, i+4, ..., n\}$, we define:

$$\begin{aligned} &t^{*^m}_j\left(\boldsymbol{z}^m_j, t^*_{j-1} \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j-1)}, \boldsymbol{t}_{i-1}, t_i, t^*_{(i+1):(j-2)}, \boldsymbol{y}_{i-1}\right) \\ &= \underset{t_j}{\operatorname{argmin}}\, \Psi\left(t_j \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j)}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}^*_{(i+1):(j-1)}, \boldsymbol{y}_{i-1}\right). \end{aligned} \tag{5.19}$$

For the mth trajectory, we obtain a pseudo-design with n patients where the ith treatment is 1, as well as a pseudo-design where the ith patient receives treatment -1. We denote the objective function of the two designs as follows:

$$\Psi\left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i = 1, \boldsymbol{t}^{*^m}_{(i+1):(n-1)}, \boldsymbol{y}_{i-1}\right), \tag{5.20}$$

$$\Psi\left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i = -1, \boldsymbol{t}^{*^m}_{(i+1):(n-1)}, \boldsymbol{y}_{i-1}\right). \tag{5.21}$$

We define the average objective function for $i = n_0 + 1, ..., n-1$ across the M designs, assuming, firstly, that $t_i = 1$, and secondly, that $t_i = -1$, as:

$$\overline{\Psi}(t_i = 1) = \frac{1}{M} \sum_{m=1}^{M} \Psi \left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i = 1, \boldsymbol{t}^{*m}_{(i+1):(n-1)}, \boldsymbol{y}_{i-1} \right), \tag{5.22}$$

$$\overline{\Psi}(t_i = -1) = \frac{1}{M} \sum_{m=1}^{M} \Psi \left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i = -1, \boldsymbol{t}^{*m}_{(i+1):(n-1)}, \boldsymbol{y}_{i-1} \right). \tag{5.23}$$

For $i = n$, we do not generate any future covariates so we have:

$$\overline{\Psi}(t_i = t) = \Psi \left(t_n = t \mid \boldsymbol{Z}_n, \boldsymbol{t}_{n-1}, \boldsymbol{y}_{n-1} \right), \tag{5.24}$$

for $t \in \{-1, 1\}$.

We sample t_i from the set $\{-1, 1\}$ where the probability of selecting 1 is given by

$$\frac{\overline{\Psi}(t_i = 1)^{-1}}{\overline{\Psi}(t_i = 1)^{-1} + \overline{\Psi}(t_i = -1)^{-1}}. \tag{5.25}$$

We then observe the response y_i and refit the model to obtain $\hat{\boldsymbol{\beta}}_i$.

Pseudocode is provided in Algorithm 6.

Algorithm 6 Pseudo-nonmyopic sequential design for binary response: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$, the distributions of the covariates $f_{z_1}, ..., f_{z_k}$ and size of initial design n_0

function PSEUDONONMY($\boldsymbol{Z}_n, f_{z_1}, ..., f_{z_k}, n_0$)

Initialization
Construct initial design $\boldsymbol{X}_{n_0}$: using the exchange algorithm assuming that $\boldsymbol{\beta} = \mathbf{0}$.
Observe responses $\boldsymbol{y}_{n_0} = \{y_1, y_2, ..., y_{n_0}\}$.
Fit the model $\boldsymbol{y}_{n_0} \sim glm(\boldsymbol{X}_{n_0}, link = logit)$ to obtain the MLE $\hat{\boldsymbol{\beta}}_0$.

for i in 1 to n **do**
Observe $\boldsymbol{z}_i = (z_{i,1}, z_{i,2}, ..., z_{i,k})^\intercal$

for t_i in $\{-1, 1\}$ **do**

for m in 1 to M **do**
Generate $\boldsymbol{Z}^m_{(i+1):n}$ using the assumed covariate distributions $f_{z_1}, ..., f_{z_k}$.

Allocate treatments $t^{*^m}_{i+1}, t^{*^m}_{i+2}, ..., t^{*^m}_n$ along the trajectory:

$$t^{*^m}_{i+1}(\boldsymbol{z}^m_{i+1}, t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}) = \underset{t_{i+1}}{\operatorname{argmin}} \Psi\left(t_{i+1} \mid \boldsymbol{Z}_i, \boldsymbol{z}^m_{i+1}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{y}_{i-1}\right). \tag{5.26}$$

For $j \in \{i+2, i+4, ..., n\}$:

$$t^{*^m}_j(\boldsymbol{z}^m_j, t^*_{j-1} \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j-1)}, \boldsymbol{t}_{i-1}, t_i, t^*_{(i+1):(j-2)}, \boldsymbol{y}_{i-1}) = \underset{t_j}{\operatorname{argmin}} \Psi\left(t_j \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):(j)}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}^*_{(i+1):(j-1)}, \boldsymbol{y}_{i-1}\right). \tag{5.27}$$

end for
Define $\overline{\Psi}(t_i) = \frac{1}{M} \sum_{m=1}^{M} \Psi\left(t_n \mid \boldsymbol{Z}_i, \boldsymbol{Z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{t}^{*^m}_{(i+1):(n-1)}, \boldsymbol{y}_{i-1}\right)$
end for
Sample t_i from $\{-1, 1\}$ where probability of selecting 1 is given by

$$\frac{\overline{\Psi}(t_i = 1)^{-1}}{\overline{\Psi}(t_i = 1)^{-1} + \overline{\Psi}(t_i = -1)^{-1}}. \tag{5.28}$$

Observe response y_i.
$\boldsymbol{X}_i = [\mathbf{1}\ \boldsymbol{Z}_i\ \boldsymbol{t}_i]$
Refit the model $\boldsymbol{y}_i \sim glm(\boldsymbol{X}_i, link = logit)$ to obtain the MLE $\hat{\boldsymbol{\beta}}_i$.
end for
return $\boldsymbol{X}_n$

end function

5.2.1 Simulations

We run two simulations to compare the myopic approach to the pseudo-nonmyopic approach in the case of the logistic model, for different values of M. Similarly to Section 5.1.1, in the first simulation, we consider the case of a single dynamic covariate with no interaction between treatment and covariate. In the second simulation, we have two covariates, one static and one dynamic and we assume that there are treatment-covariate interactions.

Reducing variability

Similarly to Section 4.2.5, we need to make sure that sources of variability are controlled as much as possible for the logistic model case, where the incorporation of responses adds a greater amount of variability than for the linear model. We wish to take measures so that differences between the results for the myopic and pseudo-nonmyopic approaches are likely to be attributable to the differences in the treatment allocation approach. We make sure that simulations have the same initial design; the initial design is constructed with the exchange algorithm to allocate treatments to 10 units, under the assumption that $\boldsymbol{\beta}$ is a vector of zeros. We fit the models using the `R` function `bayesglm`, with Cauchy prior distribution with center zero and scale given by 2.5 for both the treatment and covariate parameters.

In order to generate responses, we generate deviates u_i from the Unif$(0, 1)$ distribution and set the response y_i as

$$y_i = \begin{cases} 1 & \text{if } u_i \geq \pi_i \\ 0 & \text{if } u_i < \pi_i \end{cases}, \tag{5.29}$$

for $i \in 1, ..., n$. We keep the deviates u_i the same for simulations which compare the same objective function, so that we can ensure that the data generating mechanism is the same.

Choice of $\boldsymbol{\beta}$

We choose values of $\boldsymbol{\beta}$ which lead to values of π_i which are close to zero or one. This allows for more clear-cut choices for treatment allocation. If we choose $\boldsymbol{\beta}$ such that π_i are close to 0.5, we would introduce more noise in the results. Tables 5.1 and 5.2 show the possible values that π_i can take for the two simulations that we consider.

We note that, in practice, in cases where π_i is close to zero or one, the treatment groups are very well separated and a simple design would suffice to establish which treatment is more

effective. However, supposing that we wish to estimate the effects of all parameters with minimum possible variance there may be merit in taking an optimal design approach, and perhaps with a pseudo-nonmyopic outlook.

Set-up

The structure of the simulations are as follows:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) 100 deviates from a Unif$(0, 1)$ distribution are generated for the response.

(c) An initial design with 10 units is constructed using the exchange algorithm with D-optimality as the objective function.

(d) The three following sequential designs are constructed using the covariates, random deviates for the responses, and initial design in part (a):

- A myopic D-optimal design.
- A pseudo-nonmyopic D-optimal design with $M = 10$, and the correct covariate distribution assumed.
- A pseudo-nonmyopic D-optimal design with $M = 100$, and the correct covariate distribution assumed.

(e) Designs are evaluated using the performance measure Ψ_D at each sample size between 10 and 100, inclusive. The true values of the parameters are used to calculate Ψ_D.

II (a)-(e) above is repeated 20 times to obtain a distribution of the performance measure for each sample size.

Example 1

In the first example, we have one binary covariate z. It is dynamic with a distribution given by $\mathbb{P}(z_i = 1) = 0.01i$. The model is given by $y_i \sim \text{Bernoulli}(\pi_i)$ where

$$\text{logit}\left(\frac{\pi_i}{1-\pi_i}\right) = \beta_0 + \beta_1 z_i + \beta_2 t_i. \tag{5.30}$$

We set $\boldsymbol{\beta} = \left(0, 1, 2\right)^\intercal$ which leads to possible values of π given in Table 5.1. We note that the probabilities are not close to 0.5.

Table 5.1: Possible values of π_i when $\boldsymbol{\beta} = \begin{pmatrix} 0, 1, 2 \end{pmatrix}^\intercal$

π_i when $\boldsymbol{\beta} = \begin{pmatrix} 0, 1, 2 \end{pmatrix}^\intercal$		
z_i	t_i	π_i
-1	-1	0.047
1	-1	0.2689
-1	1	0.7311
1	1	0.9526

In Figure 5.3, we see the estimates of $\boldsymbol{\beta}$ for the myopic approach, the pseudo-nonmyopic approach with $M = 10$ and with $M = 100$. We observe that the plots looks very similar across the three methods. The variability of the estimates reduces with sample size for the intercept and the coefficient of treatment. The median of the distributions converge to their true value after a sample size of approximately 40.

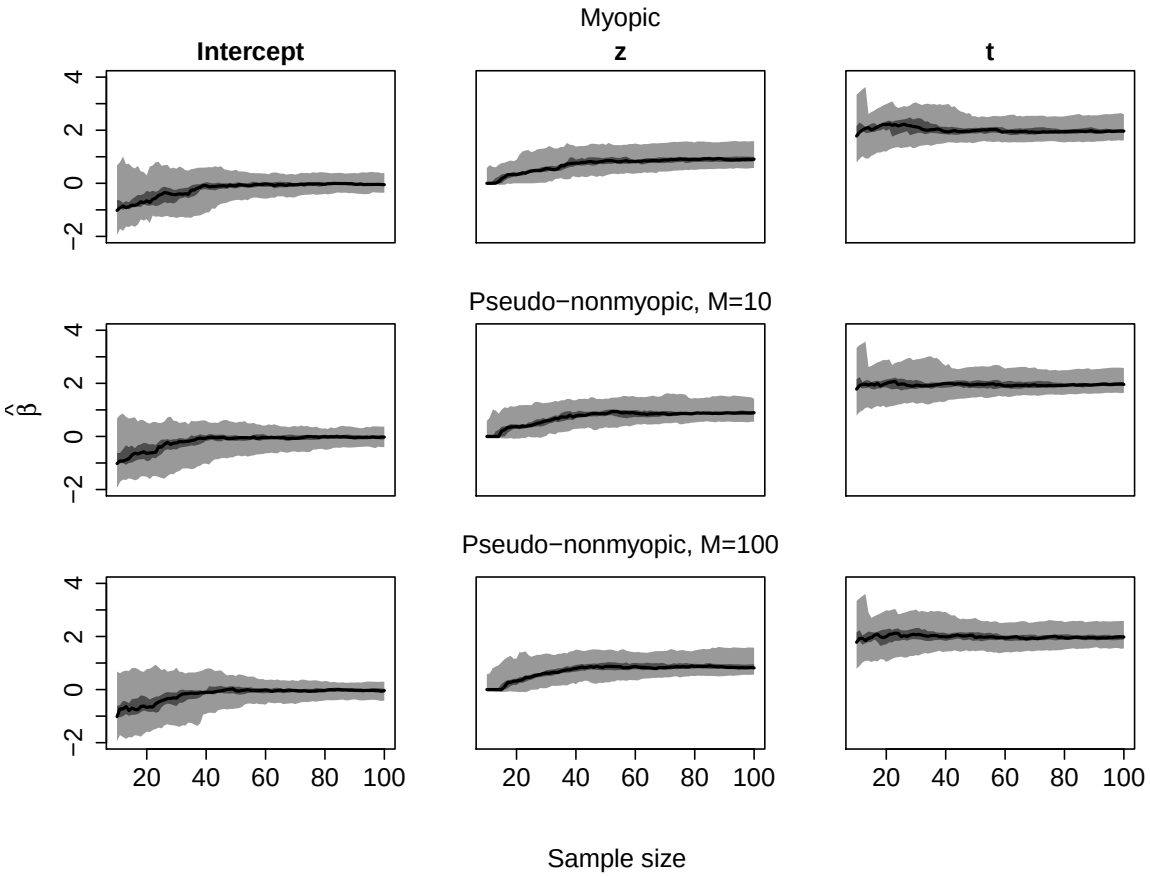

Figure 5.3: Parameter estimates given by the myopic approach, pseudo-nonmyopic approach with $M = 10$ and pseudo-nonmyopic approach with $M = 100$ for a logistic model with one dynamic covariate. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution

In Figure 5.4, the top row displays the values of Ψ_D evaluated at each sample size. This appears to be similar across all methods with slightly higher variation observed for the pseudo-nonmyopic approach with $M = 10$. In all three cases, the value of the objective function drops after a few initial patients and stabilizes after around 30 patients. The bottom row shows the relative D-efficiencies (see Equation (4.34)) of the pseudo-nonmyopic approaches, compared to the myopic approach. We see that, initially, they have equal efficiency, but then the myopic approach appears to be slightly more efficient. We note that the distributions of efficiencies are skewed; there appears to be a number of extreme points where the myopic approach is much more efficient than the pseudo-nonmyopic approach. This is partly due to the fact that the efficiency is bounded below by zero, but unbounded above.

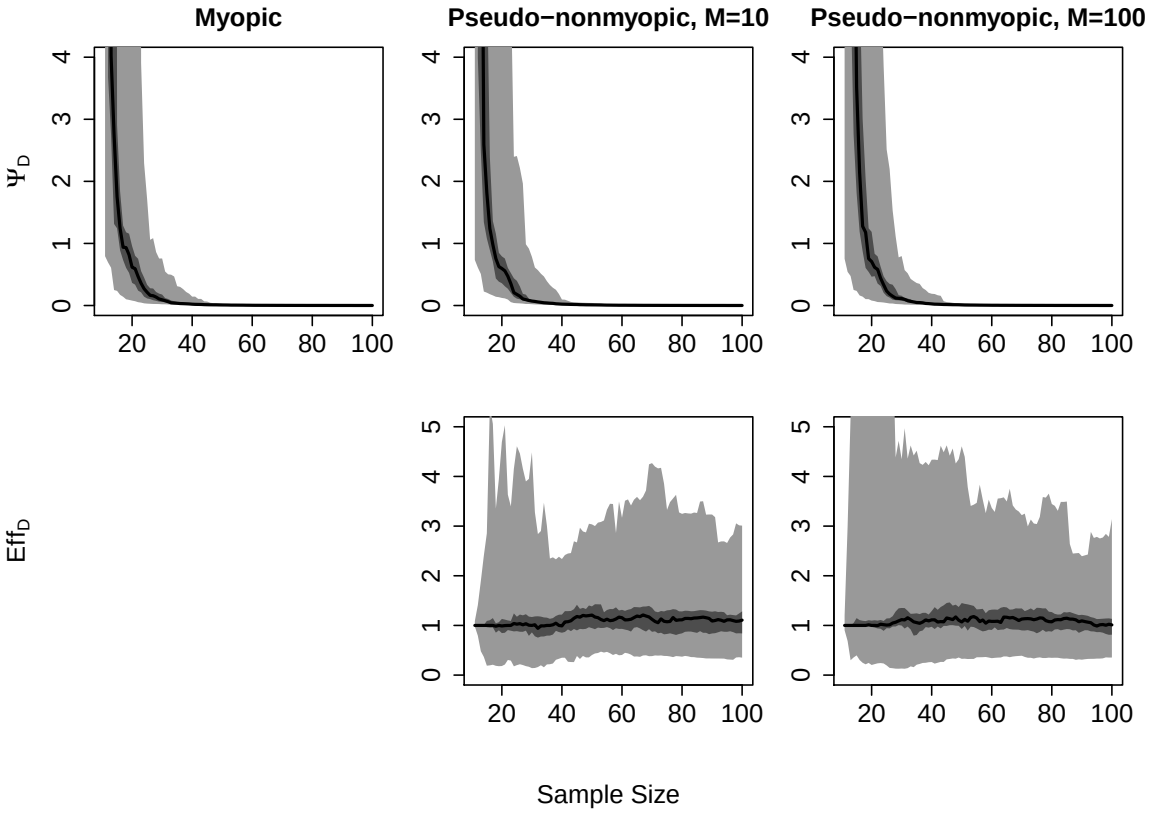

Figure 5.4: Top row: D-optimality against sample size for designs for a logistic model with one dynamic covariate. Bottom row: relative D-optimality against sample size for designs for a logistic model with one dynamic covariate. Values below 1 indicate that the pseudononmyopic approach is more beneficial than the myopic approach.

Example 3 Two covariates plus interactions

In the second example, we include interactions. In Chapters 6 and 7 which focus on personalized medicine, interactions will play a central role. This example has two binary covariates. The first is dynamic with distribution given by $\mathbb{P}(z_{i,1} = 1) = 0.01i$, and the second is static with distribution given by $\mathbb{P}(z_{i,2} = 1) = 0.3$. The model is given by $y_i \sim$ Bernoulli(π_i) where:

$$\text{logit}\left(\frac{\pi_i}{1-\pi_i}\right) = \beta_0 + z_{i,1}\beta_1 + z_{i,2}\beta_2 + t_i\beta_3 + z_{i,1}t_i\beta_4 + z_{i,2}t_i\beta_5. \tag{5.31}$$

We set $\boldsymbol{\beta} = \left(1, 2, 1, -2, 5, -4\right)^\intercal$. The possible values of π are given in Table 5.2; we observe that most are close to zero or one.

Table 5.2: Possible values of π_i when $\boldsymbol{\beta} = \left(1, 2, 1, -2, 5, -4\right)^\intercal$

π_i when $\boldsymbol{\beta} = (1, 2, 1, -2, 5, -4)^\intercal$			
$z_{1,i}$	$z_{2,i}$	t_i	π_i
−1	−1	−1	0.7311
1	−1	−1	0.0067
−1	1	−1	1.0000
1	1	−1	0.9933
−1	−1	1	0.0067
1	−1	1	0.9999
−1	1	1	0.0000
1	1	1	0.9526

Figure 5.5 shows the estimates of $\boldsymbol{\beta}$ for the three approaches. Again, the the estimates are similar across the three approaches. We observe that the larger values of $\boldsymbol{\beta}$ are not estimated as well as the smaller values. The two interaction terms, with values 5 and −4, are both underestimated. This shrinkage effect is attributable to the fact that we are using `bayesglm` with prior distributions.

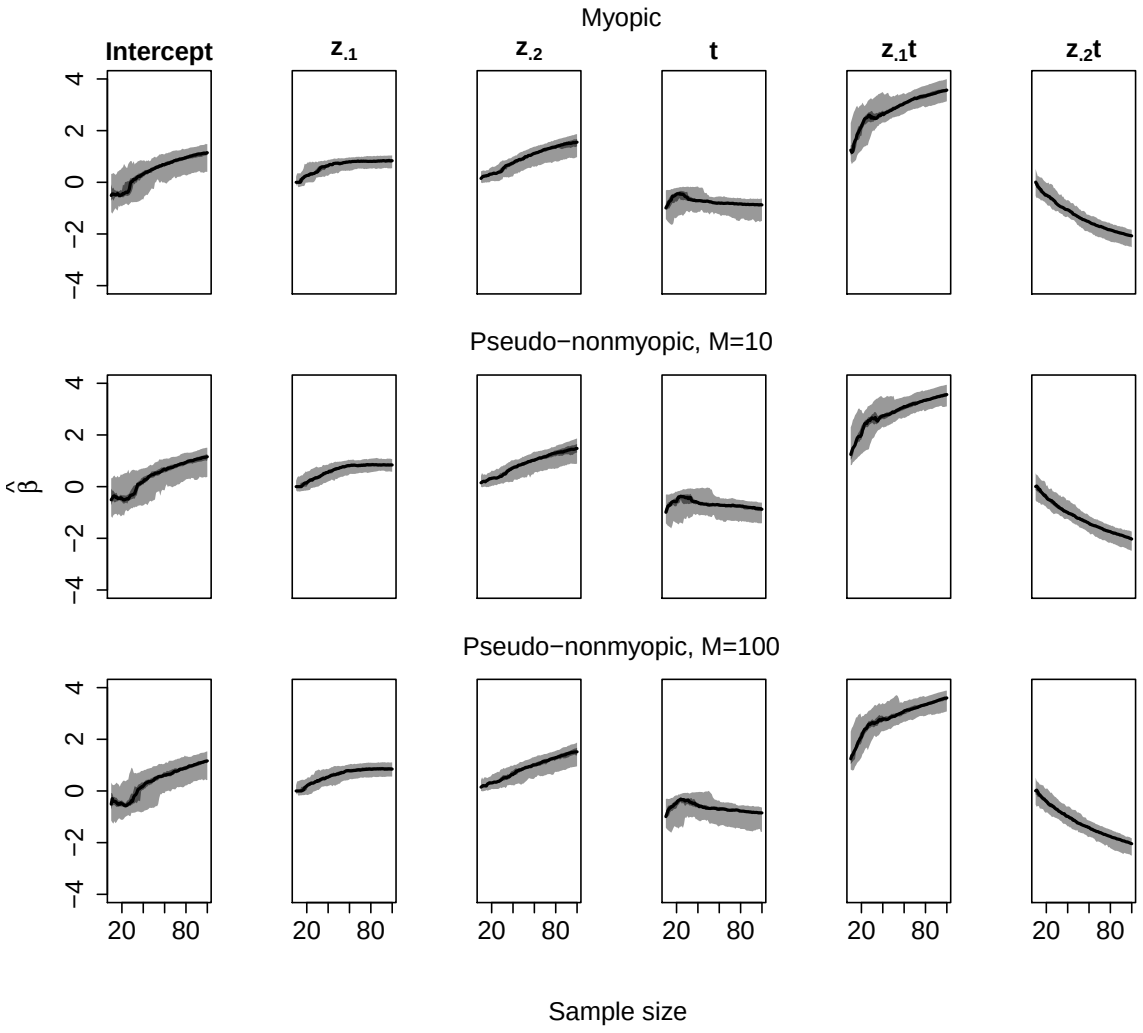

Figure 5.5: Parameter estimates given by the myopic approach, pseudo-nonmyopic approach with $M = 10$ and pseudo-nonmyopic approach with $M = 100$ for a logistic model with two covariates.

In Figure 5.6, we see in the top row displays that the values of Ψ_D evaluated at each sample size are similar across all methods. There is slightly lower variability in the plot for the myopic approach compared to the pseudo-nonmyopic approach. In the bottom row, we observe that, initially, the pseudo-nonmyopic approach and myopic approach have equal efficiency, but then the myopic approach becomes more efficient. After around sample size of 40, the two approaches appear to perform similarly. It appears that M does not affect the results in any noticeable way.

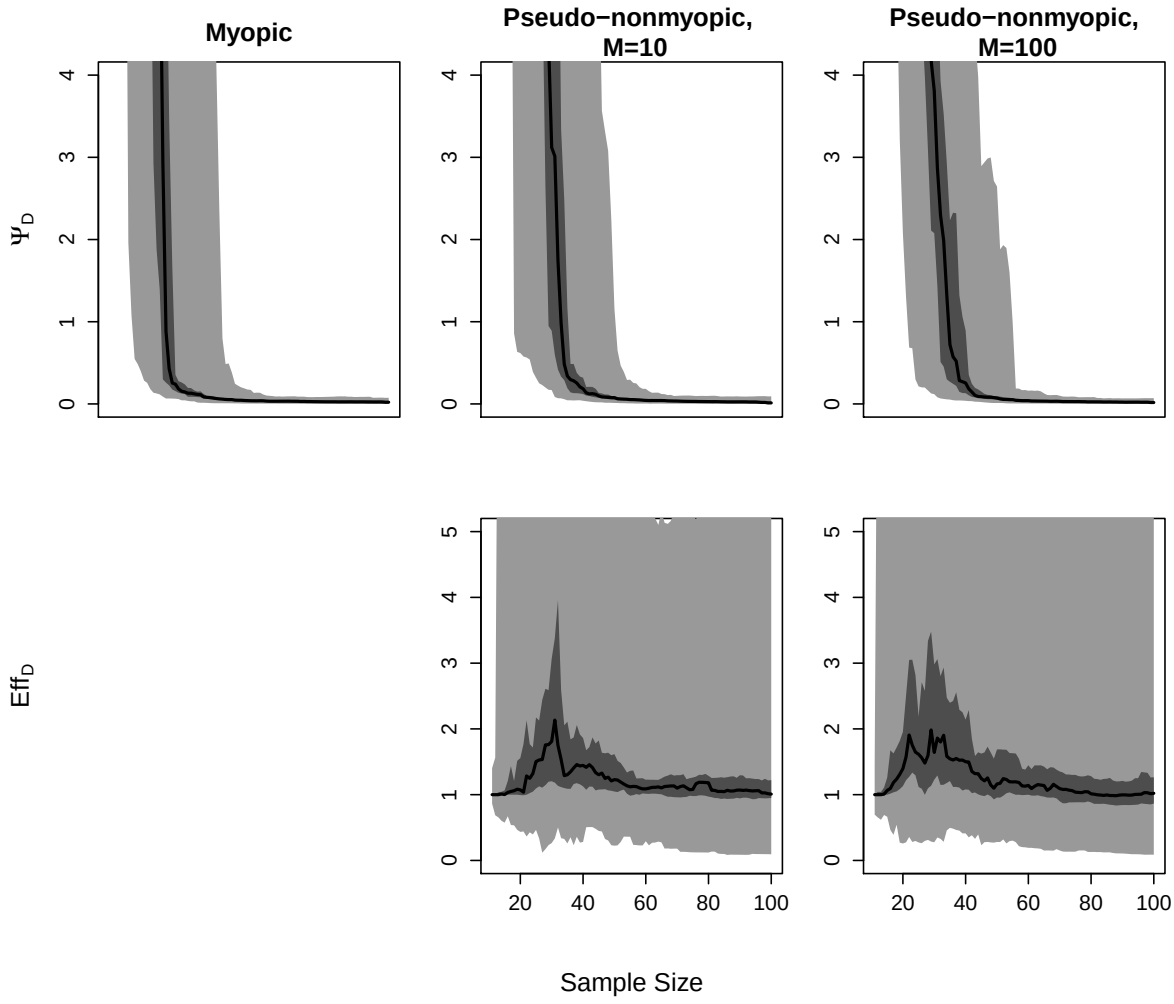

Figure 5.6: Top row: D-optimality against sample size for designs for a logistic model with two covariates. Bottom row: relative D-optimality against sample size for designs for a logistic model with two covariates. Values below 1 indicate that the pseudononmyopic approach is more beneficial than the myopic approach.

The conclusions we draw for the pseudo-nonmyopic approach for the logistic model are similar to those of the linear model; we find that the myopic approach is more efficient, and M does not appear to affect the performance of the pseudo-nonmyopic approach in the examples we have considered. As shown in Table 5.3, a larger value of M results in longer computational time, so we choose a small value of $M = 20$ in simulations using the pseudo-nonmyopic approach in Sections 5.3.3 and 7.2.1. We also see that there is more variability in the values of the objective function for the pseudo-nonmyopic approach compared to the myopic approach, which is expected as the pseudo-nonmyopic approach takes into account future possible patients.

Table 5.3: Running time in seconds for constructing a design with one covariate and 100 patients. Initial design is one patient for the linear model case, and 10 patients for the logistic model case. We compare the myopic approach ($M = 0$) against the pseudo-nonmyopic approach with M set to 10, 50, and 100. Simulations were performed on a machine with a 2GHz processor and 64 GB of memory.

M	**Linear Model Binary Treatment**		**Logistic Model Binary Treatment**		**Logistic Model Continuous Treatment**	
	Static	**Dynamic**	**Static**	**Dynamic**	**Static**	**Dynamic**
0	0.041	0.034	0.849	0.97	0.744	0.78
10	22.989	28.25	273.56	286.144	506.462	$>$ 60 hours
50	114.905	140.59	1335.804	1387.487	2441.124	$>$ 60 hours
100	229.303	282.716	2630.497	2738.561	4334.633	$>$ 60 hours

5.3 Continuous treatment

So far in this thesis, we have only considered the case of a binary treatment where we compare a new treatment with a control. However, extending this to a continuous treatment has important applications in clinical trials; dose-finding is an essential part of phase I trials where it is important to find an acceptable drug dosage which does not cause serious side-effects (Pocock, 2013, p.3). For the nonmyopic approach, computations involving recursive formulae such as Equations (4.6) or (4.46) are time consuming when the treatment is simply binary. To extend this for a continuous treatment, we would need to use Monte Carlo approximations which would further add computational complexity to the problem; we do not explore this problem in this thesis. However, for the pseudo-nonmyopic approach, extension to a continuous treatment case is relatively straightforward.

We note that, in dose-finding, continuous treatments effects are typically not modeled in a linear fashion in practice. There is a need to balance the two aims of having a high enough dose for a desired response, and a low enough dose so avoid adverse effects, so a nonlinear approach is needed (Yu et al., 2016). This is outside of the scope of this thesis.

5.3.1 Optimal design for continuous treatments

We assume without loss of generality that the treatments take values in the constrained region $[-1, 1] \in \mathbb{R}$. To allocate a continuous treatment sequentially using optimal design with a myopic outlook, we need to select a treatment t_i for the ith patient to minimize some objective function $\Psi\left(t_i \mid \boldsymbol{Z}_{i-1}, \boldsymbol{z}_i, \boldsymbol{t}_{i-1}\right)$ for the linear model case or $\Psi\left(t_i \mid \boldsymbol{Z}_{i-1}, \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}\right)$ for the logistic model case. This is a one-dimensional optimization problem which we solve using a search algorithm. We use Brent's method which is a hybrid between Newton and bisec-

tion; this can be implemented in `R` using the function `optim()` with option `method="Brent"`.

For a non-sequential approach where all n treatments are allocated simultaneously, the optimal treatment vector $\boldsymbol{t} = \left(t_1, ..., t_n\right)^\intercal$ is found to minimize the objective function $\Psi\left(\boldsymbol{X}\right)$ where $\boldsymbol{X}$ is a design matrix with $\boldsymbol{t}$ as its treatment column. For this multi-parameter optimization problem, we use the limited-memory quasi-Newton code for bound-constrained optimization which is implemented using `optim()` with option `method="L-BFGS-B"`.

In Chapter 2, we allocate treatments according to a probability such as Equation (2.37) to make a decision that is optimal but also mitigates the chance of selection bias. Here, in the case of a continuous treatment, we allocate the treatment deterministically.

5.3.2 Pseudo-nonmyopic approach for continuous treatments

We describe the pseudo-nonmyopic approach when there is a continuous treatment. We describe only the case for the logistic model; the linear model is analogous to the logistic model case, but the responses y_i do not enter the equations below as they are not needed to compute the objective function.

Logistic model

To allocate a treatment for patient i for $i \in \{1, 2, ..., n-1\}$, we observe $\boldsymbol{z}_i$. As before, we generate M possible trajectories, $\boldsymbol{Z}^1_{(i+1):n}, \boldsymbol{Z}^2_{(i+1):n}, ..., \boldsymbol{Z}^m_{(i+1):n}$. Then, for the mth trajectory, $m \in \{1, ..., M\}$, we need to perform a nested optimization; given some treatment t_i, we find the vector of optimal treatments ${\boldsymbol{t}^{*m}}_{(i+1):n}$ for the patients in the trajectory; this is a multi-parameter optimization which requires a search algorithm:

$$\begin{aligned} &\boldsymbol{t}^{*^m}_{(i+1):n}\left(\boldsymbol{z}^m_{(i+1):n}, t_i \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_i\right) \\ &= \underset{\boldsymbol{t}_{(i+1):n} \in [-1,1]^{n-i-1}}{\operatorname{argmin}} \Psi\left(\boldsymbol{t}_{(i+1):n} \mid \boldsymbol{Z}_i, \boldsymbol{z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, t_i, \boldsymbol{y}_i\right). \end{aligned} \tag{5.32}$$

We denote the quantity in Equation (5.32) by $\boldsymbol{t}^{*^m}_{(i+1):n}$ to simplify notation. We then need to search over the space $[-1, 1]$ to find the optimal choice of treatment t^m_i for the objective function at the end of the experiment:

$$t_i{}^m \left(\boldsymbol{z}^m_{(i+1):n}, \boldsymbol{t}^{*^m}_{(i+1):n} \mid \boldsymbol{Z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_i \right) = \underset{t_i \in [-1,1]}{\operatorname{argmin}} \Psi \left(t_i \mid \boldsymbol{Z}_i, \boldsymbol{z}^m_{(i+1):n}, \boldsymbol{t}_{i-1}, \boldsymbol{t}^{*^m}_{(i+1):n}, \boldsymbol{y}_i \right). \quad (5.33)$$

We compute the average of the M optimal treatments for patient i, and this is the treatment value that we assign:

$$t_i = \frac{1}{M} \sum_{m=1}^{M} t_i{}^m. \quad (5.34)$$

To allocate a treatment for patient n, there is no need to generate future patients in the trajectory; the assigned treatment is given by

$$t_n = \underset{t_n \in [-1,1]}{\operatorname{argmin}} \Psi \left(t_n \mid \boldsymbol{Z}_n, \boldsymbol{t}_{n-1} \right). \quad (5.35)$$

We note that, choosing the treatment which minimizes the average of the objective functions across the M pseudo-designs is computationally too expensive; thus we opt to select the treatment which the average of the optimal treatments as in Equation (5.34). Further, we note that the average is not the only summary statistic that could be chosen in Equation (5.34); this is discussed further in Section 5.4.

5.3.3 Simulations

We conduct the same two simulation studies as in Section 5.2.1 for logistic regression but we now assume a single continuous treatment in both cases. We also reduce the number of patients in each experiment; we consider only 50 patients instead of 100. Searching over a continuous space is more computationally intensive than the binary treatment case that we considered previously.

Example 1

In the first simulation, we have one dynamic covariate such that $\mathbb{P}(z_i = 1) = 0.02i$. For the pseudo-nonmyopic approach, we set $M = 20$. All other settings are kept the same as in the logistic model case.

In Figure 5.7, we plot the estimates of $\boldsymbol{\beta}$ for the case where we have one dynamic covariate. We observe that the estimates of $\boldsymbol{\beta}$ are similar for the myopic and pseudo-nonmyopic

approach with $M = 20$ for the continuous treatment case. We observe greater variability than in the binary treatment case in Figure 5.3.

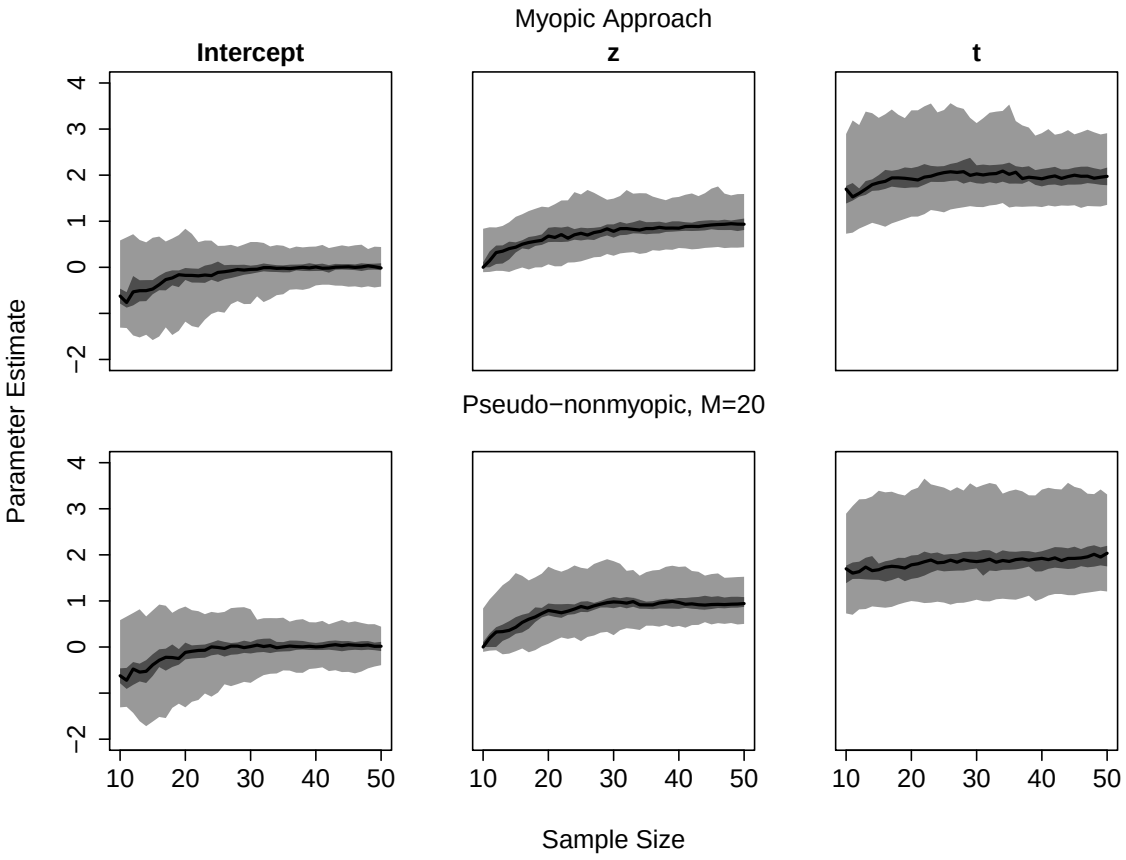

Figure 5.7: Parameter estimates given by the myopic approach and pseudo-nonmyopic approach with $M = 20$ for a logistic model with one dynamic covariate and a continuous treatment.

In Figure 5.8, we display boxplots to show the distribution of the treatment selected across the 20 simulations at each value of sample size between 1 and 50. The boxplots are the same for the first 10 patients for the myopic and pseudo-nonmyopic cases since the initial 10-patient designs are the same for the two methods. In the initial design, we assume that the true value of the parameters are zero, so the treatments selected are either -1 or 1. After the initial design, the selected treatments are spread across the interval $[-1, 1]$. We observe greater variability in the myopic case. This may be explained by the fact that the parameter estimates are less variable for the myopic approach, leading to choices of treatment that are closer to the extremes. The medians of the boxplots appear to be scattered in a random way for the myopic approach. In the pseudo-nonmyopic approach, the medians are generally greater than zero. In the final 10 patients, we observe that more treatments below zero are

allocated.

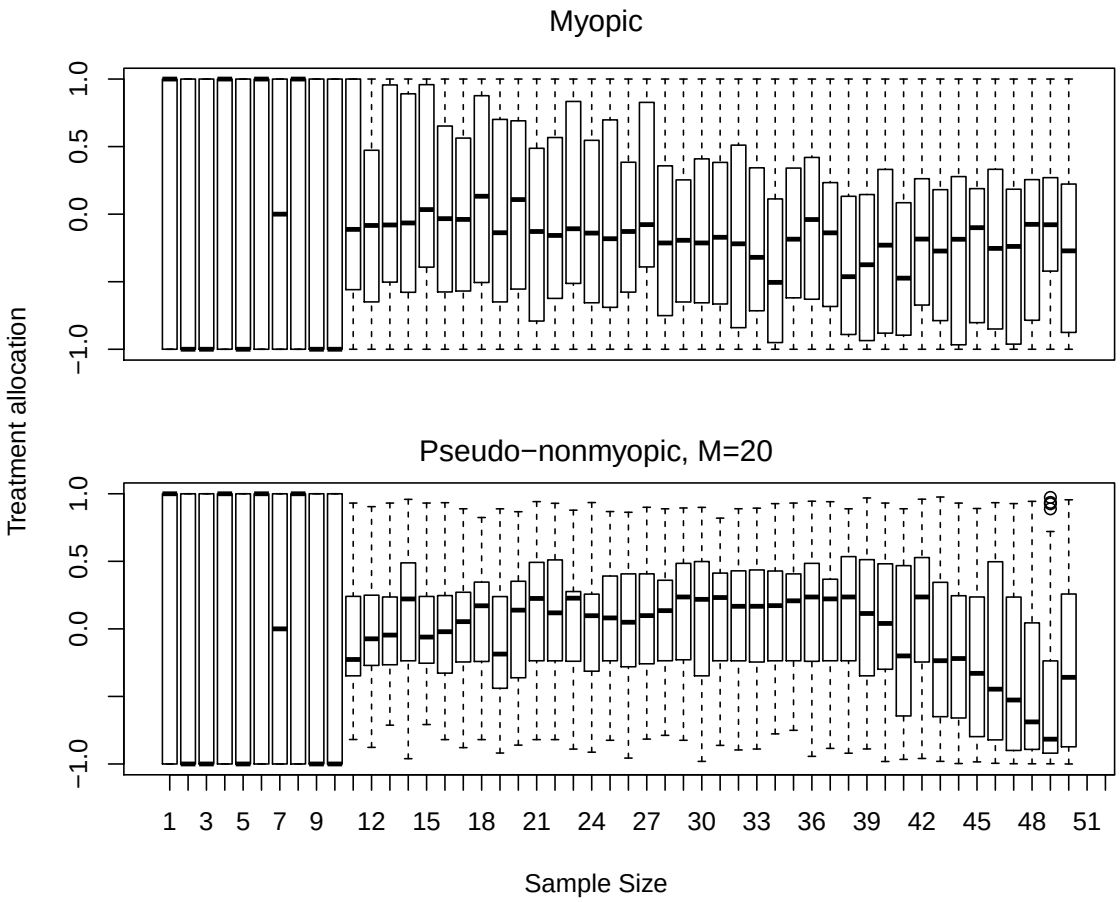

Figure 5.8: Boxplots display the distribution of the allocated treatments at each sample size for the myopic approach and pseudo-nonmyopic approach with $M = 20$ for a logistic model with one dynamic covariate and a continuous treatment.

The plots displaying the D-optimal objective function and relative D-efficiency for the continuous case below in Figure 5.9 are similar to that of the binary case in Figure 5.4. We observe that the the myopic approach is slightly more efficient than the pseudo-nonmyopic approach.

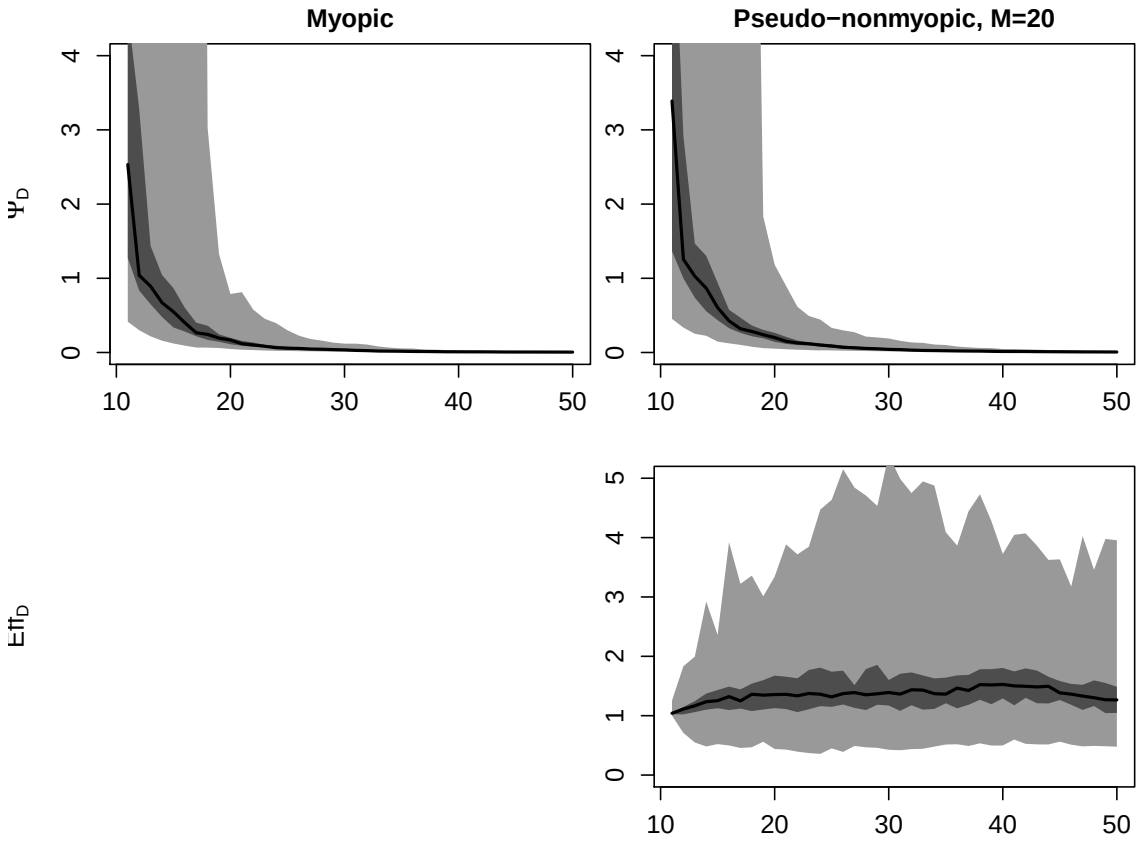

Figure 5.9: Optimality and D-efficiency given by the myopic approach and pseudo-nonmyopic approach with $M = 20$ for a logistic model with one dynamic covariate and a continuous treatment.

Example 2

In the second simulation study, we have two covariates with distributions given by $\mathbb{P}(z_{i,1} = 1) = 0.2i$ and $\mathbb{P}(z_{i,2} = 1) = 0.3$. Figure 5.10 displays the distribution of the estimates of $\boldsymbol{\beta}$. The distributions appear similar for the two approaches. In particular, the two interactions terms are shrunk towards zero in both cases. The shrinkage is slightly more extreme in the pseudo-nonmyopic case.

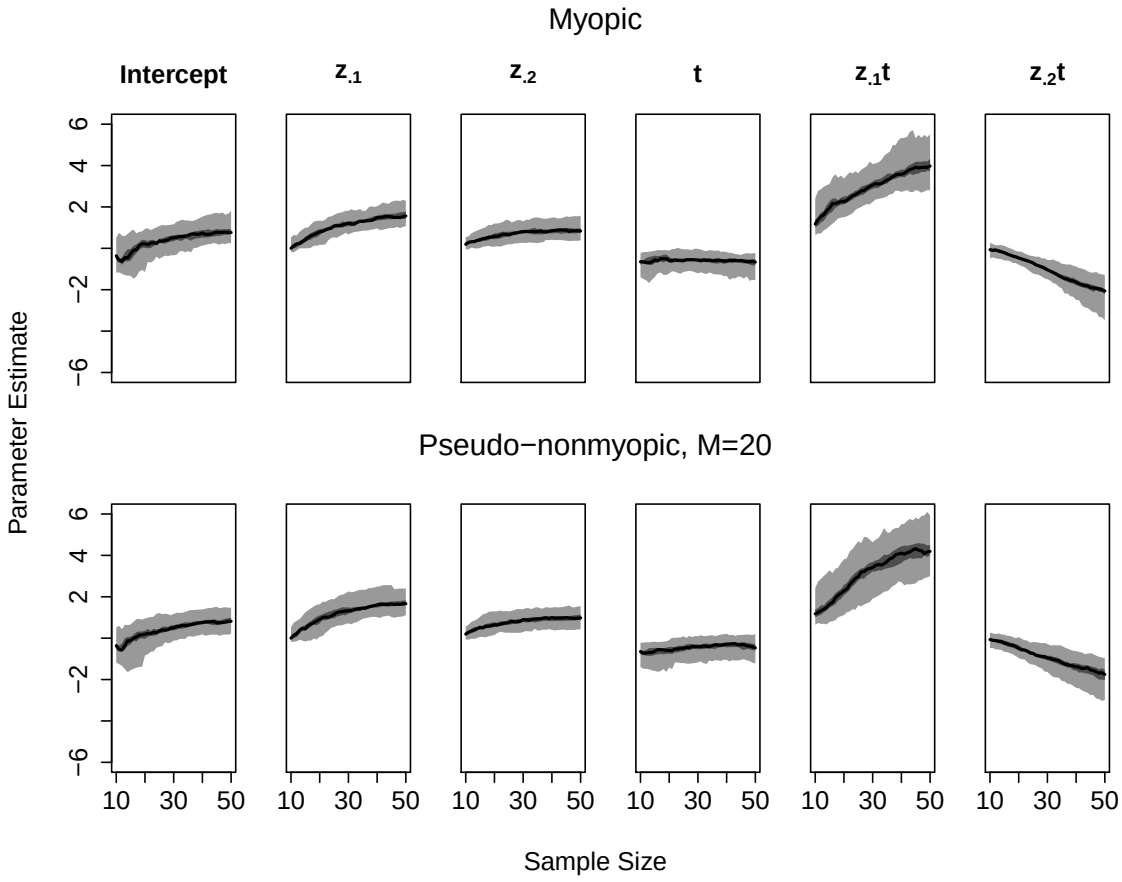

Figure 5.10: Parameter estimates given by the myopic approach, pseudo-nonmyopic approach with $M = 10$ and pseudo-nonmyopic approach with $M = 100$ for a logistic model with two covariates.

Figure 5.11 the distributions of the treatment selected across the 20 simulations at each value of sample size. Similarly to the case with one covariate, we observe that the boxplots are the same for the initial designs for the myopic and pseudo-nonmyopic cases, and after the initial design, we observe greater variability in the myopic case. We do not observe any clear patterns in the allocation of treatments in this example.

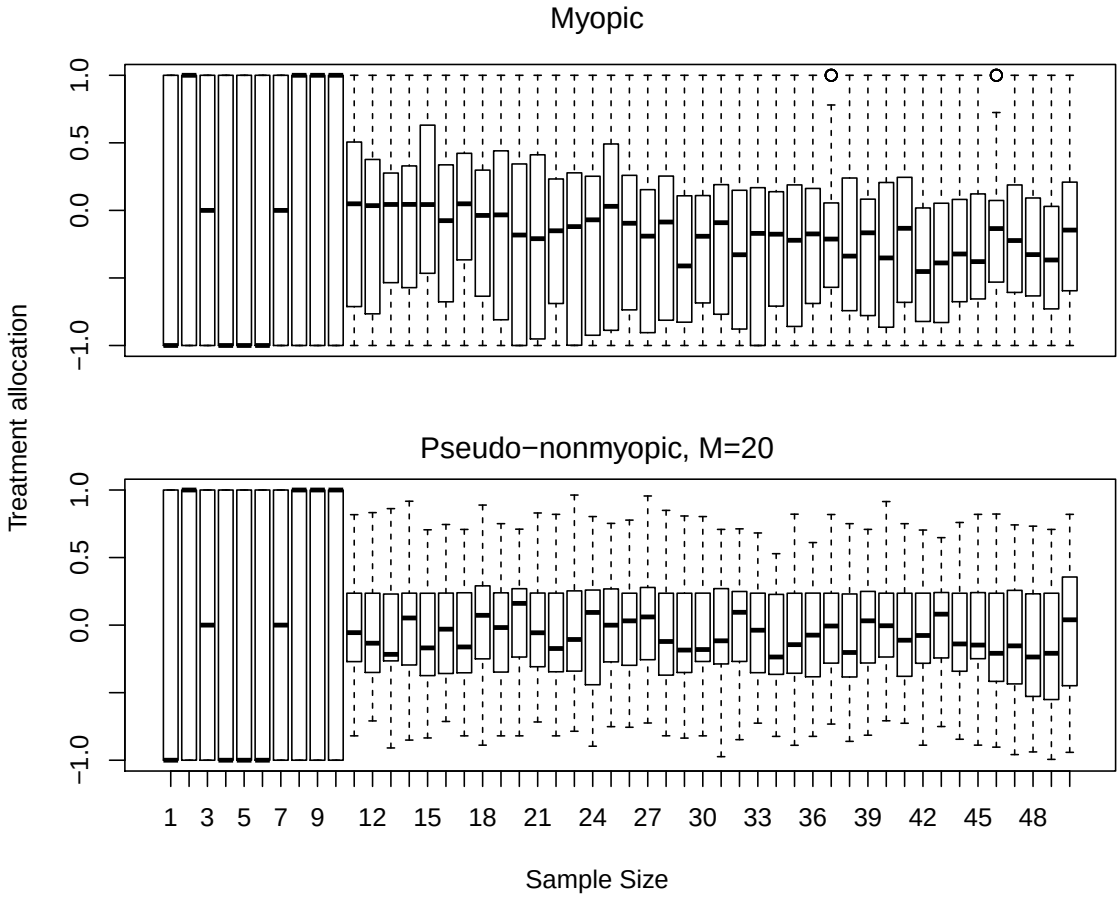

Figure 5.11: Boxplots display the distribution of the allocated treatments at each sample size for the myopic approach and pseudo-nonmyopic approach with $M = 20$ for a logistic model with two covariates.

Figure 5.12 displays the distributions of the D-optimal objective function of the two approaches, as well as the relative D-efficiency of the pseudo-nonmyopic approach against the myopic approach. We see that, initially, for sample size smaller than 20, the pseudo-nonmyopic approach is slightly more efficient than the myopic approach. If sample size increases beyond that, the myopic approach appears to be increasingly more efficient. This is likely to be caused by the fact that the myopic approach is able to achieve estimates of the interactions that are closer to the true value than the pseudo-nonmyopic approach. Since the objective function is computed using the true values of the parameters, the myopic approach will have more efficient values of objective function.

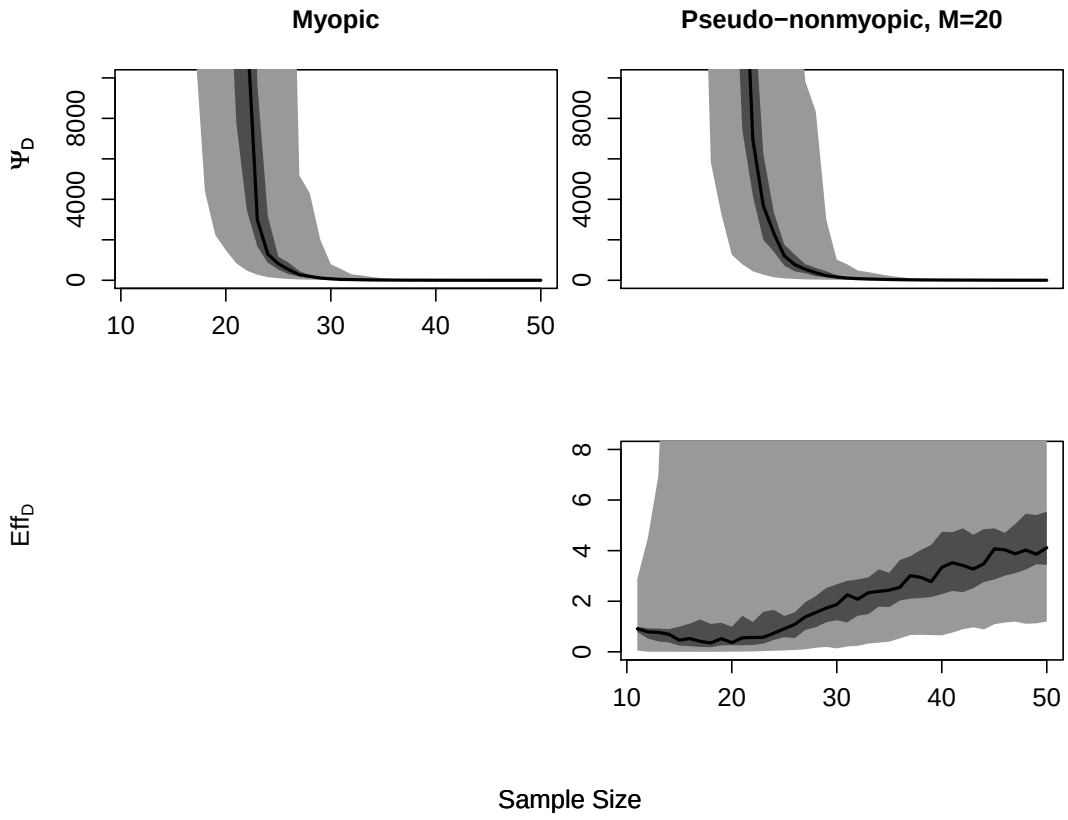

Figure 5.12: Top row: D-optimality against sample size for designs for a logistic model with two covariates. Bottom row: relative D-efficiency against sample size for designs for a logistic model with two covariates. Values below 1 indicate that the pseudononmyopic approach is more beneficial than the myopic approach.

In Chapter 9, we provide an overview of simulation results from all parts of the thesis and the overall conclusions that we draw from them.

5.4 Conclusion

In this chapter, we introduced a novel approach to sequential treatment allocation which can take into account the impact of future possible decisions without the need for computationally expensive recursive formulae that are an integral part of the fully nonmyopic approach. The pseudo-nonmyopic approach essentially approximates a criterion that is similar to that of the nonmyopic approach. Further, it can be fairly straightforwardly applied to the case of a continuous treatment, which is more difficult to do for the for the nonmyopic approach.

Results from our simulations show that, in the case of a linear model, the myopic and pseudo-nonmyopic approaches appear to produce fairly similar estimates of the parameters in both the one-covariate and also the two-covariate case. The myopic approach is slightly more efficient. We did not see any difference in results when the number of trajectories M in the pseudo-nonmyopic approach is set to 10, 50, or 100.

In the case of a logistic model, we find that the myopic and pseudo-nonmyopic approaches perform similarly if there is one covariate, and the myopic approach is more efficient when there are two covariates.

Extending the logistic model case for a continuous treatment, our main finding is that, in the case with two covariates and large interactions, the myopic approach is better able to estimate the interactions. This means that the myopic approach appears to be more efficient, especially for larger sample sizes.

We described in this chapter how to allocate treatments $i+1$ to n in the pseudo-design sequentially. However, we note that it is also possible to allocate these treatments using the exchange algorithm instead. This means that a random allocation of treatments is decided for the future patients, and the best treatment is selected patient-by-patient from patient $i+1$ up to n, repeating until two passes result in identical designs. This is then repeated a number of times using different starting designs; see Algorithm 2. Using this approach, the quantity that we wish to approximate is the the minimum value of the following integral:

$$\int \Psi\left(\boldsymbol{t}_{(i+1):n} \mid \boldsymbol{z}_{1:i}, \boldsymbol{z}_{(i+1):n}, \boldsymbol{t}_{1:i}\right) f_{\boldsymbol{z}_{(i+1):n}} d\boldsymbol{z}_{(i+1):n}, \tag{5.36}$$

where we define $f_{\boldsymbol{z}_{(i+1):n}}$ as the joint distribution function of the covariates $\boldsymbol{z}_{i+1}$ up to $\boldsymbol{z}_n$.

Our simulations showed no noticeable difference in efficiency when the treatments were selected sequentially vs via the exchange algorithm, so we chose to focus on the sequential approach as this is less computationally expensive.

Lastly, we note one extension that may be interesting further work. In the continuous treatment approach, we allocate the treatment which is the average of the optimal treatments for the pseudo-designs, as computed in Equation (5.34). By taking the average, we assign a treatment which is what we would typically expect out of the optimal decisions made for the pseudo-designs. It is possible to use some other summary statistic here; for example, a minimax decision rule would mean that treatments are assigned which minimize

the loss for potential worst-case pseudo-design. It may be interesting to compare differences in the pseudo-nonmyopic designs that are constructed with different summary statistics.

Chapter 6

An application in personalized medicine

In this chapter, we consider how to adapt the sequential treatment allocation methods introduced in Chapter 2 based on optimal design theory to design trials for personalized medicine. Personalized medicine, also referred to as stratified or precision medicine, is concerned with treatments that are tailored to patients' genetic information (Senn, 2016). For example, some cancers can now be characterized at the molecular level, and treatments can be targeted at biomarkers, which are specific genetic or biological mechanisms (Kaplan, 2015). There is a need for clinical trials to be able to identify and validate effective treatment-biomarker combinations. This can lead to a large number of interaction effects that need to be estimated.

Lee and Wason (2019) propose how data from a Phase II trial can be used to design a Phase III trial which seeks to estimate parameters associated with the effective treatment-biomarker combinations as precisely as possible. The proposed design for Phase III trial is a weighted L-optimal design, where the weights are calculated using data from a Phase II trial.

We adapt this problem for a sequential setting where we wish to obtain an adaptive treatment allocation scheme based on the weighted L-optimality criterion which seeks to estimate important treatment-biomarker interactions with low variance. The weights are not constructed using data from a previous trial; instead, they are learned from the current trial sequentially. We begin with an overview of personalized medicine, we then illustrate the methodology and show some simulation results.

6.1 Clinical trials for Personalized Medicine

We provide an overview of the types of designs that are used in clinical trials, and particularly, those relevant for personalized medicine. A classical design refers to an experiment which does not consider biomarkers and tests the effect of one treatment factor in the general population. This has been the setting that we have considered up until this chapter. Multi-arm multi-stage (MAMS) trials allow for multiple treatments to be tested at the same time, as well as for treatments to be dropped during the trial if interim analysis shows that they are ineffective (Royston et al., 2003). A MAMS trial which is currently ongoing in the UK is the STAMPEDE trial for treatments for prostate cancer, which began in 2005 (Kaplan, 2015). In a stratified or personalized design, patients are recruited from the general population, and they are subsequently stratified into biomarker-defined subgroups. Each biomarker has a linked treatment and the effect of that treatment is tested within that subgroup. We will consider this type of trial in this chapter. Finally, an enrichment design refers to an experiment where patients' biomarker status are first established and only those who are biomarker positive participate in the trial (Ondra et al., 2016b). The motivation behind enrichment trials is usually to validate that the biomarker has a significant impact on the treatment effectiveness. They may also be used if it is unethical to test the treatment on biomarker-negative patients due to, for example, potential side effects (Antoniou et al., 2016). We display the four types of trial designs visually in Figure 6.1.

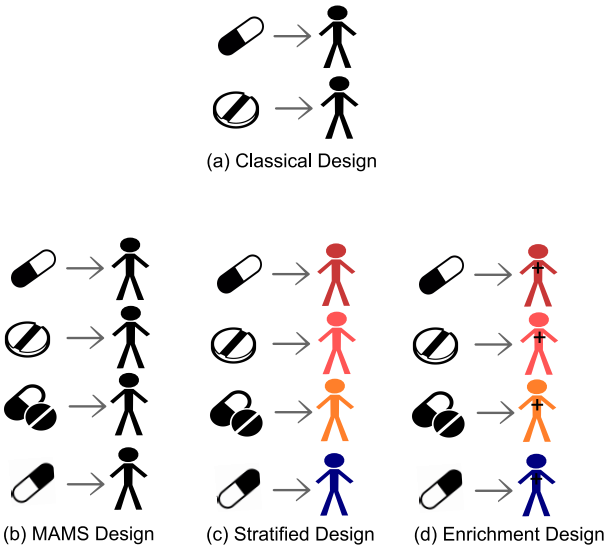

Figure 6.1: Types of designs for clinical trials: (a) a classical design testing one treatment against a control, (b) a MAMS design with multiple treatments, (c) a stratified/personalized design where patients are stratified into biomarker-defined groups and biomarker-linked treatments are tested, (d) an enrichment group where only biomarker-positive patients are entered in the trial.

Further, there are two types of clinical trials which incorporate genetic characteristics: umbrella trials and basket trials. An umbrella trial is a stratified design for cancer characterized by the same types of cells (histology) where independent randomized controlled trials are run for each biomarker-treatment subgroup; an example is the FOCUS4 clinical trial investigating treatments for colon cancer which began in the UK in 2004 and is ongoing (Kaplan, 2015). Basket trials test a single treatment in several different groups of people who have the same biomarker, but have different histologies (Renfro and Sargent, 2017). It is possible that the presence of the genetic biomarker is a stronger predictor of the effectiveness of the drug, rather than the histology of the tumour; the basket trial is designed with this in mind. Umbrella and basket trials are illustrated in Figure 6.2.

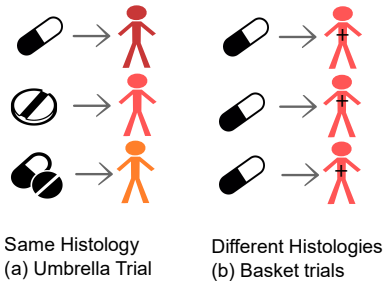

Figure 6.2: More designs for clinical trials: (a) an umbrella trial where patients have the same histology and independent trials are carried out for each biomarker-treatment combination, (b) a basket trial where patients have different histologies, but have the same biomarker and the same treatment is tested.

Antoniou et al. (2016) carried out a literature review of non-adaptive biomarker-linked designs for clinical trials and listed the main types, their advantages and limitations. Antoniou et al. (2017) undertook a similar review of the literature for adaptive biomarker-linked trials, where accumulating information is used to change part of the study. Usually, this means that treatment arms are dropped, or the probability of treatment assignment changes over time. Antoniou et al. (2017) listed both classical and Bayesian approaches which have been used. Renfro and Sargent (2017) and Parmar et al. (2017) provided overviews of umbrella, basket and multi-stage trials. Freidlin et al. (2010) provided a summary of issues related to designing personalized trials, including compliance issues and the case where biomarker status may be missing.

Minimization has been used in some personalized trials, such as the STAMPEDE trial, as a treatment allocation scheme (Sydes et al., 2012). The play-the-winner rule is a response-adaptive treatment allocation scheme that has also been used in clinical trials, though it does not in general include covariates (Yao and Wei, 1996). The play-the-winner rule is appropriate for binary treatments and works by sampling treatments from an "urn". Suppose

that, at the start of the trial, the urn contains one of each treatment, and one of them is randomly selected for the first patient. Then, assuming a binary response, if the patient responds well to the given treatment, then n_1 more of the same treatment are added to the urn. If the patient does not respond well, then n_2 of the other treatment are added to the urn. This is repeated for each patient. Bandyopadhyay and Biswas (1999) and Zhang et al. (2016) consider modified versions of the play-the-winner rule which incorporate prognostic factors.

Bayesian adaptive randomization is an approach gaining popularity in research for clinical trial methodology (Zhang et al., 2016). Originally presented by Thall and Wathen (2007), this approach is able to incorporate a physician's preference for one treatment over another based on previous data or experience. Treatment allocation probabilities are updated according to accruing data in the trial. The method has been extended to allow for the presence of biomarkers and other covariates by Zhang et al. (2016) and Wason et al. (2015).

Recently, some work has been undertaken on approaching personalized clinical trials from an optimal design framework. Ondra et al. (2016b) used a decision theoretic framework, where utility functions are constructed to find optimal sample sizes and also to determine whether testing should be performed in a particular subgroup or in the general population. These utility functions incorporated such factors as the costs of running the trial and the expected effect of the treatment. Zhang et al. (2007) were the first to consider optimal response-adaptive designs for survival trials, focusing on how to minimize the total hazard. Sverdlov et al. (2011) proposed a D_A-optimal approach to estimate most precisely the treatment effect under power constraints, as well as a method of treatment allocation involving non-linear programming which is designed to maximize power. Ethical constraints were also incorporated, such as minimizing the expected total hazard, or minimizing the number of patients who receive sub-optimal treatments.

Lee and Wason (2019) proposed an optimal design approach to designing Phase III biomarker-driven trials. The idea is that data from a Phase II trial can be used to inform which hypotheses should be tested in a Phase III trial. They proposed a treatment allocation scheme for the Phase III trial based on a weighted L-optimal design. They showed that the design helps to ensure that false null hypotheses are rejected with high power.

We discuss two important statistical topics that arise in sequential trials for personalized medicine: the issue of multiplicity, and the the issue of subgroup selection. Although these topics are outside the scope of what we investigate in our simulations, we mention briefly some key findings in the literature.

6.1.1 Multiplicity

In the setting of sequential clinical trials, multiple hypothesis tests are conducted. The first source of multiplicity comes from repeatedly testing data which accumulates over time (Pocock, 1982). For ethical reasons, it is desirable to stop a trial once there is sufficient evidence of treatment difference or lack thereof, so it is necessary to conduct intermittent analyses. The second source of multiplicity comes from testing hypotheses for different subgroups. Both aspects result in a need to adjust significance levels of the tests.

Repeated hypothesis testing as data accumulates can be conducted either after each patient response or after a batch of patients respond to treatment. This results in the problem that, even if the null hypothesis is true, running the trial for a sufficiently long time can result in a test statistic which appears significant (Armitage et al., 2001, p.496). One common solution is to choose a nominal significance level for each test which is stricter than the overall significance level; the nominal significance level can be found by numerical integration given a specified overall significance level and power of the test. Pocock (1982) investigated batch sequential designs where stricter nominal significance levels are used. Results from simulations show that for the case of a continuous response and two treatments, it is unlikely that having more than five batches will lead to a smaller expected sample size if the alternative hypothesis is true. One exception is if there is an extremely large treatment effect. In this case, frequent testing can lead to early stopping of the trial. They also considered approaches where the nominal significance level is allowed to vary. In an approach described by O'Brien and Fleming (1979), nominal significance levels were set to be monotonically increasing with the significance level of the final test set to the overall significance level. Pocock (1982) found from simulations with two and five batches where high power is required, that there is little advantage to having varying significance levels as opposed to having a constant nominal significance level. Armitage et al. (1969) showed simulation results for the probability of the trial stopping at or before the nth observation, for given constant nominal significance levels and assuming that tests are conducted after each response. They looked at cases where the response is normal, binomial and exponential.

Hypothesis testing within subgroups also results in multiplicity; in particular, Tanniou et al. (2016) described how subgroup analysis is often conducted in a way which can be misleading, and their results are sometimes mistakenly interpreted as being confirmatory rather than exploratory. Moyé and Deswal (2001) described methodology of selecting type I error rates in advance in order to calculate treatment effect within a prespecified subgroup which has a confirmatory interpretation.

6.1.2 Subgroup selection

A major challenge in personalized clinical trials is developing methodologies to select appropriate biomarker subgroups (Maitland and Schilsky, 2011). Sargent et al. (2005) remarked that there is a lack of systematic methodology for validating suitability of biomarker subgroups for clinical trials, which can lead to biomarkers being included without evidence of being predictive of treatment effect, or effective biomarkers being introduced slowly into clinical settings. A number of clinical trial designs have been developed specifically to validate the biomarker's effectiveness in predicting treatment effect (Sargent et al., 2005). A literature review on methods for confirming targeted subgroups for clinical trials is provided by Ondra et al. (2016a). Zhang et al. (2018) proposed methodology for enrichment designs where subgroups are selected based on accruing data in the trial.

6.2 Method

We now describe the method for treatment allocation for when there is interest in estimating effects of treatment-biomarker interactions in the context of personalized medicine.

6.2.1 Model and set-up

In previous chapters, we assumed that there was one binary treatment. We now generalize this and assume there are t binary treatments, which have levels coded by 1 (new treatment) and 0 (control), as this is the conventional form of coding treatment in medical statistics. Biomarkers are represented by binary covariates, and they are assumed to be independent. The biomarkers also have levels coded by 1 (patient has the biomarker) and 0 (patient does not have the biomarker). For k biomarkers and t treatments, there are $R = 2^k$ subgroups in total and $2^k \times t$ hypotheses to test. In the case of treatment-linked biomarkers, treatments are usually assigned only to specific subgroups for whom the treatment is suspected to be more effective for. For the moment, we ignore this aspect and assume that all treatments may be applied to all subgroups.

We consider the case where the response is continuous. We assume a linear model for the response:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{6.1}$$

where $\boldsymbol{X}$ includes a column for the intercept, columns for each biomarker, columns for each treatment, and columns for each treatment - biomarker interaction; there are $1 + k + t + tk$

columns. The vector $\boldsymbol{\beta}$ has length q and we assume that $\boldsymbol{\epsilon} \sim N(0, \boldsymbol{I}_n\sigma^2)$.

We follow the notation by Lee and Wason (2019). In each of the $R = 2^k$ subgroups, we wish to test the null hypothesis, $\boldsymbol{c}_r^\intercal\boldsymbol{\beta} \geq \tau_r$, against the alternative hypothesis, $\boldsymbol{c}_r^\intercal\boldsymbol{\beta} < \tau_r$, where $\tau_r \leq 0$ is the minimum uninteresting threshold. A negative parameter value is deemed interesting since it is interpreted as reducing the symptoms of the disease. The vector $\boldsymbol{c}_r$ is an indicator vector for the linear combination of parameters of interest in subgroup r, for $r \in \{1, ..., R\}$. Typically, $\boldsymbol{c}_r$ will have value one for the entry corresponding to the treatment effect in $\boldsymbol{\beta}$ and potentially also for entries corresponding to interaction effects of interest, and zeros otherwise.

We define the probability P_r as

$$P_r = P(\boldsymbol{c}_r^\intercal\hat{\boldsymbol{\beta}} < \tau_r) \text{ for all } r \in \{1, ..., R\}. \tag{6.2}$$

In order to compute P_r, we note that, under the model in Equation (6.1), if σ^2 is known, $\boldsymbol{c}_r^\intercal\hat{\boldsymbol{\beta}}$ has distribution given by $N\left(\boldsymbol{c}_r^\intercal\boldsymbol{\beta}, \sigma^2\boldsymbol{c}_r^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c}_r\right)$. If σ^2 is unknown, $\boldsymbol{c}_r^\intercal\hat{\boldsymbol{\beta}}$ has a t-distribution with degrees of freedom $n - q$, with location and scale parameters given by $\boldsymbol{c}_r^\intercal\boldsymbol{\beta}$ and $\sigma^2\boldsymbol{c}_r^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c}_r$, respectively: $t_{n-q}\left(\boldsymbol{c}_r^\intercal\boldsymbol{\beta}, \sigma^2\boldsymbol{c}_r^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c}_r\right)$.

Lee and Wason (2019) define weights w_r as

$$w_r = \begin{cases} P_r \text{ if } P_r \geq \kappa \\ 0 \text{ otherwise} \end{cases} \quad \text{for all } r \in \{1, ..., R\}, \tag{6.3}$$

where κ is a chosen threshold. In our simulations, we set κ to zero, so we have $w_r = P_r$.

As we are interested in minimizing the variance of linear combinations of parameters of the model, the L-optimal objective function is an appropriate choice for an optimality criterion (Atkinson et al., 2007, p. 142). Further, since some linear combinations are deemed more important than others, we consider weighting this criterion (Morgan and Wang, 2010). Given a design matrix $\boldsymbol{X}$ and vectors $\boldsymbol{c}_r$, $1 \leq r \leq R$, we define the weighted L-optimal objective function:

$$\Psi_L(\boldsymbol{X}, \boldsymbol{c}_r) = \sum_{r=1}^{R} w_r\boldsymbol{c}_r^\intercal(\boldsymbol{X}^\intercal\boldsymbol{X})^{-1}\boldsymbol{c}_r. \tag{6.4}$$

We wish to construct $\boldsymbol{X}$ to minimize $\Psi_L(\boldsymbol{X}, \boldsymbol{c}_r)$.

As a comparison, in our simulations in Section 6.3.3, we also consider the D_A-optimal design:

$$\Psi_{D_A}(\boldsymbol{X}) = \det\left(\boldsymbol{A}^\intercal(\boldsymbol{X}^\intercal\boldsymbol{X})^{-1}\boldsymbol{A}\right), \tag{6.5}$$

where the matrix $\boldsymbol{A}$ contains vectors $\boldsymbol{c_r}$ as columns for some values of $r \in \{1, ..., R\}$ so that $\boldsymbol{A}$ has full column rank (Atkinson et al., 2007, p.137). We denote by b the number of columns in $\boldsymbol{A}$. The objective function for a weighted D_A-optimal design is as follows:

$$\Psi_{D_Aw}(\boldsymbol{X}, \boldsymbol{c_r}) = \prod_{r=1}^{R}\left(\boldsymbol{c}_r^\intercal(\boldsymbol{X}^\intercal\boldsymbol{X})^{-1}\boldsymbol{c}_r\right)^{w_r}. \tag{6.6}$$

The weighted D_A-optimal criterion can be thought of as a geometric mean of the variances of the linear combinations of interest, whereas the weighted L-optimal criterion can be thought of as an arithmetic mean. In addition, we also show in our simulations in Section 6.3.3 how minimization fares against these optimal design approaches.

We can also define weighted L- and D_A-efficiencies of a design $\boldsymbol{X}$ relative to another design $\boldsymbol{X}^*$ as

$$\text{Eff}_L = \frac{\Psi_L(\boldsymbol{X}^*, \boldsymbol{c}_r)}{\Psi_L(\boldsymbol{X}, \boldsymbol{c}_r)}, \tag{6.7}$$

$$\text{Eff}_{D_A} = \left\{\frac{\Psi_{D_A}(\boldsymbol{X}^*, \boldsymbol{c}_r)}{\Psi_{D_A}(\boldsymbol{X}, \boldsymbol{c}_r)}\right\}^{1/b}, \tag{6.8}$$

where b is the number of rows in matrix $\boldsymbol{A}$.

6.2.2 The sequential implementation

In Section 2.6.2, we discussed a method proposed by Atkinson (1982) for allocating treatments sequentially to patients in order to construct an approximately D_A-optimal design. We adapt this method to obtain a weighted L-optimal design sequentially. For the moment, we focus on a myopic approach, which we will then extend to a nonmyopic approach in Chapter 7. We assume that we have one binary treatment.

We begin with an initial design with n_0 patients, which is constructed using an exchange algorithm with the L-optimal objective function with equal weights. We need n_0 to be sufficiently large to be relatively confident that parameter values can be estimated. Problems with estimation can occur if $n_0 < q$, if there is exact collinearity in the columns of the

design matrix, or if not all combinations of treatments and covariates are represented by the sample in the initial design. To avoid problems with estimation when fitting the model even with a small sample size, we choose to use a Bayesian linear model with a diffuse normal-inverse-gamma prior. We assume that responses are observed immediately after treatments are assigned. From the initial treatments and responses, we obtain estimates of $\hat{\boldsymbol{\beta}}$, $\hat{\sigma}$ and we can also estimate the weights w_r. In Section 6.2.3, we describe a fully Bayesian approach to constructing the design.

We denote by $\boldsymbol{T}_{i-1}$ the $(i-1) \times k$ matrix where each column is the vector $\boldsymbol{t}_{i-1}$:

$$\boldsymbol{T}_{i-1} = \left[\underbrace{\boldsymbol{t}_{i-1} \ \ldots \ \boldsymbol{t}_{i-1}}_{\text{k times}} \right]. \tag{6.9}$$

Given that the treatments $\boldsymbol{t}_{i-1}$ have been assigned, the responses for the previous patients $\boldsymbol{y}_{i-1}$ have been observed and the biomarkers for all patients up until the current patient $\boldsymbol{Z}_i$ have been observed, we define the design matrix as follows:

$$\boldsymbol{X}(t_i) = \begin{bmatrix} \mathbf{1} & \boldsymbol{Z}_{i-1} & \boldsymbol{t}_{i-1} & \boldsymbol{Z}_{i-1} \circ \boldsymbol{T}_{i-1} \\ 1 & \boldsymbol{z}_i^\intercal & t_i & t_i \boldsymbol{z}_i^\intercal \end{bmatrix}, \tag{6.10}$$

where $\circ$ denotes the Hadamard product. We express the weighted L-optimal design criterion defined previously for a static design matrix in Equation (6.4) as

$$\Psi_L(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1}) = \sum_{r=1}^{R} w_{r,i-1} \left(\boldsymbol{c}_r^\intercal \left(\boldsymbol{X}(t_i)^\intercal \boldsymbol{X}(t_i) \right)^{-1} \boldsymbol{c}_r \right). \tag{6.11}$$

A treatment is chosen by sampling from all of the possible treatments, where the sampling weight for treatment m, for $m \in \{1, 2, ..., t\}$, is given by

$$\frac{\Psi_L(t_i = j \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})^{-1}}{\sum_{m=1}^{t} \Psi_L(t_i = m \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})^{-1}}. \tag{6.12}$$

Once the treatment is chosen, the design matrix is updated, a new response is observed, and we update $\hat{\boldsymbol{\beta}}$, $\hat{\sigma}$ and w_r. The pseudocode is provided in Algorithm 7.

Algorithm 7 Sequential weighted L-optimal design: function returns a design matrix given covariate values for n patients, $\boldsymbol{Z}_n$, and the number of patients in initial design n_0

function SEQWL($\boldsymbol{Z}_n, n_0$)

Initialization
Construct initial design $\boldsymbol{X}_{n_0}$: use the exchange algorithm with L-optimality with equal weights.
Observe responses $y_1, y_2, ..., y_{n_0}$
Fit the model $\boldsymbol{y}_{n_0} \sim lm(\boldsymbol{X}_{n_0})$ to obtain $\hat{\boldsymbol{\beta}}_0$ and $\hat{\sigma}_0$.
Calculate $w_{r_0} = P(\boldsymbol{c}_r^\intercal \hat{\boldsymbol{\beta}}_0 < \tau_r)$ for all r.

for i in $n_0 + 1$ to n **do**
    Observe $z_{i,1}, ..., z_{i,k}$
    Calculate $\Psi_L(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$ for each treatment.
    Sample treatment for patient i where probability of treatment j is given in Equation (6.12).
    Append new row containing covariates for patient i, t_i and their interactions to the design matrix to form $\boldsymbol{X}_i$.
    Observe response y_i.
    Refit model and $\boldsymbol{y}_i \sim lm(\boldsymbol{X}_i)$ update the parameter estimates $\hat{\boldsymbol{\beta}}_i$ and $\hat{\sigma}_i$.
    Update weights w_{r_i} for all r.
end for

return $\boldsymbol{X}$

end function

6.2.3 A Bayesian approach

It is helpful in the case where we have a small initial sample size or when we have rare biomarkers, to use the Bayesian linear model so that we do not have problems with estimation. In this section we describe a fully Bayesian approach where we choose a prior distribution for $(\boldsymbol{\beta}, \sigma^2)$, e.g. normal inverse gamma distribution, and calculate the posterior probability $w_r = P(\boldsymbol{c}_r^\intercal \boldsymbol{\beta} < \tau_r | \boldsymbol{y})$.

Suppose we have a normal inverse gamma distribution for $(\boldsymbol{\beta}, \sigma^2)$. Using notation similar to O'Hagan (1994, p.307), we have:

$$\boldsymbol{\beta} \mid \sigma^2, \boldsymbol{m}, \boldsymbol{V}^{-1} \sim \mathrm{N}(\boldsymbol{m}, \sigma^2 \boldsymbol{V}^{-1}), \tag{6.13}$$

$$\sigma^2 \mid a, d \sim \Gamma^{-1}(a, d), \tag{6.14}$$

with hyperparameters a, d, $\boldsymbol{m}$, and $\boldsymbol{V}$. The q-vector $\boldsymbol{m}$ is the prior mean of $\boldsymbol{\beta}$ given σ^2, and the $n \times n$ matrix $\boldsymbol{V}$ is the prior precision of $\boldsymbol{\beta}$ given σ^2. The joint prior is given by:

$$\begin{aligned} & f(\boldsymbol{\beta}, \sigma^2 \mid \boldsymbol{m}, \boldsymbol{V}, a, d) \\ &= \mid \boldsymbol{V} \mid^{-1} (2\pi)^{-q/2} \frac{d^a}{\Gamma(a)} \left(\frac{1}{\sigma^2}\right)^{k/2+a+1} \exp\left(-\frac{2d + (\boldsymbol{\beta} - \boldsymbol{m})^\intercal \boldsymbol{V}^{-1} (\boldsymbol{\beta} - \boldsymbol{m})}{2\sigma^2}\right). \end{aligned}$$

The posterior distribution of $(\boldsymbol{\beta}, \sigma^2)$ is a normal inverse gamma distribution, $\text{NIG}(a^*, d^*, \boldsymbol{m}^*, \boldsymbol{V}^*)$, where the parameter values are given below (O'Hagan, 1994, p.308):

$$\begin{aligned} a^* &= a + \boldsymbol{m}^\intercal \boldsymbol{V}^{-1} \boldsymbol{m} + \boldsymbol{y}^\intercal \boldsymbol{y} - (\boldsymbol{m}^*)^\intercal (\boldsymbol{V}^*)^{-1} \boldsymbol{m}^*, \\ d^* &= d + n, \\ \boldsymbol{V}^* &= (\boldsymbol{V}^{-1} + \boldsymbol{X}^\intercal \boldsymbol{X})^{-1}, \\ \boldsymbol{m}^* &= (\boldsymbol{V}^{-1} + \boldsymbol{X}^\intercal \boldsymbol{X})^{-1} (\boldsymbol{V}^{-1} \boldsymbol{m} + \boldsymbol{X}^\intercal \boldsymbol{y}). \end{aligned}$$

We estimate the parameters under the Bayesian approach by taking the posterior mode of the marginal distributions. The marginal distribution of $\boldsymbol{\beta}$ is given by the multivariate t distribution with degrees of freedom d^* and posterior location parameter $\boldsymbol{m^*}$ and posterior shape parameter $\boldsymbol{a^* V^*}$. The posterior mode is given by $\boldsymbol{m^*}$. The marginal distribution of σ^2 is an inverse gamma distribution with hyperparameters a^* and d^*. The posterior mode is given by $\sqrt{\frac{a^*}{d^*-2}}$.

In the Bayesian setting, obtaining an estimate for w_r or P_r as defined by Equations (6.3) and (6.2) is more cumbersome. The marginal distribution of $\boldsymbol{\beta}$ a multivariate t distribution, and linear combinations of t distributions do not have a simple analytic form. Thus, it is necessary to generate h simulations of $\boldsymbol{\beta}$ and compute the proportion of those h deviates for which we have $\boldsymbol{c}_r^\intercal \boldsymbol{\beta} < \tau_r$, where we set $h = 100$. Given this estimate of w_r, the weighted L- and D_A-optimal objective functions are defined in the Bayesian approach as in Equations (6.4) and (6.6).

6.3 Simulations

We wish to construct weighted L-optimal designs sequentially for some simple scenarios so that we can observe how well the parameters are being estimated and examine some characteristics of the designs. At each sample size between n_0 and n, we consider the following performance measures:

1. The estimates $\hat{\boldsymbol{\beta}}$.

2. The estimates $\hat{\sigma}$.
3. The weights w_r.
4. The proportion of patients assigned to the new treatment ($t_i = 1$).
5. The weighted L-optimality of the design.

In addition, we are interested in the following measures which we explain below:

6. The theoretical value of the power of the hypothesis test in each subgroup.
7. The empirical value of the power of the hypothesis test in each subgroup.
8. The empirical value of 1-specificity of the hypothesis test in each subgroup.

We are interested in the power of the hypothesis test. If the true value of σ^2 is known, $\boldsymbol{c_r}^\intercal\hat{\boldsymbol{\beta}}$ has distribution given by $N\left(\boldsymbol{c_r}^\intercal\boldsymbol{\beta}, \sigma^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}\right)$. Assuming that σ^2 is unknown, we have

$$\boldsymbol{c_r}^\intercal\hat{\boldsymbol{\beta}} \sim t_{n-q}\left(\boldsymbol{c_r}^\intercal\boldsymbol{\beta}, \sigma^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}\right). \tag{6.15}$$

Our test statistic is given by

$$\frac{\boldsymbol{c_r}^\intercal\hat{\boldsymbol{\beta}} - \tau_r}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}}, \tag{6.16}$$

and the p-value of the hypothesis test is given by the probability that the test statistic is in the critical region:

$$\begin{aligned}
&P\left(\frac{\boldsymbol{c_r}^\intercal\hat{\boldsymbol{\beta}} - \tau_r}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}} < t_{n-q,\alpha}\right) \\
&= P\left(\frac{\boldsymbol{c_r}^\intercal\hat{\boldsymbol{\beta}} - \boldsymbol{c_r}^\intercal\boldsymbol{\beta}}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}} < t_{n-q,\alpha} + \frac{\tau_r - \boldsymbol{c_r}^\intercal\boldsymbol{\beta}}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}}\right) \\
&= P\left(t_{n-q} < t_{n-q,\alpha} + \frac{\tau_r - \boldsymbol{c_r}^\intercal\boldsymbol{\beta}}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}}\right) \\
&= T_{n-q}\left(t_{n-q,\alpha} + \frac{\tau_r - \boldsymbol{c_r}^\intercal\boldsymbol{\beta}}{\sqrt{\hat{\sigma}^2\boldsymbol{c_r}^\intercal\left(\boldsymbol{X}^\intercal\boldsymbol{X}\right)^{-1}\boldsymbol{c_r}}}\right),
\end{aligned}$$

where T_{n-q} is the cumulative distribution function of the t_{n-q}-distribution and $t_{n-q,\alpha}$ is the α quantile of the t_{n-q}-distribution. We calculate this quantity in subgroups in which

the alternative hypothesis is true; this gives us a theoretical quantity of the power of the hypothesis test. We also find the proportion of simulations in which the test statistic in Equation (6.16) is smaller than the critical value $t_{n-q,\alpha}$. In subgroups in which the null hypothesis is true, this proportion is equivalent to $1 - specificity$, and in subgroups in which the alternative hypothesis is true, this proportion is equivalent to an empirical value of the power of the hypotheses test.

We note that, in a Bayesian framework, power is equivalent to *sensitivity* (Sharma et al., 2009).

6.3.1 An example with one biomarker

We consider the case where there is one binary biomarker z, with levels coded by 0 and 1. This defines two subgroups: the subgroup with $z_i = 0$ and the subgroup with $z_i = 1$. We assume there is one binary treatment, and we wish to test its effectiveness in the two subgroups. The ith row of the design matrix for patient i is $\{1 \quad z_i \quad t_i \quad z_i t_i\}$.

The vectors $\boldsymbol{c}_r$ are as follows,

$$\boldsymbol{c}_1 = (0, 0, 1, 0) \tag{6.17}$$

$$\boldsymbol{c}_2 = (0, 0, 1, 1)\,, \tag{6.18}$$

where $\boldsymbol{c}_1$ corresponds to the effectiveness of the treatment for the subgroup where $z_i = 0$, and $\boldsymbol{c}_2$ is that of the subgroup where $z_i = 1$.

We consider a simple simulation with the sequential weighted L-optimal design described in Section 6.2.2 with pseudocode given by Algorithm 7 where we choose a normal-inverse-gamma prior $\boldsymbol{m} = \boldsymbol{0}$, $\boldsymbol{V} = \frac{1}{100}\boldsymbol{I}_4$, $a = 2, d = 2$. This means that we have set the prior distribution of $\boldsymbol{\beta}$ given σ^2 to be centered around zero with a high variance, and the prior distribution of σ^2 is relatively flat. Essentially, our prior is diffuse. Other settings for the simulation are given below:

Table 6.1: Settings for a simulation in a personalized medicine setting with one biomarker

Biomarkers	$z \sim$ Bernoulli(0.3)
True $\boldsymbol{\beta}$	all entries (for intercept, $\boldsymbol{t}$, $\boldsymbol{z}$, $\boldsymbol{tz}$) are zero
True σ	1.5
τ_r	-1 for all r
κ	0
Initial sample size	15
Number of simulations	100

The structure of our simulations are as follows:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) The weighted L-optimal criterion, given by Equation (6.4), is used to generate designs sequentially with settings given in Table 6.1.

(c) Designs are evaluated using the performance measures given at the start of Section 6.3 at each sample size.

II (a)-(c) above is repeated 100 times to obtain a distribution of the performance measure for each sample size.

Figures 6.3 to 6.8 display how the distribution of a number of performance measures change as sample size increases. The performance measures considered include the parameter estimates $\hat{\boldsymbol{\beta}}$, $\hat{\sigma}$, the weights w_r, the power of the hypothesis tests, the proportion of patients assigned to the new treatment ($t_i = 1$) and the weighted L-optimality of the design.

We observe in Figure 6.3 that the distribution of the estimates of the coefficients are all centered around zero. The variance appears to be highest for the interaction term, and lowest for the intercept. Figure 6.4 shows that the estimates of σ are also distributed around the true value. Overall, it seems that the design is estimating the parameters in a way that we expect.

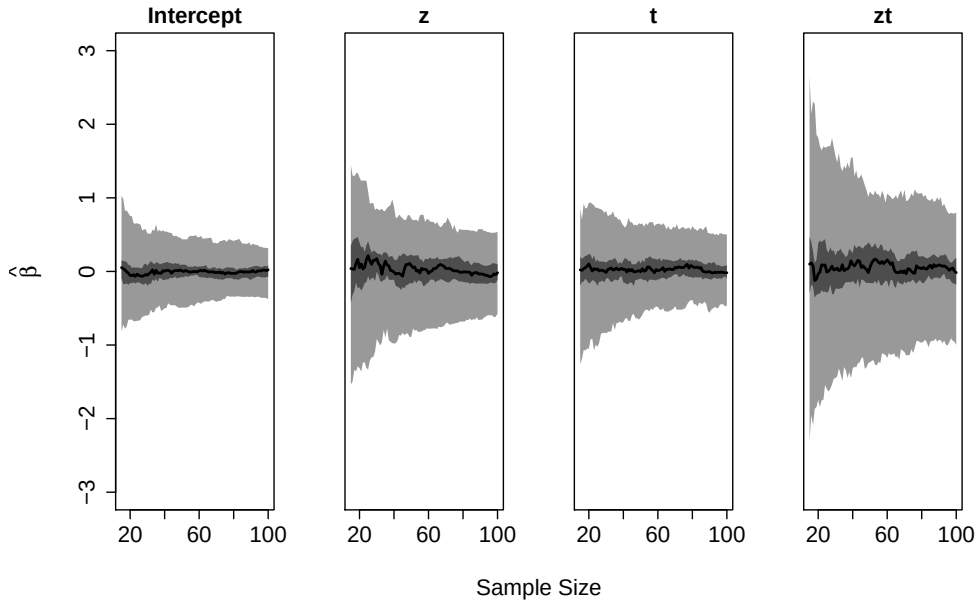

Figure 6.3: Distribution of $\hat{\boldsymbol{\beta}}$ vs sample size for the weighted L-optimal design for the linear model case with one biomarker. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

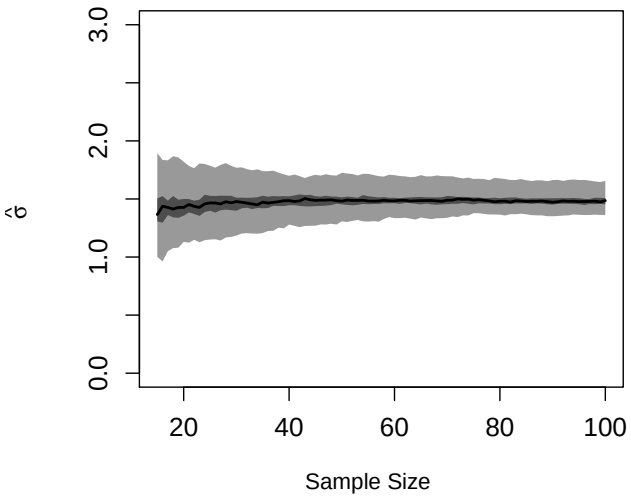

Figure 6.4: Distribution of $\hat{\sigma}$ vs sample size for the weighted L-optimal design for the linear model case with one biomarker.

In Figure 6.5, we plot the distribution of the weights for the two subgroups. Since the treatment is not effective in either subgroup, we expect the true value of the weights to be zero for both subgroups. We observe that as sample size increases, the estimates of

the weights do appear to converge to zero. This convergence occurs more quickly with the subgroup where z takes the value zero, since our simulation settings are such that more patients with this profile are generated.

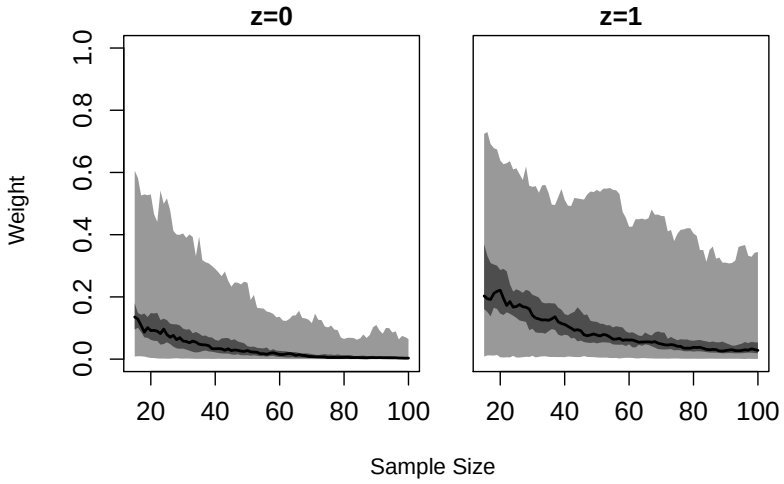

Figure 6.5: Distribution of w_r vs sample size for the weighted L-optimal design for the linear model case with one biomarker. This is shown for subgroups $(z = 0)$ and $(z = 1)$.

In Figure 6.6, we find the proportion of patients allocated to treatment $t = 1$ in each subgroup at each sample size, and we plot the distribution of these proportions against the total sample size. We observe that the distributions are centered around 0.5, so the weighted L-optimal design leads to a balanced design under the settings chosen in the simulations. There is greater variance in the subgroup where $z = 1$ since this biomarker profile is more rare.

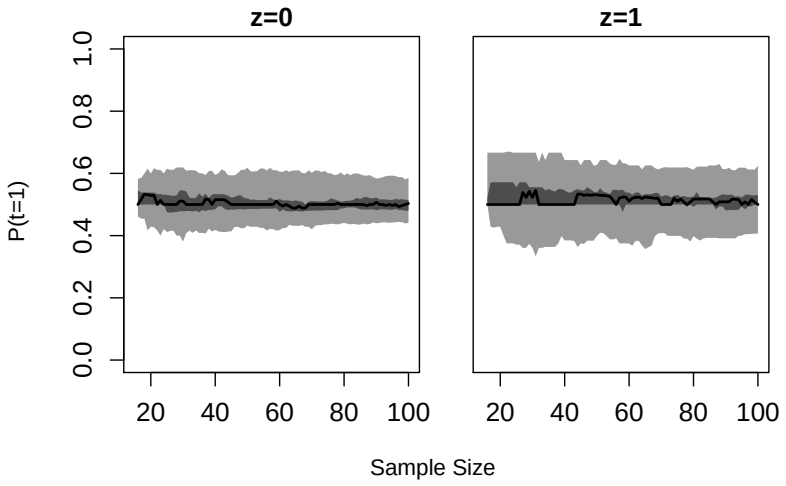

Figure 6.6: Distribution of the proportion of new treatment in each subgroup vs total sample size for the weighted L-optimal design for the linear model case with one biomarker.

In our example, the treatment is not effective so the null hypothesis is true. We plot 1-specificity for the test for each sample size in Figure 6.7 for both subgroups.

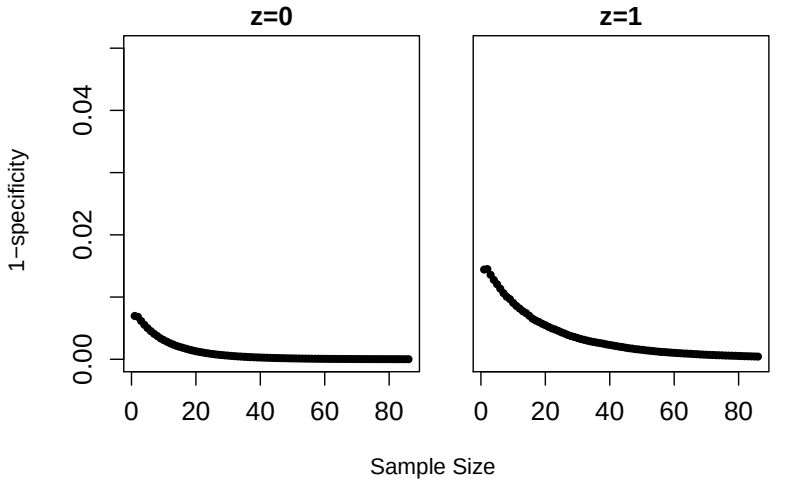

Figure 6.7: 1-specificity for the hypothesis test vs sample size for subgroups ($z = 0$) (left panel) and ($z = 1$) (right panel) for the weighted L-optimal design for the linear model case with one biomarker. The proportion of correct null hypotheses that have been rejected is given for each sample size.

We plot the distribution of Ψ_L calculated at each value of sample size. We observe the values of the objective functions decrease as sample size increases; as our objective function is weighted, this decreasing trend is due to a combination of more information accruing as

well as weights becoming smaller.

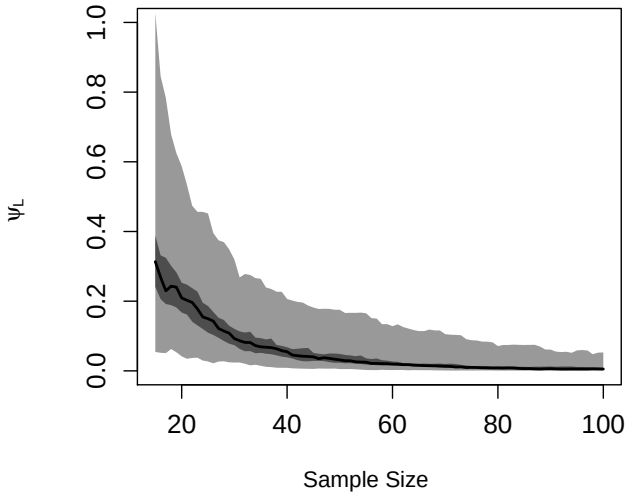

Figure 6.8: Distribution of Ψ_L vs sample size for the weighted L-optimal design for the linear model case with one biomarker.

In Appendix F, we provide plots for simulations that we ran for two more cases, where settings for the simulation are as in Table 6.1, except for the following changes:

1. The coefficient of $\boldsymbol{t}$ is set to -3. This means that the treatment is effective in both subgroups so instead of plotting 1-specificity in the two subgroups, we plot the power of the hypotheses tests instead. The plots are given in Appendix F.1
2. The coefficients of $\boldsymbol{t}$ and $\boldsymbol{tz}$ are both set to -3. This means that the treatment is effective overall, but is particularly effective for patients who have a positive biomarker status. For the first subgroup, we plot 1-specificity for the hypothesis tests and for the second subgroup, we plot the power. The plots are given in Appendix F.2

6.3.2 Comparing the normal inverse gamma prior, the Cauchy prior, and the fully Bayesian approach

Now, we run a simulation to compare three approaches of constructing the weighted L-optimal design: using the normal inverse gamma prior, using the Cauchy prior, and using the fully Bayesian approach. We now consider a slightly more interesting case where there are two binary biomarkers. Subgroups are denoted by two-tuples of the form $(z_{\cdot 1}, z_{\cdot 2})$, where $z_{\cdot 1}$ is the level of the first biomarker, and $z_{\cdot 2}$ is the level of the second biomarker. There are four subgroups:

Subgroup 1: $(0, 0)$
Subgroup 2: $(1, 0)$
Subgroup 3: $(0, 1)$
Subgroup 4: $(1, 1)$.

We assume one treatment, and we wish to test its effectiveness in each of the four subgroups.

For patient i with biomarker values $(z_{i,1}, z_{i,2})$ and assigned treatment t_i, the corresponding row of the design matrix $\boldsymbol{X}$ is given by:

$$\boldsymbol{x}_i = \begin{pmatrix} 1 & z_{i,1} & z_{i,2} & t_i & z_{i,1}t_i & z_{i,2}t_i \end{pmatrix}^\intercal . \tag{6.19}$$

We are interested in testing the effectiveness of treatment in each of the four subgroups. The vectors $\boldsymbol{c}_r$ are as follows:

$$\boldsymbol{c}_1 = (0, 0, 0, 1, 0, 0) \tag{6.20}$$
$$\boldsymbol{c}_2 = (0, 0, 0, 1, 1, 0) \tag{6.21}$$
$$\boldsymbol{c}_3 = (0, 0, 0, 1, 0, 1) \tag{6.22}$$
$$\boldsymbol{c}_4 = (0, 0, 0, 1, 1, 1)\,, \tag{6.23}$$

and correspond to testing the effectiveness of treatment in subgroups 1, 2, 3 and 4, respectively.

For the treatment allocation methods, we first consider the sequential weighted L-optimal design described in Section 6.2.2 with pseudocode given by Algorithm 7. We choose a diffuse normal-inverse-gamma prior, as in the previous simulation in Section 6.3.1, with $\boldsymbol{m} = \boldsymbol{0}$, $\boldsymbol{V} = \frac{1}{100}\boldsymbol{I}_6$, $a = 2, d = 2$. Secondly, we use this method with a Cauchy prior with center zero and scale 2.5. Thirdly, we use the fully Bayesian approach with the normal-inverse-gamma distribution with parameters $\boldsymbol{m} = \boldsymbol{0}$, $\boldsymbol{V} = \frac{1}{100}\boldsymbol{I}_6$, $a = 2, d = 2$, as described in Section 6.2.3. In the fully Bayesian approach, the weights $\boldsymbol{w}_r$ are computed using the posterior distribution. Other settings for the simulation are given in Table 6.2:

Table 6.2: Settings for a simulation comparing weighted L-optimal designs constructed with the NIG prior, Cauchy prior, and the fully Bayesian approach

Biomarkers	$z_{i,1}$ ∼Bernoulli(0.5), $z_{i,2}$ ∼Bernoulli(0.7)
True $\boldsymbol{\beta}$	all zero except for the entry corresponding to $\boldsymbol{z}_{\cdot 1}\boldsymbol{t}$, which is -2
True σ	1.5
τ_r	-1 for all r
κ	0
Initial sample size	15
Number of simulations	1000

In Figure 6.9, we display the distribution of $\hat{\boldsymbol{\beta}}$ under the NIG prior and Cauchy prior, and for the fully Bayesian approach, we display the posterior mode of $\boldsymbol{\beta}$. The plots appear similar for the three methods. For the estimates of the coefficients for the intercept, $\boldsymbol{t}$ and $\boldsymbol{z}_{\cdot 2}\boldsymbol{t}$, the variability is slightly reduced for small sample sizes under the Cauchy prior.

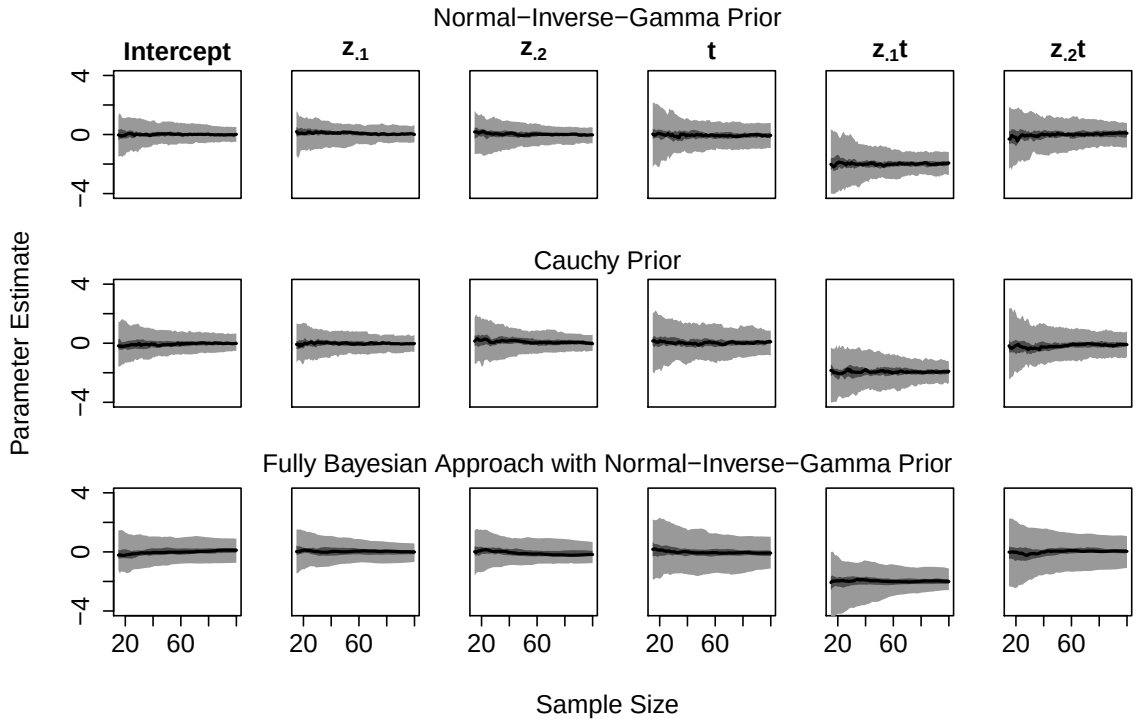

Figure 6.9: Distribution of $\hat{\boldsymbol{\beta}}$ ($\boldsymbol{m}^*$ for the Bayesian case). The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

The distribution of estimates of σ for the first two methods and the posterior mode, $\sqrt{\frac{a^*}{d^*-2}}$, for the Bayesian method, are given in Figure 6.10. There seems to be no noticeable difference

between the normal-inverse-gamma prior and the Cauchy prior. The Bayesian approach seems to lead to estimates that are biased downwards for small sample sizes.

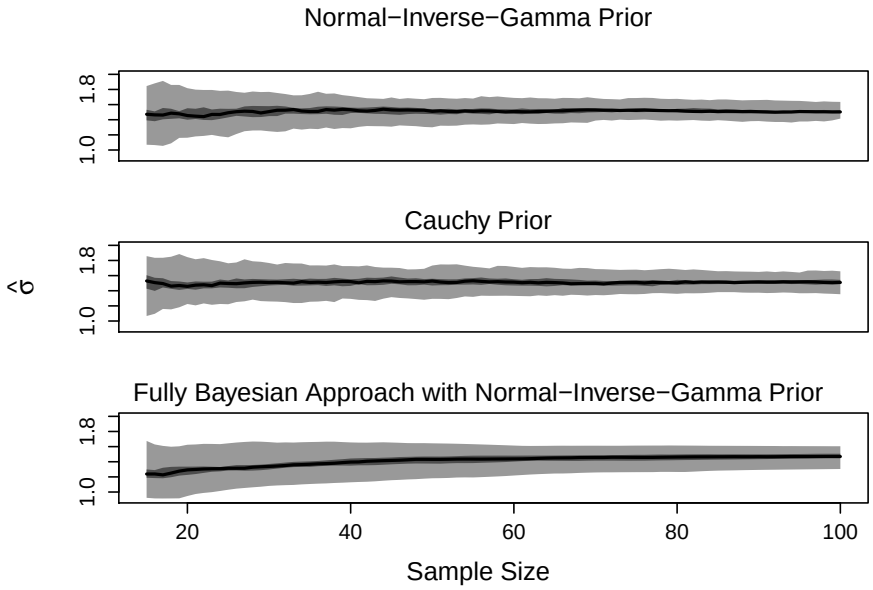

Figure 6.10: Distribution of $\hat{\sigma}$. The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

We observe in Figure 6.11 that the Bayesian method results in the weights reaching their true value much faster than the other two methods. It appears that constructing the weights using the posterior distribution, as in the Bayesian method, can be more effective than computing the weights as in Equation (6.3), but leads to greater variability.

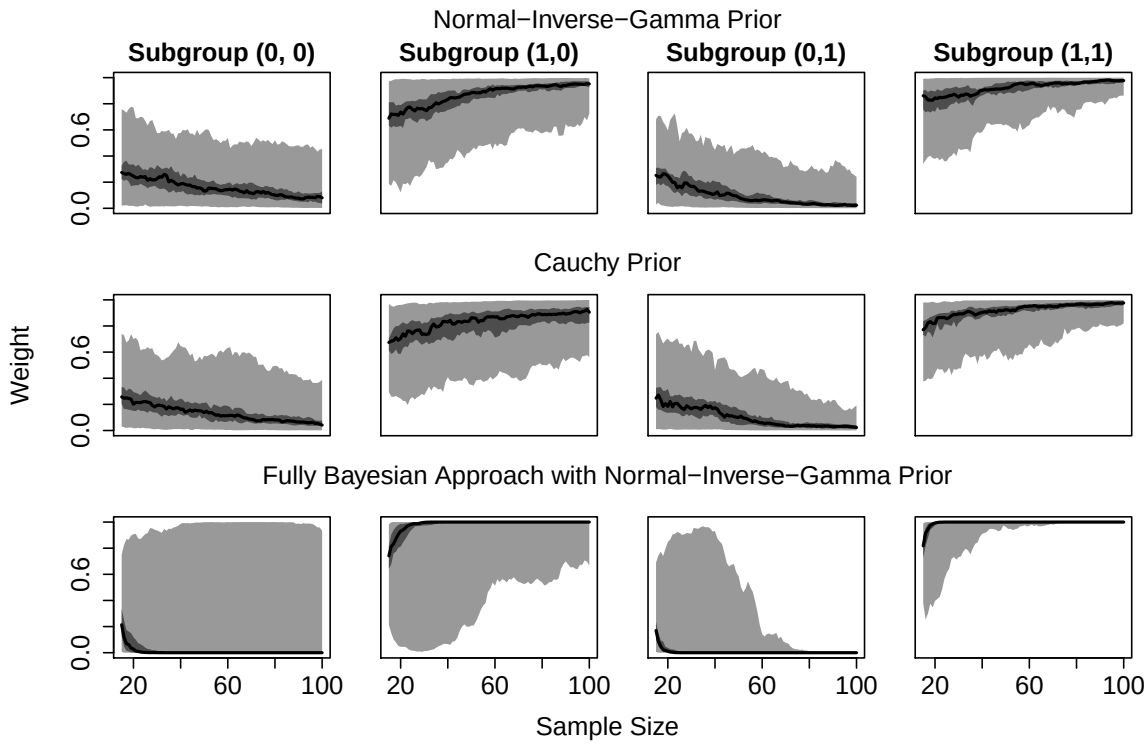

Figure 6.11: Distribution of w_r. The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

We consider the treatment balance in Figure 6.12. It appears that subgroups $(0, 0)$ and $(1, 0)$ are more imbalanced than the other two groups under all three approaches, particularly for small sample sizes. This is due to the fact that the biomarker $z_{.2}$ takes the value zero in these subgroups, which is rarer. The three approaches produce similar results.

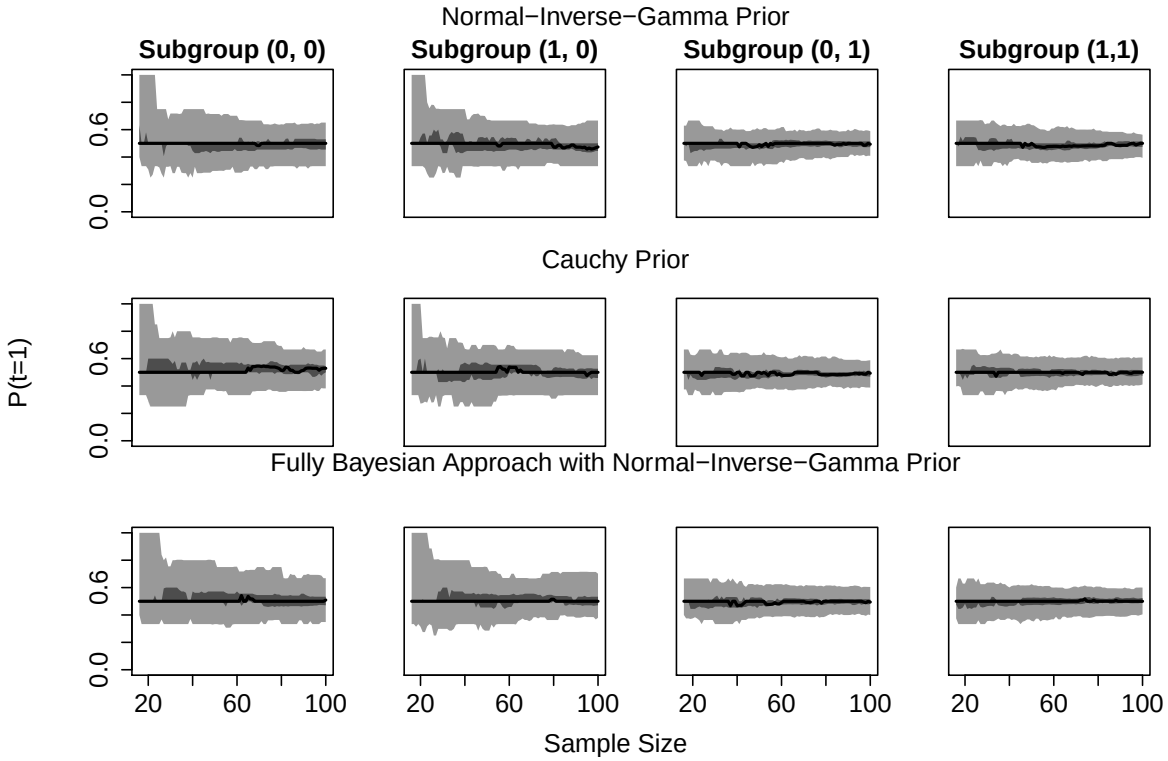

Figure 6.12: Distribution of the proportion of new treatment in each subgroup. The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

We plot the distribution of the theoretical value of the power (sensitivity in the Bayesian case) in Figure 6.13. We display the power for all three methods for subgroups (1, 0) and (1, 1) in which the null hypothesis is false. We observe that the slope of the curve is higher for the subgroup (1, 1) which is due to the fact that the second biomarker takes the more common value of one. In the subgroup (1, 0), it takes the rare value of zero so the sample size is smaller and power is reduced. The empirical values of the power of the hypothesis tests for all three methods are shown in Figure 6.14, which roughly appear to lie along the median of the theoretical distribution plots.

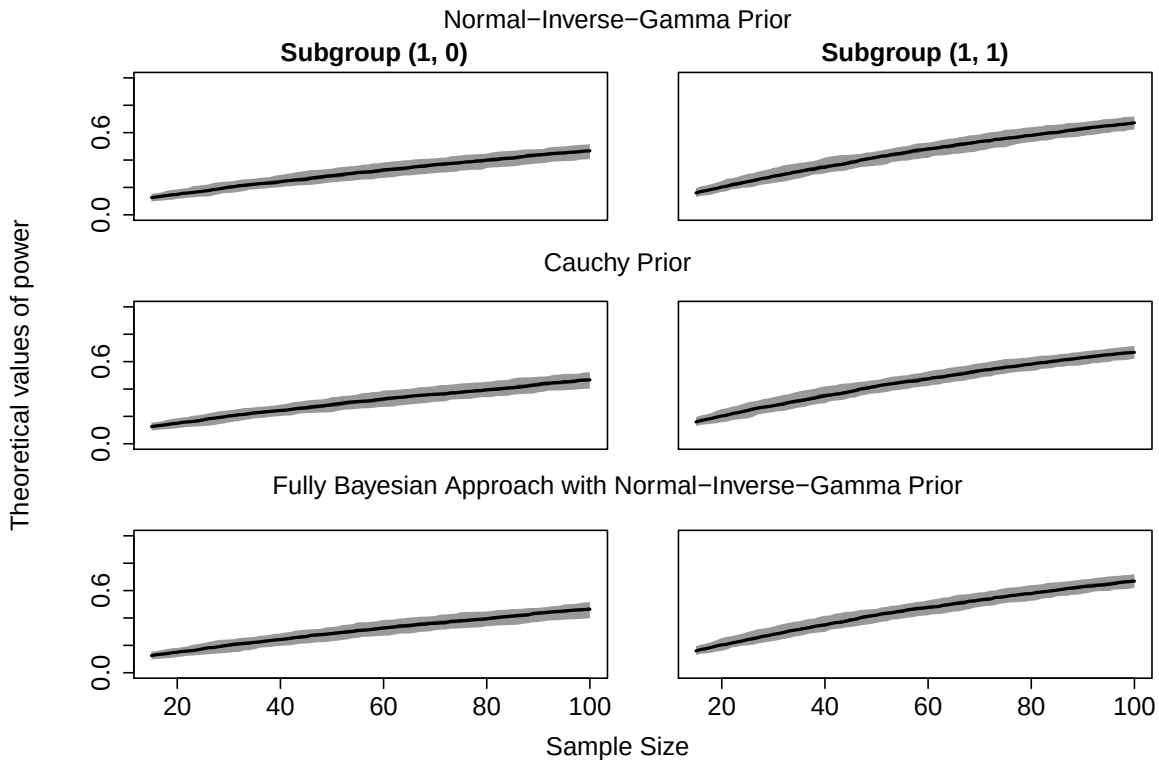

Figure 6.13: Distribution of the theoretical value of the power of the hypothesis test, or sensitivity in the Bayesian case. The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

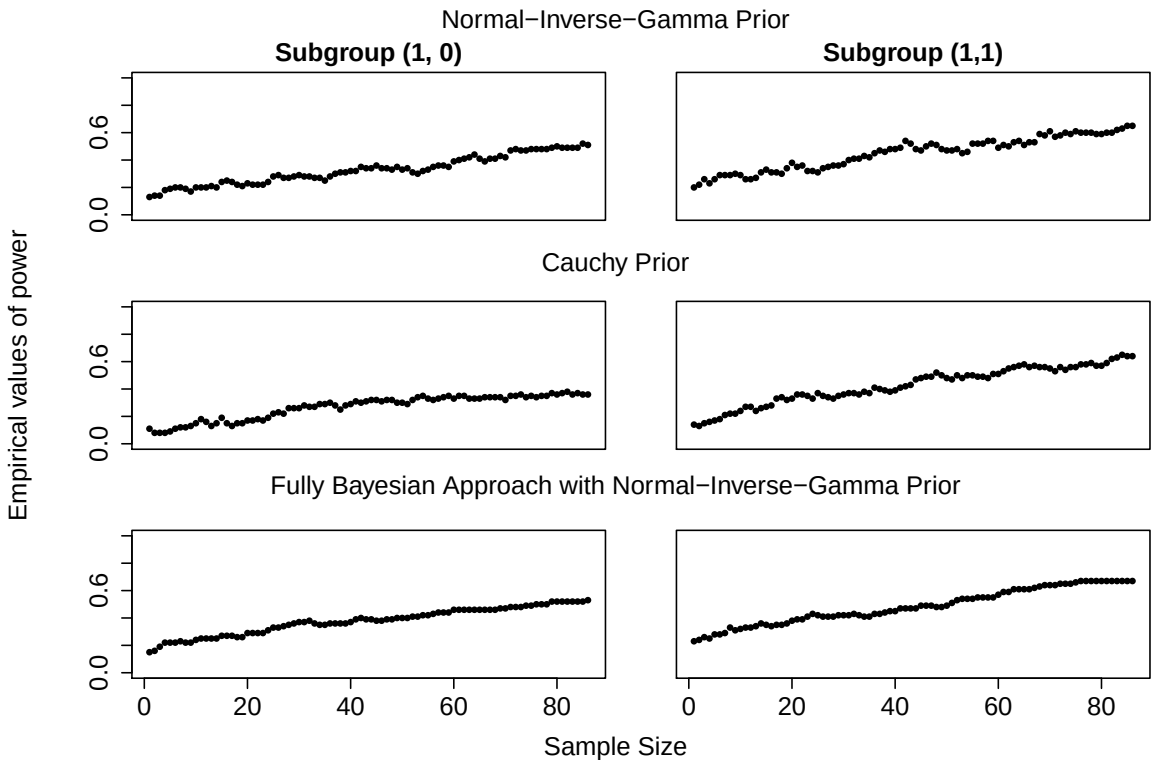

Figure 6.14: Plots of the empirical value of the power of the hypothesis test for subgroups (1,0) and (1,1), or 1− specificity in the Bayesian case. The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

In Figure 6.15, we plot 1-specificity for the subgroups (0, 0) and (1, 1). We observe that it decreases as sample size increases.

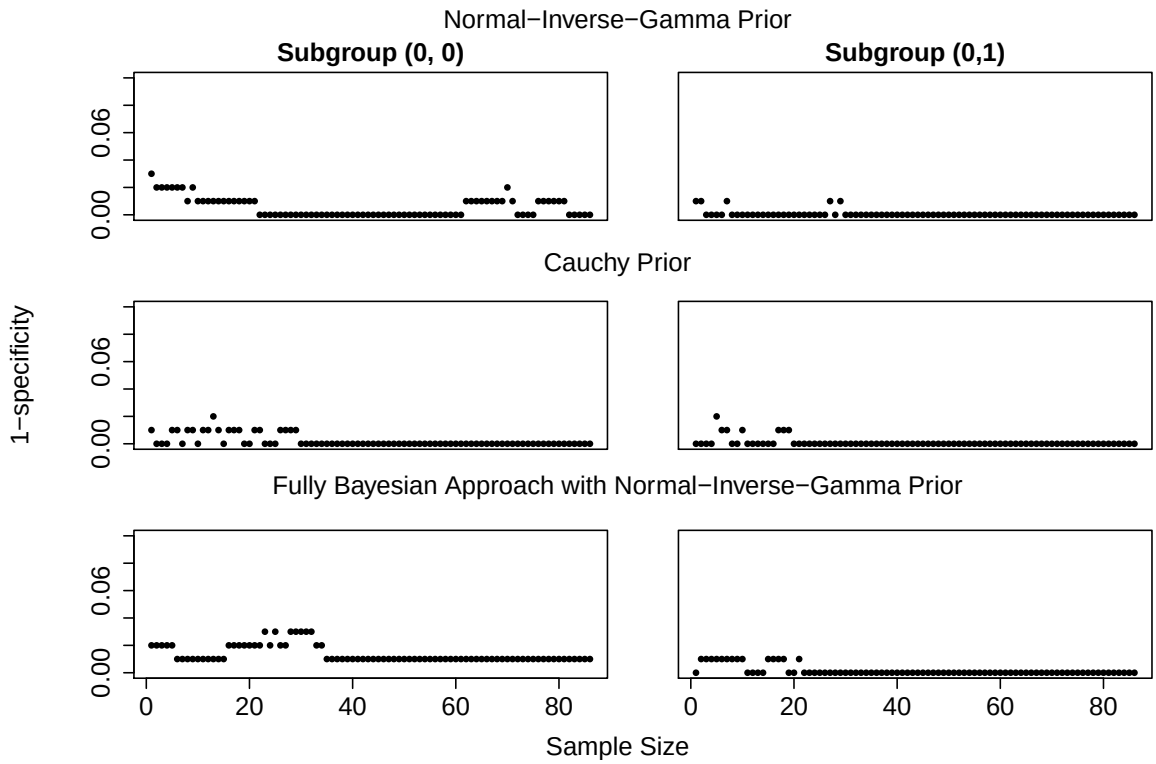

Figure 6.15: 1-specificity for subgroups (0,0) and (0,1). The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

In Figure 6.16, we observe that the value of the weighted L-optimal objective function is lower for the design constructed using the Bayesian approach. This is likely to be due to the fact that the weights reach their true value more quickly under the Bayesian approach, as shown in Figure 6.11.

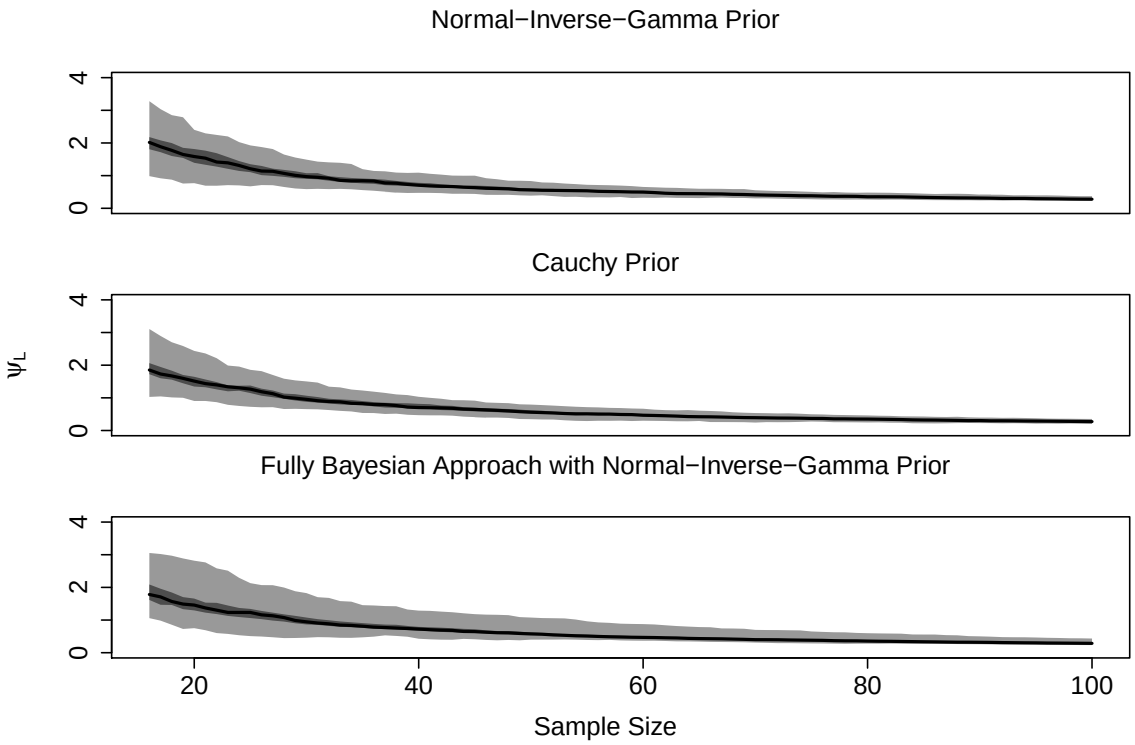

Figure 6.16: Distribution of Ψ_L sample size for subgroups (1,0) and (1,1). The top row is when the normal-inverse-gamma prior is used, the middle row is when the Cauchy prior is used, and the bottom row is when the fully Bayesian approach is used with a normal-inverse-gamma prior.

Overall, this simulation demonstrates that the two priors lead to very similar results. We note that the fully Bayesian approach allows weights to reach their true value more quickly, and, as a consequence, lower values of the weighted L-optimal objective function can be achieved.

6.3.3 Comparing minimization, D_A-optimality, weighted D_A-optimality, weighted L-optimality

We now compare the weighted L-optimal design against other common approaches: the classic form of minimization with $p = \frac{2}{3}$ (see Section 2.5), D_A-optimality (see Equation 6.5), and weighted D_A-optimality (Equation 6.6). For the D_A-optimality criterion, the matrix $\boldsymbol{A}$ is given by $\boldsymbol{A} = \left(\boldsymbol{c_1}^\intercal \, \boldsymbol{c_2}^\intercal \, \boldsymbol{c_3}^\intercal\right)$. Since $\boldsymbol{c_4}$ is a linear combination of the vectors $\boldsymbol{c_1}, \boldsymbol{c_2}$, and $\boldsymbol{c_3}$, it cannot be included in $\boldsymbol{A}$ because $\boldsymbol{A}$ needs to have full column rank. We use the Cauchy prior with center zero and scale 2.5 for the weighted L-optimal approach. The settings for this simulation are as in Table 6.2. The weighted L-optimal design uses an exchange algorithm

with the L-optimality criterion with equal weights to construct the initial design with 15 patients. The D_A- and weighted D_A-optimal designs use an exchange algorithm with the D_A-optimality criterion to construct the initial design.

It appears in Figures 6.17, 6.18 and 6.19 that the four approaches lead to similar estimates of $\hat{\boldsymbol{\beta}}$, $\hat{\sigma}$ and similar values of w_r. The variability of the estimates also appears to be very similar.

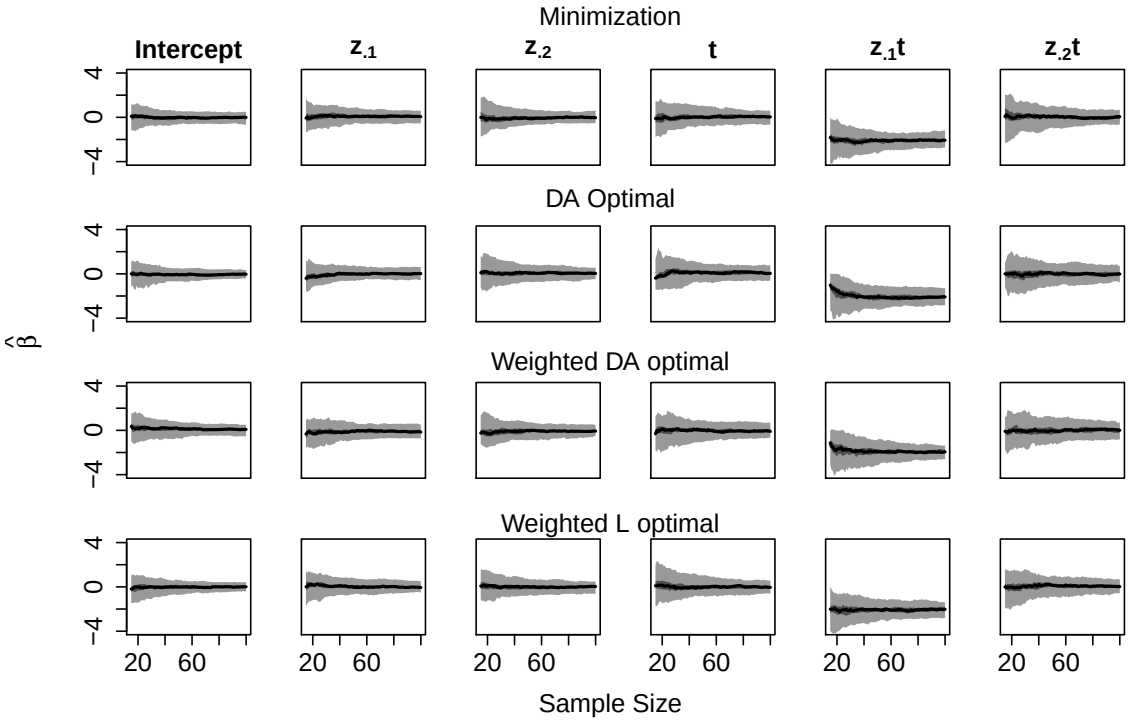

Figure 6.17: Distribution of $\hat{\boldsymbol{\beta}}$ for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

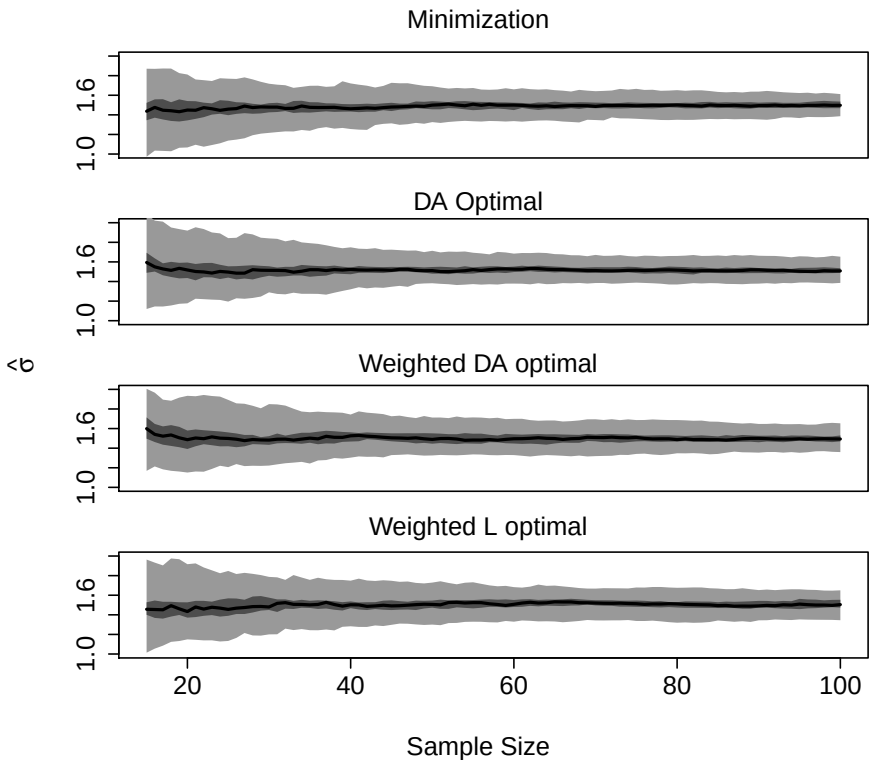

Figure 6.18: Distribution of $\hat{\sigma}$ for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row).

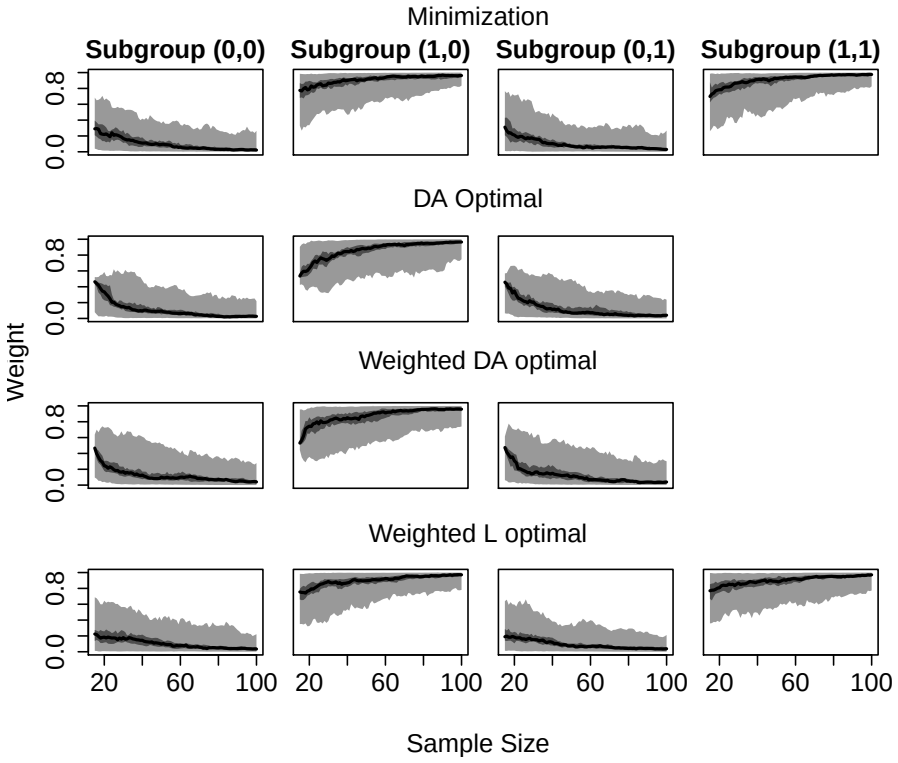

Figure 6.19: Distribution of w_r for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row).

Further, we observe in Figure 6.20 that the distribution of the theoretical value of the power of the hypothesis test for the two subgroups $(1, 0)$ and $(1, 1)$ is extremely similar for all four methods. This is also true for the plots of the empirical value of the power for these subgroups, as well as 1-specificity for the two subgroups $(0, 0)$ and $(0, 1)$, as shown in Figures 6.21 and 6.22.

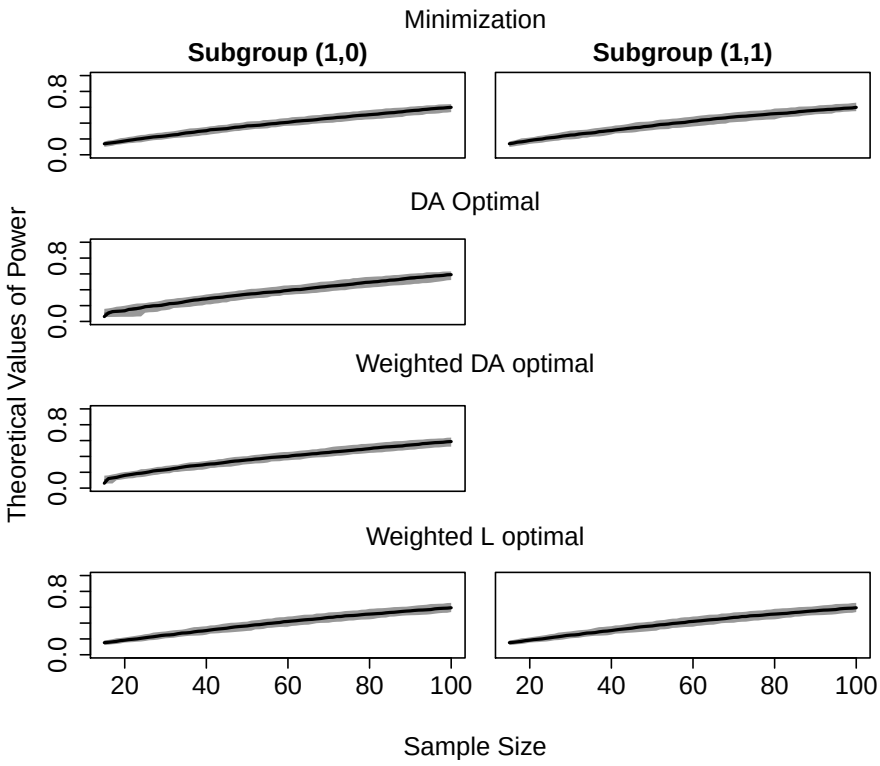

Figure 6.20: Distribution of the theoretical value of the power of the hypothesis test for subgroups (1,0) and (1,1) for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row).

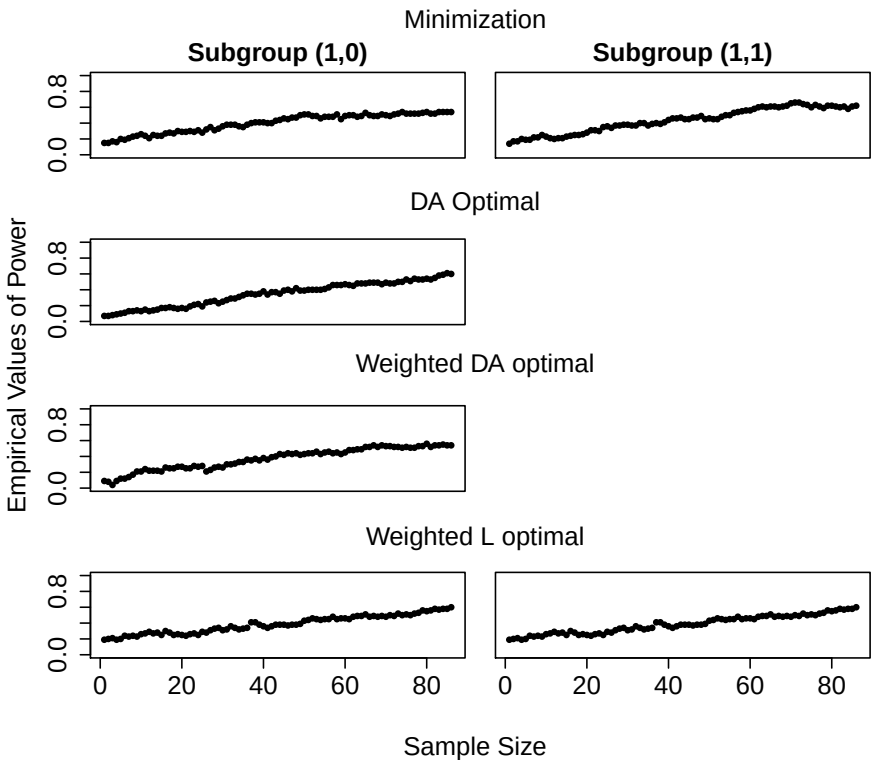

Figure 6.21: The empirical value of the power of the hypothesis test for subgroups (1,0) and (1,1) for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row).

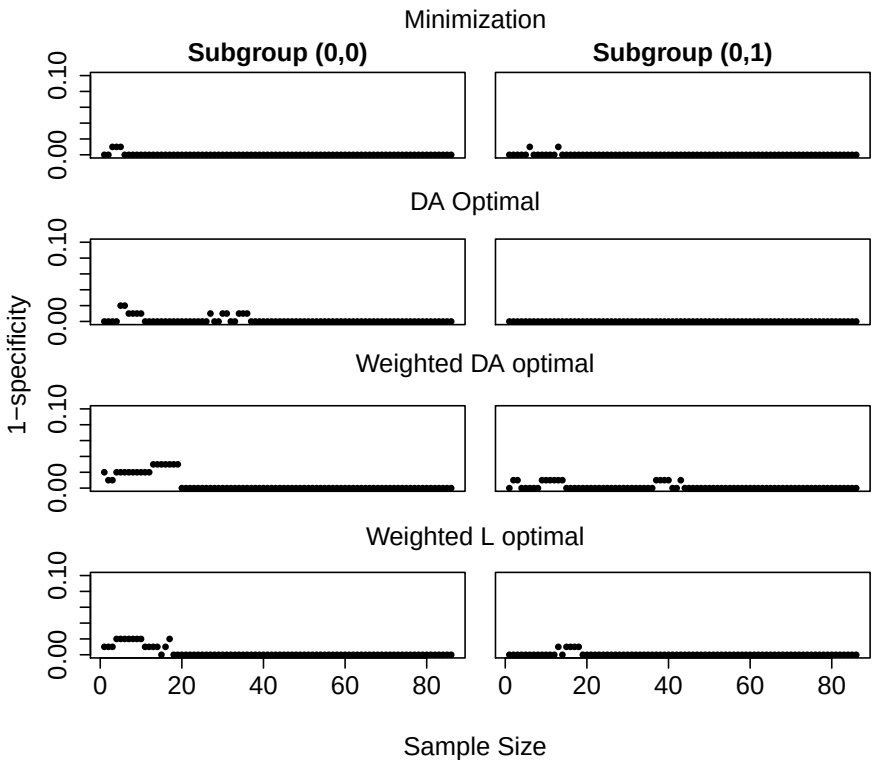

Figure 6.22: 1-specificity for the hypothesis test for subgroups (1,0) and (1,1) for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row).

We observe in Figure 6.23 that, overall, the treatment balance is better for the weighted L-optimal design rather than the other designs. This is particularly true for small sample sizes. For large sample sizes, surprisingly, it seems that minimization leads to better treatment balance in the four subgroups.

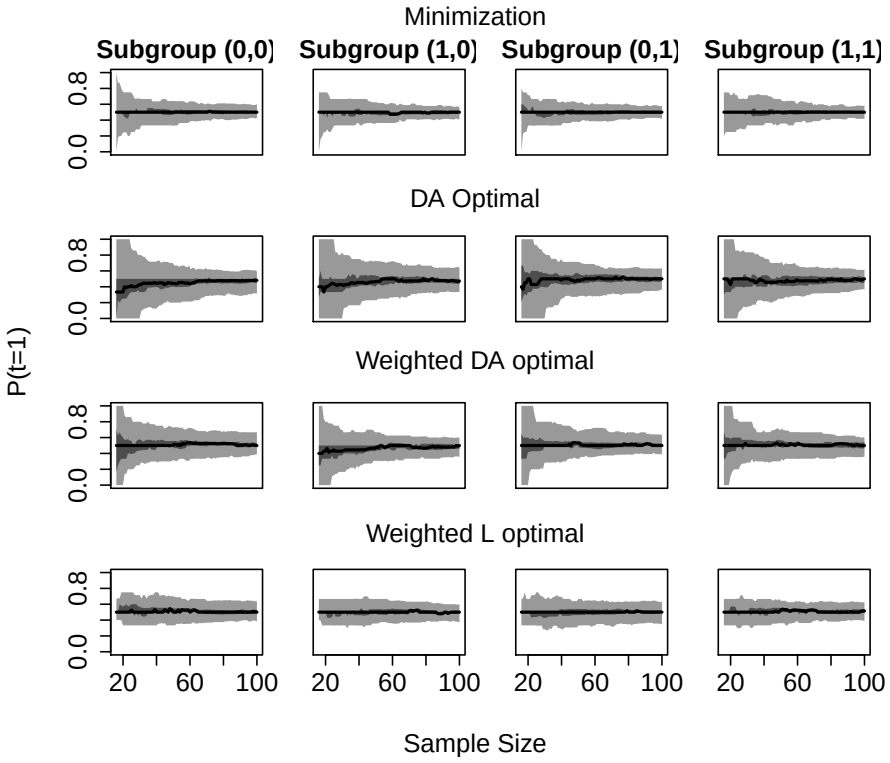

Figure 6.23: Distribution of the proportion of new treatment in each subgroup for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row)

In Figure 6.24 we consider three measures of optimality: D_A-optimality, weighted D_A-optimality and weighted L-optimality. We observe that the three optimal design-based designs have very similar performance. It is slightly surprising that the weighted L-optimal design performs slightly less well than the D_A-optimal and the weighted D_A-optimal design under weighted L-optimality; further, the weighted L-optimal design performs slightly better than the D_A-optimal and the weighted D_A-optimal designs under weighted D-optimality. Perhaps unsurprisingly, minimization performs less well than the three optimal design-based approaches, with much higher values of the objective function initially and also greater variability.

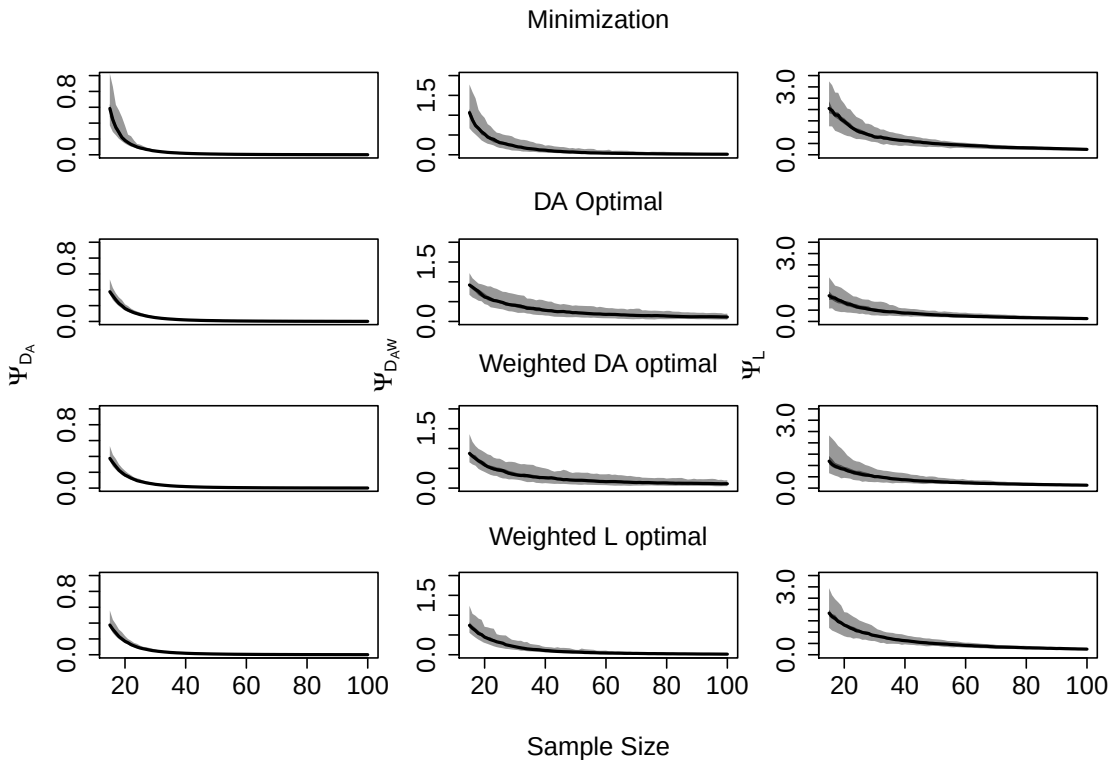

Figure 6.24: Distribution of optimality for minimization (top row), D_A-optimal design (second row), weighted D_A-optimal design (third row) and weighted L-optimal design (bottom row)

Overall, we observe that under the settings of this simulation, the four approaches lead to similar parameter estimates and we note that the optimal design based approaches lead to more efficient designs than minimization.

6.4 Conclusion

In this chapter, we considered how to design Phase III clinical trials using a sequential optimal design approach where the aim is to estimate specific linear combinations of parameters within the model. A weighted L-optimal design is used to identify effective treatment-biomarker combinations. Since separation is an issue for early stages of the design, we examined two options to deal with this: firstly, by simply putting prior distributions on the model parameters and secondly, by using a fully Bayesian approach where the posterior distribution of the parameters is used to define the weights. Simulations demonstrated that the fully Bayesian approach has a slight advantage in that weights appear to reach their true value more quickly.

We ran simulations to compare the weighted L-optimal design against minimization, D_A- and weighted D_A-optimal designs. We found that all four methods lead to similar parameter estimates, but the optimal design based approaches lead to more efficient designs than minimization. In the next chapter, we investigate whether the nonmyopic approach is beneficial in this context of personalized medicine.

Chapter 7

Nonmyopic approaches in personalized medicine

In Chapter 6, we introduced the notion of personalized medicine and demonstrated methodology for designing experiments where interest lies in estimating important interactions between treatments and biomarkers with low variance. In this chapter, we show how the nonmyopic approach introduced in Chapter 4 and the pseudo-nonmyopic approach described in Chapter 5 can be applied in the context of personalized medicine. In our simulations, we focus on the logistic model case; for the nonmyopic approach, we look at the case of a binary treatment, and in the pseudo-nonmyopic approach, we consider the case of a continuous treatment.

7.1 Nonmyopic approach for personalized medicine

As described in Section 4.2.2, if the response is binary rather than continuous, the information matrix is $\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X}$ where $\boldsymbol{W}$ is a matrix contingent on the parameters of the model. For a binary GLM with the logistic link function, $\boldsymbol{W}$ is a diagonal matrix with the ith entry equal to $\hat{\pi}_i(1-\hat{\pi}_i)$ where $\hat{\pi}_i = \frac{\exp \eta_i}{1+\exp \eta_i}$, where η_i is the linear predictor. The weights for the optimality criterion are generated as for the case of a continuous response in Equation (6.3); we are using an asymptotic result on the approximate normality of $\boldsymbol{c}_r^\intercal \hat{\boldsymbol{\beta}}$. Given a design matrix $\boldsymbol{X}$, vectors $\boldsymbol{c}_r$, $1 \leq r \leq R$, and values for the parameters $\boldsymbol{\beta}$, we have that the weighted L-optimal objective function in the case of a binary response is

$$\Psi_L(\boldsymbol{X}, \boldsymbol{c}_r, \boldsymbol{\beta}) = \sum_{r=1}^{R} w_r \boldsymbol{c}_r^\intercal (\boldsymbol{X}^\intercal \boldsymbol{W} \boldsymbol{X})^{-1} \boldsymbol{c}_r. \tag{7.1}$$

The relative weighted L-efficiency of a design $\boldsymbol{X}$ relative to another design $\boldsymbol{X}^*$, with

parameter values $\boldsymbol{\beta}$ in the case of a binary response, is

$$\text{Eff}_L = \frac{\Psi_L\left(\boldsymbol{X}^*, \boldsymbol{c}_r, \boldsymbol{\beta}\right)}{\Psi_L\left(\boldsymbol{X}, \boldsymbol{c}_r, \boldsymbol{\beta}\right)}. \tag{7.2}$$

We note that, in our simulations, we use the true parameter values to compute Ψ_L in order to evaluate Eff_L.

The sequential nonmyopic approach for personalized medicine is equivalent to that described in Section 4.2.4, where the optimality criterion Ψ is the weighted L-optimality criterion. In calculating $\Psi_L(t_i \mid \boldsymbol{z}_i, \boldsymbol{t}_{i-1}, \boldsymbol{y}_{i-1})$, there is an extra step of calculating the weights $w_{r_{i-1}}$ which depend on $\boldsymbol{y}_{i-1}$ through $\hat{\boldsymbol{\beta}}$.

In the current literature, there is some work where dynamic programming has been used in designing personalized trials. Zhang et al. (2016) considered a problem where treatments are allocated sequentially using optimal Bayesian adaptive randomization for the first m patients. Then, based on the results of this trial, future patients are given the treatment that is deemed best for their biomarker subgroup. Backwards induction is used to solve this problem. This approach is nonmyopic in the sense that the design of the first m patients is constructed with the anticipation that future patients after the trial will be given treatments based on the results. So far, no work has been done on designing personalized clinical trials where each step of the construction of sequential trial with n patients involves a nonmyopic optimization.

7.1.1 Simulations

We run a simulation to compare the performance of the nonmyopic approach with horizon 1 and horizon 3 against the myopic approach. We wish to see whether there is an advantage to using the nonmyopic algorithm over the myopic algorithm. We consider an example where two biomarkers are independently generated from Bernoulli(0.5) distributions. We use the Cauchy prior which is recommended by Gelman et al. (2008) as it effectively alleviates problems related to separation when the response is binary. We assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = -2z_{i,1}t_i$. The settings of the simulation are given in Table 7.1. We assume the true distribution of the biomarkers is known in the nonmyopic case. Our simulation has the following structure:

I (a) 100 patients are assumed and their covariates are generated from a specified distribution.

(b) 100 deviates from a Unif$(0, 1)$ distribution are generated for the response.

(c) An initial design with 15 units is constructed using the exchange algorithm with weighted L-optimality as the objective function, given by Equation (7.1).

(d) The following three sequential designs are constructed using the covariates, random deviates for the responses, and initial design:

- A myopic weighted L-optimal design.
- A nonmyopic weighted L-optimal design with horizon $N = 1$.
- A nonmyopic weighted L-optimal design with horizon $N = 3$.

(e) Designs are evaluated using the performance measure Ψ_L, given by Equation (7.1), at each sample size between 15 and 100, inclusive. The true values of the parameters are used to calculate Ψ_L.

II (a)-(e) above is repeated 20 times to obtain a distribution of the performance measure for each sample size.

Table 7.1: Settings for a simulation comparing myopic and nonmyopic approaches for the weighted L-optimal approach

Biomarkers	$z_{\cdot 1} \sim$ Bernoulli(0.3), $z_{\cdot 2} \sim$ Bernoulli(0.5)
True $\boldsymbol{\beta}$	all zero except for the entry corresponding to $\boldsymbol{z}_{\cdot 1}\boldsymbol{t}$, which is -2
τ_r	-1 for all r
κ	0
Initial sample size	15
Number of simulations	100
Design Criterion	Weighted L-optimality
Prior	Cauchy with center 0 and scale 2.5

In Figures 7.1 and 7.2, we compare the estimates across the three methods for the parameters and the weights, respectively. We observe that there is very little difference in the distributions of the estimates of $\boldsymbol{\beta}$. The weights for the subgroups where $z_{i,1} = 1$ go to one, and the weights for the subgroups where $z_{i,1} = 0$ go to zero, as expected. There appears to be little difference between the weights for the nonmyopic and myopic approaches.

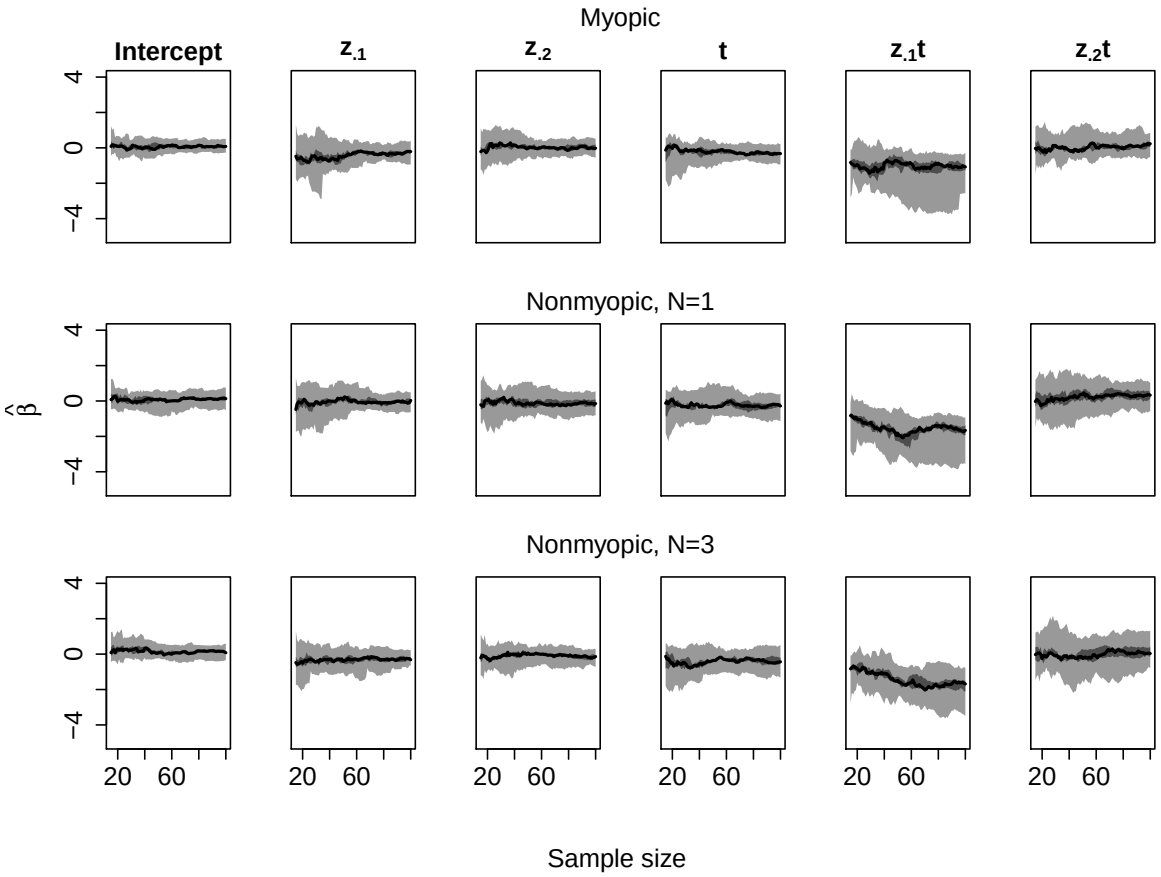

Figure 7.1: Distributions of $\hat{\boldsymbol{\beta}}$ for weighted L-optimal designs for the logistic model under the Cauchy prior are plotted against sample size. We show the myopic approach (top row), nonmyopic approach with horizon 1 (middle row) and nonmyopic approach with horizon 3 (bottom row). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

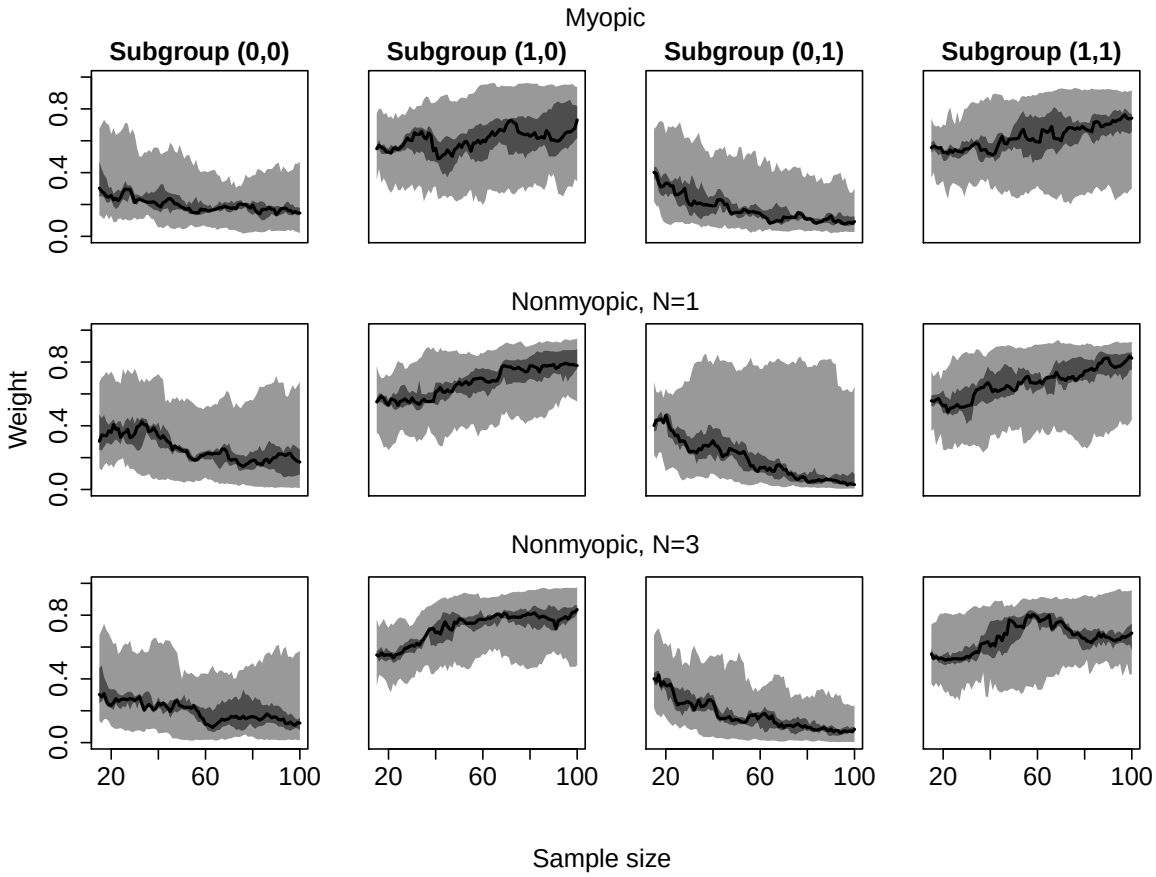

Figure 7.2: Distribution of w_r vs sample size under the Cauchy prior for the myopic approach (top row), nonmyopic approach with horizon 1 (middle row) and nonmyopic approach with horizon 2 (bottom row).

We plot the distribution of Ψ_L against sample size for the myopic and nonmyopic approaches in Figure 7.3. We observe that there appears to be very little difference between the plots of Ψ_L in the three approaches, and further, the plots of relative efficiency suggest that they perform similarly.

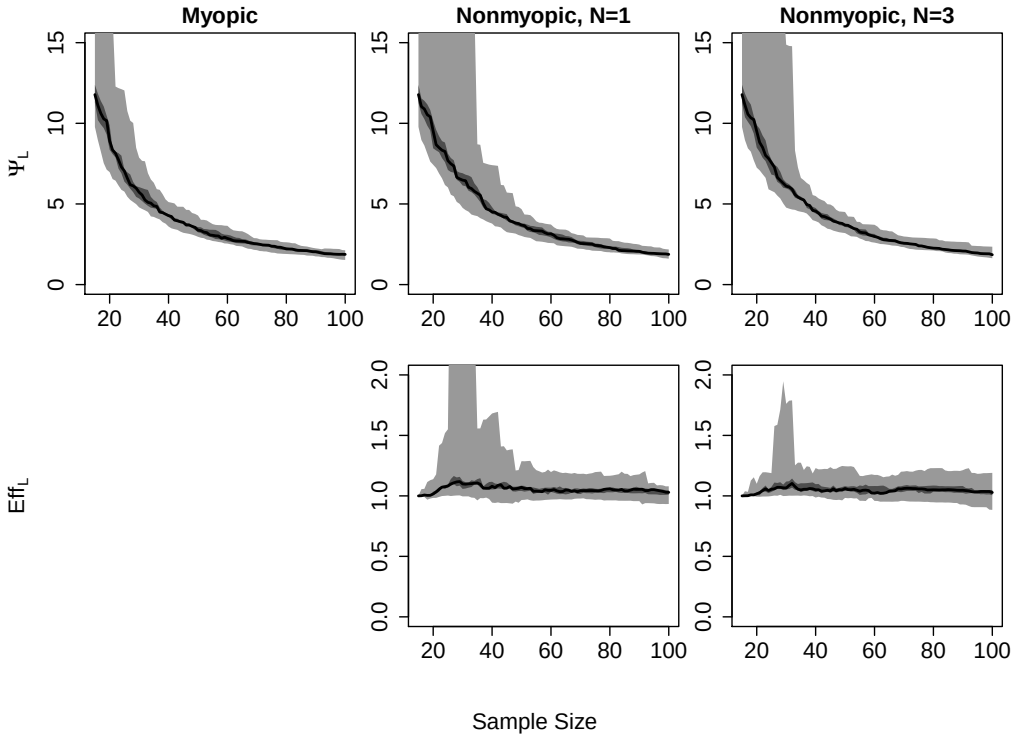

Figure 7.3: In the top row, we have the distribution of Ψ_L vs sample size under the Cauchy prior for the myopic approach (left column), nonmyopic approach with horizon 1 (middle column) and nonmyopic approach with horizon 2 (right column). In the bottom row, we have the relative efficiencies of the nonmyopic approach with horizon 1 against the myopic approach (left) and the nonmyopic approach with horizon 3 against the myopic approach (right).

We repeat this simulation with the same settings as in Table 7.1, except we use a normal-inverse-gamma prior with $\boldsymbol{m} = \mathbf{0}$, $\boldsymbol{V} = \frac{1}{100}\boldsymbol{I}_6$, $a = 2$, and $d = 2$ is used instead of a Cauchy prior. The results are shown in Appendix G. In both the simulation with the Cauchy prior and the normal-inverse-gamma prior, we observe that there is no benefit to using the nonmyopic approach over the myopic approach, even when we know the true distribution of the biomarkers. We note that the plots show greater variability when the normal-inverse-gamma prior is used.

7.2 Pseudo-nonmyopic approach for personalized medicine

We now turn our attention to the pseudo-nonmyopic approach for personalized medicine. The algorithm for the logistic model case is as described in Section 6 where the optimality criterion Ψ is the weighted L-optimality criterion and there is an extra step of updating

the weights after refitting the model in the algorithm. In the case of a binary treatment, simulations revealed that, much like for the nonmyopic approach, the pseudo-nonmyopic approach performs very similarly to the myopic approach. This is true even in cases where we have dynamic covariates. We therefore focus on results in the continuous treatment case in the following section.

7.2.1 Simulations

Example 1

In this simulation, we compare the performance of the pseudo-nonmyopic approach with $M = 20$, the myopic approach and a non-sequential design where the true values of the parameters are assumed to be known. This allows us to see whether there is any relative advantage of the pseudo-nonmyopic approach over the myopic approach, and we can also compare these two types of designs against a design where we have perfect knowledge of all covariate values and parameter values. We consider an example where two biomarkers are independently generated from a Bernoulli(0.5) and Bernoulli(0.7) distributions, respectively. We assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = -2z_{i,1}t_i$. The settings of the simulation are given in Table 7.2. We assume true distribution of the biomarkers is known in the pseudo-nonmyopic case, and the structure of the simulation is analogous to that shown in Section 7.1.1.

Table 7.2: Comparing myopic and pseudo-nonmyopic approaches for the weighted L-optimal approach

Biomarkers	$z_{.1}$ ~Bernoulli(0.5), $z_{.2}$ ~Bernoulli(0.7)
True $\boldsymbol{\beta}$	all zero except for the entry corresponding to $\boldsymbol{z}_{.1}\boldsymbol{t}$, which is -2
True σ	1.5
τ_r	-1 for all r
κ	0
Initial sample size	5
Total sample size	50
M	20
Number of simulations	100

In Figure 7.4, we observe that the myopic and pseudo-nonmyopic approaches lead to estimates of the parameter values that are close to the true value. We note that there is more variability in the estimates produced by the pseudo-nonmyopic approach for the intercept, the coefficient of $\boldsymbol{t}$ and the coefficient of $\boldsymbol{z}_{.1}\boldsymbol{t}$.

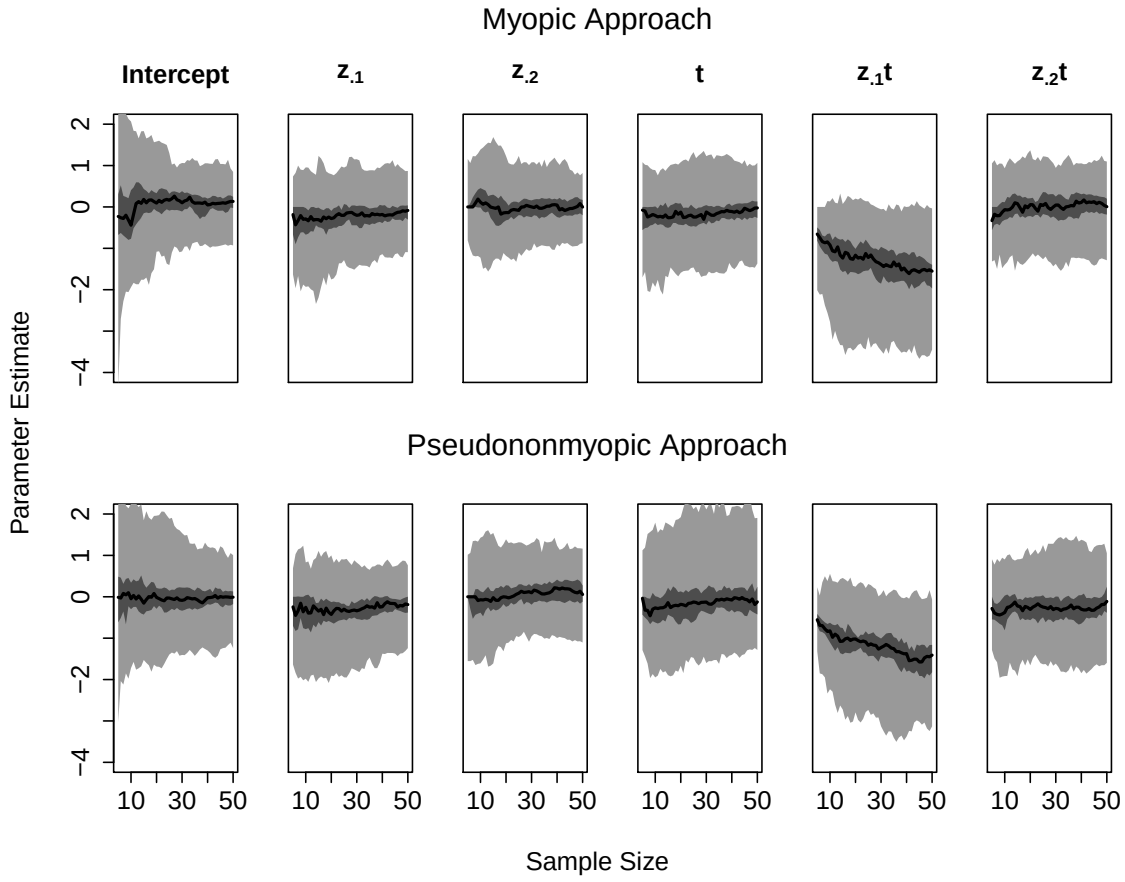

Figure 7.4: Distributions of $\hat{\beta}$ for designs for the logistic model with a continuous treatment are plotted against sample size for the myopic and pseudo-nonmyopic approaches. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

We observe in Figure 7.5 that the myopic approach leads estimates of the weights that are closer, on average, to the true values than the pseudo-nonmyopic approach.

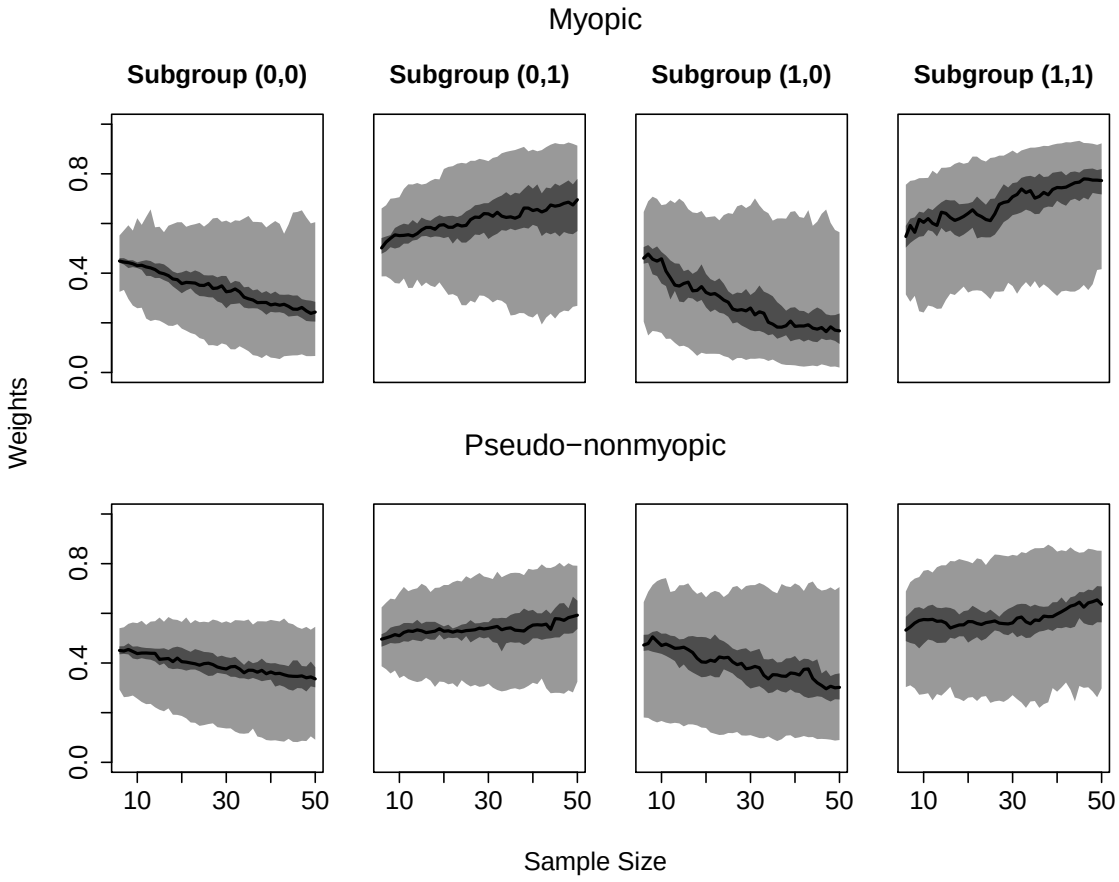

Figure 7.5: Distributions of w_r for designs for the logistic model with a continuous treatment are plotted against sample size for the myopic and pseudo-nonmyopic approaches.

In Figure 7.6, we display boxplots of the allocated treatment at each sample size for the three allocation methods: the nonsequential design, the myopic approach and the pseudo-nonmyopic approach. For the nonsequential design, the boxplots have a peculiar appearance because all treatments lie at the extremes of one or zero. Proportionately more patients are given the treatment value of one. For the myopic approach, all allocated treatments are again either zero or one, but there appears to be more zeros than for the nonsequential design. The fact that the assigned treatments are at the extremes is likely to be due to the fact that all but one of the parameters are set to zero (Atkinson and Woods, 2015). In contrast, for the pseudo-nonmyopic approach, apart from the initial design in which treatments are zero or one, we find that most treatment values lie between 0.4 and 0.6.

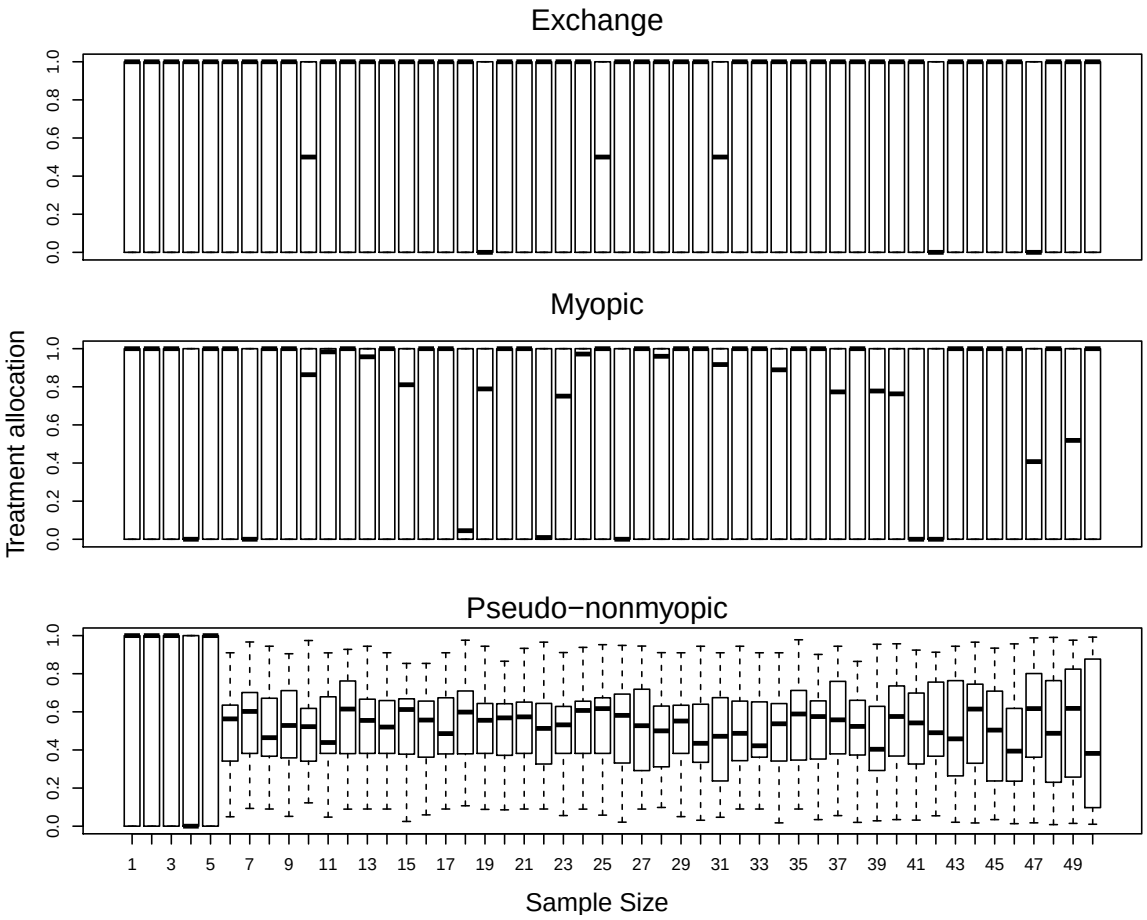

Figure 7.6: Boxplots display the distribution of the allocated treatments at each sample size for the exchange, myopic and pseudo-nonmyopic approaches for the logistic model case with a continuous treatment.

In Figure 7.7, we compare the distributions of the weighted L-optimality of the three approaches in the top row. The nonsequential designs have more variable values of weighted L-optimality compared to the myopic approach initially, but they later appear to have similar performance. We also see that the myopic approach has much lower variability and has a sharper downward slope than the pseudo-nonmyopic approach; it is better able to reach more efficient designs. The bottom row shows the relative efficiency of the nonsequential designs against the myopic designs, as well the relative optimality of the pseudo-nonmyopic designs against the myopic desigsn. We observe that the myopic designs and the nonsequential designs have similar values of efficiency after sample size of 10. When sample size is greater than 10, the myopic approach is two to three times more efficient than the pseudo-nonmyopic approach.

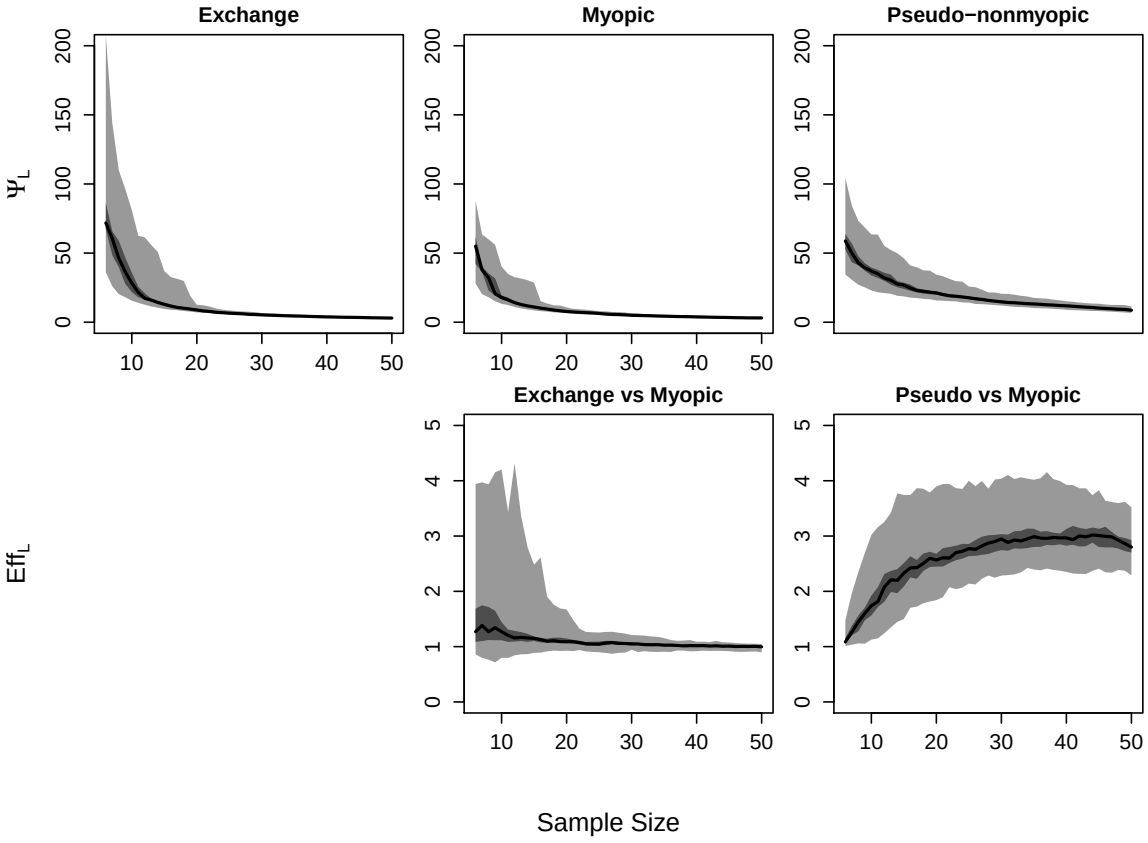

Figure 7.7: Distribution of weighted L-optimality is plotted against sample size for the exchange algorithm (top left), myopic approach (top middle) and the pseudo-nonmyopic approach (top right) for the logistic model case with a continuous treatment. The relative weighted L-efficiency of the exchange algorithm vs the myopic approach is plotted against sample size on the bottom left, and the pseudo-nonmyopic approach vs the myopic approach is on the bottom right.

Example 2

We now consider a simulation study where all settings are the same as the previous one, except for two changes. Firstly, we assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = -3z_{i,1} + 5z_{i,2} + t_i - 20z_{i,1}t_i + 3z_{i,2}t_i$. Secondly, $z_{\cdot 1}$ has a Bernoulli(0.1) distribution and $z_{\cdot 2}$ has a Bernoulli(0.5) distribution. We purposefully selected an example with a large interaction effect and chose the biomarker associated with that interaction to be rare. The simulation settings are shown in Table 7.3.

Table 7.3: Comparing myopic and pseudo-nonmyopic approaches for the weighted L-optimal approach with a larger interaction

Biomarkers	$z_{.1}$ ∼Bernoulli(0.1), $z_{.2}$ ∼Bernoulli(0.5)
True $\boldsymbol{\beta}$	$(0, -3, 5, 1, -20, 3)$
True σ	1.5
τ_r	-1 for all r
κ	0
Initial sample size	5
Total sample size	50
M	20
Number of simulations	100

We see very clearly in Figure 7.8 that the median of the estimates of the coefficient for $\boldsymbol{z}_{.1}\boldsymbol{t}$ is much closer to the true value for the pseudo-nonmyopic approach than the myopic approach. In both approaches, the estimates are shrunk towards zero, but it is more extreme for the myopic approach. We observe, again, more variability in the estimates for the pseudo-nonmyopic approach.

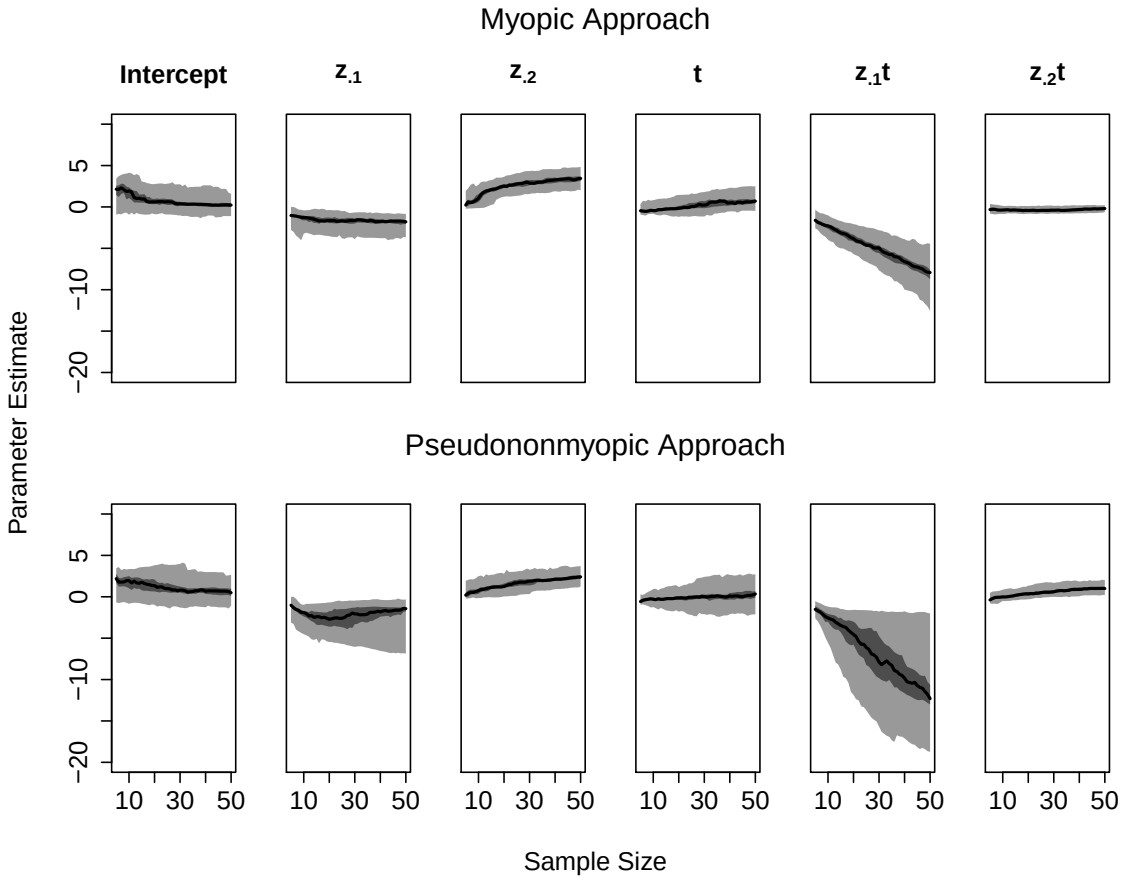

Figure 7.8: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model with a continuous treatment and large interaction and rare biomarker are plotted against sample size for the myopic and pseudo-nonmyopic approaches. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

In Figure 7.9, we see much greater variability in the estimates of the weights for the pseudo-nonmyopic approach than the myopic approach; this is particularly true for the subgroups $(1, 0)$ and $(1, 1)$. This is likely to be due to the fact that the covariate value $z_{i,1} = 1$ is rare.

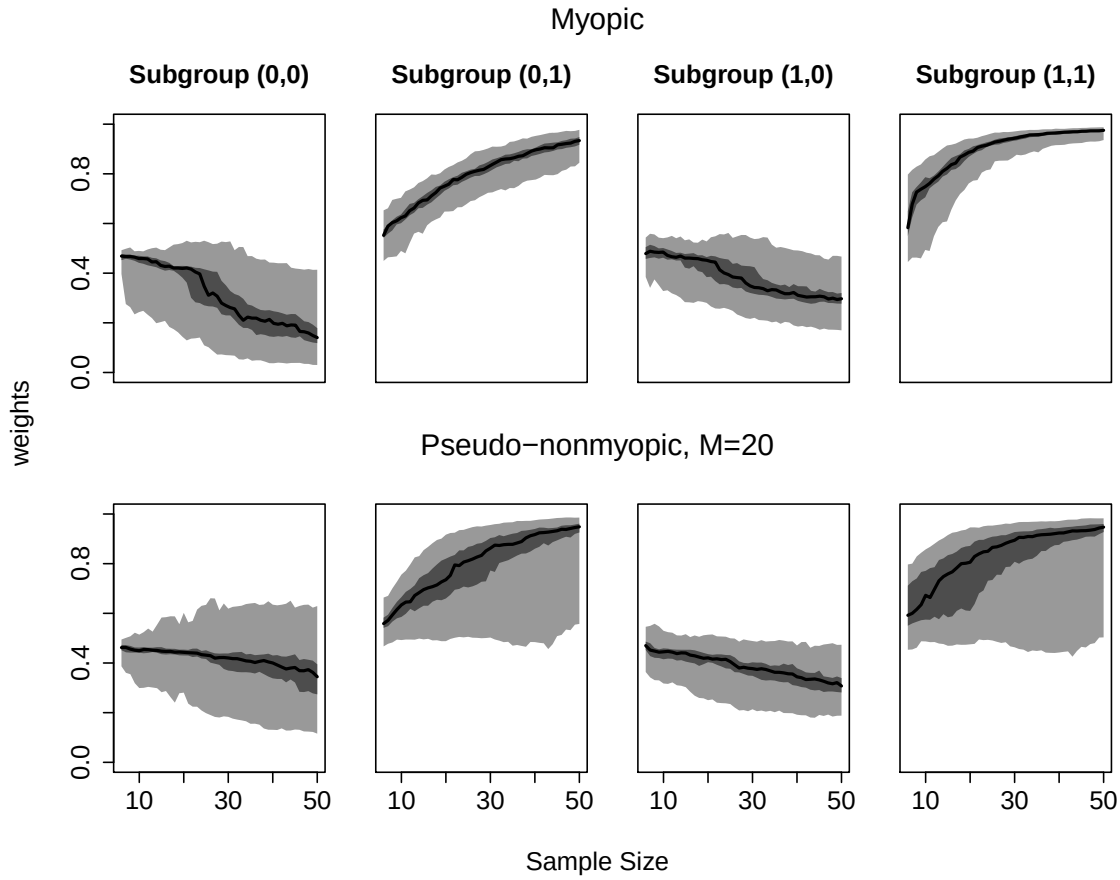

Figure 7.9: Distributions of w_r for designs for the logistic model with a continuous treatment and large interaction and rare biomarker are plotted against sample size for the myopic and pseudo-nonmyopic approaches.

In Figure 7.10, we see that the allocated treatments in the nonsequential design are mostly centered around 0.3, which is quite different from what we observed in the previous example in Figure 7.10; this is due to the values of the true parameters being far from zero in this example. For the myopic approach, we observe again that most treatments are zeros and ones, whereas for the pseudo-nonmyopic approach, the treatments can lie between the extremes and most are centered around 0.5.

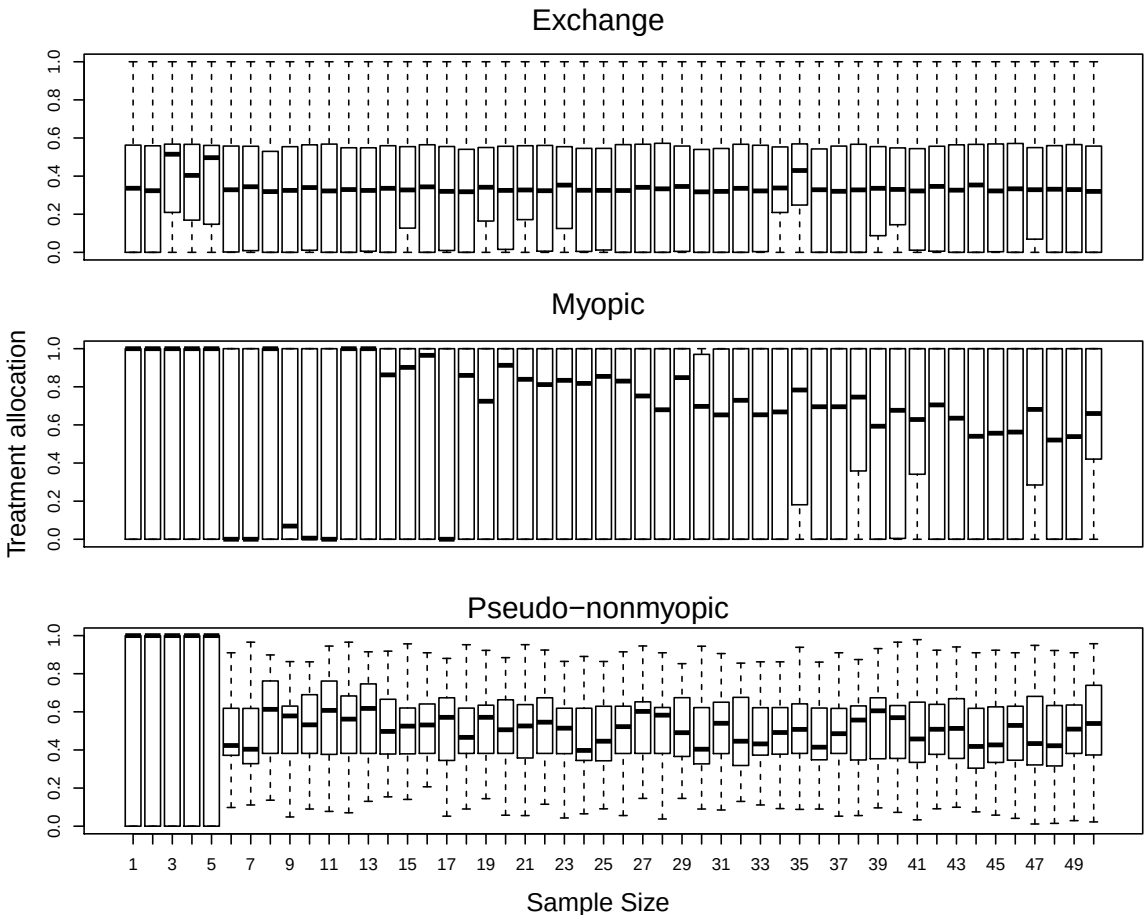

Figure 7.10: Boxplots display the distribution of the allocated treatments at each sample size for exchange, myopic and pseudo-nonmyopic approaches for the logistic model with a continuous treatment and large interaction and rare biomarker.

In Figure 7.11, we observe that the designs constructed using the exchange algorithm are the most efficient. Since we assume the true values of the parameters in the exchange algorithm, and for the other approaches, the estimates are biased towards zero, we expect the nonsequential designs to be most efficient. We observe also that the pseudo-nonmyopic approach is more efficient than the myopic approach for most of the distribution, for all values of sample size considered in this simulation.

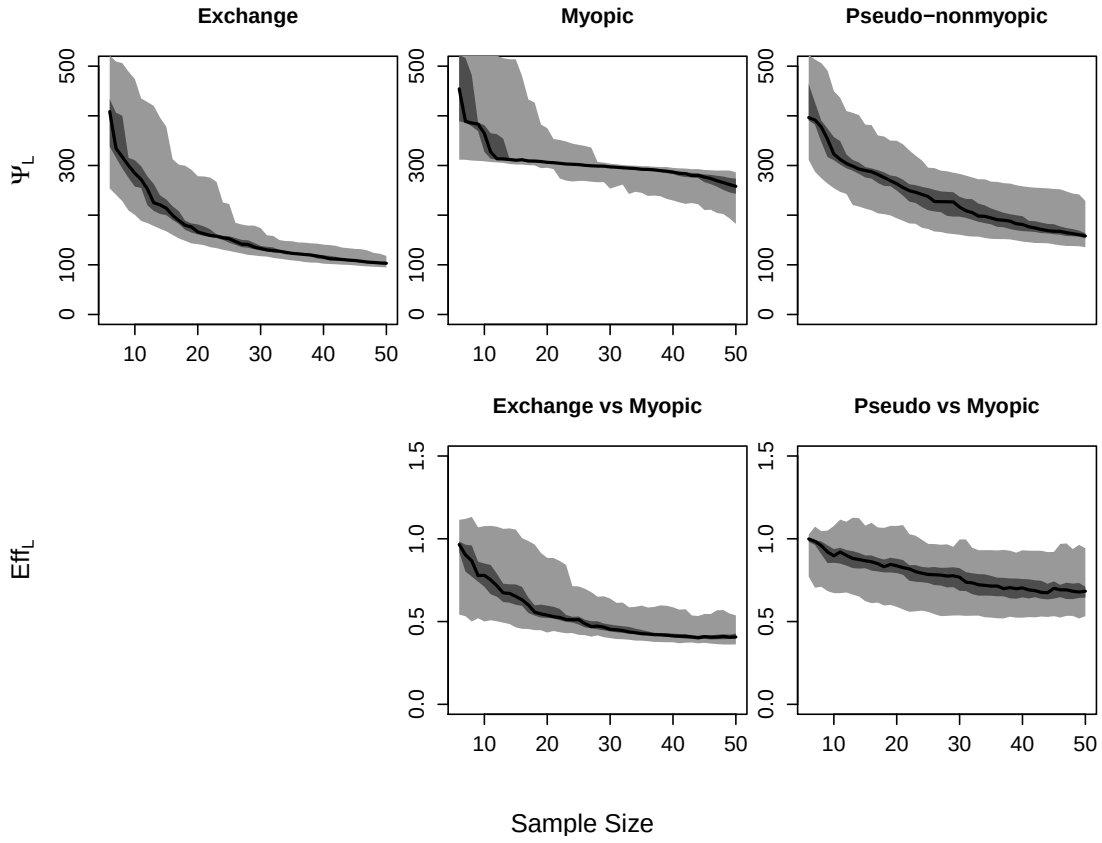

Figure 7.11: Distribution of weighted L-optimality is plotted against sample size for the exchange algorithm (top left) myopic approach (top middle) and the pseudo-nonmyopic approach (top right) for the logistic model case with a continuous treatment and large interaction and rare biomarker. The relative weighted L-efficiency of the exchange algorithm against the myopic approach is plotted against sample size on the bottom left and we have the pseudo-nonmyopic approach vs the myopic approach on the bottom right.

We observed in this example that, in the case where we have a large interaction and the biomarker associated with that interaction is rare, the pseudo-nonmyopic approach offers benefit compared to the myopic approach. It is better able to estimate the large interaction and the resulting designs are more efficient. In Appendix I, we look at the same example as above, but the biomarkers are not rare; $\boldsymbol{z}_1$ has a Bernoulli(0.5) distribution and $\boldsymbol{z}_2$ has a Bernoulli(0.7) distribution; the pseudo-nonmyopic approach still offers a slight benefit over the myopic approach, but the difference is not as pronounced.

7.3 Conclusions

In this chapter, we looked into how the three model-based approaches to treatment allocation considered in this thesis compare in the context of personalized medicine: the myopic approach, the non-myopic approach and the pseudo-nonmyopic approach. Our simulation studies show that the nonmyopic approach performs very similarly to the myopic approach for the binary response and binary treatment case. For the pseudo-nonmyopic approach, we found that:

- the myopic approach is more efficient than the pseudo-nonmyopic approach in general for the binary treatment case.
- the myopic approach is more efficient than the pseudo-nonmyopic approach for the continuous treatment case if interactions are not large.
- the pseudo-nonmyopic approach is more efficient than the myopic approach for the continuous treatment case if there is a large interaction term. This is even more pronounced when the biomarker associated with the large interaction is rare.

We observe in the second simulation study of the pseudo-nonmyopic approach in this chapter that the parameter estimates are biased towards zero for both the myopic and pseudo-nonmyopic approaches. In order to investigate this further, we did some preliminary work on designs which are optimal for mean-squared error. However, when the true coefficients are unknown, such as in a real experiment, it is difficult to obtain an accurate estimate of the bias matrix. This is an area which requires further investigation and may be appropriate for future work.

Chapter 8

Vignette of R package

This chapter is a vignette of the R package created as part of this thesis, which can be accessed here: https://github.com/mst1g15/biasedcoin. The package contains functions which implement the sequential treatment allocation algorithms described in the thesis. We provide simple examples with small sample sizes to illustrate how the functions work.

8.1 Myopic approaches

To illustrate the functions in the package to implement myopic approaches to treatment allocation, we begin by generating binary covariate values for 10 patients, and storing them in a dataframe. We then create a design where the treatments are randomly allocated using `rand()`:

```
covar <- as.data.frame(sample(c(-1,1), 10, replace=T))
colnames(covar) <- "covar"
rand.D <- rand(covar)
rand.D
#>    1 covar tmt
#> 1  1     1   1
#> 2  1    -1  -1
#> 3  1     1  -1
#> 4  1     1   1
#> 5  1    -1   1
#> 6  1     1   1
#> 7  1    -1  -1
#> 8  1    -1   1
#> 9  1    -1   1
```

```
#> 10 1    -1   1
```

We can also create a design based on Efron's biased coin with `efron()`, or the classic form of minimization with `min.classic()`. We set $p = \frac{2}{3}$:

```
efron.D <- efron(covar, p=2/3)
min.D <- min.classic(covar, p=2/3)
```

For a single continuous convariate, we generate designs using minimization with the Efron, K-S or Max.imb imbalance measures:

```
cont.covar <- as.data.frame(runif(10))
efron.cont.D <- min.gen(cont.covar, imb="efron", p=2/3)
ks.D <- min.gen(cont.covar, imb="ks", p=2/3)
maximb.D <- min.gen(cont.covar, imb="max.imb", p=2/3)
maximb.D
#>    1      covar tmt
#> 1  1 0.57712495  -1
#> 2  1 0.17155262   1
#> 3  1 0.04256604   1
#> 4  1 0.13985023  -1
#> 5  1 0.97035001   1
#> 6  1 0.06374479   1
#> 7  1 0.03342408   1
#> 8  1 0.68122202  -1
#> 9  1 0.75750067   1
#> 10 1 0.88481663  -1
```

We now look at functions that construct optimal design-based designs. Here, we construct a D-optimal design using the exchange algorithm. We set the number of repetitions from different starting designs `k` to 2 and we do not assume any covariate-treatment interactions (`int=NULL`).

```
coord.D <- coordex(covar, k=2, lossfunc=calc.D, int=NULL)
```

The function `calc.D` is used here to construct a D-optimal design here, but we can also use `calc.DA` for an D_A-optimal design, where an extra argument `A` is included to indicate the subsets of parameters of interest. The functions `calc.A` and `calc.G` are used for A- and G-optimal designs, respectively.

We can also construct an approximately D-optimal designs sequentially with `atkins()`. We choose not to begin with a pre-constructed initial design by setting the number of patients in the initial design `init` to 1 and setting `same.start=NULL`. Having a pre-constructed

initial design can be useful when comparing different designs. We allocate treatments with a probability based on a ratio of values from objective functions such as in Equation (2.37) (`stoc=T`), rather than deterministically (`stoc=F`):

```
seq.D <- atkins(covar, lossfunc=calc.D, int=NULL, init=1,
                same.start=NULL, stoc=T)
seq.D
#> $D
#>    1 covar tmt
#> 1  1     1   1
#> 2  1    -1   1
#> 3  1     1  -1
#> 4  1     1  -1
#> 5  1    -1   1
#> 6  1     1  -1
#> 7  1    -1  -1
#> 8  1    -1  -1
#> 9  1    -1   1
#> 10 1    -1  -1
#>
#> $opt
#> [1] 1.249991e+04 6.249906e-02 3.124961e-02 1.562484e-02 1.041657e-02
#> [6] 3.676452e-03 2.232134e-03 1.602558e-03 1.157404e-03
#>
#> $loss.m
#> [1] 1.249991e+04 6.249906e-02 3.124961e-02 8.928514e-03 1.041657e-02
#> [6] 3.676452e-03 2.232134e-03 1.602558e-03 1.157404e-03
#>
#> $loss.p
#> [1] 1.249991e+04 6.249965e+03 3.124961e-02 1.562484e-02 7.812441e-03
#> [6] 6.944387e-03 2.604155e-03 1.602558e-03 1.249996e-03
```

The output of `atkins()` has more than just the design matrix; the value of the objective function for each sample size is given, as well as the value of the objective function evaluated when $t_i = 1$ and $t_i = -1$ at each point in the design.

We can evaluate this design based on other optimality criteria:

```
calc.D(rand.D)
#> [1] 0.001249996
calc.DA(rand.D, A=t(c(0, 0, 1)))
```

```
#> [1] 0.1199998
calc.A(rand.D)
#> [1] 0.3499995
calc.G(rand.D)
#> [1] 0.5199991
```

8.2 Nonmyopic approach

8.2.1 Linear model

We now show how to implement the non-myopic algorithm for the linear model case. We use the same set of binary covariates as above. We set the number of patients in the initial design `init=1`. The argument `z.probs` is the assumed distribution of the covariate:

- If `z.probs` is a scalar, this means that there is one static covariate and the scalar is the probability that the covariate takes value 1.
- If `z.probs` is a vector, this means that there is one dynamic covariate where the ith element is the probability that the ith value of the covariate is equal to 1.
- If `z.probs` is a matrix, this means that there is two or more dynamic covariates. The (i, j)th element is the probability that the ith value of covariate j is equal to 1.

The argument `k` is the number of repetitions from different starting designs for constructing the initial design, and `N` is set to an integer value for the horizon. An optimality criterion is specified in `lossfunc`; we choose D-optimality in this example. Similarly to the myopic algorithms, the treatments can be allocated deterministically (`stoc=F`) or with a stochastic element (`stoc=T`). The output is the design matrix and the optimality at each point of the design between $init$ and n:

```
z.probs=0.5
linear.nonmyop(covar, init=1, z.probs=z.probs, k=NULL,
               N=1, int=NULL, lossfunc=calc.D, stoc=F)
#> $D
#>   intercept covar tmt
#> 1         1     1  -1
#> 2         1    -1   1
#> 3         1     1   1
#> 4         1     1  -1
#> 5         1    -1  -1
#> 6         1     1   1
```

```
#> 7         1    -1  -1
#> 8         1    -1   1
#> 9         1    -1  -1
#>
#> $opt
#> [1] 1.249991e+04 6.249906e-02 3.124961e-02 8.928514e-03 5.208305e-03
#> [6] 3.124986e-03 1.953118e-03 1.420450e-03
```

If the distribution of the covariate is unknown, an empirical approximation can be used at every step of the algorithm. This is done with the option `z.probs="learn"`. If one wishes to use D_A-optimality as the criterion, the vector `A` can be added as an extra argument:

```
linear.nonmyop(covar, init=1, z.probs="learn", k=NULL,
               N=1, int=NULL, lossfunc=calc.DA, stoc=F,
               A=t(c(0, 0, 1)))
```

8.2.2 Logistic model

We now consider the case of logistic regression. When comparing across different designs, it can be helpful to start with the same initial design and to generate responses using Equation (4.48). We show how an initial D-optimal design with five units can be constructed using the exchange algorithm assuming that all parameters are set to zero and assuming no treatment-covariate interactions:

```
Des <-logit.coord(covar[1:5, ], beta=c(0, 0, 0), k=2,
                  int=NULL,  lossfunc=calc.y.D)
```

We have used the function `calc.y.D` to construct a D-optimal design in the logistic regression case. Alternatively, `calc.y.DA`, `calc.y.A` or `calc.y.G` could be used for D_A-, A- or G-optimal designs, respectively.

Now, we can construct a myopic D-optimal design for the logistic regression case using the initial design `Des`. We also generate random uniform numbers `u` in order to simulate responses. For logistic regression, there are a few additional required arguments; we need to include the true parameter values `true.beta`, whether the objective function should be computed using the true value of the parameters (`true.bvcov=T`) or using the estimates at the current point in the algorithm (`true.bvcov=F`), and whether the model fitting is performed with priors (`bayes=T`) or not (`bayes=F`):

```
u <- runif(10,0,1)

desD0 <- logit.des(covar, true.beta=c(0,1,1), init=5, int=NULL,
```

```
                  lossfunc=calc.y.D, same.start=Des,
                  stoc=F, bayes=T, u=u, true.bvcov=T)
desD0
#> $D
#>    rep.1..n. covar tmt
#> 1          1     1  -1
#> 2          1    -1  -1
#> 3          1     1  -1
#> 4          1     1   1
#> 5          1    -1   1
#> 6          1     1   1
#> 7          1    -1   1
#> 8          1    -1  -1
#> 9          1    -1   1
#> 10         1    -1  -1
#>
#> $y
#>  [1] 0 0 1 1 1 1 0 0 1 0
#>
#> $all.beta
#>           (Intercept) D[, -1]covar D[, -1]tmt
#> all.beta 9.412891e-01    0.3871869  1.5706776
#> beta     1.047621e+00    0.4182632  1.6851147
#> beta     1.001289e-01    0.8115854  0.9811492
#> beta     1.667600e-17    0.9413737  1.1198958
#> beta     2.406404e-01    0.8062223  1.2919086
#> beta     1.812436e-01    0.9055358  1.4080177
#>
#> $beta
#>  (Intercept) D[, -1]covar   D[, -1]tmt
#>    0.1812436    0.9055358    1.4080177
#>
#> $all.probs
#> [1] 0.4740768 0.4591522 0.5057856 0.4814327 0.5030691
#>
#> $yprop.tot
#> [1] 0.6666667 0.5714286 0.5000000 0.5555556 0.5000000
#>
#> $opt
```

```
#> [1] 1.1178682 0.6200678 0.4192049 0.2939608 0.2235769
#>
#> $loss.p
#> [1] 1.1178682 0.6200678 0.4290200 0.2939608 0.2263386
#>
#> $loss.m
#> [1] 1.2401214 0.7303947 0.4192049 0.3166351 0.2235769
```

The output shows the design matrix, and the generated responses. Further, at each point of the design between the initial design and the final patient the following quantities are computed:

- Estimate of model parameters.
- Probability of treatment assignment.
- Proportion of units with response 1.
- Objective function.
- Objective function if unit i is assigned treatment 1.
- Objective function if unti i is assigned treatment -1.

For the nonmyopic approach for logistic regression, two additional arguments are needed, `z.probs` and `N`:

```
desD1 <-logit.nonmy(covar,  true.beta=c(0,1,1), init=5, z.probs=0.5,
                    N=1, int=NULL, lossfunc=calc.y.D, same.start=Des,
                    stoc=F, bayes=T, u=u, true.bvcov=T)
```

8.3 Pseudononmyopic approach

For the pseudo-nonmyopic approach for the linear model, we specify the number of trajectories to simulate in the future with the argument `sim`.

```
pseudo.sim10 <- simfuture(covar, sim=10, z.probs=z.probs, int=NULL,
                          lossfunc=calc.D)
pseudo.sim10
#> $design
#>       [,1] [,2] [,3]
#>  [1,]    1    1    1
#>  [2,]    1   -1   -1
#>  [3,]    1    1    1
```

```
#>  [4,]    1    1   -1
#>  [5,]    1   -1    1
#>  [6,]    1    1    1
#>  [7,]    1   -1    1
#>  [8,]    1   -1    1
#>  [9,]    1   -1    1
#> [10,]    1   -1   -1
#>
#> $opt
#>  [1] 3.333322e+09 1.249991e+04 6.249965e+03 3.124961e-02 8.928514e-03
#>  [6] 6.249963e-03 3.676452e-03 2.604155e-03 2.016120e-03 1.249996e-03
#>
#> $loss.m
#>  [1] 0.001177079 0.001114580 0.001062497 0.001114580 0.001104163
#>  [6] 0.001166663 0.001094904 0.001076385 0.001087959 0.001249996
#>
#> $loss.p
#>  [1] 0.001083330 0.001083330 0.001174475 0.001392356 0.001104163
#>  [6] 0.001094904 0.001083330 0.001099534 0.001249996 0.001644730
```

For the pseudo-nonmyopic approach for the logistic model, there are additional arguments that have already been introduced: the true model parameters `true.beta`, the number of repetitions from different starting designs for the exchange algorithm to generate an initial design `k`, whether an initial design is already specified `same.start`, whether the treatment is assigned with a stochastic element `stoc`, whether priors are used in the model fitting `bayes`, a vector of random uniform numbers for generating the response `u` and whether the optimality criterion is computed using the true parameter values or the estimates `true.bvcov`:

```
true.beta <- c(0, 1, 1)
init=5
design.sim.10 <- simfuture.logis(covar, true.beta, init, k=2, sim=10,
                                 z.probs=z.probs, int=NULL,
                                 lossfunc=calc.y.D, same.start=NULL,
                                 stoc=T, bayes=T,  u=u,
                                 true.bvcov=T)
```

8.3.1 Continuous treatment

We now consider the case of the logistic model with continuous treatment. Suppose that we wish to have an initial design with five units. We select the treatments using the optimization

routine `optim()` as follows, where we assume that the parameter values are all set to zero. The function is selecting treatments in the interval $[-1, 1]$ which minimize the D-optimality function written for the case when multiple treatment values are simultaneously optimized, `Dopt.y.t.init()`. We then append the selected treatments to a column for the intercept and a column for the covariates to construct the design matrix:

```
t <-optim(runif(init, min = -1, max = 1), Dopt.y.t.init,
          method="L-BFGS-B", lower=-1, upper=1,
          z=as.numeric(covar[1:init,]), int=NULL, beta=c(0, 0, 0),
          epsilon=0.00001)$par
cont.des <-cbind(1, V1=covar[1:5,], tmt=t)
```

Using the initial design `cont.des`, we construct a myopic D-optimal design and a pseudo-nonmyopic D-optimal design with the number of trajectories set to 10. Here, the optimality function `Dopt.y.t()` is for the continuous treatment case when only one treatment is optimized at a time:

```
myop.cont <-logit.cont(covar, true.beta, init, int=NULL,
                       lossfunc=Dopt.y.t,  same.start=cont.des,
                       bayes=T, u=u, true.bvcov=NULL)
myop.cont
#> $D
#>    V1 V1.1 tmt
#> 1   1    1   1
#> 2   1    1  -1
#> 3   1   -1  -1
#> 4   1   -1   1
#> 5   1   -1  -1
#> 6   1   -1   1
#> 7   1   -1  -1
#> 8   1   -1   1
#> 9   1   -1  -1
#> 10  1    1   1
#>
#> $y
#>  [1] 0 0 0 1 0 1 0 0 0 1
#>
#> $all.beta
#>          (Intercept)  D[, -1]V1 D[, -1]tmt
#> all.beta   -1.581620 -0.6714105  0.9461478
```

```
#> beta       -1.354708 -0.9417476  1.3052552
#> beta       -1.545341 -0.9205978  1.5006143
#> beta       -1.745482 -0.6371323  1.0833715
#> beta       -1.864677 -0.6238266  1.2020117
#> beta       -1.274832 -0.1114414  1.3959728
#>
#> $beta
#> (Intercept)   D[, -1]V1  D[, -1]tmt
#>  -1.2748322  -0.1114414   1.3959728
#>
#> $yprop.tot
#> [1] 0.3333333 0.2857143 0.2500000 0.2222222 0.3000000
#>
#> $opt
#> [1] 2.1839737 1.5744443 1.4840709 1.0571803 0.7062811
pseudnon.cont <- simfuture.logis.cont(covar, true.beta, init=5,
                     k=2, sim=20, z.probs, int=NULL,
                     lossfunc=Dopt.y.t, same.start=cont.des,
                     bayes=T,  u=u, true.bvcov=NULL)
pseudnon.cont
#> $D
#>          V1        tmt
#>  [1,] 1  1  1.0000000
#>  [2,] 1  1 -1.0000000
#>  [3,] 1 -1 -1.0000000
#>  [4,] 1 -1  1.0000000
#>  [5,] 1 -1 -1.0000000
#>  [6,] 1 -1  0.5278612
#>  [7,] 1 -1  0.9293569
#>  [8,] 1 -1 -0.2360708
#>  [9,] 1 -1 -0.2260180
#> [10,] 1  1 -0.2370483
#>
#> $y
#>  [1] 0 0 0 1 0 0 0 0 0 1
#>
#> $beta
#>           (Intercept)  as.matrix(D[, -1])V1 as.matrix(D[, -1])tmt
#>            -1.2896394             0.3504995             0.5079986
```

```
#>
#> $all.betas
#>           (Intercept) as.matrix(D[, -1])V1 as.matrix(D[, -1])tmt
#> all.betas   -1.581620           -0.6714105              0.9461478
#> beta        -1.912799           -0.5044194              0.8313030
#> beta        -2.145080           -0.4014743              0.6719523
#> beta        -2.327405           -0.3750006              0.8198003
#> beta        -2.489750           -0.3587483              0.9670611
#> beta        -1.289639            0.3504995              0.5079986
#>
#> $yprop.tot
#> [1] 0.1666667 0.1428571 0.1250000 0.1111111 0.2000000
#>
#> $opt
#> [1] 2.608524 2.819743 3.587260 4.357316 4.340369
```

8.4 Personalized medicine

To design experiments for personalized medicine, a number of settings have to be chosen, as described in Section 6.2. In the example below, we have two static covariates. The first covariate has a Bernoulli(0.3) distribution and the second covariate has Bernoulli(0.5) distribution. We set the threshold to be -1 for all four linear combinations of interest and κ is set to 0. The prior scale is set to 100. We choose the true parameter values to be zero except for the interaction between the first covariate and treatment. We use the weighted L-optimal criterion to select the treatment; this is the function `calc.logit.wL`. We generate the data for 10 patients using the function `gencov()`, where the argument `code=T` indicates that we are using the $(0, 1)$ coding instead of $(-1, 1)$ coding for binary treatments.

```
z.probs <- c(0.3, 0.5)
init=5
threshold <- rep(-1, 4)
kappa=0
true.beta <- c(0, 0, 0, 0, -2, 0)
prior.scale=100
cr.lossfunc=calc.logit.wL

#generate data
covar <- gencov(z.probs, n.r=10, code=T)
```

We construct an intial design using the exchange algorithm with the number of random starts set to two and assuming that all parameters values are equal to zero. We then obtain 10 random numbers for generating the responses and construct a myopic design:

```
k=2
#initial design
med.D <-cr.logit.coord(as.data.frame(covar[1:init,]), beta=rep(0,6),
                    threshold, kappa, cr.lossfunc=calc.logit.wL,
                    k, wt=NULL)
u <- runif(10)
#myopic design
myop <- cr.logit.des(covar, true.beta, threshold, kappa, init,
                     cr.lossfunc, k, wt=T, int=T, prior.scale=100,
                     same.start=med.D, rand.start=NULL, stoc=F,
                     bayes=T,  u=u, prior.default=T,
                     true.bvcov=T)
myop
#> $D
#>    rep.1..n. covar.V1 covar.V2 tmt covar.V1.1 covar.V2.1
#> 1          1        1        1   1          1          1
#> 2          1        0        1   1          0          1
#> 3          1        1        1   0          0          0
#> 4          1        0        1   0          0          0
#> 5          1        0        0   1          0          0
#> 6          1        1        0   0          0          0
#> 7          1        0        1   1          0          1
#> 8          1        0        0   0          0          0
#> 9          1        0        1   0          0          0
#> 10         1        0        1   1          0          1
#>
#> $y
#> 1 2 3 4 5
#> 0 0 0 1 0 0 0 1 1 0
#>
#> $all.beta
#>          (Intercept) D[, -1]covar.V1 D[, -1]covar.V2 D[, -1]tmt
#> all.beta  -0.4516261       -1.384843     0.610366049  -1.514374
#> beta      -0.7985549       -1.614261     0.894124105  -1.354218
#> beta      -0.8056438       -1.561382     0.827834903  -1.524451
#> beta       0.8068893       -2.460256     0.004850823  -2.457159
```

```
#> beta       1.1814783       -2.941211     0.232304326  -2.907058
#> beta       1.2105222       -2.941316     0.189504449  -3.086601
#>          D[, -1]covar.V1 D[, -1]covar.V2
#> all.beta    -0.121243035      -0.6822250
#> beta        -0.103877132      -0.7087699
#> beta        -0.067543821      -0.8951602
#> beta        -0.018807396      -0.8109399
#> beta        -0.012971166      -0.8910774
#> beta        -0.004244081      -1.0039260
#>
#> $all.wr
#>           [,1]      [,2]      [,3]      [,4]
#> [1,] 0.5383291 0.5331493 0.6375480 0.5940264
#> [2,] 0.5329429 0.5266102 0.6245403 0.5799878
#> [3,] 0.5479203 0.5321682 0.6799909 0.5924758
#> [4,] 0.6851391 0.5711845 0.7698418 0.6131536
#> [5,] 0.7303195 0.5865671 0.8335958 0.6293927
#> [6,] 0.7425796 0.5904409 0.8682042 0.6364270
#>
#> $beta
#>     (Intercept) D[, -1]covar.V1 D[, -1]covar.V2      D[, -1]tmt
#>     1.210522211    -2.941315561     0.189504449    -3.086600919
#> D[, -1]covar.V1 D[, -1]covar.V2
#>    -0.004244081    -1.003926010
#>
#> $yprop
#> [1] 0.2000000 0.1666667 0.1428571 0.2500000 0.3333333 0.3000000
#>
#> $tmtprop
#> [1] 0.6000000 0.5000000 0.5714286 0.5000000 0.4444444 0.5000000
#>
#> $all.lopt
#> [1] 42.88224 23.89727 22.35003 19.01539 18.31301 17.79933
#>
#> $loss.p
#> [1] 42.23968 24.81373 23.96394 23.14424 22.88740
#>
#> $loss.m
#> [1] 26.97869 24.97893 21.73986 22.73942 23.16425
```

The ouput provides the design matrix, the responses, and the following quantities for each unit between *init* and n:

- Estimates of the model parameters.
- Estimates of the weights.
- Proportion of units with response 1.
- Probability of treatment assignment.
- Optimality.
- Optimality if unit i is assigned treatment 1.
- Optimality if unti i is assigned treatment -1.

For a nonmyopic design for personalized medicine, we specify additional arguments `z.probs` and `N`, and for a pseudo-nonmyopic design, we specify additional arguments `z.probs` and `sim`:

```
nonmy1 <- logit.Lbnon(covar, true.beta,  threshold,kappa, init,
                      z.probs, N=1, prior.scale, same.start=med.D,
                      stoc=T, bayes=T, u=u,  true.bvcov=T)
nonmy1
#> $D
#>    rep.1..n. covar.V1 covar.V2 tmt covar.V1.1 covar.V2.1
#> 1          1        1        1   1          1          1
#> 2          1        0        1   1          0          1
#> 3          1        0        1   0          0          0
#> 4          1        0        0   1          0          0
#> 5          1        1        0   0          0          0
#> 6          1        0        1   0          0          0
#> 7          1        1        0   0          0          0
#> 8          1        0        1   1          0          1
#> 9          1        1        1   1          1          1
#> 10         1        0        0   1          0          0
#>
#> $y
#>  [1] 1 1 1 0 1 0 1 0 0 1
#>
#> $all.beta
#>          (Intercept) design[, -1]covar.V1 design[, -1]covar.V2
#> all.beta   0.8027954           0.9537492           1.4980099
#> beta      -0.2878227           1.5435617           0.5337140
```

```
#> beta      -0.1126728            1.9688283            0.4075259
#> beta      -0.1423335            2.1703495            0.2591197
#> beta       0.3680976            1.1699257           -0.1191595
#> beta       0.8931736            0.8702034           -0.5898950
#>          design[, -1]tmt design[, -1]covar.V1 design[, -1]covar.V2
#> all.beta      -1.2882865             0.2100298            0.88443941
#> beta          -0.4019993             0.2213949            1.48615757
#> beta          -0.4956859             0.1618019            1.46970546
#> beta          -0.7667745             0.4016875            0.55193107
#> beta          -1.1061759            -0.5201627            0.43076536
#> beta          -0.5898950            -0.4486288            0.05831601
#>
#> $all.wr
#>           [,1]      [,2]      [,3]      [,4]
#> [1,] 0.5250915 0.5054764 0.4421918 0.4485065
#> [2,] 0.4465471 0.4410512 0.2197166 0.3459816
#> [3,] 0.4548373 0.4544799 0.2297898 0.3627202
#> [4,] 0.4790369 0.4456447 0.3439974 0.4118782
#> [5,] 0.5096669 0.5751461 0.4348516 0.5170890
#> [6,] 0.4617505 0.5050592 0.4058917 0.4982697
#>
#> $beta
#>          (Intercept) design[, -1]covar.V1 design[, -1]covar.V2
#>           0.89317356           0.87020340          -0.58989501
#>      design[, -1]tmt design[, -1]covar.V1 design[, -1]covar.V2
#>          -0.58989501          -0.44862877           0.05831601
#>
#> $yprop
#> [1] 0.8000000 0.6666667 0.7142857 0.6250000 0.5555556 0.6000000
#>
#> $tmtprop
#> [1] 0.6000000 0.5000000 0.4285714 0.5000000 0.5555556 0.6000000
#>
#> $all.lopt
#> [1] 36.33154 34.84562 33.37969 32.30242 28.95182 27.51782
pseudnon <-cr.simfuture.logis(covar, true.beta, threshold, kappa,
                              init, int=T, cr.lossfunc=calc.logit.wL,
                              k, wt=T,  sim=10, z.probs,
                              same.start=med.D, stoc=F, bayes=T, u=u,
```

```
                              prior.default=T, true.bvcov=T)
pseudnon
#> $D
#>       rep(1, n) covar.V1 covar.V2 tmt covar.V1 covar.V2
#> 1             1        1        1   1        1        1
#> 2             1        0        1   1        0        1
#> 3             1        0        1   0        0        0
#> 4             1        0        0   1        0        0
#> 5             1        1        0   0        0        0
#> new.d         1        0        1   1        0        1
#> new.d         1        1        0   1        1        0
#> new.d         1        0        1   0        0        0
#> new.d         1        1        1   0        0        0
#> new.d         1        0        0   0        0        0
#>
#> $y
#> 1 2 3 4 5
#> 1 1 1 0 1 1 1 0 0 0
#>
#> $all.betas
#>           (Intercept) D[, -1]covar.V1 D[, -1]covar.V2  D[, -1]tmt
#> all.betas  0.80279536       0.9537492       1.4980099 -1.28828652
#> beta       0.81130405       0.9120125       1.6722489 -1.21738697
#> beta       0.81158628       1.3338382       1.5442393 -0.93740698
#> beta      -0.30673664       1.8838579       0.5806413 -0.10397307
#> beta       0.04560457       0.3623769      -0.3328009 -0.06345723
#> beta      -0.73795372       0.7282190       0.1559534  0.33605659
#>           D[, -1]covar.V1 D[, -1]covar.V2
#> all.betas       0.2100298       0.8844394
#> beta            0.1463624       1.0686841
#> beta            0.6887546       0.9211908
#> beta            0.6065763       1.6931354
#> beta            1.1558203       2.1888505
#> beta            1.1401543       2.0750946
#>
#> $all.wr
#>           [,1]      [,2]       [,3]       [,4]
#> [1,] 0.5250915 0.5054764 0.44219184 0.44850648
#> [2,] 0.5187983 0.5047251 0.41692698 0.43982405
```

```
#> [3,] 0.4944715 0.4256129 0.40192020 0.38813517
#> [4,] 0.4185510 0.3451212 0.14236028 0.27543233
#> [5,] 0.3819065 0.2421883 0.09244249 0.10683454
#> [6,] 0.2915271 0.1923016 0.07202091 0.09731824
#>
#> $beta
#>     (Intercept) D[, -1]covar.V1 D[, -1]covar.V2      D[, -1]tmt
#>      -0.7379537       0.7282190       0.1559534       0.3360566
#> D[, -1]covar.V1 D[, -1]covar.V2
#>       1.1401543       2.0750946
#>
#> $yprop
#> [1] 0.8000000 0.8333333 0.8571429 0.7500000 0.6666667 0.6000000
#>
#> $tmtprop
#> [1] 0.6000000 0.6666667 0.7142857 0.6250000 0.5555556 0.5000000
#>
#> $all.lopt
#> [1] 36.68298 35.26751 29.01025 27.70550 15.62064 13.09603
#>
#> $all.dawopt
#> [1] 253.36673 215.92678 148.82077 121.81010  52.50876  36.08903
#>
#> $allt.lopt
#> [1] 36.68298 35.26751 29.01025 27.70550 15.62064 13.09603
#>
#> $not.lopt
#> [1] 52.275105 55.143863 36.667052 21.596137  7.658392  4.877878
#>
#> $loss.p
#> [1] 15.754118 15.806702 15.514552 12.871152  6.554679
#>
#> $loss.m
#> [1] 16.42159 16.49943 15.35843  8.58600  5.20490
```

Chapter 9

Conclusion

In this thesis, we reviewed methods in the literature for sequential design of experiments with covariates. We focused in particular on model-based approaches where decisions for allocating treatment are made based on optimal design criteria and developed a novel pseudo-nonmyopic approach for optimal design taking into account the covariate values of future possible patients. We also considered how these approaches apply in the context of personalized medicine. In this chapter, we summarize the results in Section 9.1, outline the limitations of our work in Section 9.2 and provide scope for future work in Section 9.3.

9.1 Discussion

In Chapter 2, we introduced myopic covariate-adaptive approaches to treatment allocation and compared them using simulations. We found that the sequential optimal design method offers improvement in efficiency over other approaches such as randomization, Efron's biased coin and minimization for covariate balance and loss, particularly when sample sizes are small. This was demonstrated for both continuous and binary responses. Further, the Atkinson approach is able to take into account dependence between covariates while other approaches are not able to, and it is flexible in the sense that both discrete and continuous covariates can be considered simultaneously; for minimization, separate approaches have been developed for discrete and continuous covariates.

We then extended Atkinson's method in two ways in Chapter 4: firstly, it was placed in a nonmyopic framework for linear responses where any optimality criterion can be used. We focused in particular on D, D_A and G-optimality. Secondly, we allowed for binary responses which we again placed in a nonmyopic framework. In our simulations, we observed that for the linear model case, there appears to a slight benefit in using the nonmyopic

method for the D-optimality case when there are one or two binary covariates and the true covariate distribution is known. We also observed a slight benefit in the G-optimal case. We consider only horizon up to five due to the computational complexity of the problem. For the logistic model case, we observed no benefit to using the nonmyopic approach over the myopic approach.

We developed a novel methodology called the pseudo-nonmyopic approach in Chapter 5 which is still able to take into account future possible patients, but is less computationally expensive than the nonmyopic approach. We further extended this method to be applicable for a continuous treatment, which is too computationally intensive to be considered at this moment for the nonmyopic approach. Simulations in this chapter showed that the pseudo-nonmyopic approach performs similarly to the myopic approach for the linear and logistic model case with a binary treatment. For a continuous treatment, we found one case with a large interaction term where the pseudo-nonmyopic approach is more efficient than the myopic approach for small sample sizes, but not for sample sizes greater than 20.

In Chapter 6, we introduced the idea of personalized medicine and how the optimal design approach can be taken to construct designs which aim to estimate the effects of important treatment-biomarker combinations. We compared the performance of the myopic weighted L-optimal design under three different settings: using the normal-inverse-gamma prior distribution, using the Cauchy prior distribution, and using the fully Bayesian approach with a normal-inverse-gamma distribution. We found that the fully Bayesian approach has a slight advantage in that the weights reach their true value quicker and leads to slightly more efficient designs. We then considered how the weighted L-optimal design fares against minimization, D_A-optimal and weighted D_A-optimal designs. We found that they provide surprisingly similar results in terms of estimates of the parameters, the weights, and the power and 1-specificity of the hypothesis tests. Minimization is slightly less efficient than the three optimal design-based approaches, particularly for small sample sizes.

The main themes of the thesis are culminated in Chapter 7 where we apply the nonmyopic and pseudo-nonmyopic approaches to personalized medicine. Our simulation results echo the results of Chapter 4 in showing that the non-myopic approach is not more efficient than the myopic approach in the context of personalized medicine. In our simulations comparing the pseudo-nonmyopic approach with the myopic approach for personalized medicine, we see that, given a large interaction term and particularly when that interaction term involves a rare covariate, the pseudo-nonmyopic approach can offer the potential to be more efficient than the myopic approach.

Finally, Chapter 8 is a vignette of the `R` package with functions that implement the approaches to sequential allocation of treatment described in this thesis.

9.2 Limitations

9.2.1 Computational time

A major limitation of our implementation of the nonmyopic approach is that it is extremely computationally expensive. We conducted simulations with horizon up to five for the linear model case, and horizon up to three for the logistic model case. We found that the nonmyopic approach does not offer improvement compared to the myopic case, but this may be because we were not able to make the horizon large enough to see any substantive improvements. The pseudo-nonmyopic approach was also time-consuming and we chose the total sample size to be smaller than in other simulations ($n = 50$) in Chapter 7. By writing some of the code in C instead of `R`, the myopic and pseudo-nonmyopic approaches would have faster running times. They would be more convenient to use and we would be able to explore larger values for the horizon and sample size in simulations.

9.2.2 Applicability to clinical trials

There are a number of limitations to our work in its ability to be directly applicable to clinical trials. Firstly, we assume responses are measured immediately after treatments are given to patients. This is not a realistic assumption so some method to allow for a delay between treatment allocation and response could be useful. One modification would be to allow for the method to be batch sequential; instead of allocating treatments to one patient at a time, a group of patients may be given optimal treatments by using the exchange algorithm. It is also possible to incorporate delay in adaptive designs. Hardwick et al. (2006) achieve this by assuming that patients arrive according to a Poisson process, and they consider three types of designs: a two-armed bandit rule which incorporates delay, a delayed play-the-winner rule and an adaptive hyperopic rule which makes decisions based on the number of patients yet to be treated in the trial, as well as the observed responses and covariates so far. Their key conclusion is that the performance of the bandit rule with delay is nearly as efficient as the case where responses are observed immediately after treatment, unless the delay rate is many orders of magnitude greater than the arrival rate of the patients.

Secondly, we do not consider toxicity in our work. We assume that the treatment which leads to a better response is the more desirable treatment, but it is possible that such a treatment has unsafe toxicity levels (Rosenberger, 1999). In our algorithms for treatment assignment,

if the optimality criterion is equal for treatment $t_i = 1$ and $t_i = 0$, we would assign the treatment at random. In clinical trials, this is less likely to happen as relative efficiency of the treatments need to be considered in conjunction with relative toxicity (Simon, 1977). In general, Rosenberger (1999) recommended that adaptive designs should be considered after previous experiments have been able to establish low toxicity of the treatments.

A further limitation of our work is that we arbitrarily assume in all of our simulations that we have 100 patients in the trial. In clinical trials, there are stopping rules that determine when the trial should terminate (Stallard et al., 2001). See Whitehead (1993) for a frequentist perspective and Berry (1989) and Freedman and Spiegelhalter (1989) for a Bayesian perspective on stopping rules in interim analysis. Including this element into our designs would mean that our methodology is more generally applicable to clinical trials. Further, we may be able to make statements about relative numbers of patients and costs required in order to detect a significant difference in treatment effect for each method.

9.3 Future work

In this section, we describe a number of ways in which our current work can be extended. We then consider how the nonmyopic approach can be adopted for CARAEE designs.

9.3.1 Extending current work

Our optimality criteria have been chosen with parameter estimation as our primary goal. In the example of personalized medicine, we assume that we wish to estimate certain linear combinations of parameters as precisely as possible so that we can identify the effective treatment-biomarker combinations. However, this problem of identifying the treatment-biomarker combinations can be framed as a model selection problem. With this outlook, the T-optimal objective function may be appropriate since it is designed to discriminate between competing models (Atkinson and Fedorov, 1975).

The non-myopic algorithm considers only the case where the response and treatments are binary. Natural extensions include allowing for more complex treatment structures, such as factorial designs, or allowing for a continuous response. Computing the expected objective function for a continuous response would require Monte-Carlo simulation. As explained in Section 9.2.1, extending our framework for the non-myopic approach to allow for a more general response will require greater computational efficiency in our algorithms. This is also true of the pseudo-nonmyopic approach.

9.3.2 CARAEE designs

In the optimality criteria that we have considered in the thesis, the response of the patient is either not included (minimization, optimal design methods for the linear model case), or they are included in order to update parameter estimates (optimal design methods for the logistic model case, weighted L-optimal design). The response has not been used in order to inform treatment allocation based on the efficacy of the treatment. Covariate adjusted response-adaptive designs based on efficiency and ethics (CARAEE) aim to optimize a utility function which takes into account the number of patients who receive the more effective treatment. We did some preliminary work on CARAEE designs. Here, our optimality criterion is a function which has a component for efficiency and a component for ethics, as well as a tuning parameter which allows the practitioner to decide which aim is more important. The CARA (covariate adjusted response adaptive) design and RAR (response adaptive randomization) design are special cases of the CARAEE design.

In Appendix J, we describe in more detail the CARAEE designs that we implemented. Some simulation results showed that the CARAEE designs, as expected, produce designs that are less efficient but more ethical than myopic D-optimal designs. Investigation is needed to verify suitable choices of the tuning parameter, and to check whether our choices of measuring efficiency and ethics are suitable. In particular, one concern is that they are not measured on different scales. An interesting area of future work is extending the CARAEE design for a nonmyopic or pseudo-nonmyopic framework; we have seen in Chapters 4 and 5 how to calculate the expected objective for the efficiency, $d_i(t)$, from a nonmyopic and pseudo-nonmyopic framework, respectively. How to incorporate future possible responses into the ethics component and combine it with the efficiency component into a sensible objective function is a question worth investigating and would further make the work more applicable to clinical trials.

Appendix A

Proof of Result 1

We show that, for the optimality criterion $\Psi = \log \left| \boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right|$, the derivative with respect to the information matrix $\frac{\partial \Psi}{\partial \boldsymbol{M}}$ is given by $\boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1}$. The proof of this result relies on a number of definitions and results that we outline below.

Suppose that $\boldsymbol{M}_1$ and $\boldsymbol{M}_2$ are information matrices. The Frechet derivative of Ψ at $\boldsymbol{M}_1$ in the direction of $\boldsymbol{M}_2$ is given by:

$$F_\Psi(\boldsymbol{M}_1, \boldsymbol{M}_2) = \lim_{\epsilon \to 0^+} \frac{1}{\epsilon} \left[\Psi \left((1-\epsilon)\boldsymbol{M}_1 + \epsilon \boldsymbol{M}_2 \right) \right) - \Psi(\boldsymbol{M}_1) \right]. \tag{A.1}$$

Intuitively, this derivative measures how much the design criterion Ψ changes when the design matrix $\boldsymbol{M}_1$ is nudged slightly in the direction of $\boldsymbol{M}_2$. In order to compute the Frechet derivative, it is often easier to first derive the Gateaux derivative of Ψ at $\boldsymbol{M}_1$ in the direction of $\boldsymbol{M}_2$, given by:

$$G_\Psi(\boldsymbol{M}_1, \boldsymbol{M}_2) = \lim_{\epsilon \to 0^+} \frac{1}{\epsilon} \left[\Psi(\boldsymbol{M}_1 + \epsilon \boldsymbol{M}_2) - \Psi(\boldsymbol{M}_1) \right], \tag{A.2}$$

and to use the relationship between the two derivatives:

$$F_\Psi(\boldsymbol{M}_1, \boldsymbol{M}_2) = G_\Psi(\boldsymbol{M}_1, \boldsymbol{M}_2 - \boldsymbol{M}_1). \tag{A.3}$$

Now, the derivative $\frac{\partial \Psi}{\partial \boldsymbol{M}}$ that we wish to obtain can be observed after noting that the Frechet derivative with $\boldsymbol{M}_1 = \boldsymbol{M}$ and $\boldsymbol{M}_2 = \boldsymbol{x}\boldsymbol{x}^\intercal$ can be written the form (Atkinson et al., 2007, p.136):

$$F_\Psi(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal) = -\operatorname{tr} \boldsymbol{M} \frac{\partial \Psi}{\partial \boldsymbol{M}} + \boldsymbol{x}^\intercal \frac{\partial \Psi}{\partial \boldsymbol{M}} \boldsymbol{x}. \tag{A.4}$$

Finally, we also rely on two matrix identities. The first is the Woodbury matrix identity, where we have, for conformable matrices $\boldsymbol{A}, \boldsymbol{U}, \boldsymbol{C}, \boldsymbol{V}$:

$$(\boldsymbol{A} + \boldsymbol{U}\boldsymbol{C}\boldsymbol{V})^{-1} = \boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}\boldsymbol{U}\left(\boldsymbol{C}^{-1} + \boldsymbol{V}\boldsymbol{A}^{-1}\boldsymbol{U}\right)^{-1}\boldsymbol{V}\boldsymbol{A}^{-1}. \tag{A.5}$$

The second is a property of the determinant. For conformable matrices $\boldsymbol{I}$ and $\boldsymbol{H}$:

$$|\boldsymbol{I} + \epsilon\boldsymbol{H}| = 1 + \epsilon \operatorname{tr}(\boldsymbol{H}) + O(\epsilon^2), \tag{A.6}$$

where $O(\epsilon^2)$ indicates that the quantity is bounded above by some constant on the order of ϵ^2, for small ϵ, which does not depend on $\boldsymbol{H}$.

We now have the necessary tools to prove the following result:

Result 1. *For the optimality criterion* $\Psi = \log\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|$, *we have that the derivative with respect to the information matrix is given by:* $\frac{\partial\Psi}{\partial\boldsymbol{M}} = \boldsymbol{M}^{-1}\boldsymbol{A}\left(\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}$.

Proof. The Gateau derivative of $\Psi = \log\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|$ at $\boldsymbol{M}$ in the direction of $\boldsymbol{x}\boldsymbol{x}^\intercal$ is given by:

$$\begin{aligned}
G_\Psi(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal) &= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\left[\Psi(\boldsymbol{M} + \epsilon\boldsymbol{x}\boldsymbol{x}^\intercal) - \Psi(\boldsymbol{M})\right] \\
&= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\left[\log\left|\boldsymbol{A}^\intercal(\boldsymbol{M} + \epsilon\boldsymbol{x}\boldsymbol{x}^\intercal)^{-1}\boldsymbol{A}\right| - \log\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|\right] \\
&= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\log\left(\frac{\left|\boldsymbol{A}^\intercal(\boldsymbol{M} + \epsilon\boldsymbol{x}\boldsymbol{x}^\intercal)^{-1}\boldsymbol{A}\right|}{\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|}\right) \\
&= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\log\left|\boldsymbol{A}^\intercal(\boldsymbol{M} + \epsilon\boldsymbol{x}\boldsymbol{x}^\intercal)^{-1}\boldsymbol{A}\right|\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|^{-1}.
\end{aligned}$$

By the Woodbury identity (Equation A.5), we have:

$$\begin{aligned}
G_\Psi(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal) &= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\log\left|\boldsymbol{A}^\intercal\left(\boldsymbol{M}^{-1} - \boldsymbol{M}^{-1}\epsilon\boldsymbol{x}\left(\boldsymbol{I}^{-1} + \boldsymbol{x}^\intercal\boldsymbol{M}^{-1}\epsilon\boldsymbol{x}\right)^{-1}\boldsymbol{x}^\intercal\boldsymbol{M}^{-1}\right)\boldsymbol{A}\right|\left|\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\right|^{-1} \\
&= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\log\left|\boldsymbol{I}_s - \epsilon\boldsymbol{A}^\intercal\boldsymbol{M}^{-1}\boldsymbol{x}\left(\boldsymbol{I}_s^{-1} + \boldsymbol{x}^\intercal\boldsymbol{M}^{-1}\epsilon\boldsymbol{x}\right)^{-1}\boldsymbol{x}^\intercal\boldsymbol{M}^{-1}\boldsymbol{A}\left(\boldsymbol{A}^\intercal\boldsymbol{M}_1^{-1}\boldsymbol{A}\right)^{-1}\right|
\end{aligned}$$

By the property of determinants given in Equation (A.6), we have:

$$\begin{aligned} G_\Psi(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal) &= \lim_{\epsilon \to 0^+} \frac{1}{\epsilon} \log \Bigg(\boldsymbol{I}_s - \epsilon \operatorname{tr} \Big(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x} \left(\boldsymbol{I}_s^{-1} + \boldsymbol{x}^\intercal \boldsymbol{M}^{-1} \epsilon \boldsymbol{x} \right)^{-1} \boldsymbol{x}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right)^{-1} \\ &\qquad + O(\epsilon^2) \Big) \Bigg) \\ &= s - \operatorname{tr} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x}\boldsymbol{x}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \right)^{-1} \right) \\ &\propto - \operatorname{tr} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x}\boldsymbol{x}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \right)^{-1} \right). \end{aligned}$$

The Frechet derivative of $\Psi = \log \left| \boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right|$ at $\boldsymbol{M}$ in the direction of $\boldsymbol{x}\boldsymbol{x}^\intercal$ is given by the relation to the Gateaux derivative in Equation (A.3):

$$\begin{aligned} F_\Psi(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal) &= G_\Psi \left(\boldsymbol{M}, \boldsymbol{x}\boldsymbol{x}^\intercal - \boldsymbol{M} \right) \\ &= - \operatorname{tr} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \left(\boldsymbol{x}\boldsymbol{x}^\intercal - \boldsymbol{M} \right) \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \right)^{-1} \right) \\ &= - \operatorname{tr} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \left(\boldsymbol{x}\boldsymbol{x}^\intercal \right) \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \right)^{-1} \right) + \operatorname{tr} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{M} \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}_1^{-1} \boldsymbol{A} \right)^{-1} \right) \\ &= \operatorname{tr}(\boldsymbol{M} \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1}) - \boldsymbol{x}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{x} \\ &\propto - \operatorname{tr} \left(\boldsymbol{M} \frac{\partial \Psi}{\partial \boldsymbol{M}} \right) + \boldsymbol{x}^\intercal \frac{\partial \Psi}{\partial \boldsymbol{M}} \boldsymbol{x}. \end{aligned}$$

Thus by Equation (A.4), we have that $\frac{\partial \Psi}{\partial \boldsymbol{M}} = \boldsymbol{M}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^\intercal \boldsymbol{M}^{-1} \boldsymbol{A} \right)^{-1} \boldsymbol{A}^\intercal \boldsymbol{M}^{-1}$.

□

Appendix B

Proof of Result 2

Result 2. *Senn et al. (2010) claim that the Atkinson approach is equivalent to minimizing* $\mathcal{L}_i$.

Proof. We wish to show that there is a correspondence between $d(\boldsymbol{z}_i, t)$ given in Equation (2.36) and the quantity $\mathcal{L}_i$ such that maximizing the first implies minimizing the second and vice versa.

We first show that $F_{i-1} \geq 0$ where $F_{i-1} = i - 1 - \mathcal{L}_{i-1}$. We can write $\mathcal{L}_{i-1}$ as

$$\mathcal{L}_{i-1} = \boldsymbol{t}_{i-1}^{\intercal} \boldsymbol{P}_{i-1} \boldsymbol{t}_{i-1},$$

where $\boldsymbol{P}_{i-1}$ is the symmetric matrix $\overline{\boldsymbol{Z}}_{i-1}(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1})^{-1}\overline{\boldsymbol{Z}}_{i-1}^{\intercal}$. We denote the largest eigenvalue of $\boldsymbol{P}_{i-1}$ by λ_{max}. Since $\boldsymbol{P}_{i-1}$ is a projection matrix, its eigenvalues are zero or one. Therefore, we have:

$$\mathcal{L}_{i-1} < \lambda_{max} \boldsymbol{t}_{i-1}^{\intercal} \boldsymbol{t}_{i-1} < \boldsymbol{t}_{i-1}^{\intercal} \boldsymbol{t}_{i-1} < i - 1,$$

and we conclude that F_{i-1} is positive.

Next, we consider $d(\boldsymbol{z}_i, t)$. Below, the information matrix $\boldsymbol{M}_{i-1}$ is based on the design with the $i-1$ patients in the study so far, and $\boldsymbol{x} = (1 \;\; \boldsymbol{z}_i^{\intercal} \;\; t)^{\intercal}$. We assume that interest lies only in the treatment effect α, so we have that $\boldsymbol{A} = (0 \;\; 0 \;\; 0 \;\; ... \;\; 1)^{\intercal}$. We denote by $\overline{\boldsymbol{z}_i}$ the vector $(1 \;\; \boldsymbol{z}_i^{\intercal})^{\intercal}$ and we denote by $\overline{\boldsymbol{Z}_i}$ the matrix $\left(\mathbf{1} \;\; \boldsymbol{z}_{i,1}, \boldsymbol{z}_{i,2}, ..., \boldsymbol{z}_{i,k}\right)$.

$$d(\boldsymbol{z}_i, t) = \boldsymbol{x}^{\intercal} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \left(\boldsymbol{A}^{\intercal} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \right)^{-1} \boldsymbol{A}^{\intercal} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= \boldsymbol{x}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \left(i - 1 - \mathcal{L}_{i-1}\right) \boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= F_{i-1} \boldsymbol{x}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{A} \boldsymbol{A}^\intercal \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= F_{i-1} \boldsymbol{x}^\intercal \begin{pmatrix} \left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} - \mathcal{L}_{i-1}\right)^{-1} & -(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} F_{i-1}^{-1} \\ -F_{i-1}^{-1} \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} & F_{i-1}^{-1} \end{pmatrix} \begin{pmatrix} \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{1} \end{pmatrix} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= F_{i-1} \boldsymbol{x}^\intercal \begin{pmatrix} \boldsymbol{0} & -(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} F_{i-1}^{-1} \\ \boldsymbol{0} & F_{i-1}^{-1} \end{pmatrix} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= \boldsymbol{x}^\intercal \begin{pmatrix} \boldsymbol{0} & -(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \\ \boldsymbol{0} & 1 \end{pmatrix} \boldsymbol{M}_{i-1}^{-1} \boldsymbol{x}$$

$$= F_{i-1}^{-1} \boldsymbol{x}^\intercal \begin{pmatrix} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} & -(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \\ -\boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} & 1 \end{pmatrix} \boldsymbol{x}$$

$$= F_{i-1}^{-1} \begin{pmatrix} \boldsymbol{z}_i^\intercal (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} - t \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \\ -\overline{\boldsymbol{z}}_i^\intercal (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + t \end{pmatrix}^\intercal \boldsymbol{x}$$

$$\begin{aligned} = F_{i-1}^{-1} \Big\{ & \overline{\boldsymbol{z}}_i^\intercal (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{z}}_i \\ & + -t \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_j)^{-1} \overline{\boldsymbol{z}}_i - t \overline{\boldsymbol{z}}_i^\intercal (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + 1 \Big\} \end{aligned}$$

$$\begin{aligned} = F_{i-1}^{-1} \Big\{ & 1 - 2t \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1\intercal} \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{z}}_i + \\ & \overline{\boldsymbol{z}}_i^\intercal (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{z}}_i \Big\} \end{aligned}$$

Since we are interested in the choice of t which maximizes this quantity, we can disregard the terms that do not involve t. Since F_{i-1}^{-1} is positive, we wish to maximize:

$$-2t \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} (\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1})^{-1} \overline{\boldsymbol{z}}_i. \tag{B.1}$$

Next, we consider $\mathcal{L}_i$. The loss after i trials can be obtained by updating the loss after $i-1$

trials, as given in the formula below (Senn et al., 2010).

$$\mathcal{L}_i = \left(\boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1} + t\overline{\boldsymbol{z}}_i\right)\left(\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1} - \frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\right)\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + \overline{\boldsymbol{z}}_i t\right) \tag{B.2}$$

Expanding the right hand side, we obtain:

$$\begin{aligned}
\mathcal{L}_i = {} & \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\boldsymbol{z}_i t \\
& - \boldsymbol{t}^\intercal \overline{\boldsymbol{Z}}_{i-1}\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \\
& - \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{z}}_i t \\
& + t\overline{\boldsymbol{z}}_i\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + t\overline{\boldsymbol{z}}_i\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i t \\
& - t\overline{\boldsymbol{z}}_i\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{z}}_i t \\
& - t\overline{\boldsymbol{z}}_i\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1}
\end{aligned}$$

$$\begin{aligned}
= {} & \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} + 2t\boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i \\
& - \boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{Z}}_{i-1}^\intercal \boldsymbol{t}_{i-1} \\
& - 2\boldsymbol{t}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{z}}_i t + \overline{\boldsymbol{z}}_i\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i \\
& - \overline{\boldsymbol{z}}_i\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^\intercal\left(\overline{\boldsymbol{Z}}_{i-1}^\intercal \overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{z}}_i
\end{aligned}$$

Again, since we are only interested in the choice of t which minimizes the above quantity, the terms which we focus on are:

$$2t\boldsymbol{t}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i - 2t\boldsymbol{t}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\frac{\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}}{1+\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\overline{\boldsymbol{z}}_i. \quad \text{(B.3)}$$

We use partial fractions to re-write the second term:

$$2t\boldsymbol{t}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i - 2t\left(\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\boldsymbol{t}_{i-1} - \frac{\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\boldsymbol{t}_{i-1}}{1+\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}\right). \quad \text{(B.4)}$$

This simplifies to ive us the following term. We see that the choice of t which minimizes the expression below maximizes the expression $-2t\boldsymbol{t}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1})^{-1}\overline{\boldsymbol{z}}_i$ from Equation (B.1).

$$2t\frac{\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\boldsymbol{t}_{i-1}}{1+\overline{\boldsymbol{z}}_i^{\intercal}\left(\overline{\boldsymbol{Z}}_{i-1}^{\intercal}\overline{\boldsymbol{Z}}_{i-1}\right)^{-1}\overline{\boldsymbol{z}}_i}. \quad \text{(B.5)}$$

□

Appendix C

Simulations for covariate-adaptive approaches

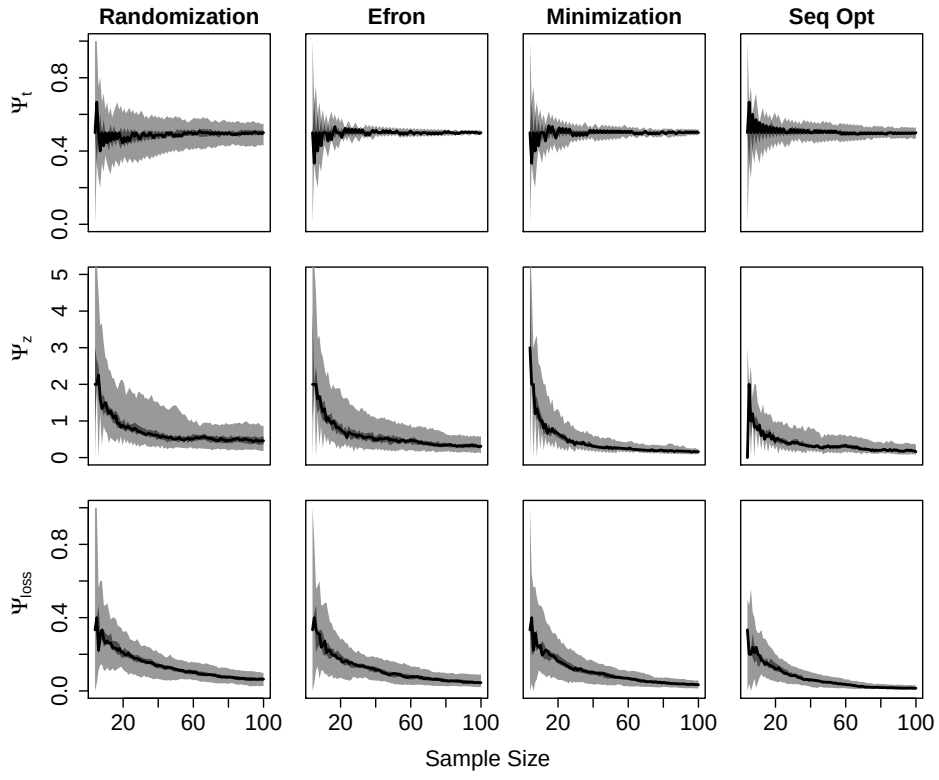

Figure C.1: Distributions of Ψ_t, Ψ_z and Ψ_{loss} as sample size increases for five Bernoulli($p = 0.8$) covariates which have correlation 0.8, based on 100 simulations. Four allocation methods are considered: Randomization, Efron's biased coin, the classic form of minimization, and the sequential optimal design method. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

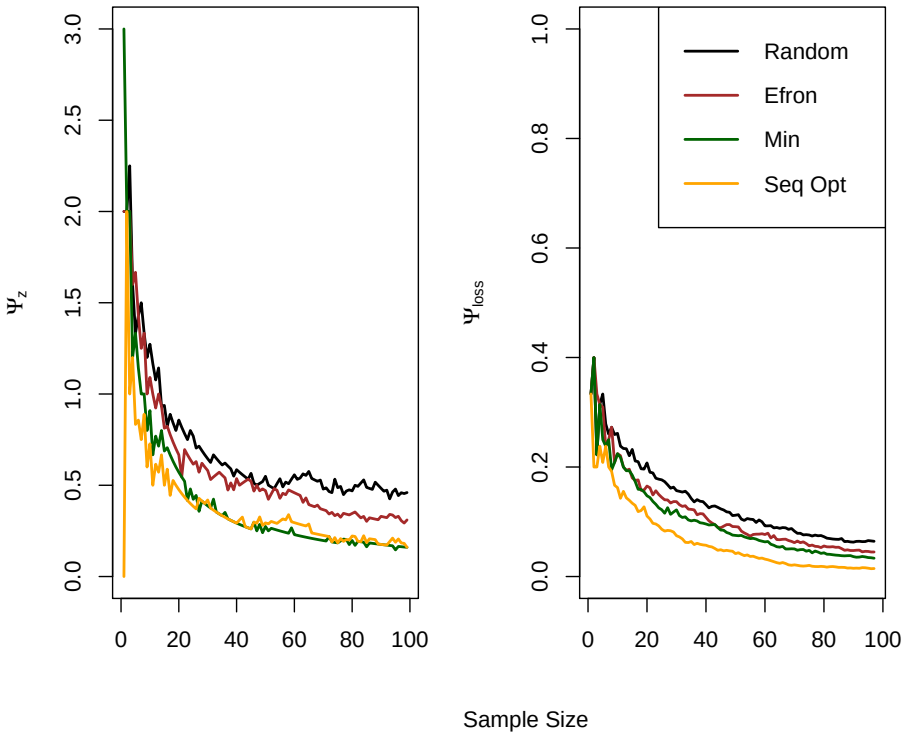

Figure C.2: Medians of the distributions of Ψ_z and Ψ_{loss} as sample size increases for five Bernoulli($p = 0.8$) covariates which have correlation 0.8, based on 100 simulations. Four allocation methods are considered: Randomization (black), Efron's biased coin (brown), the classic form of minimization (green), and the sequential optimal design method (yellow).

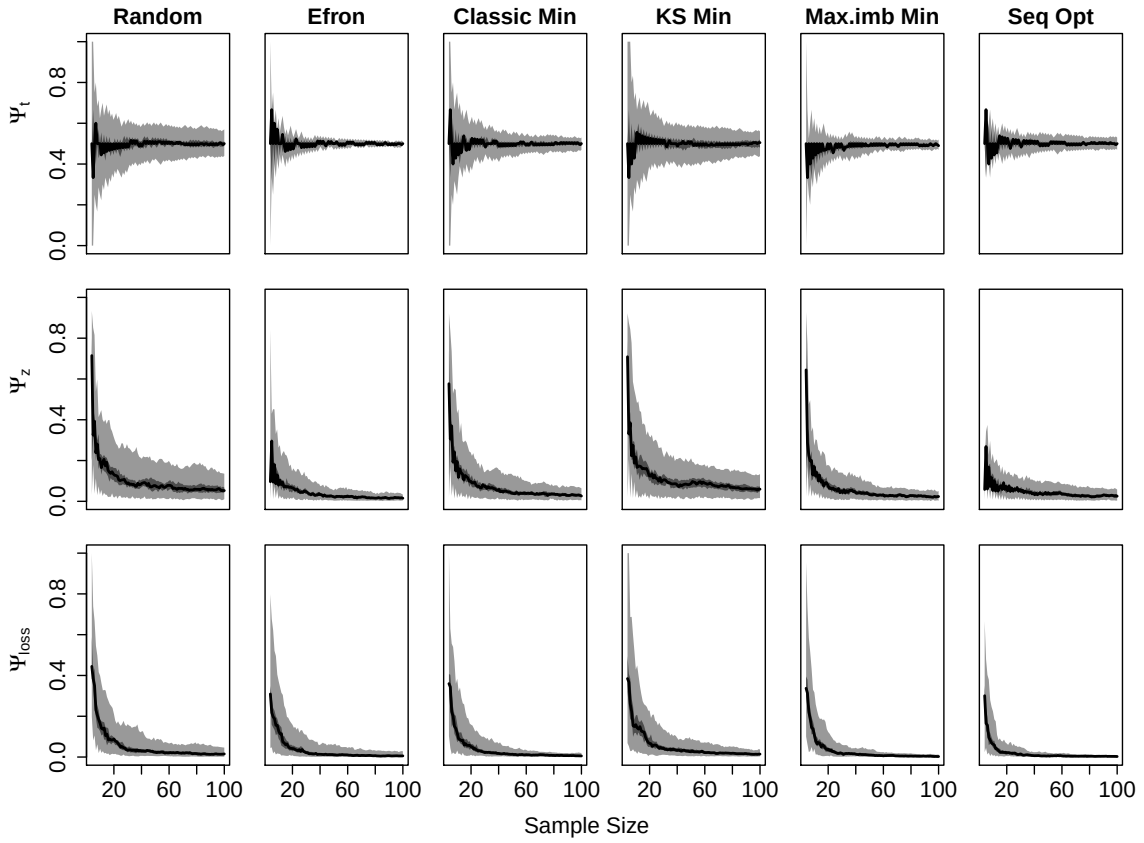

Figure C.3: Distributions of Ψ_t, Ψ_z and Ψ_{loss} as sample size increases for one Beta($\alpha = 5, \beta = 1$) covariate, based on 100 simulations. Six allocation methods are considered: Randomization, Efron's biased coin, the classic form of minimization, minimization using the K-S measure, minimization using the Max.imb measure, and the sequential optimal design method. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

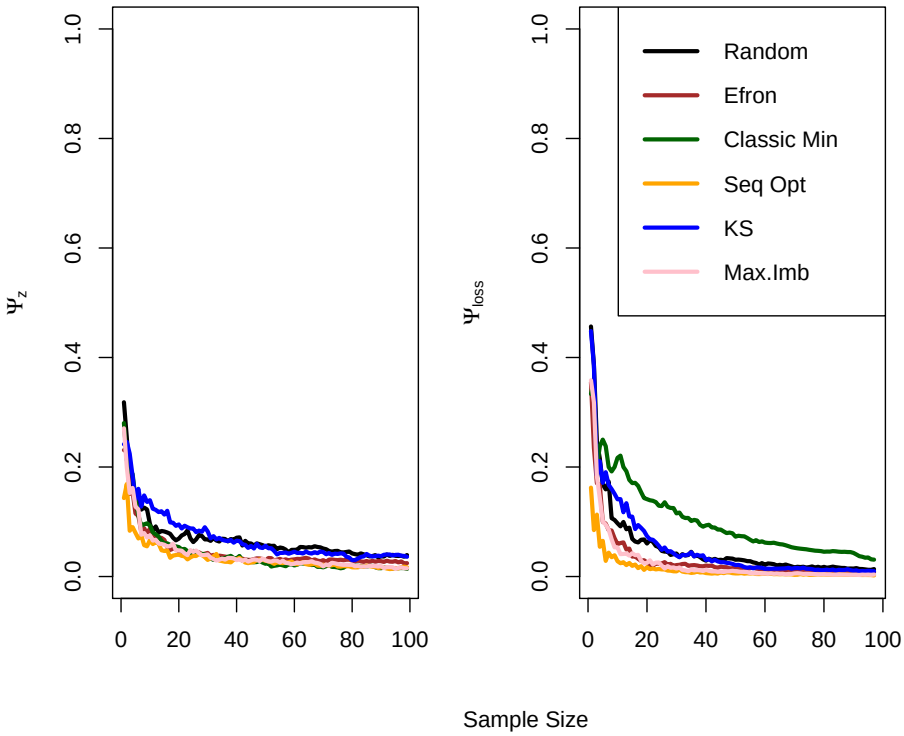

Figure C.4: Medians of the distributions of Ψ_z and Ψ_{loss} as sample size increases for one Beta($\alpha = 5, \beta = 1$) covariate, based on 100 simulations.Six allocation methods are considered: Randomization (black), Efron's biased coin (brown), the classic form of minimization (green), and the sequential optimal design method (yellow), the K-S form of minimization (blue) and the Max.imb form of minimization (pink).

Appendix D

Nonmyopic approach for linear model

D.1 Efficiencies for the one covariate case

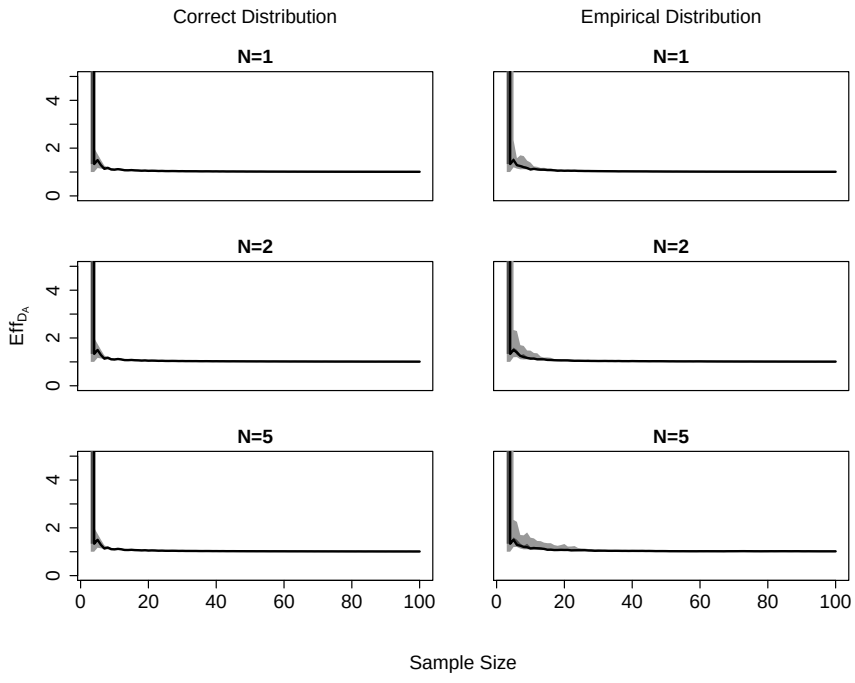

Figure D.1: Distributions of the efficiencies of nonmyopic D_A-optimal designs for the linear model from Figure 4.1 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

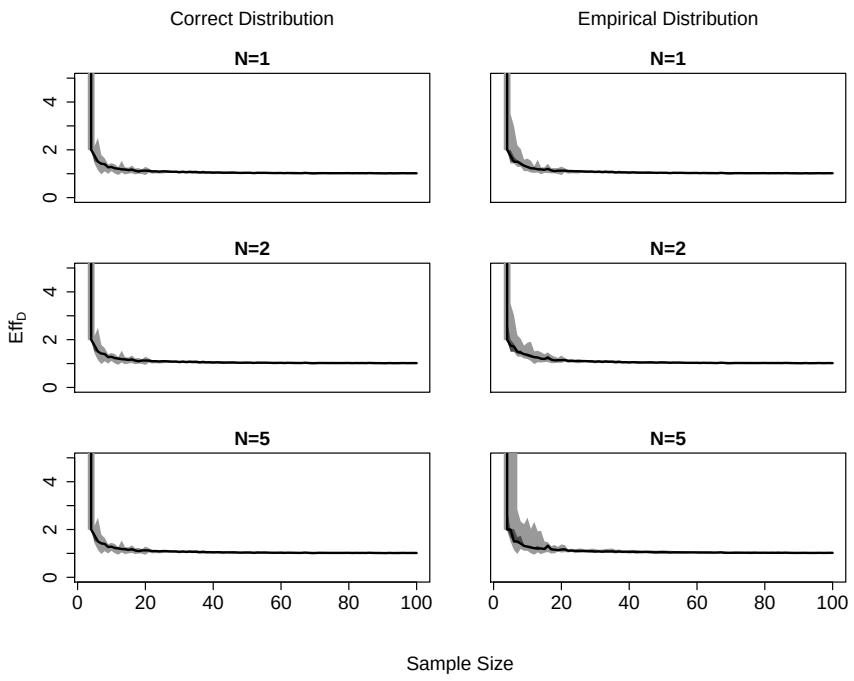

Figure D.2: Distributions of the efficiencies of nonmyopic D-optimal designs for the linear model from Figure 4.2 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

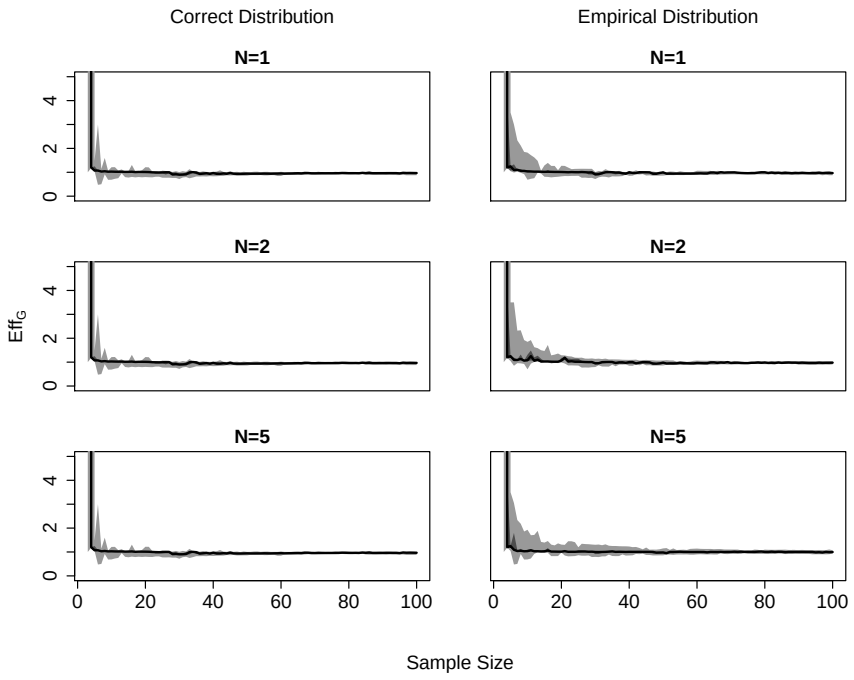

Figure D.3: Distributions of the efficiencies of nonmyopic G-optimal designs for the linear model from Figure 4.3 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

D.2 Results for two covariates

This simulation investigates the efficiency of the myopic and non-myopic approaches when two covariates are recorded in the experiment and included in the model. We follow the same structure for the simulation study as described in section 4.1.4, but we generate two independent covariates, both with Bernoulli(0.5) distribution.

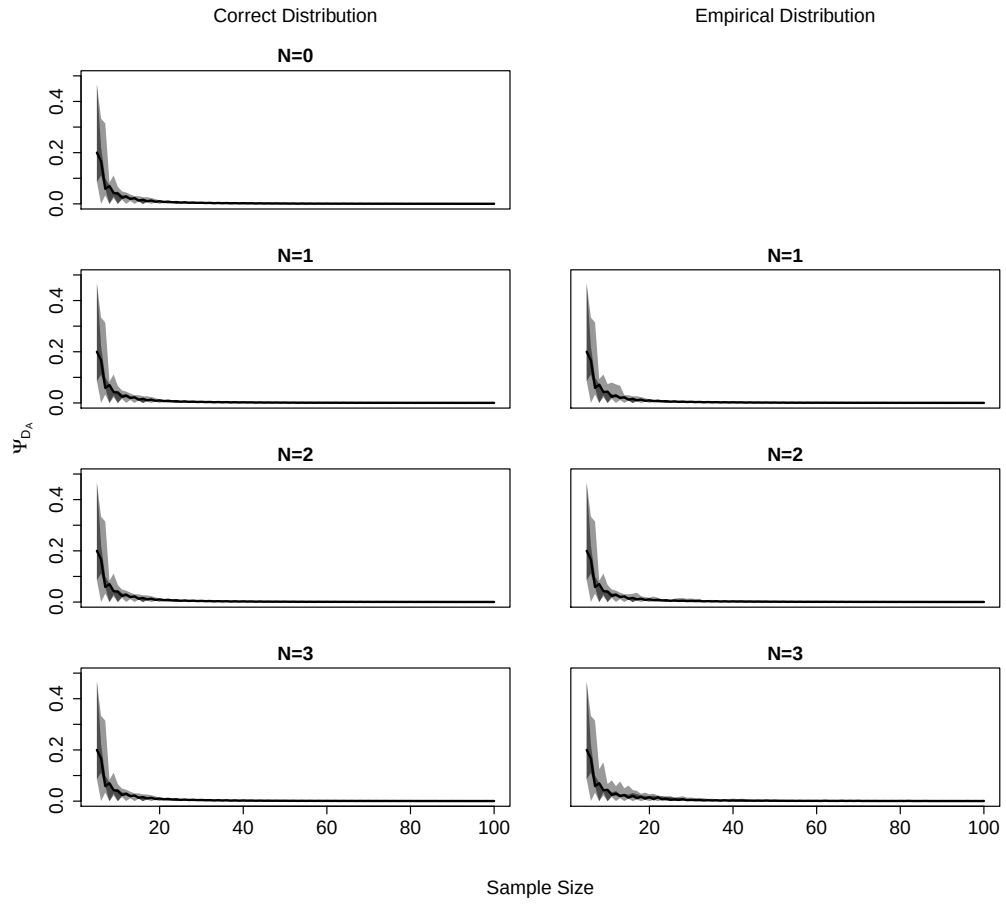

Figure D.4: Distributions of Ψ_{D_A} for designs for the linear model for two covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the non-myopic approach to constructing D_A-optimal designs with horizon $N = 1$, 2 and 3. For the non-myopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

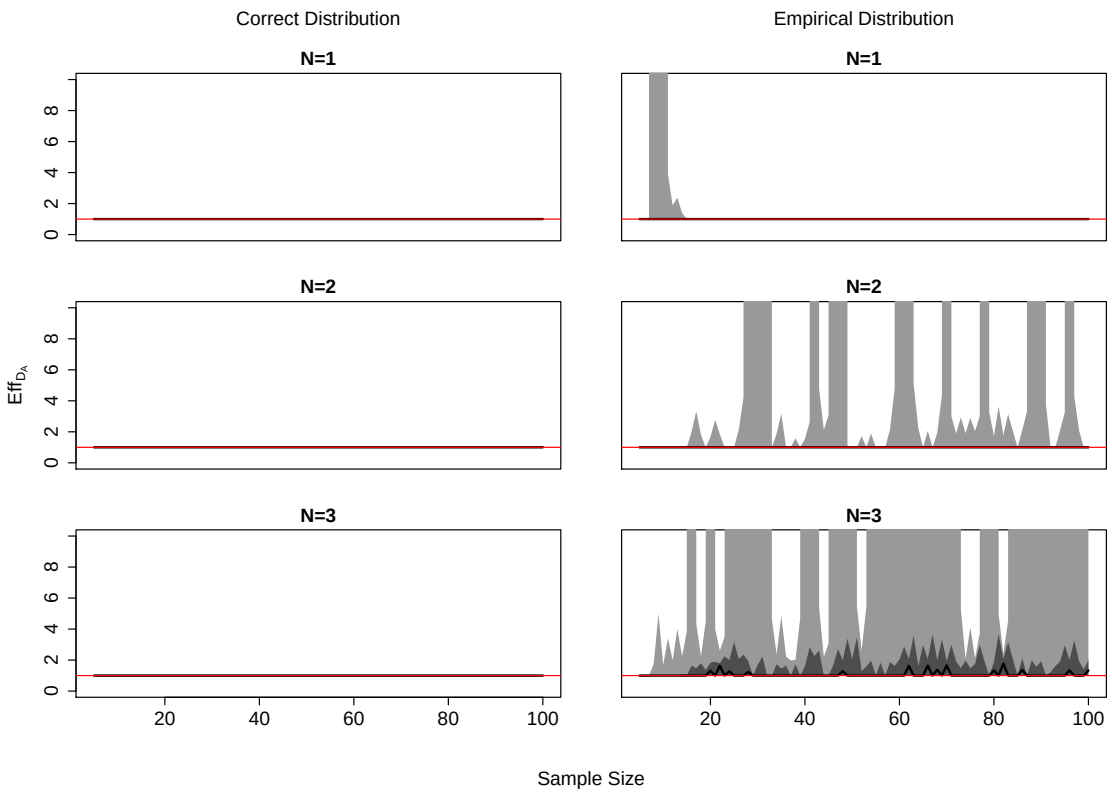

Figure D.5: Distributions of the efficiencies of nonmyopic D_A-optimal designs for the linear model for two covariates from Figure D.4 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

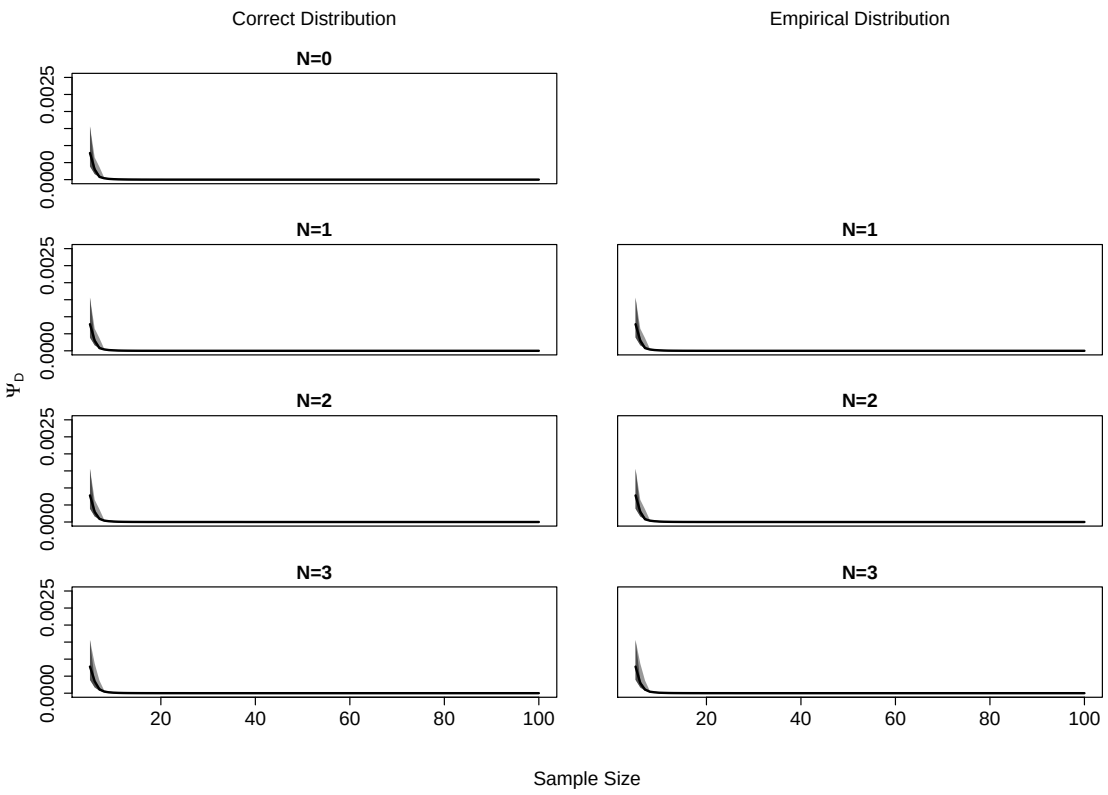

Figure D.6: Distributions of Ψ_D for designs for the linear model for two covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the non-myopic approach to constructing D-optimal designs with horizon $N = 1$, 2 and 3. For the non-myopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

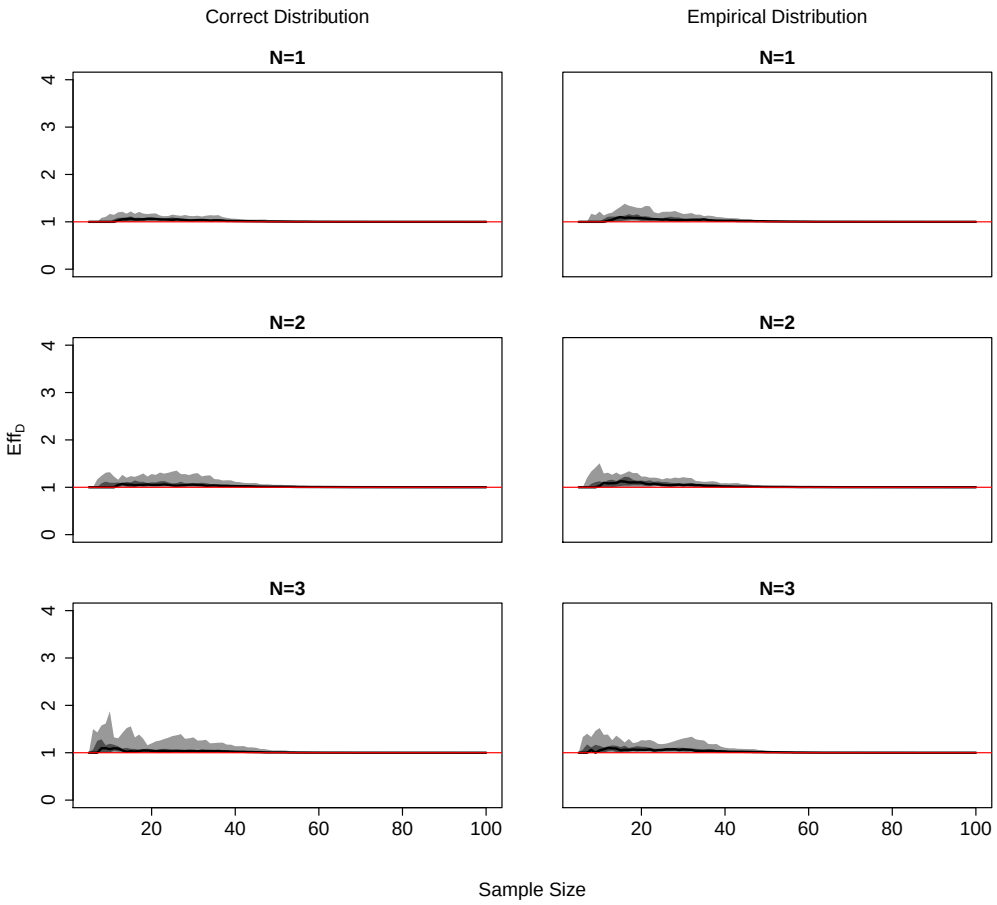

Figure D.7: Distributions of the efficiencies of nonmyopic D-optimal designs for the linear model for two covariates from Figure D.6 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

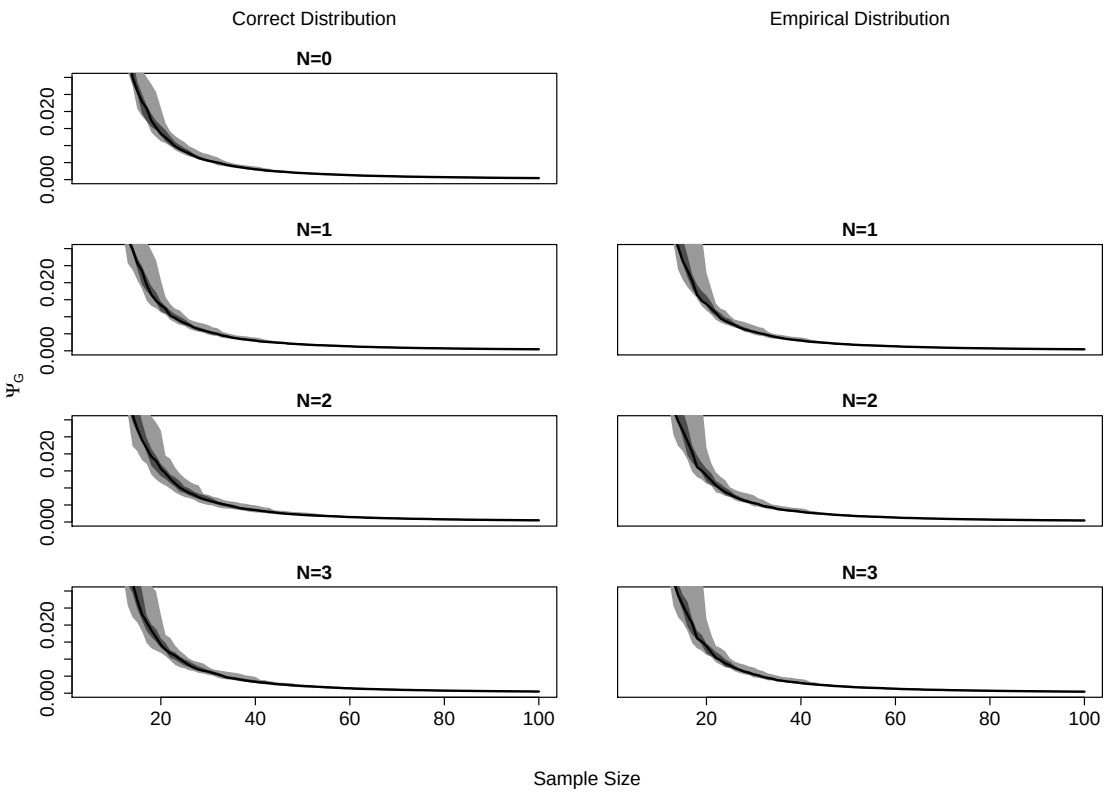

Figure D.8: Distributions of Ψ_G for designs for the linear model for two covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the non-myopic approach to constructing G-optimal designs with horizon $N = 1$, 2 and 3. For the non-myopic approach, we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used.

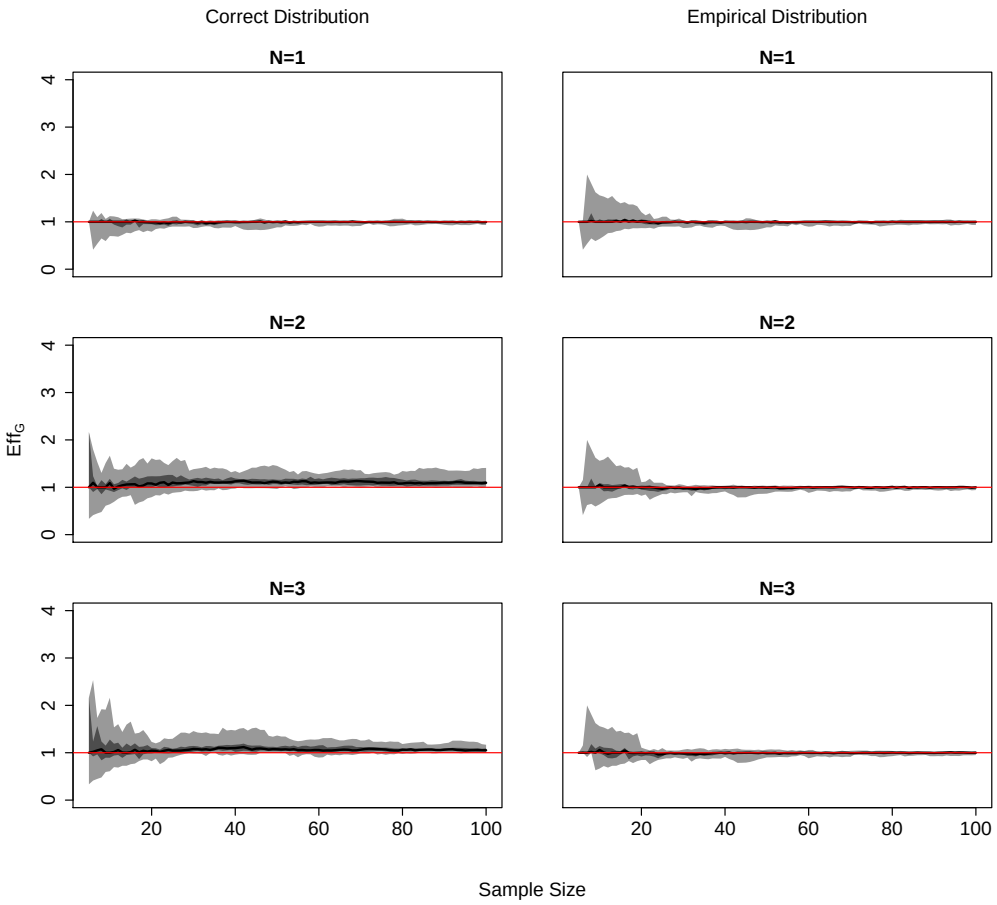

Figure D.9: Distributions of the efficiencies of nonmyopic G-optimal designs for the linear model for two covariates from Figure D.8 are shown, relative to the myopic design. Values of the efficiency under 1 indicate that the non-myopic design is more efficient than the myopic.

Appendix E

Nonmyopic approach for the logistic model

E.1 One dynamic covariate

We compare myopic and non-myopic designs for logistic regression where we have one dynamic covariate in this simulation. We generate the covariate $\boldsymbol{z}$ from a Bernoulli(p_i) distribution, where p_i is the ith element of $\boldsymbol{p}$:

$$\boldsymbol{p} = \left[\underbrace{0.1, ..., 0.1}_{25 \text{ times}}, \underbrace{0.9, ..., 0.9}_{25 \text{ times}}, \underbrace{0.1, ..., 0.1}_{25 \text{ times}}, \underbrace{0.9, ..., 0.9}_{25 \text{ times}}\right]. \tag{E.1}$$

All other settings for the simulations are as stated in Section 4.2.5.

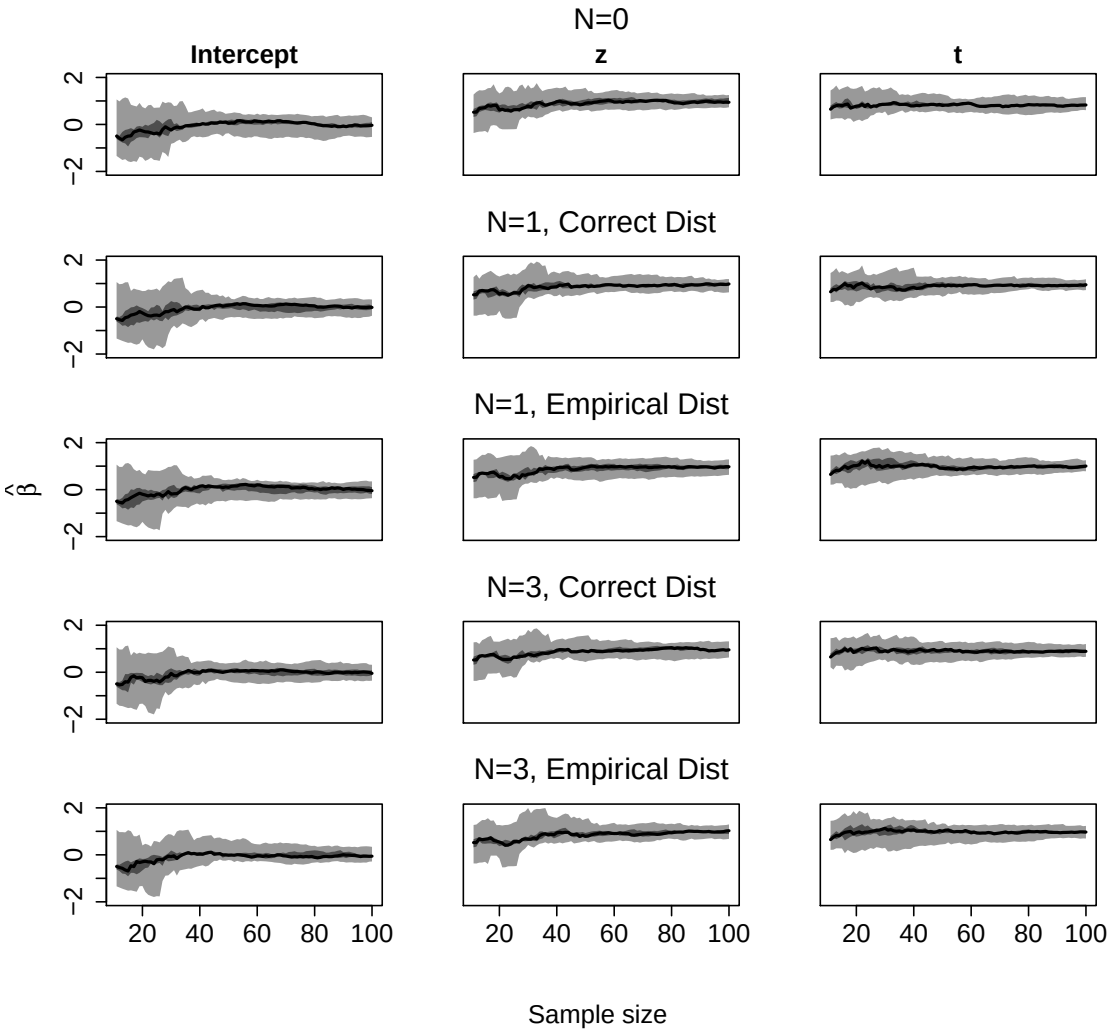

Figure E.1: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

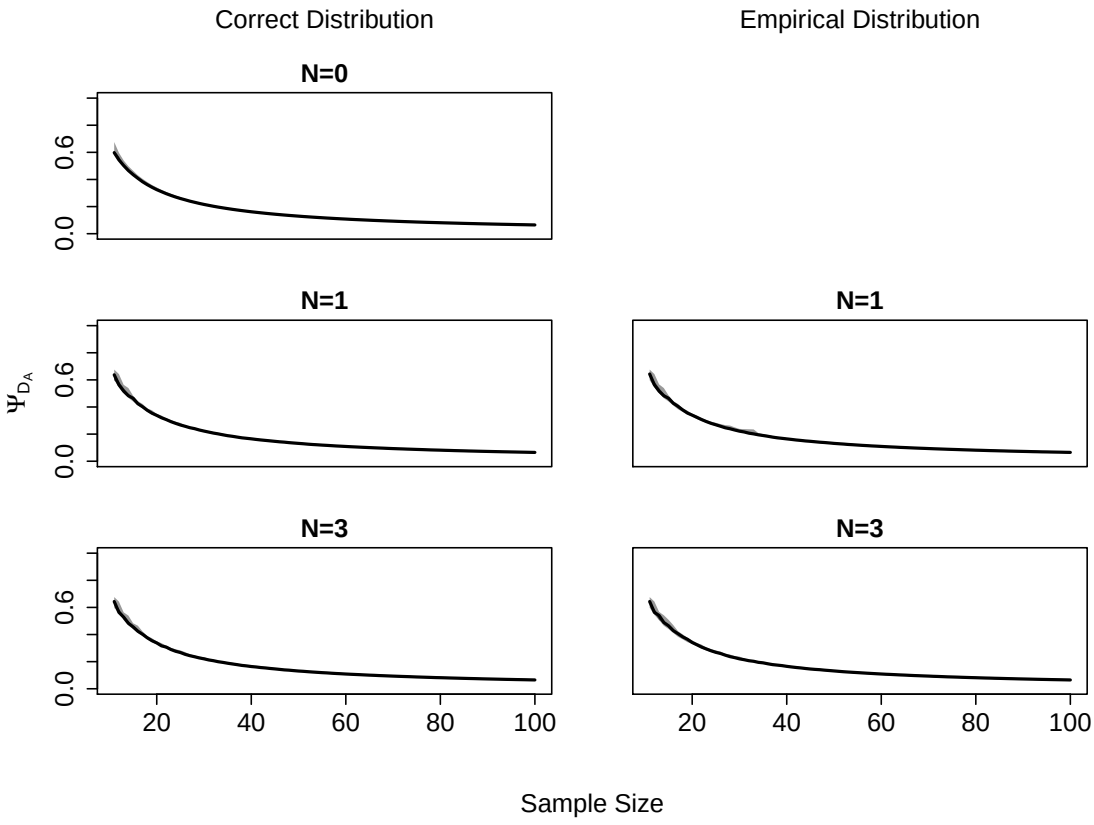

Figure E.2: Distributions of Ψ_{D_A} for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

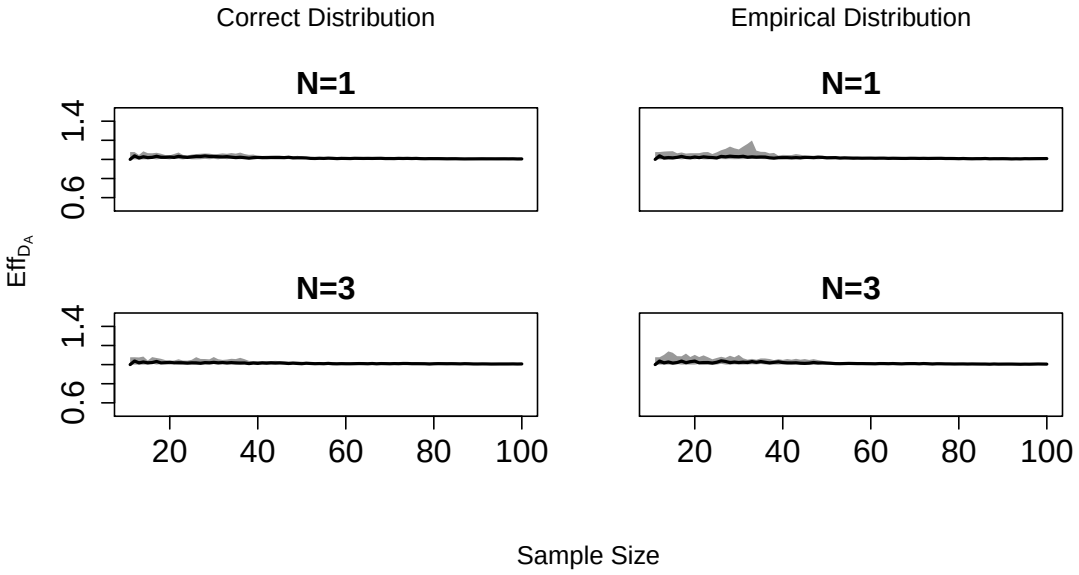

Figure E.3: Distributions of the efficiencies of the nonmyopic D_A-optimal designs against the myopic D_A-optimal designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

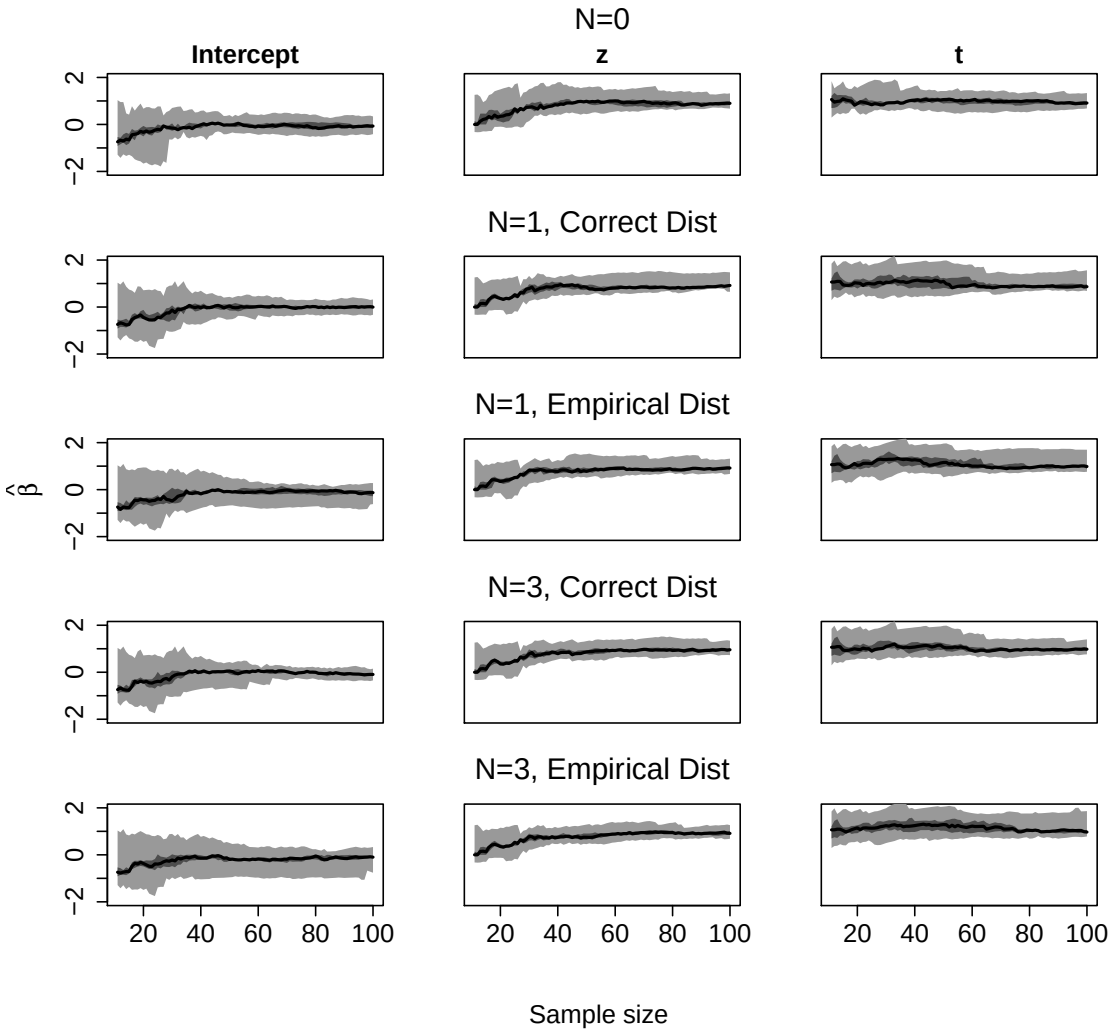

Figure E.4: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

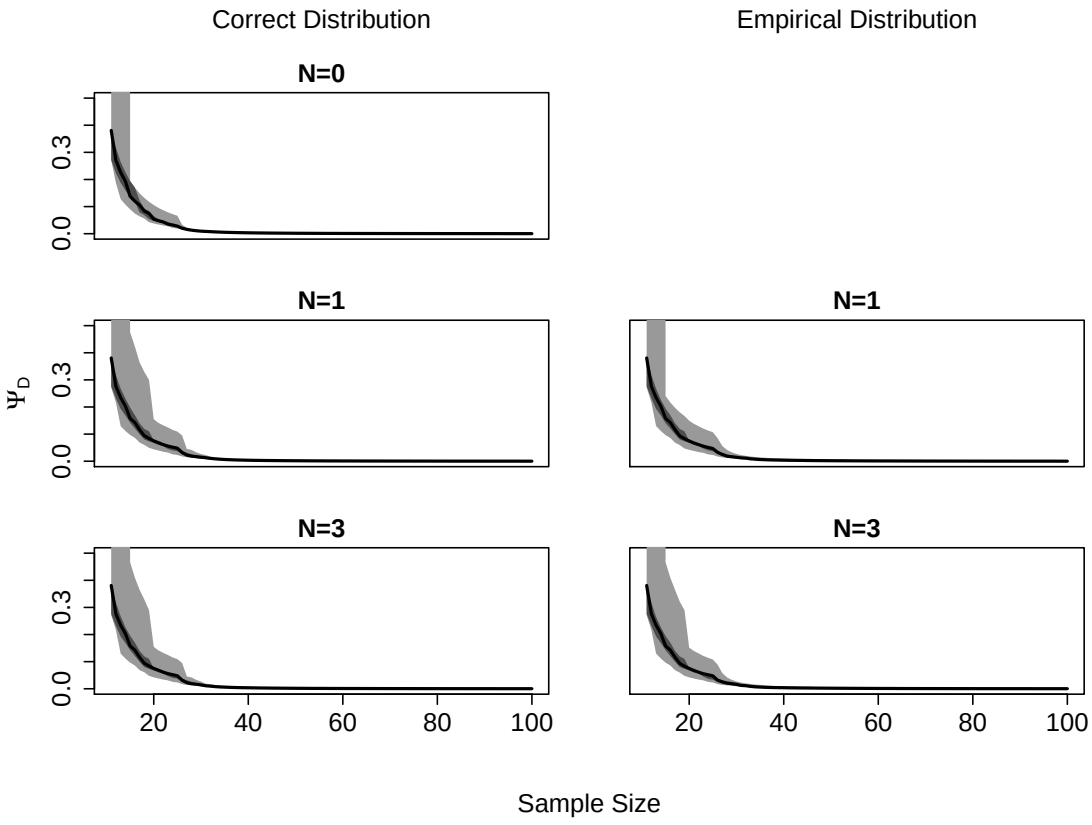

Figure E.5: Distributions of Ψ_D for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

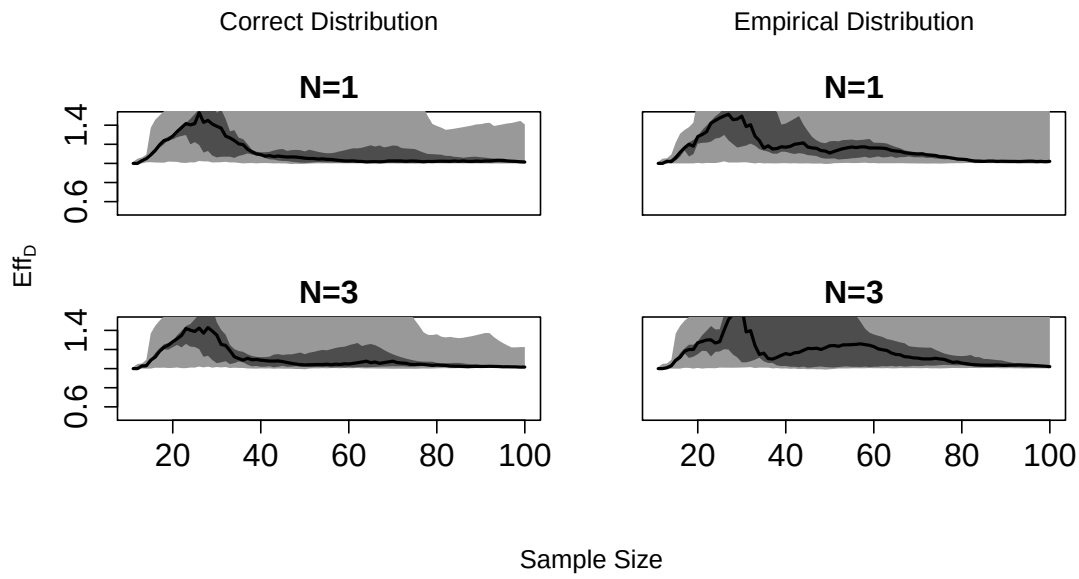

Figure E.6: Distributions of the efficiencies of the nonmyopic D-optimal designs against the myopic D-optimal designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

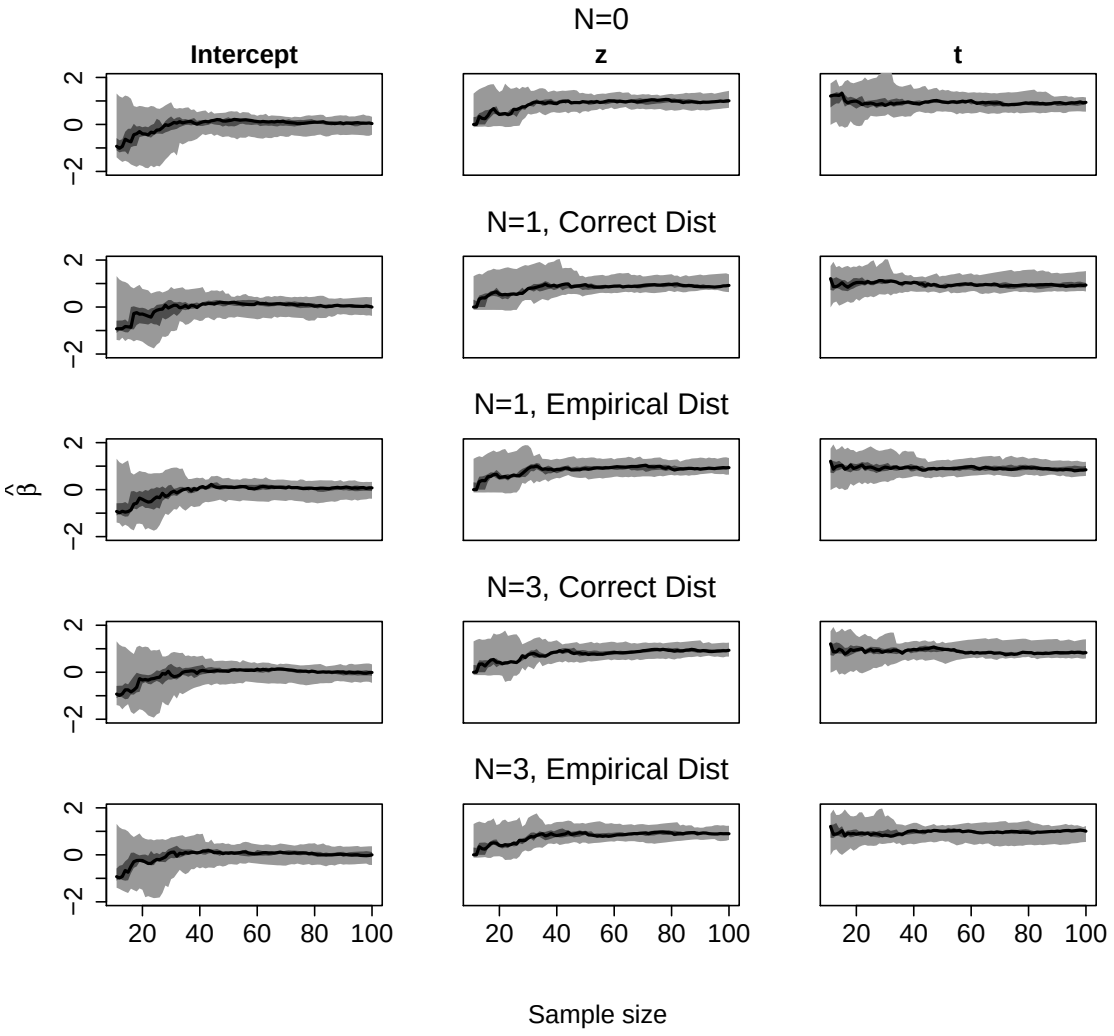

Figure E.7: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

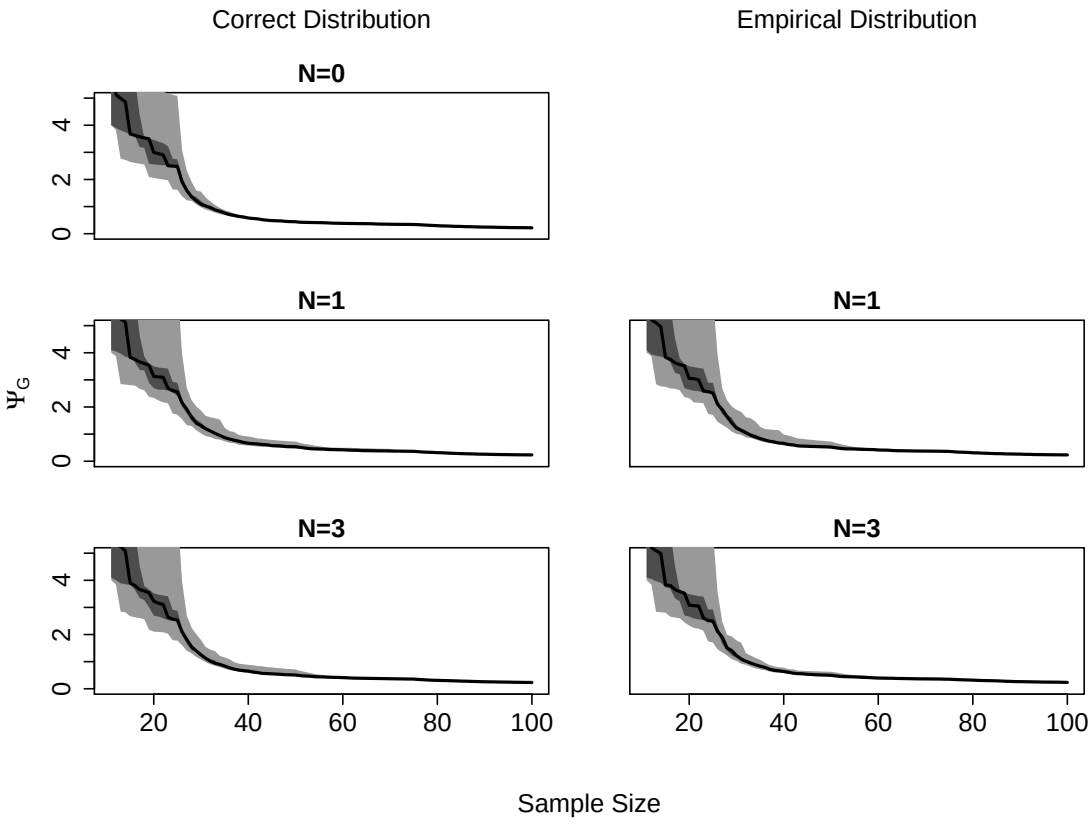

Figure E.8: Distributions of Ψ_G for designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

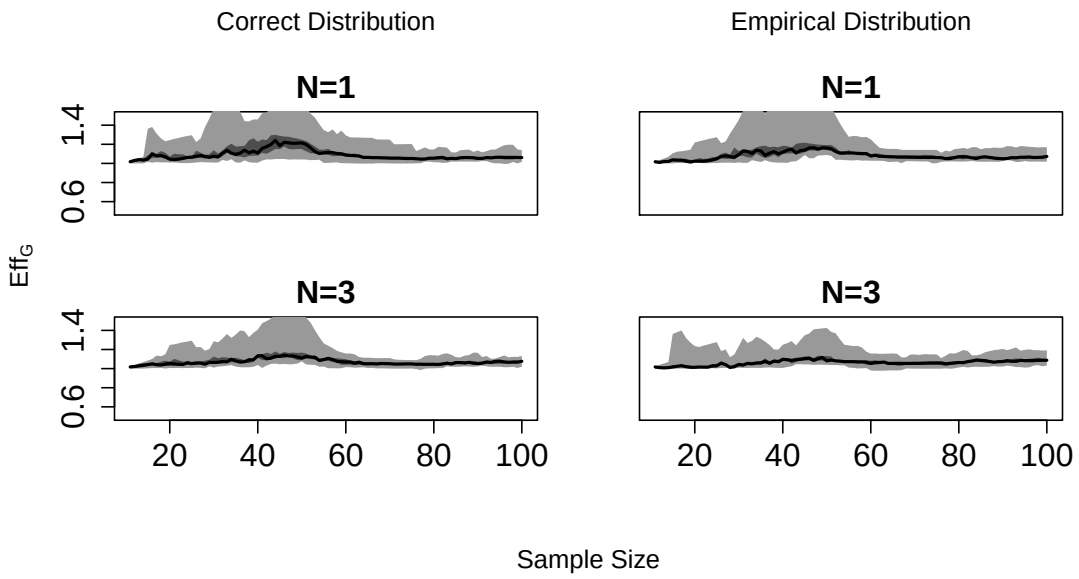

Figure E.9: Distributions of the efficiencies of the nonmyopic G-optimal designs against the myopic G-optimal designs for the logistic model for one dynamic covariate are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

E.2 Two covariates

In this simulation, we compare myopic and non-myopic designs for logistic regression where we have two static covariates and interactions between treatments and covariates are assumed. The model is given by $\boldsymbol{y} = \boldsymbol{\beta}\left(1 \;\; \boldsymbol{z}_{\cdot 1} \;\; \boldsymbol{z}_{\cdot 2} \;\; \boldsymbol{t} \;\; \boldsymbol{z}_{\cdot 1}\boldsymbol{t} \;\; \boldsymbol{z}_{\cdot 2}\boldsymbol{t}\right)$ where the true values of the parameters are $\boldsymbol{\beta} = \left(1, 2, 1, -2, 5, -4\right)^{\intercal}$. The first covariate is generated such that $\mathbb{P}(z_{1,i} = 1) = 0.5$ and the second covariate is generated such that $\mathbb{P}(z_{i,2} = 1) = 0.3$ for all i. All other settings for the simulations are as stated in Section 4.2.5.

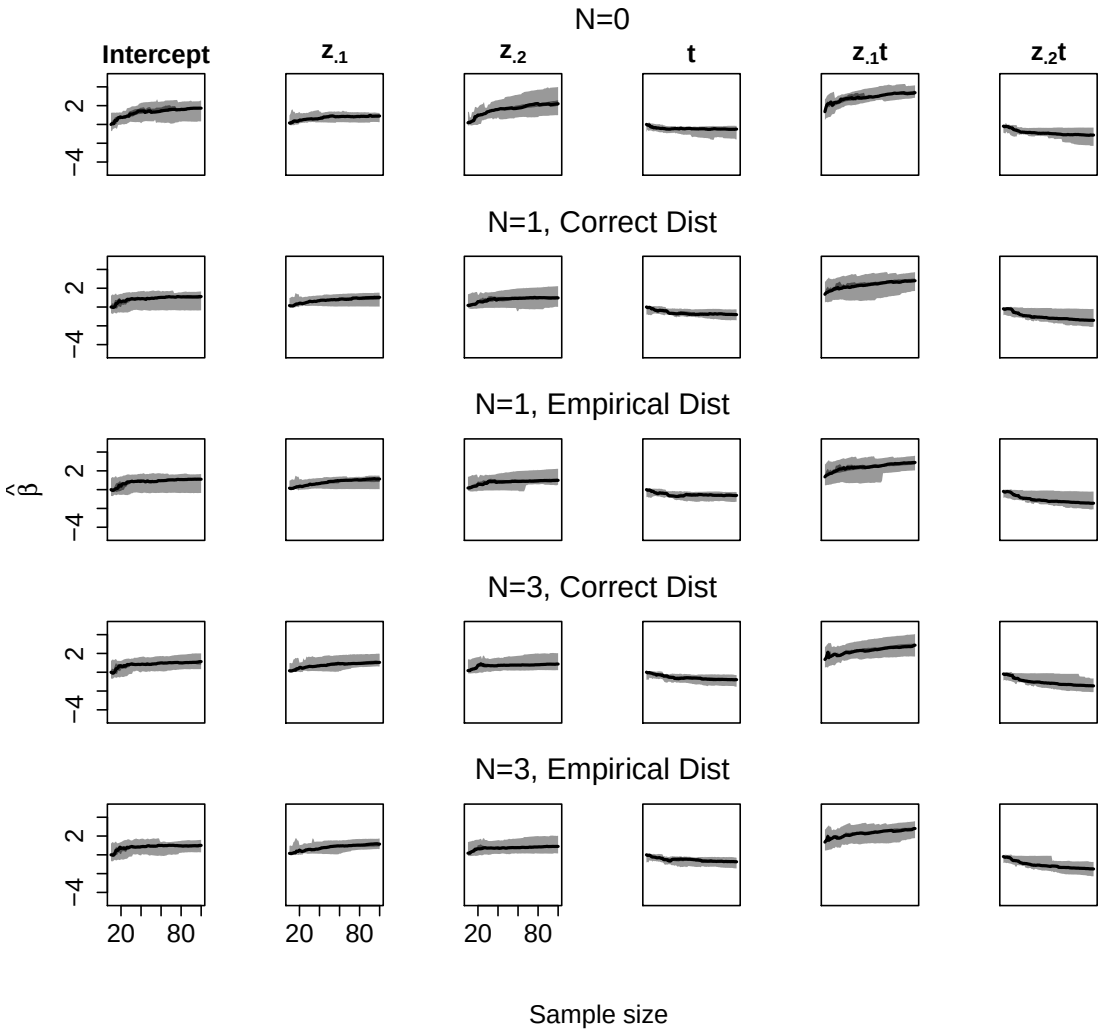

Figure E.10: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

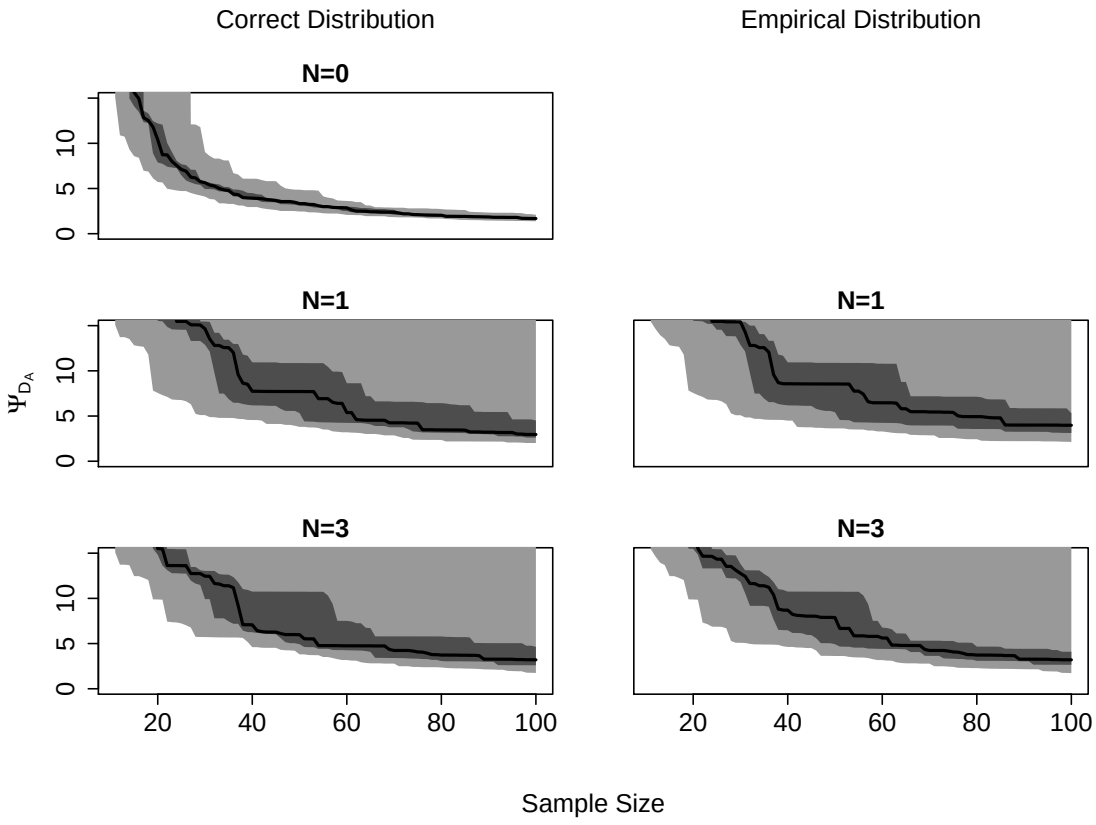

Figure E.11: Distributions of Ψ_{D_A} for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D_A-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

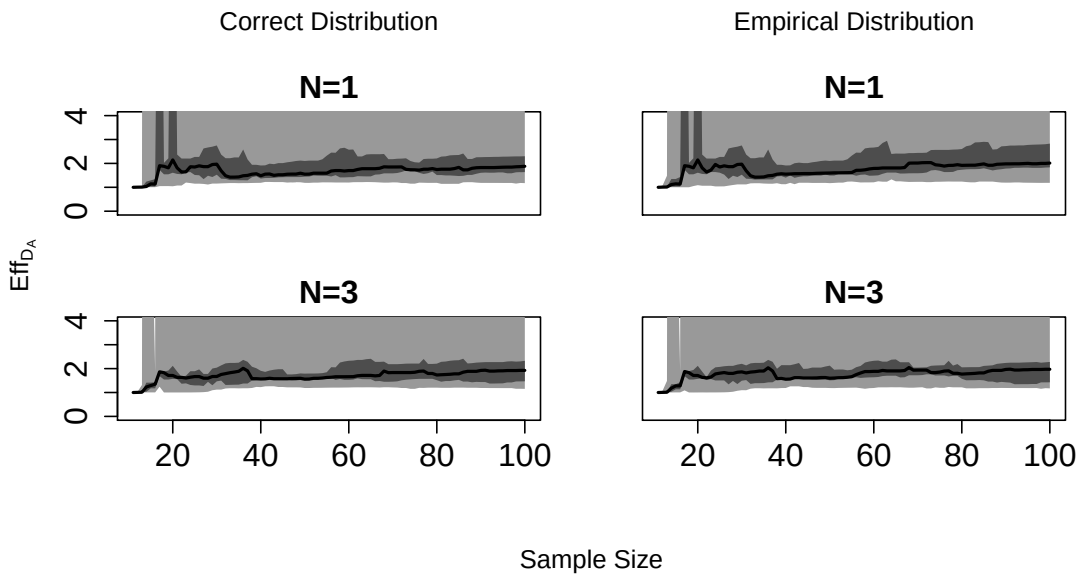

Figure E.12: Distributions of the efficiencies of the nonmyopic D_A-optimal designs against the myopic D_A-optimal designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

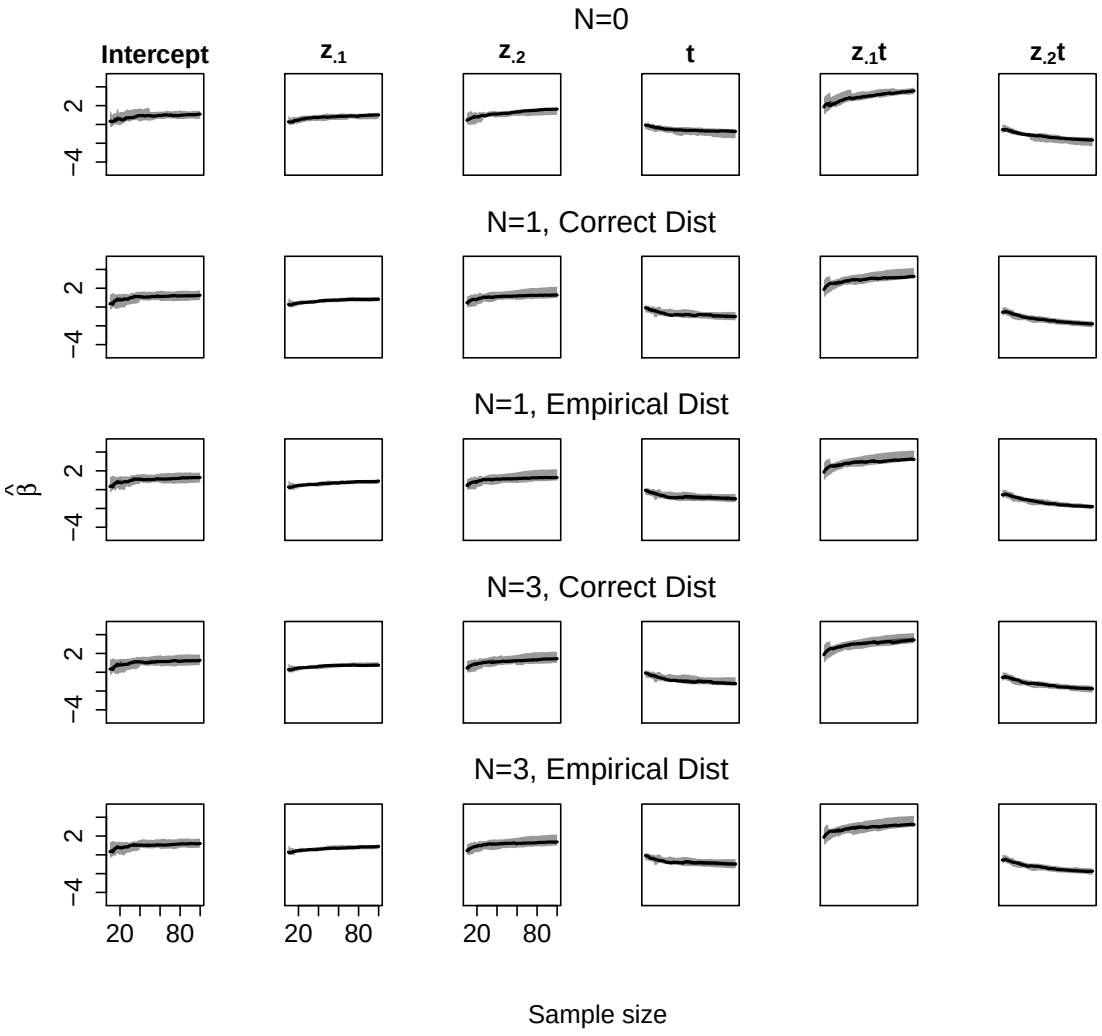

Figure E.13: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

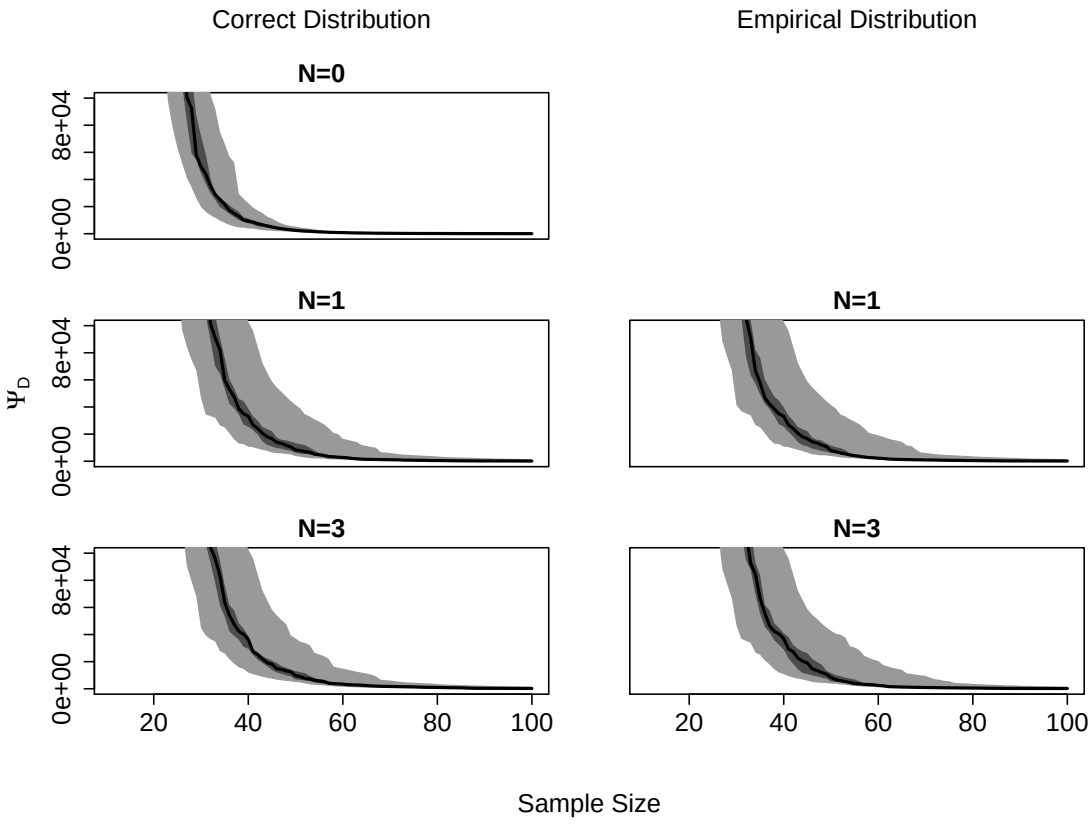

Figure E.14: Distributions of Ψ_D for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing D-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

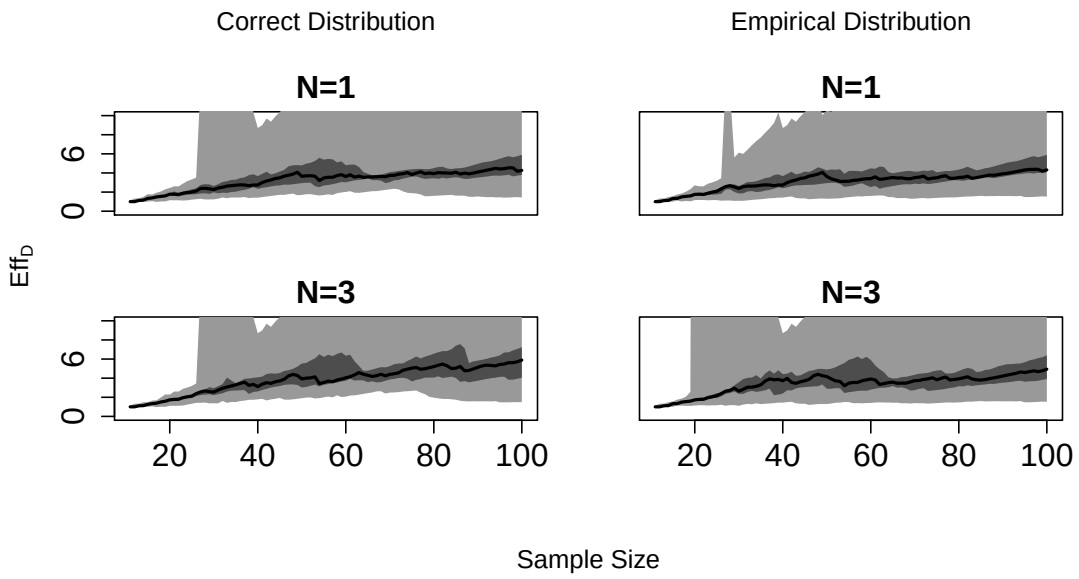

Figure E.15: Distributions of the efficiencies of the nonmyopic D-optimal designs against the myopic D-optimal designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1 and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

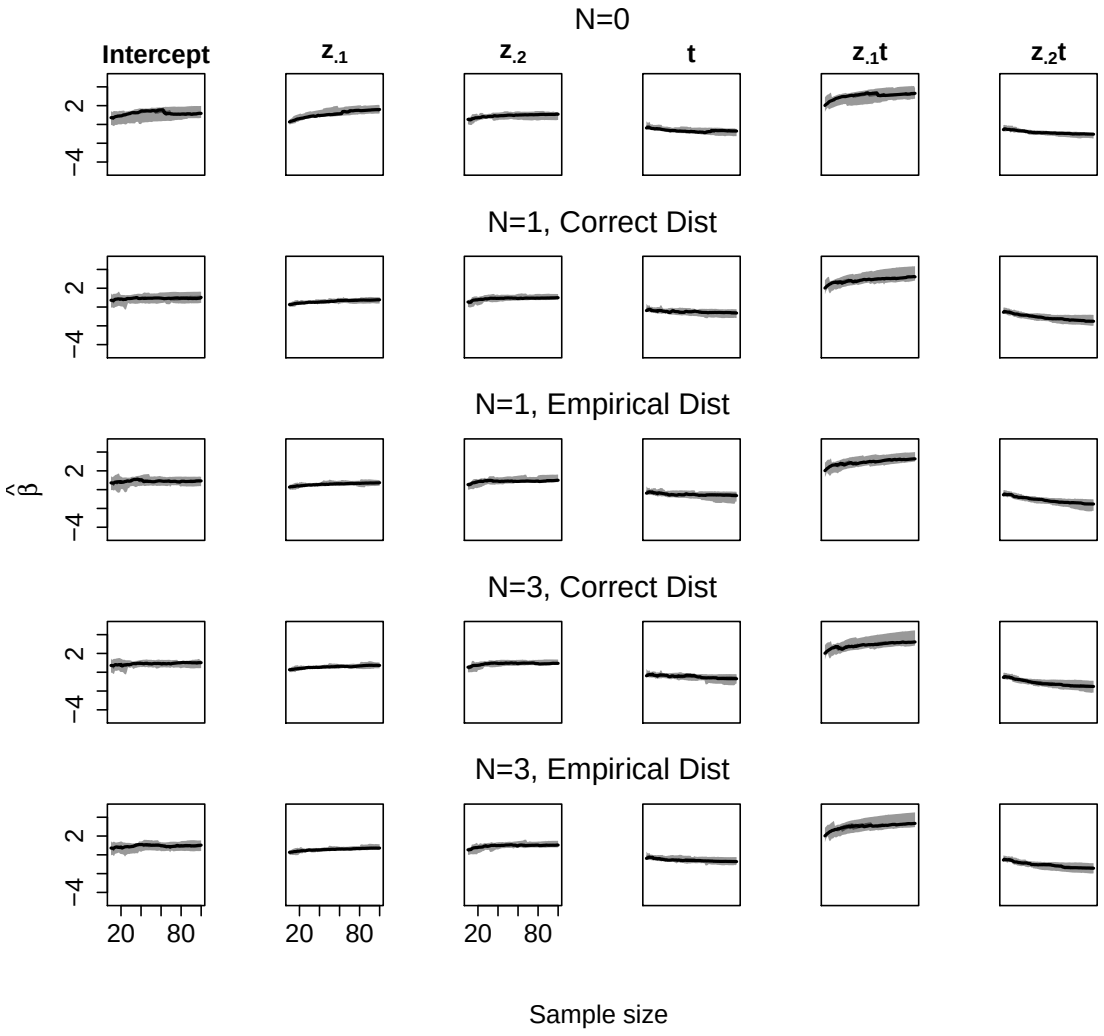

Figure E.16: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

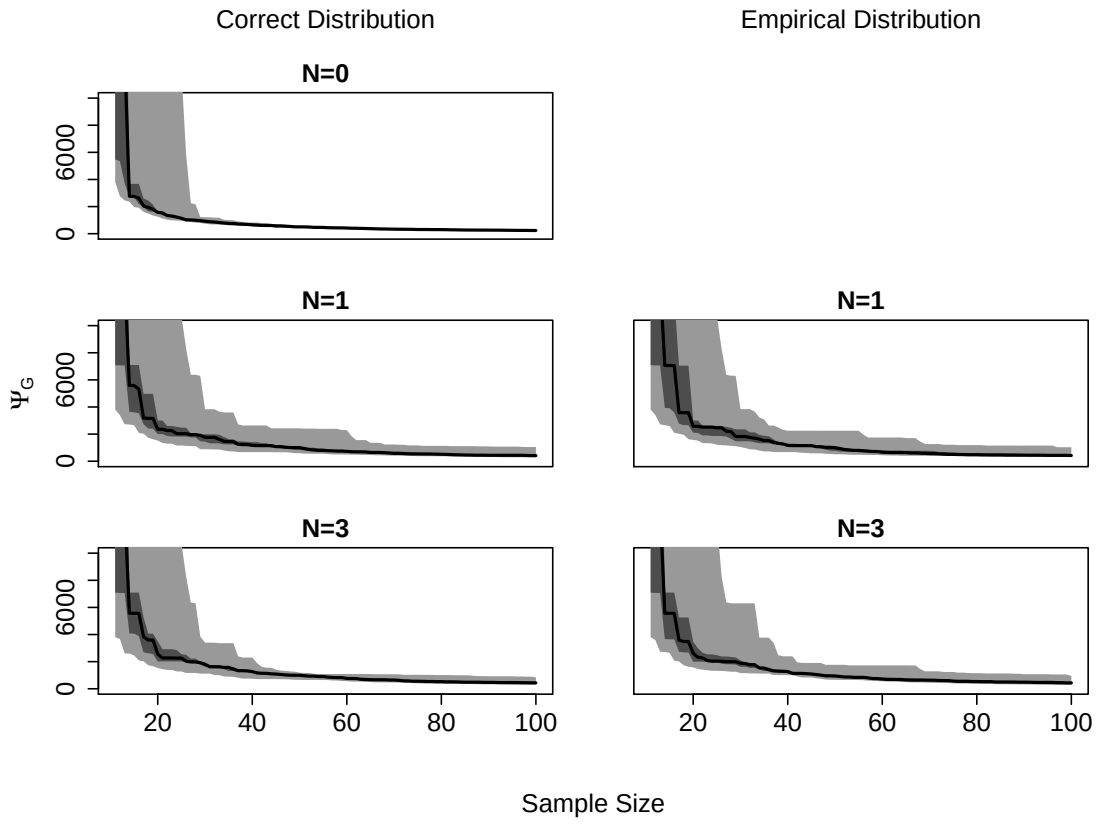

Figure E.17: Distributions of Ψ_G for designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the myopic approach ($N = 0$), as well as the nonmyopic approach to constructing G-optimal designs with horizon $N = 1$ and 3. For the nonmyopic approach, we consider both the case where the correct covariate distribution is known (left panel), and when it is unknown so the empirical covariate distribution is used (right panel).

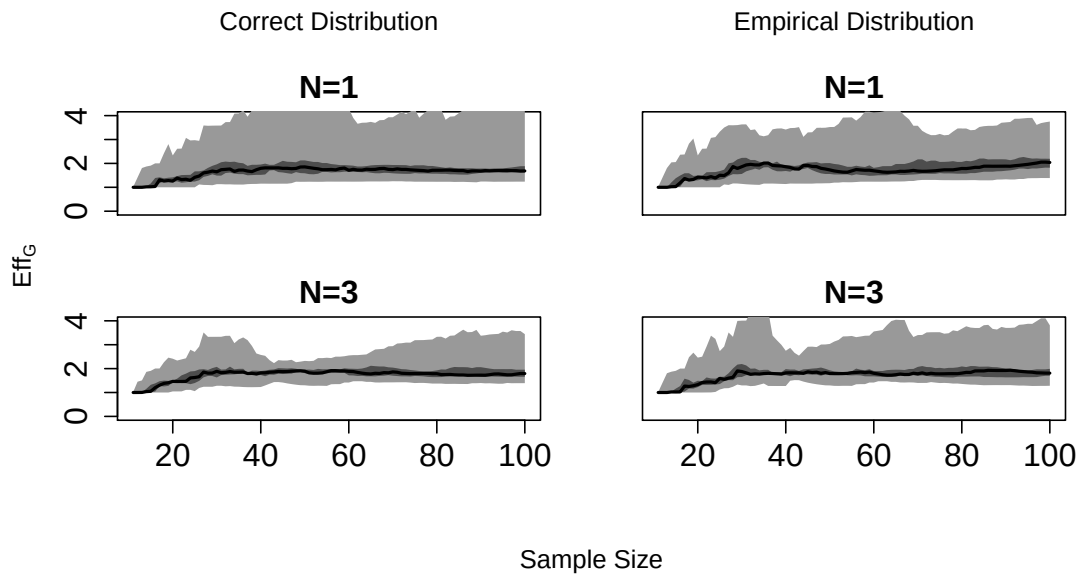

Figure E.18: Distributions of the efficiencies of the nonmyopic G-optimal designs against the myopic G-optimal designs for the logistic model for two static covariates are plotted against sample size, based on 20 simulations. We consider the efficiencies of the non-myopic approach with horizons 1, 2, and 3, with the correct and empirical distributions, against the myopic approach as the baseline.

Appendix F

Simulations for the weighted L-optimal designs

F.1 Effective Treatment

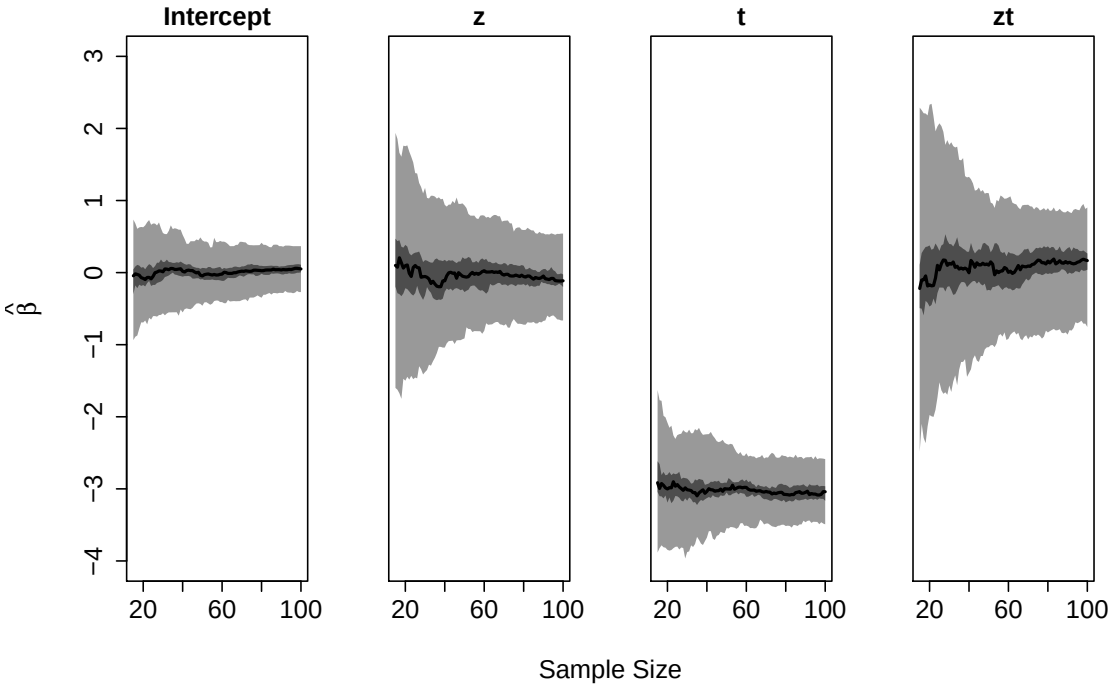

Figure F.1: Distribution of $\hat{\boldsymbol{\beta}}$ vs sample size for the weighted L-optimal design for the linear model case with effective treatment. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

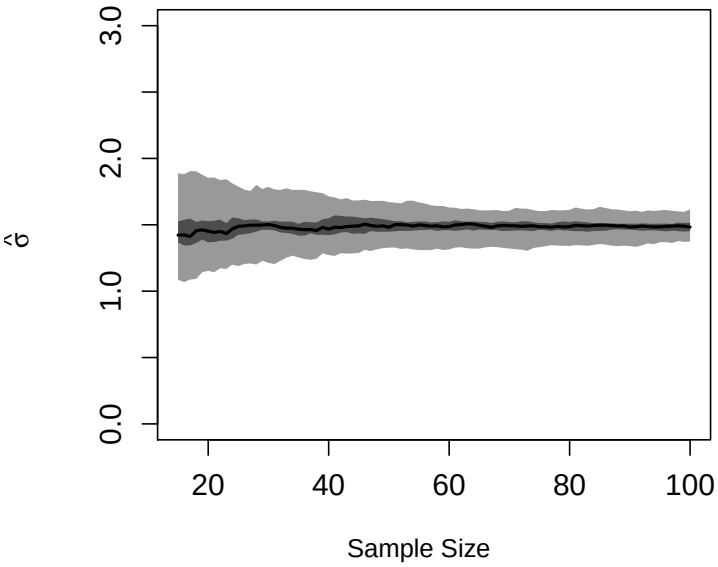

Figure F.2: Distribution of $\hat{\sigma}$ vs sample size for the weighted L-optimal design for the linear model case with effective treatment.

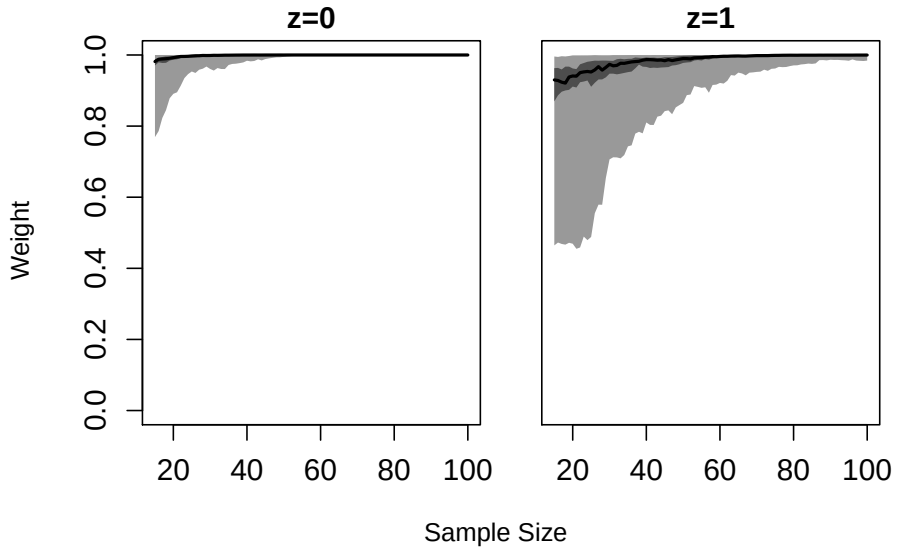

Figure F.3: Distribution of w_r vs sample size for subgroups $(z = 0)$ and $(z = 1)$ for the weighted L-optimal design for the linear model case with effective treatment.

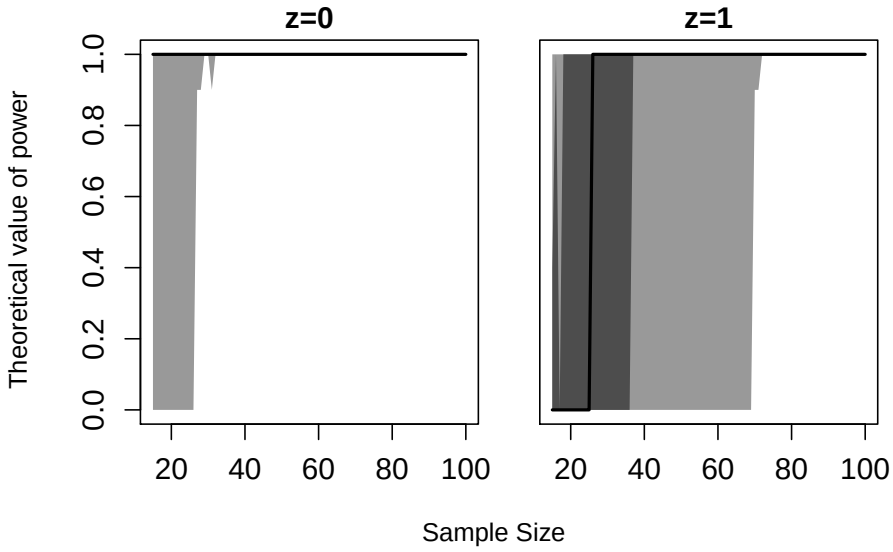

Figure F.4: Distribution of theoretical values of the the power of the hypothesis test vs sample size for subgroups ($z = 0$) and ($z = 1$) for the weighted L-optimal design for the linear model case with effective treatment.

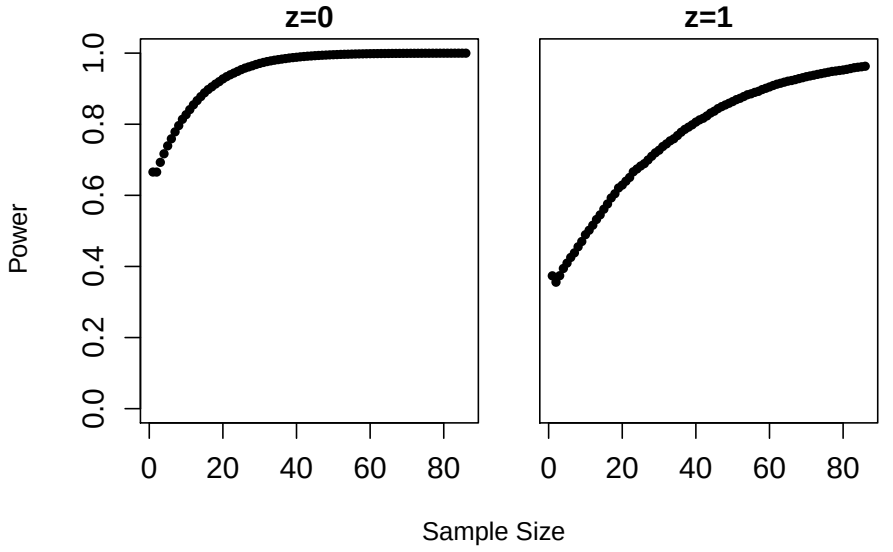

Figure F.5: Empirical values of the the power of the hypothesis test vs sample size for subgroups ($z = 0$) and ($z = 1$). Each point is the proportion of simulations in which the null hypothesis is correctly rejected.

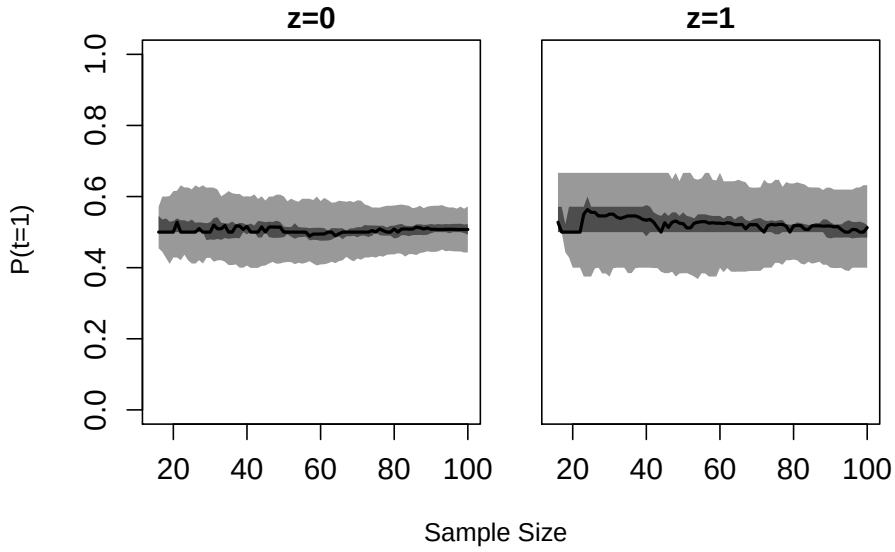

Figure F.6: Distribution of the proportion of new treatment in each subgroup vs total sample size for the weighted L-optimal design for the linear model case with effective treatment.

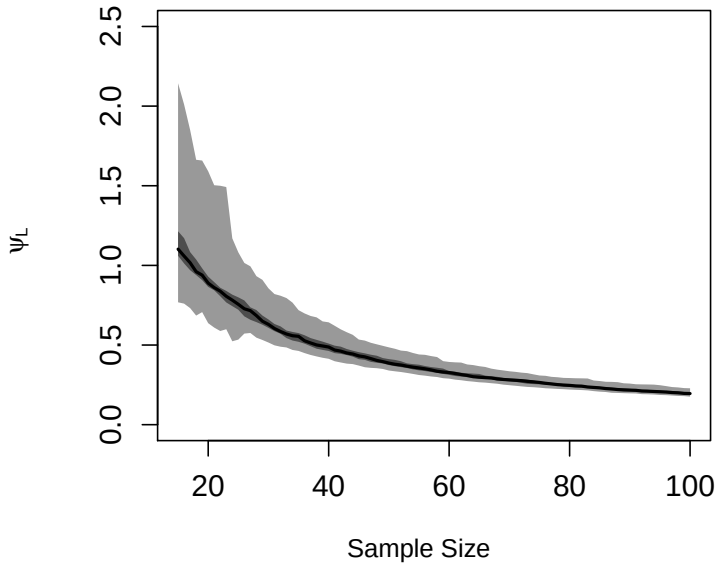

Figure F.7: Distribution of Ψ_L vs sample size for the weighted L-optimal design for the linear model case with effective treatment.

F.2 Non-zero Interaction

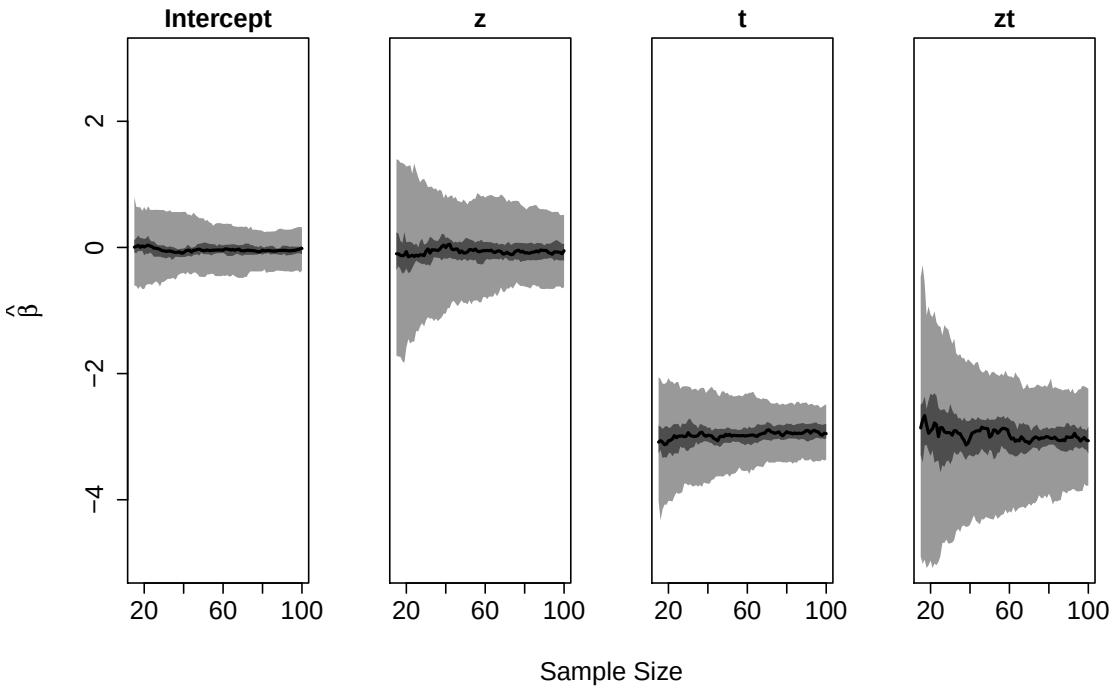

Figure F.8: Distribution of $\hat{\beta}$ vs sample size for the weighted L-optimal design for the linear model case with non-zero interaction.

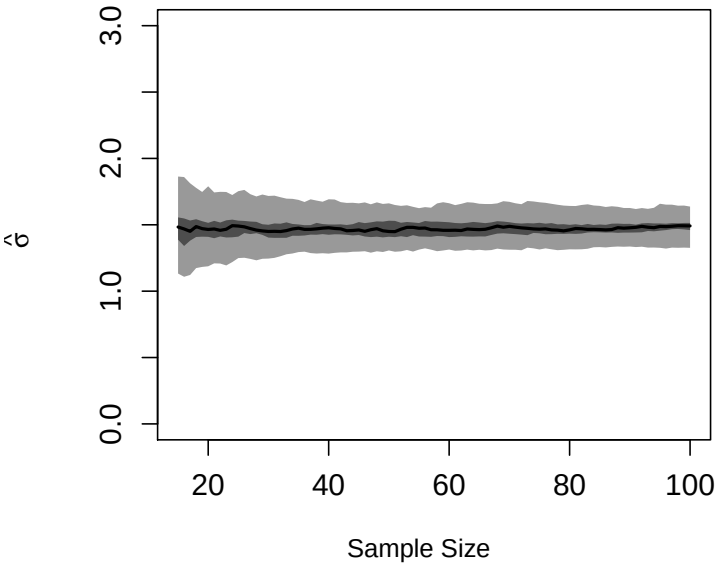

Figure F.9: Distribution of $\hat{\sigma}$ vs sample size for the weighted L-optimal design for the linear model case with non-zero interaction.

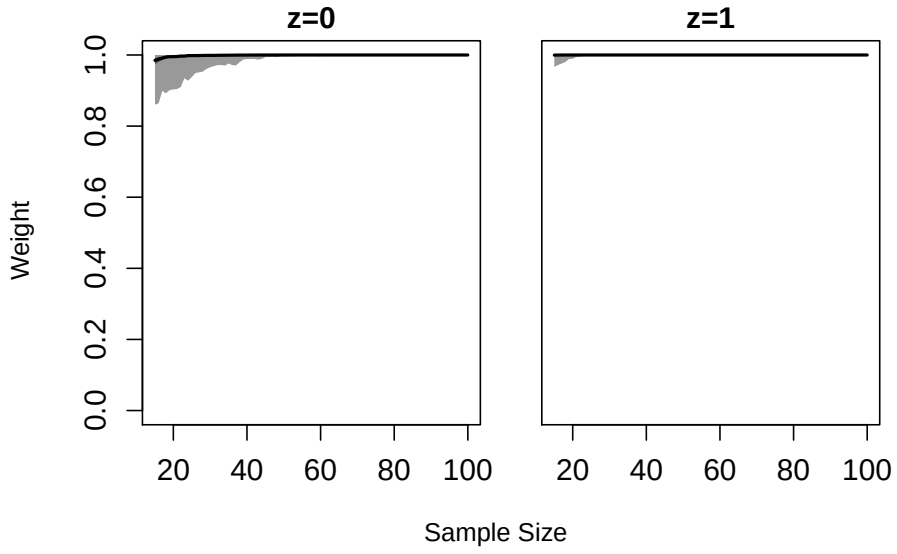

Figure F.10: Distribution of w_r vs sample size for subgroups $(z = 0)$ and $(z = 1)$ for the weighted L-optimal design for the linear model case with non-zero interaction.

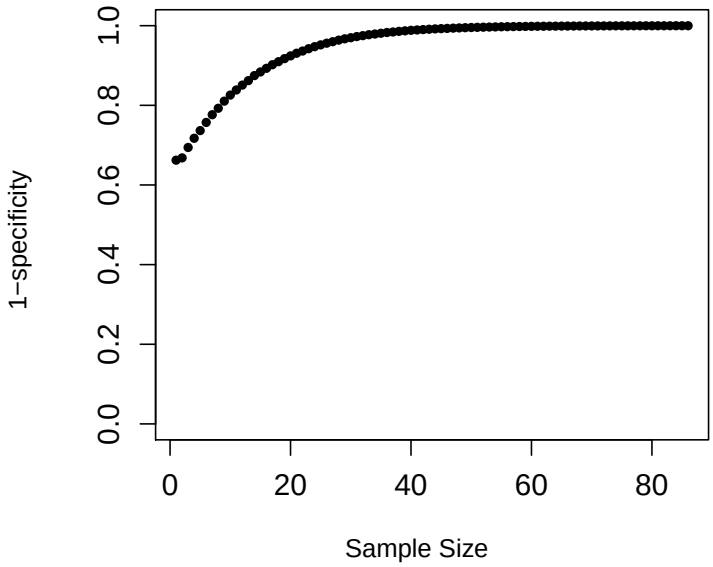

Figure F.11: 1-specificity for the hypothesis test vs sample size for subgroup ($z = 0$) for the weighted L-optimal design for the linear model case with non-zero interaction. The proportion of correct null hypotheses that have been rejected is given for each sample size.

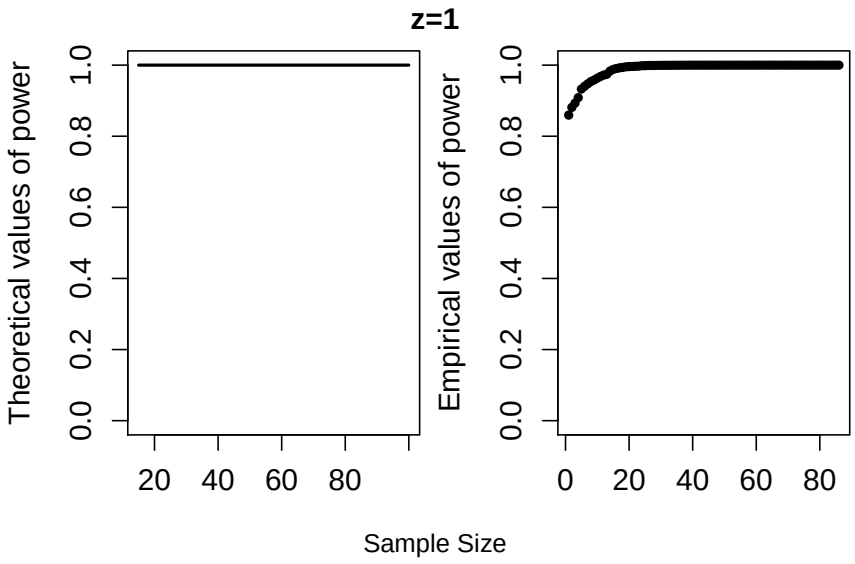

Figure F.12: Distribution of theoretical values of the the power of the hypothesis test vs sample size (left panel) and empirical values of the the power of the hypothesis test vs sample size (right panel) for the weighted L-optimal design for the linear model case with non-zero interaction.

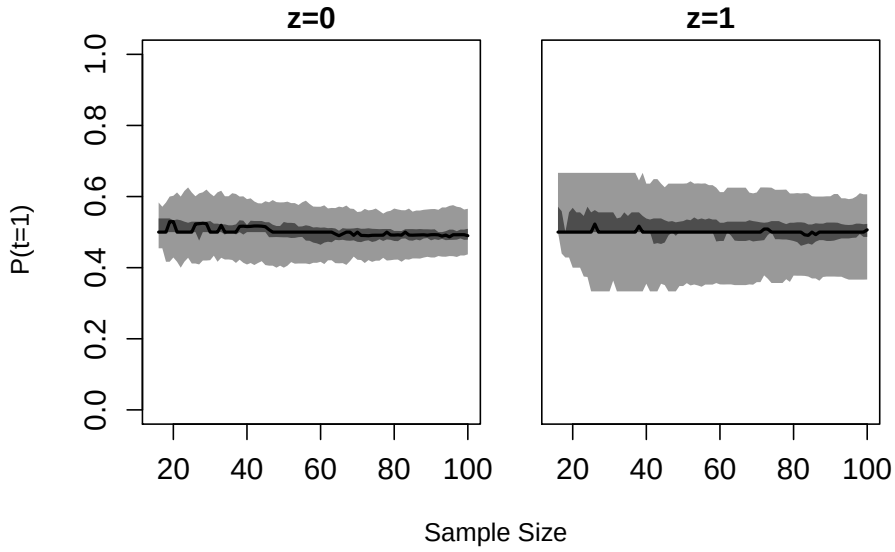

Figure F.13: Distribution of the proportion of new treatment in each subgroup vs total sample size for the weighted L-optimal design for the linear model case with non-zero interaction.

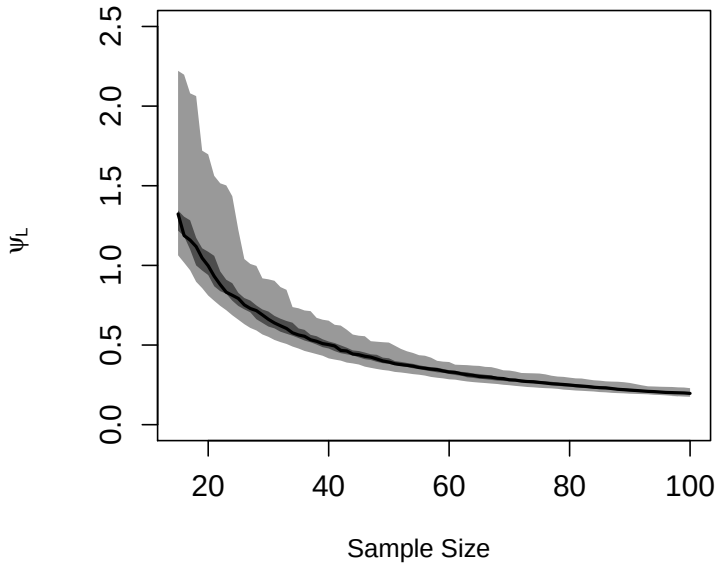

Figure F.14: Distribution of Ψ_L vs sample size for the weighted L-optimal design for the linear model case with non-zero interaction.

Appendix G

Nonmyopic weighted L-optimal design with NIG prior

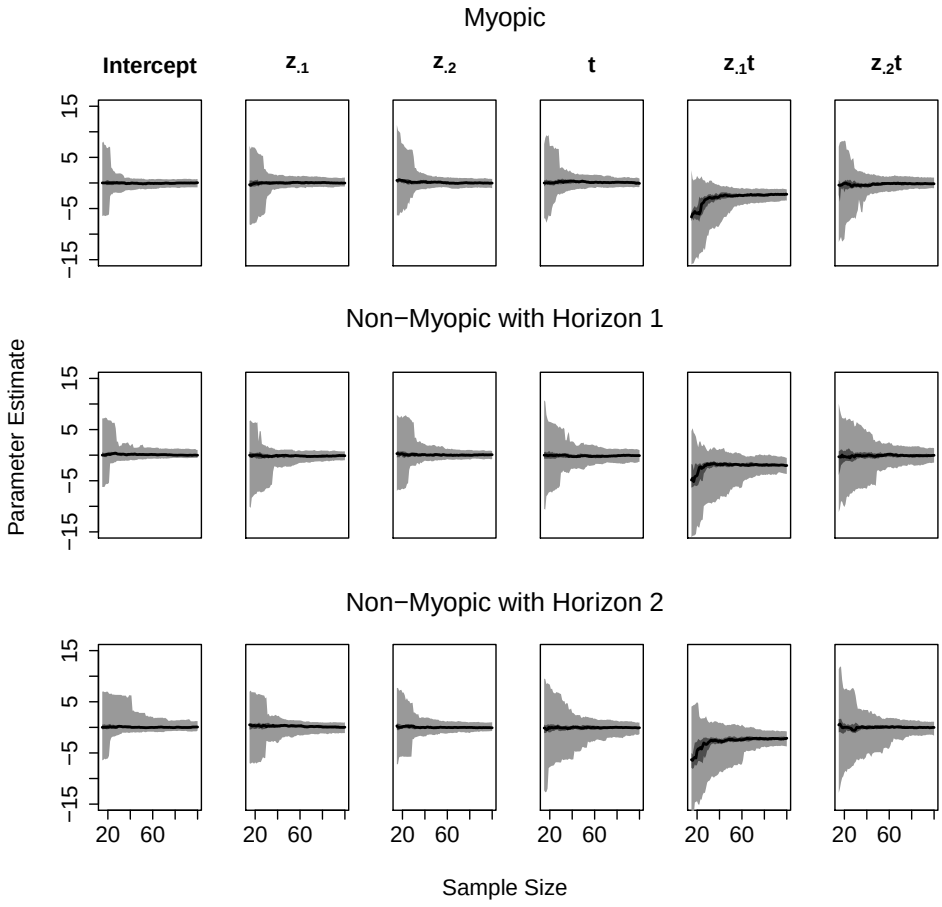

Figure G.1: Distribution of $\hat{\beta}$ vs sample size under the normal-inverse-gamma prior for the myopic approach (top row), non-myopic approach with horizon 1 (middle row) and non-myopic approach with horizon 2 (bottom row). The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

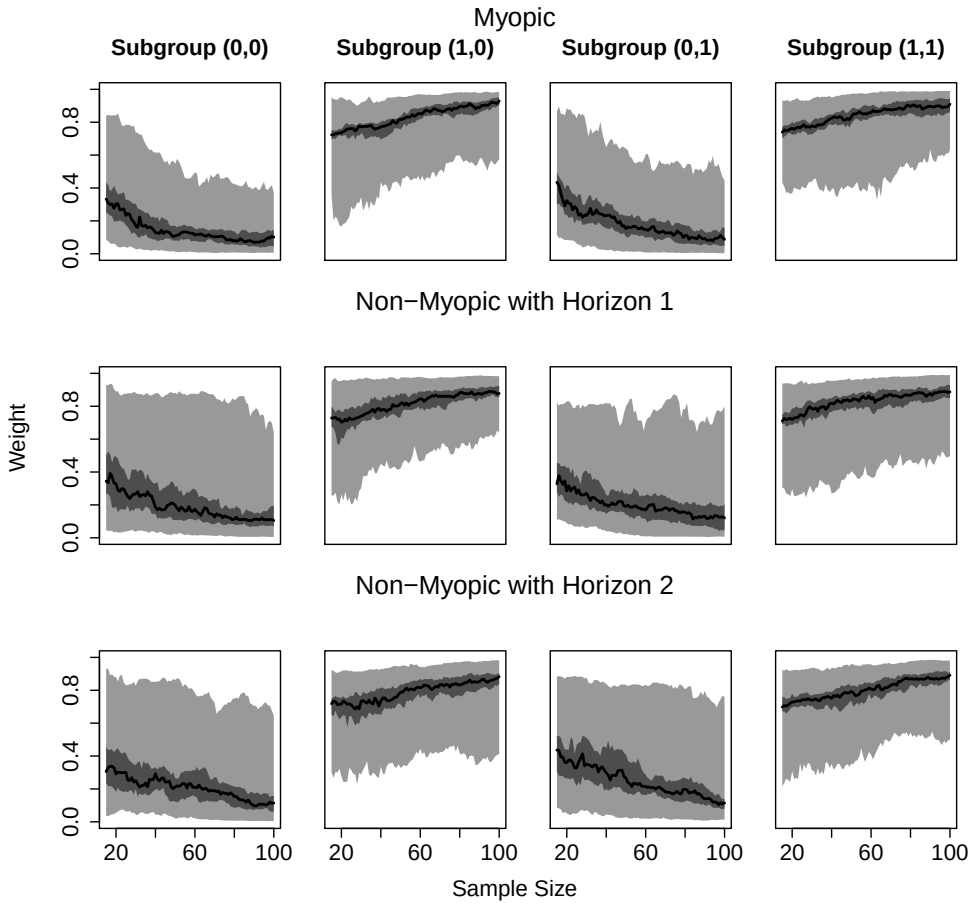

Figure G.2: Distribution of w_r vs sample size under the normal-inverse-gamma prior for the myopic approach (top row), non-myopic approach with horizon 1 (middle row) and non-myopic approach with horizon 2 (bottom row).

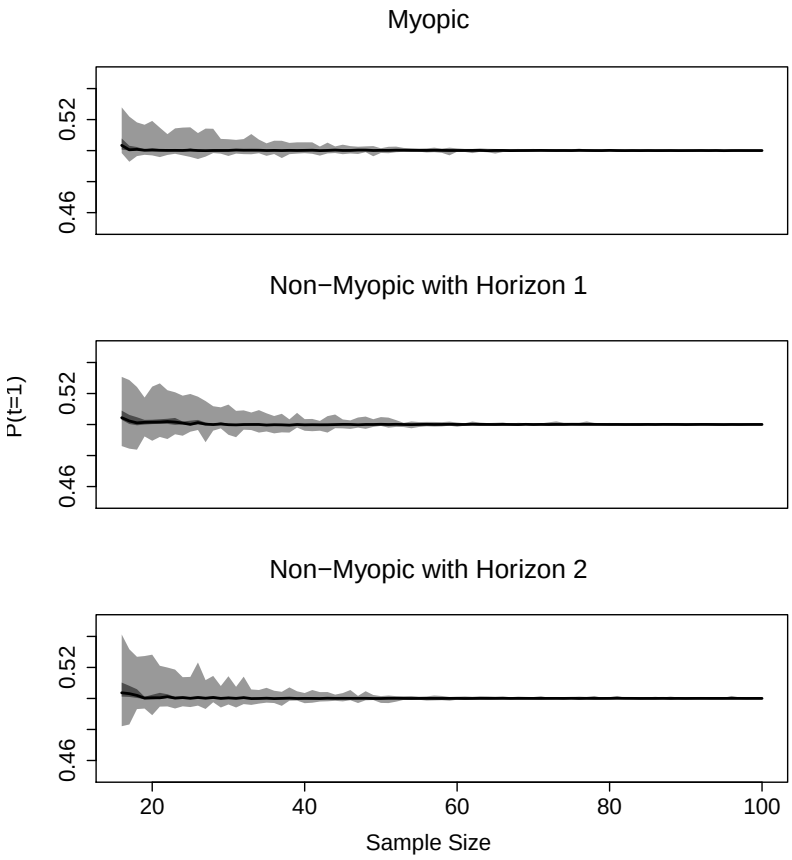

Figure G.3: Distribution of $P(t_i = 1)$ vs sample size under the normal-inverse-gamma prior for the myopic approach (top row), non-myopic approach with horizon 1 (middle row) and non-myopic approach with horizon 2 (bottom row).

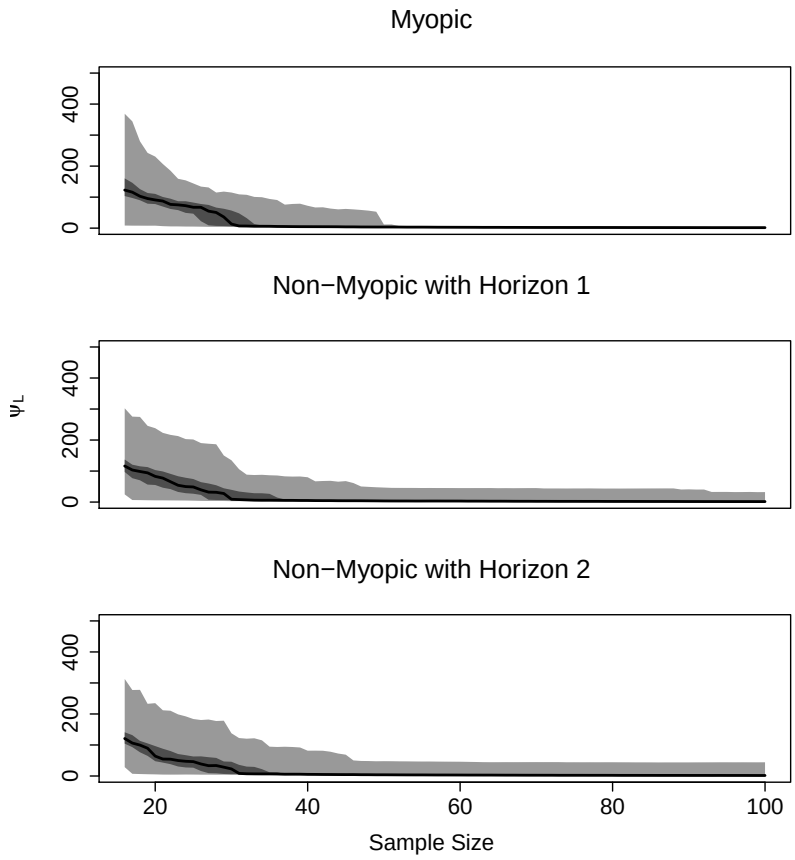

Figure G.4: Distribution of Ψ_L vs sample size under the normal-inverse-gamma prior for the myopic approach (top row), non-myopic approach with horizon 1 (middle row) and non-myopic approach with horizon 2 (bottom row)

Appendix H

Simple illustrative example

In this example, suppose that we have one binary covariate z generated by the distribution given by: $\mathbb{P}(z = 1) = 0.3$ and $\mathbb{P}(z = -1) = 0.7$. The response is binary and we assume a linear model where we have an intercept and effects for the treatment and covariate:

$$\mathbb{E}(\boldsymbol{y}) = \beta_0 + \beta_1 \boldsymbol{z} + \beta_2 \boldsymbol{t}. \tag{H.1}$$

We do not assume treatment-covariate interactions.

Suppose that we already have three patients in our design with covariate values $z_1 = 1, z_2 = -1$ and $z_3 = 1$, respectively. Their treatments are $t_1 = -1, t_2 = -1$ and $t_3 = 1$. The design matrix is as follows:

$$\begin{pmatrix} [r]1 & 1 & -1 \\ 1 & -1 & -1 \\ 1 & 1 & 1 \end{pmatrix} \tag{H.2}$$

Suppose we observe that $z_4 = -1$. We illustrate three methods of assigning treatment t_4:

H.1 Myopic approach

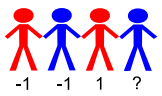

Figure H.1

As illustrated in Figure H.1, we only consider patients one up to four when we make the decision about the fourth patient. There are two possible options for the design matrix after four patients, $\boldsymbol{X}_a$ where $t_4 = 1$ and $\boldsymbol{X}_b$ where $t_4 = -1$:

$$\boldsymbol{X}_a = \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & -1 \\ 1 & 1 & 1 \\ 1 & -1 & 1 \end{pmatrix}, \boldsymbol{X}_b = \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & -1 \\ 1 & 1 & 1 \\ 1 & -1 & -1 \end{pmatrix} \tag{H.3}$$

We have that $\Psi_D(\boldsymbol{X}_a) = 0.03125$ and $\Psi_D(\boldsymbol{X}_b) = 0.015625$. If we select treatments deterministically, we select -1 to be the treatment for patient 4 since $\Psi_D(\boldsymbol{X}_b) < \Psi_D(\boldsymbol{X}_a)$. If we wish to select treatments with a stochastic element, we assign $t_4 = 1$ with probability given by:

$$\frac{\frac{1}{0.03125}}{\frac{1}{0.03125} + \frac{1}{0.015625}} = 0.3333. \tag{H.4}$$

H.2 Nonmyopic approach

With this approach, we consider what the expected loss would be after one patient in the horizon if $t_4 = 1$ versus $t_4 = -1$. As shown in Figure H.2, there are four possibilities to consider: t_4 could either be 1 or -1, and for each of these possibilities, there are two covariate values that we could potentially observe for z_5. For each of these four possible options, we determine the best choice of treatment for t_5 and note the value of D-optimality if the optimal treatment is assigned. This expected D-optimality is shown on the right hand side of Figure H.2.

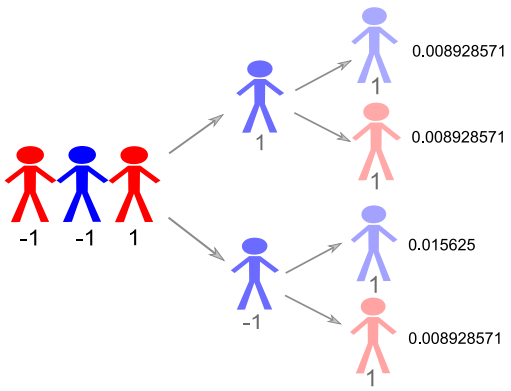

Figure H.2

Next, we compute the expected loss with one patient in the horizon for $t_4 = 1$ and $t_4 = -1$, using Equation (4.2):

$$\Phi_1(t_4 = 1 \mid \boldsymbol{z}_4, \boldsymbol{t}_3) = 0.3\,(0.00893) + 0.7\,(0.00893) = 0.00893 \tag{H.5}$$

$$\Phi_1(t_4 = -1 \mid \boldsymbol{z}_4, \boldsymbol{t}_3) = 0.3\,(0.0156) + 0.7\,(0.00893) = 0.010931 \tag{H.6}$$

We find that $\Phi_1(t_4 = 1 \mid \boldsymbol{z}_4, \boldsymbol{t}_3) < \Phi_1(t_4 = -1 \mid \boldsymbol{z}_4, \boldsymbol{t}_3)$, so we would select $t_4 = 1$ if treatment is to be chosen deterministically. If we are adding a stochastic element to the treatment allocation procedure, we select $t_4 = 1$ with probability given by:

$$\frac{\frac{1}{0.00893}}{\frac{1}{0.00893} + \frac{1}{0.010931}} = 0.5503751. \tag{H.7}$$

H.3 Pseudo-nonmyopic approach

To illustrate the pseudo-nonmyopic approach, we consider what the expected loss would be after all patients up until the end of the experiment if $t_4 = 1$ versus $t_4 = -1$. Let us suppose there is a total of $n = 7$ patients in the trial, and we choose the number of trajectories, M, to be 3. As shown in Figure H.3, supposing that $t_4 = 1$, we generate three trajectories, each with three possible covariate values for future patients. We select treatments for those three patients using the exchange algorithm, and note the value of D-optimality after seven patients for each of those three trajectories. The values of the D-optimality are given on the right hand side of Figure H.3. We then suppose that $t_4 = -1$ and repeat the same process.

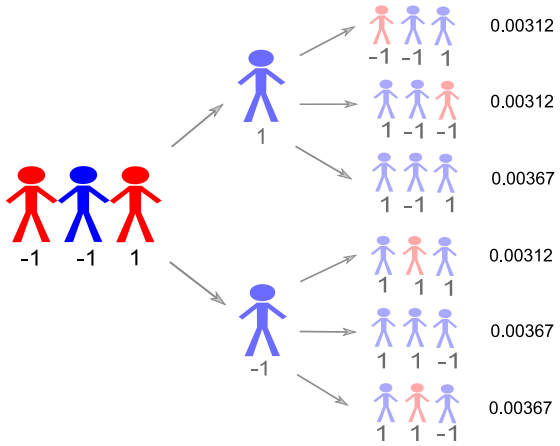

Figure H.3

Next, we compute the average objective function across the M designs, assuming $t_4 = 1$ and $t_4 = -1$, using Equation (5.9):

$$\bar{\Psi}(t_4 = 1) = \frac{1}{3}\,(0.00312 + 0.00312 + 0.00367) = 0.00330, \tag{H.8}$$

$$\bar{\Psi}(t_4 = -1) = \frac{1}{3}\left(0.00367 + 0.00367 + 0.00312\right) = 0.00349. \tag{H.9}$$

Since $\bar{\Psi}(t_4 = 1) < \bar{\Psi}(t_4 = -1)$, so we would select $t_4 = 1$ if treatment is to be chosen deterministically. If treatment is allocated stochastically, we would allocate $t_4 = 1$ with probability given by:

$$\frac{\frac{1}{0.00330}}{\frac{1}{0.00330} + \frac{1}{0.00349}} = 0.514. \tag{H.10}$$

We observe that, for this small example, the decisions for treatment assignment for patient 4 do not necessarily coincide for the three methods.

Appendix I

Pseudo-nonmyopic approach for personalized medicine

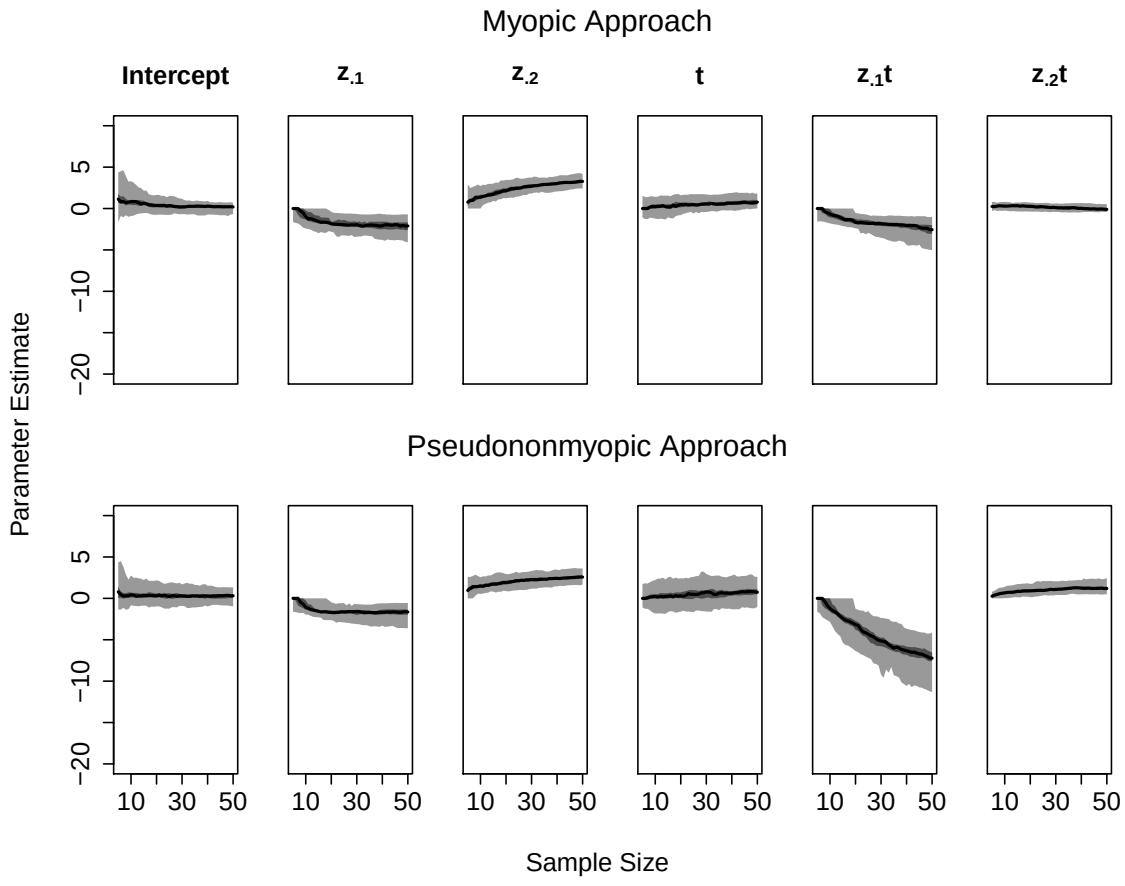

Figure I.1: Distributions of $\hat{\boldsymbol{\beta}}$ for designs for the logistic model with a continuous treatment and large interaction are plotted against sample size for the myopic and pseudo-nonmyopic approaches. The black line indicates the median, the dark grey indicates the 40th to 60th percentile, and the light grey indicates the 10th to 90th percentile of the distribution.

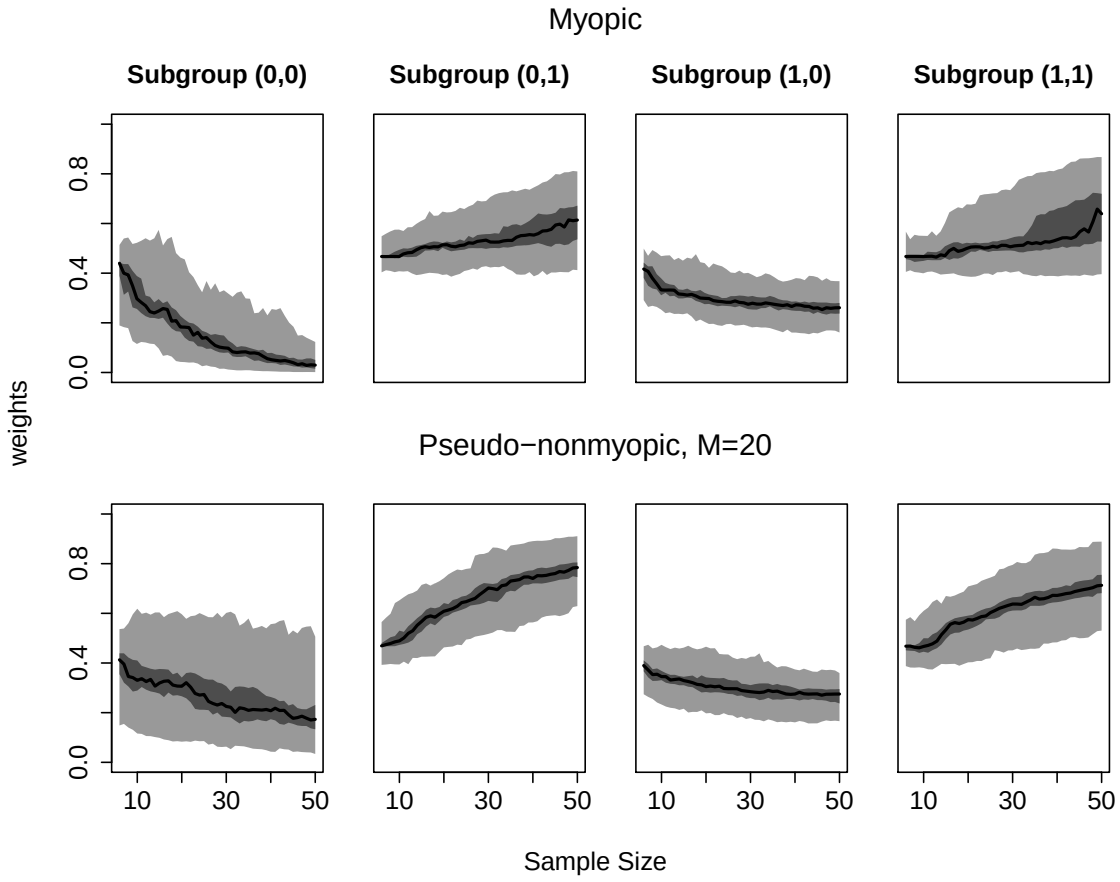

Figure I.2: Distributions of w_r for designs for the logistic model with a continuous treatment and large interaction are plotted against sample size for the myopic and pseudo-nonmyopic approaches.

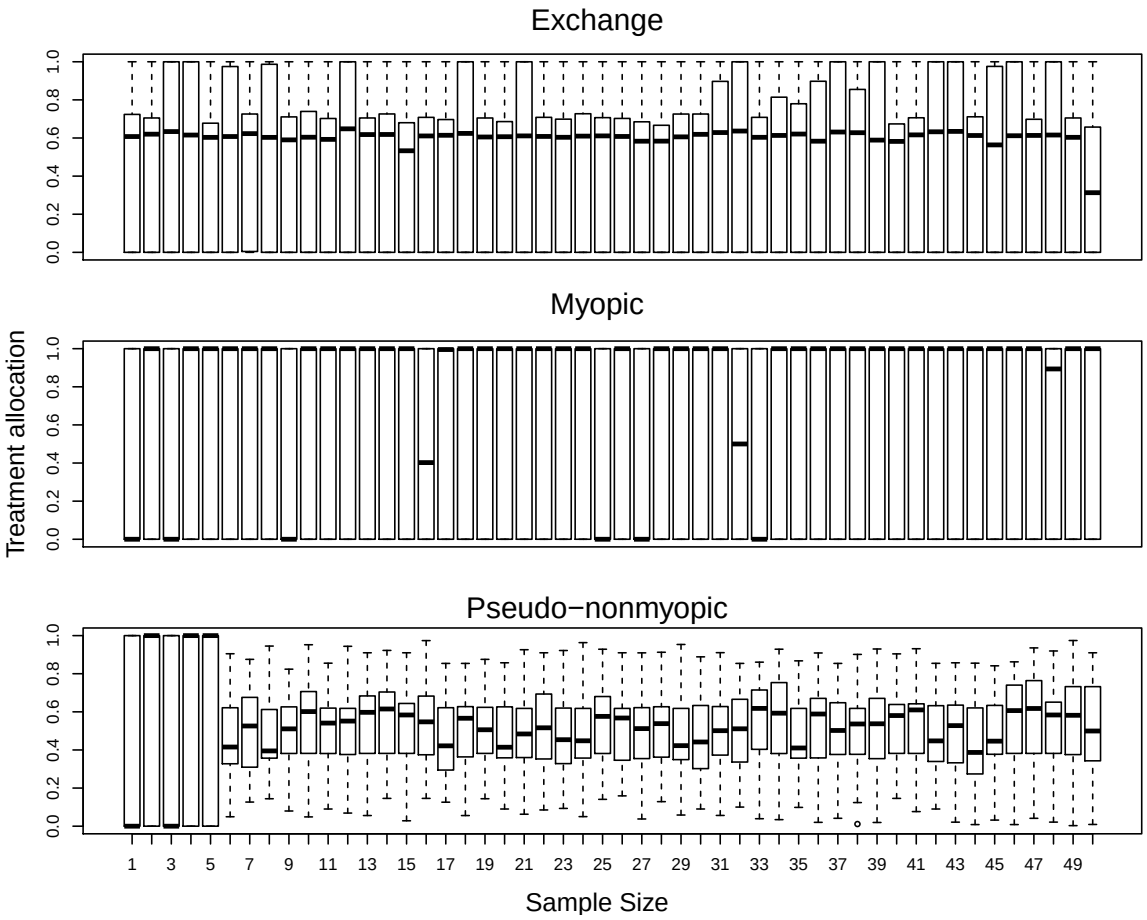

Figure I.3: Boxplots display the distribution of the allocated treatments at each sample size for the the exchange, myopic and pseudo-nonmyopic approaches for the logistic model with a continuous treatment and large interaction.

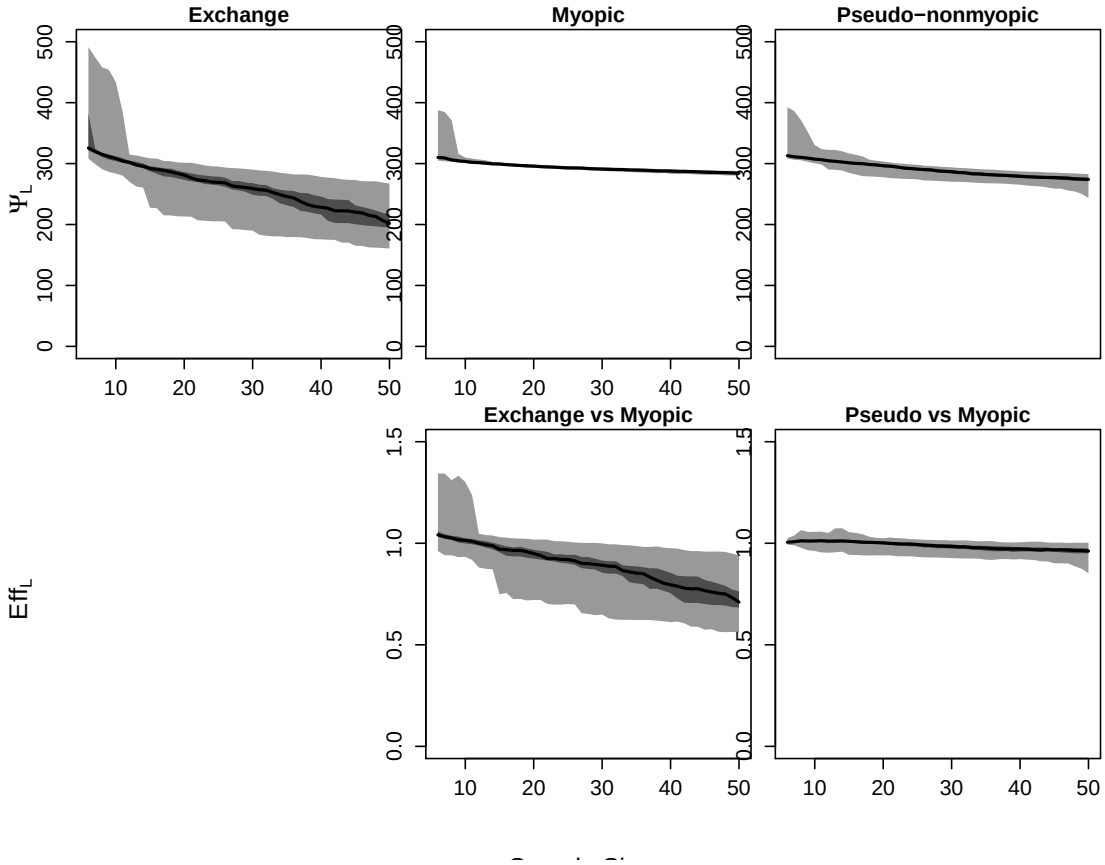

Figure I.4: Distribution of weighted L-optimality is plotted against sample size the exchange algorithm (top left), myopic approach (top middle) and the pseudo-nonmyopic approach (top right) case with a continuous treatment and large interaction. The relative weighted L-optimality of the exchange algorithm vs the myopic approach is plotted against sample size on the bottom left, and the pseudo-nonmyopic approach vs the myopic approach is on the bottom right.

Appendix J

Preliminary work on CARAEE designs

In this appendix, we describe some of the preliminary work that we did on CARAEE designs. Here, there is a need to consider both the efficiency and ethics components in making the decision about treatment allocation. We denote the efficiency and ethics measurements as $d_t(z_i, \hat{\beta}_i)$ and $e_t(z_i, \hat{\beta}_i)$, respectively. In a CARAEE design, the ith patient is assigned treatment 1 with probability given by:

$$\frac{e_1(z_i, \hat{\boldsymbol{\beta}}_i)\left(d_1(z_i, \hat{\boldsymbol{\beta}}_i)\right)^{\gamma}}{e_1(z_i, \hat{\boldsymbol{\beta}}_i)\left(d_1(z_i, \hat{\boldsymbol{\beta}}_i)\right)^{\gamma} + e_{-1}(z_i, \hat{\boldsymbol{\beta}}_i)\left(d_{-1}(z_i, \hat{\boldsymbol{\beta}}_i)\right)^{\gamma}}, \tag{J.1}$$

where $\gamma \geq$ is a tuning parameter. With $\gamma = 0$, you have a CARA (covariate adjusted response adaptive) design. With $e_1(z_i, \hat{\boldsymbol{\beta}}_i) = 1$, you have a RAR (response adaptive randomization) design. We chose D-optimality as the efficiency measure for the ith patient when they are assigned treatment t, for $t \in \{-1, 1\}$:

$$d_t(z_i, \hat{\boldsymbol{\beta}}_i) = \Psi_D(t|\tilde{z}_i, \tilde{t}_{i-1}, \tilde{y}_{i-1}), \tag{J.2}$$

which we wish to minimize. We set the ethics measure as the expected response for patient i, given their covariate values and assuming that they receive treatment t:

$$e_t(z_i, \hat{\boldsymbol{\beta}}_i) = \mathbb{E}\left(y_i|\tilde{z}_i, \tilde{t}_{i-1}, t_i = t\right), \tag{J.3}$$

which we wish to maximize, following the proposal by Hu et al. (2015). The probability of assigning treatment 1 to patient i is given by:

$$\frac{\frac{e_1(z_i,\hat{\boldsymbol{\beta}}_i)^\gamma}{d_1(z_i,\hat{\boldsymbol{\beta}}_i)}}{\frac{e_1(z_i,\hat{\boldsymbol{\beta}}_i)^\gamma}{d_1(z_i,\hat{\boldsymbol{\beta}}_i)} + \frac{e_{-1}(z_i,\hat{\boldsymbol{\beta}}_i)^\gamma}{d_{-1}(z_i,\hat{\boldsymbol{\beta}}_i)}}, \tag{J.4}$$

where $0 \leq \gamma \leq 1$. We note that $0 \leq e_t(z_i, \hat{\beta}_i) \leq 1$ and $d_t(z_i, \hat{\beta}_i) \geq 0$ for all i and $t \in \{-1, 1\}$. Some simulation results showed that the CARAEE designs for the binary response, produce designs that are less efficient but more ethical than myopic D-optimal designs, as expected. Investigation is needed to verify suitable choices of the tuning parameter, γ, and to check whether our choices of $d_t(z_i, \hat{\beta}_i)$ and $e_t(z_i, \hat{\beta}_i)$ are suitable. A major concern is that they are not measured on different scales. An interesting area of future work is extending the CARAEE design for a nonmyopic or pseudo-nonmyopic framework.

References

Abdelbasit, K. M. and Plackett, R. L. (1983) Experimental design for binary data. *Journal of the American Statistical Association*, **78**, 90–98.

Agrawal, R. (1995) Sample mean based index policies with o(log n) regret for the multi-armed bandit problem. *Advances in Applied Probability*, **27**, 1054–1078.

Agrawal, S. and Goyal, N. (2012) Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory* (eds. S. Mannor, N. Srebro and R. C. Williamson), vol. 23 of *Proceedings of Machine Learning Research*, 39.1–39.26. Edinburgh, Scotland: PMLR.

— (2013) Thompson Sampling for Contextual Bandits with Linear Payoffs. *Proceedings of the 30th International Conference on Machine Learning*, **28**.

Antoniou, M., Jorgensen, A. L. and Kolamunnage-Dona, R. (2016) Biomarker-guided adaptive trial designs in phase II and phase III: A methodological review. *PLoS ONE*, **11**, 1–30.

Antoniou, M., Kolamunnage-Dona, R. and Jorgensen, A. (2017) Biomarker-Guided Non-Adaptive Trial Designs in Phase II and Phase III: A Methodological Review. *Journal of Personalized Medicine*, **7**, 1.

Armitage, P., Berry, G. and Matthews, J. (2001) *Statistical Methods in Medical Research*. Oxford statistical science series. Wiley.

Armitage, P., Mcpherson, K. and C. Rowe, B. (1969) Repeated significance tests on accumulating data. *Journal of the Royal Statistical Society. Series A (General)*, **132**, 235–244.

Assmann, S. F., Pocock, S. J., Enos, L. E. and Kasten, L. E. (2000) Subgroup analysis and other (mis)uses of baseline data in clinical trials. *The Lancet*, **355**, 1064–1069.

Atkinson, A. C. (1982) Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika*, **69**, 61–67.

— (1999) Optimum biased-coin designs for sequential treatment allocation with covariate information. *Statistics in Medicine*, **18**, 1741–1752.

Atkinson, A. C. and Biswas, A. (2005) Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Statistics in Medicine*, **24**, 2477–2492.

Atkinson, A. C., Donev, A. N. and Tobias, R. D. (2007) *Optimum Experimental Designs, with SAS*. Oxford University Press.

Atkinson, A. C. and Fedorov, V. V. (1975) The design of experiments for discriminating between two rival models. *Biometrika*, **62**, 57–70.

Atkinson, A. C. and Woods, D. C. (2015) *Handbook of Design and Analysis of Experiments*, chap. Designs for generalized linear models. Chapman & Hall/CRC.

Bailey, R. A. (2008) *Design of Comparative Experiments*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

Bakshy, E., Park, M., Eckles, D., Park, M. and Bernstein, M. S. (2014) Designing and Deploying Online Field Experiments. *Proceedings of the 23rd International Conference on World Wide Web*.

Bandyopadhyay, U. and Biswas, A. (1999) Allocation By Randomized Play-the-Winner Rule in the Presence of Prognostic Factors. *Sankhya: The Indian Journal of Statistics, Series B (1960-2002)*, **61**, 397–412.

— (2001) Adaptive designs for normal responses with prognostic factors. *Biometrika*, **88**, 409–419.

Bartroff, J. and Lai, T. L. (2010) Approximate dynamic programming and its applications to the design of phase I cancer trials. *Statist. Sci.*, **25**, 245–257.

Berry, D. and Fristedt, B. (1979) Bernoulli One-Armed Bandits–Arbitrary Discount Sequences. *The Annals of Statistics*, **7**, 1086–1105.

Berry, D. A. (1989) Monitoring accumulating data in a clinical trial. *International Biometric Society*, **45**, 1197–1211.

Berry, D. A. and Viscusi, W. (1981) Bernoulli two-armed bandits with geometric termination. *Stochastic Processes and their Applications*, **11**, 35 – 45.

Bertsekas, D. (1976) *Dynamic Programming and Stochastic Control*. Academic Press.

Bradley, S. P., Hax, A. C. and Magnanti, T. L. (1977) *Applied Mathematical Programming*. Addison Wesley.

Bradt, R. N., Johnson, S. M. and Karlin, S. (1956) On sequential designs for maximizing the sum of n observations. *Ann. Math. Statist.*, **27**, 1060–1074.

Brezzi, B. Y. M. and Lai, T. Z. E. L. (2000) Incomplete learning from endogenous data in dynamic allocation. *The Econometric Society*, **68**, 1511–1516.

Cancer Research UK (2015) Phases of a clinical trial.

Casella, G. and Berger, R. (2002) *Statistical Inference*. Duxbury advanced series in statistics and decision sciences. Thomson Learning.

Chakraborty, B. and Moodie, E. E. (2013) *Statistical Methods for Dynamic Treatment Regimes*. Springer, New York, NY.

Chaloner, K. and Larntz, K. (1989) optimal logistic bayesian design applied to logistic regression experiments. *Journal of Statistical Planning and Inference*, **21**, 191–208.

Cheng, Y. and Berry, D. A. (2007) Optimal adaptive randomized designs for clinical trials. *Biometrika*, **94**, 673–687.

Clairon, Q., Wilson, E., Henderson, R. and Taylor, C. (2017) Adaptive Biomedical Treatment and Robust Control 1 1This work is supported by the UK Engineering & Physical Sciences Research Council (EPSRC) grant EP/M015637/1 (2015–2018). *IFAC-PapersOnLine*, **50**, 12191–12196.

Coad, D. S. (1991) Sequential tests for an unstable response variable. *Biometrika*, **78**, 113–121.

Cox, D. (1958) *Planning of experiments*. Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley.

Cox, D. R. and Reid, N. (2000) *The theory of the design of experiments*. Chapman and Hall/CRC.

Dixon, W. J. and Mood, A. M. (1948) A method for obtaining and analyzing sensitivity data. *Journal of the American Statistical Association*, **43**, 109–126.

Dror, H. A. and Steinberg, D. M. (2008) Sequential Experimental Designs for Generalized Linear Models. *Journal of the American Statistical Association*, **103**, 288–298.

Efron, B. (1971) Forcing a sequential experiment to be balanced. *Biometrika*, **58**, 403–417.

Firth, D. (1993) Bias reduction of maximum likelihood estimates. *Biometrika*, **80**, 27–38.

Freedman, L. S. and Spiegelhalter, D. J. (1989) Comparison of bayesian with group sequential methods for monitoring clinical trials. *Controlled Clinical Trials*, **10**, 357–367.

Freidlin, B., McShane, L. M. and Korn, E. L. (2010) Randomized clinical trials with biomarkers: Design issues. *Journal of the National Cancer Institute*, **102**, 152–160.

Gautier, R. and Pronzato, L. (1998) Sequential design and active control. *New Developments and Applications in Experimental Design, Lecture Notes Monograph Series*, **34**, 138–151.

Gelman, A., Jakulin, A., Pittau, M. G. and Su, Y. S. (2008) A weakly informative default prior distribution for logistic and other regression models. *Annals of Applied Statistics*, **2**, 1360–1383.

Gelman, A. and Su, Y.-S. (2016) *arm: Data Analysis Using Regression and Multilevel/HierarchicalModels*. R package version 1.9-3.

Gentle, J. E. (2003) *Random Number Generation and Monte Carlo Methods*. Springer-Verlag New York.

Gittins, J. C. and Jones, D. M. (1979) A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, **66**.

Goos, P. and Jones, B. (2011) *Optimal Design of Experiments: A Case Study Approach*. Wiley.

Greenewald, K., Tewari, A. and Murphy, S. (2017) Action Centered Contextual Bandits. *31st Conference on Neural Information Processing Systems*.

Hardwick, J., Oehmke, R. and Stout, Q. F. (2006) New adaptive designs for delayed response models. *Journal of Statistical Planning and Inference*, **136**, 1940–1955.

Hu, J., Zhu, H. and Hu, F. (2015) A Unified Family of Covariate-Adjusted Response-Adaptive Designs Based on Efficiency and Ethics. *Journal of the American Statistical Association*, **110**, 357–367.

Hu, Y. and Hu, F. (2012) Balancing treatment allocation over continuous covariates: A new imbalance measure for minimization. *Journal of Probability and Statistics*, **2012**.

Huan, X. and Marzouk, Y. M. (2016) Sequential Bayesian optimal experimental design via approximate dynamic programming. 1–34.

Imbens, G. W. and Rubin, D. B. (2015) *Causal inference for statistics, social and biomedical sciences: an introduction*. New York: Cambridge University Press.

Joseph, V. R. (2004) Efficient robbins-monro procedure for binary data. *Biometrika*, **91**, 461–470.

Kaplan, R. (2015) The FOCUS4 design for biomarker stratified trials. *Chinese Clinical Oncology*, **4**, 1–10.

Kaufmann, E., Cappe, O. and Garivier, A. (2012) On bayesian upper confidence bounds for

bandit problems. *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, PMLR*, **22**, 592–600.

Kendall, J. (2003) Designing a research project: randomised controlled trials and their principles. *Bmj*, **20**, 164–168.

Khuri, A. I., Mukherjee, B., Sinha, B. K. and Ghosh, M. (2006) Design Issues for Generalized Linear Models: A Review. *Statistical Science*, **21**, 376–399.

Kiefer, J. and Wolfowitz, J. (1960) The equivalence of two extremum problems. *Canadian Journal of Mathematics*, **12**, 363–366.

King, J. and Wong, W.-K. (2000) Minimax d-optimal designs for the logistic model. *Biometrics*, **56**, 1263–1267.

Klotz, J. H. (1978) Maximum entropy constrained balance randomization for clinical trials. *Biometric*, **34**, 283–287.

Lai, T. L. and Robbins, H. (1985) Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, **6**, 4–22.

Lee, K. M. and Wason, J. (2019) Design of experiments for a confirmatory trial of precision medicine. *Journal of Statistical Planning and Inference*, **199**, 179 – 187.

Mahajan, Adityaand Teneketzis, D. (2008) Multi-armed bandit problems. In *Foundations and Applications of Sensor Management* (eds. A. O. Hero, D. A. Castañón, D. Cochran and K. Kastella), 121–151. Springer US.

Maitland, M. L. and Schilsky, R. L. (2011) Clinical trials in the era of personalized oncology. *CA: A Cancer Journal for Clinicians*, **61**, 365–381.

Morgan, J. P. and Wang, X. (2010) Weighted optimality in designed experimentation. *Journal of the American Statistical Association*, **105**, 1566–1580.

Morgan, K. L. and Rubin, D. B. (2012) Rerandomization to improve covariate balance in experiments. *Annals of Statistics*, **40**, 1263–1282.

— (2015) Rerandomization to Balance Tiers of Covariates. *Journal of the American Statistical Association*, **110**, 1412–1421.

Moyé, L. A. and Deswal, A. (2001) Trials within trials: Confirmatory subgroup analyses in controlled clinical experiments. *Controlled Clinical Trials*, **22**, 605–619.

MRC-CTU (2014) What is a clinical trial?

Mueller, P., Berry, D. A., Grieve, A. P., Smith, M. and Krams, M. (2007) Simulation-based sequential Bayesian design. *Journal of Statistical Planning and Inference*, **137**, 3140–3150.

Murphy, S. A. (2003) Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B-Statistical Methodology*, **65**, 331–355.

Neyer, B. T. (1994) A d-optimality-based sensitivity test. *Technometrics*, **36**, 61–70.

O'Brien, P. C. and Fleming, T. R. (1979) A Multiple Testing Procedure for Clinical Trials.

O'Hagan, A. (1994) *Kendall's Advanced Theory of Statistics: Bayesian inference. vol. 2B.* No. v. 2, pt. 2 in Kendall's library of statistics. Edward Arnold.

Ondra, T., Dmitrienko, A., Friede, T., Graf, A., Miller, F., Stallard, N. and Posch, M. (2016a) Methods for identification and confirmation of targeted subgroups in clinical trials: A systematic review. *Journal of Biopharmaceutical Statistics*, **26**, 99–119.

Ondra, T., Jobjornsson, S., Beckman, R. A., Burman, C.-F., Koenig, F., Stallard, N. and Posch, M. (2019) Optimized adaptive enrichment designs. *statistical methods in medical research*, **28**, 2096–2111.

Ondra, T., Jobjörnsson, S., Beckman, R. A., Burman, C. F., König, F., Stallard, N. and Posch, M. (2016b) Optimizing trial designs for targeted therapies. *PLoS ONE*, **11**, 1–19.

Parmar, M. K., Sydes, M. R., Cafferty, F. H., Choodari-Oskooei, B., Langley, R. E., Brown, L., Phillips, P. P., Spears, M. R., Rowley, S., Kaplan, R., James, N. D., Maughan, T., Paton, N. and Royston, P. J. (2017) Testing many treatments within a single protocol over 10 years at MRC Clinical Trials Unit at UCL: Multi-arm, multi-stage platform, umbrella and basket protocols. *Clinical Trials*, **14**, 451–461.

Parmigiani, G. (2002) Measuring uncertainty in complex decision analysis models. *Statistical Methods in Medical Research*, **11**, 513–537.

Perchet, V. and Rigollet, P. (2013) The multi-armed bandit problem with covariates. *Annals of Statistics*, **41**, 693–721.

Pocock, S. and Simon, R. (1975) Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics*, **31**, 691–694.

Pocock, S. J. (1982) Interim analyses for randomized clinical trials: The group sequential approach. *Biometrics*, **38**, 153–162.

— (2013) *Clinical Trials: a practical approach.* John Wiley and Sons Ltd,.

Powell, W. B. (2009) What you should know about approximate dynamic programming. *Naval Research Logistics (NRL)*, **56**, 239–249.

Rainey, C. (2016) Dealing with separation in logistic regression models. *Political Analysis*, **24**, 339–355.

Rao, C. (1973) *Linear Statistical Inference and Its Applications*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. Wiley.

Renfro, L. A. and Sargent, D. J. (2017) Statistical controversies in clinical research: Basket trials, umbrella trials, and other master protocols: A review and examples. *Annals of Oncology*, **28**, 34–43.

Rigollet, P. and Zeevi, A. (2010) Nonparametric bandits with covariates. *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 54–66.

Robbins, H. (1952) Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc*, **58**, 527–535.

Robbins, H. and Monro, S. (1951) A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, **22**, 400–407.

Robins, J. M. (2004) Optimal Structural Nested Models for Optimal Sequential Decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics. Lecture Notes in Statistics, vol 179.* (eds. D. Lin and P. Heagerty), 189–326. Springer, New York, NY.

Rosenberger, W. F. (1999) Randomized play-the-winner clinical trials: Review and recommendations. *Controlled Clinical Trials*, **20**, 328–342.

Rosenberger, W. F. and Sverdlov, O. (2008) Handling Covariates in the Design of Clinical Trials. *Statistical Science*, **23**, 404–419.

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.

Royston, P., Parmar, M. K. and Qian, W. (2003) Novel designs for multi-arm clinical trials with survival outcomes with an application in ovarian cancer. *Statistics in Medicine*, **22**, 2239–2256.

Rubin, D. B. (1978) Bayesian Inference for Causal Effects: The Role of Randomization. *The Annals of Statistics*, **6**, 34–58.

Ryan, K. J. (2003) Estimating Expected Information Gains for Experimental Designs With Application to the Random Fatigue-Limit Model Estimating Expected Information Gains for Experimental Designs With Application to the Random Fatigue-Limit Model. *Journal of Computational and Graphical Statistics*, **12**, 585–603.

Sambucini, V. (2015) Comparison of single-arm vs. Randomized phase II clinical trials: A bayesian approach. *Journal of Biopharmaceutical Statistics*, **25**, 474–489.

Sargent, D. J., Conley, B. A., Allegra, C. and Collette, L. (2005) Clinical trial designs for

predictive marker validation in cancer treatment trials. *Journal of Clinical Oncology*, **23**, 2020–2027.

Scott, N. W., McPherson, G. C., Ramsay, C. R. and Campbell, M. K. (2002) The method of minimization for allocation to clinical trials: A review. *Controlled Clinical Trials*, **23**, 662–674.

Scott, S. L. (2010) A modern bayesian look at the multi-armed bandit. *Appl. Stoch. Model. Bus. Ind.*, **26**, 639–658.

Senn, S. (2004) Added values controversies concerning randomization and additivity in clinical trials. *Statistics in Medicine*, **23**, 3729–3753.

— (2016) Mastering variation: Variance components and personalised medicine. *Statistics in Medicine*, **35**, 966–977.

Senn, S., Anisimov, V. V. and Fedorov, V. V. (2010) Comparisons of minimization and Atkinson's algorithm. *Statistics in Medicine*, **29**, 721–730.

Sharma, D., Yadav, U. B. and Sharma, P. (2009) the concept of sensitivity and specificity in relation to two types of errors and its application in medical research. *Journal of Reliability and Statistical Studies*, **2**, 53–58.

Simon, R. (1977) Adaptive treatment assignment methods and clinical trials. *International Biometric Society*, **21**, 721–732.

Sitter, R. R. and Wu, C. F. J. (1999) Two-stage design of quantal response studies. *Biometrics*, **55**, 396–402.

Smith, A. L. and Villar, S. S. (2017) Bayesian adaptive bandit-based designs using the Gittins index for multi-armed trials with normally distributed endpoints. *Journal of Applied Statistics*, 1–25.

Stallard, N., Whitehead, J., Todd, S. and Whitehead, A. (2001) Stopping rules for phase II studies. *British Journal of Clinical Pharmacology*, **51**, 523–529.

Sverdlov, O., Tymofyeyev, Y. and Wong, W. K. (2011) Optimal response-adaptive randomized designs for multi-armed survival trials. *Statistics in Medicine*, **30**, 2890–2910.

Sydes, M. R., Parmar, M. K., Mason, M. D., Clarke, N. W., Amos, C., Anderson, J., de Bono, J., Dearnaley, D. P., Dwyer, J., Green, C., Jovic, G., Ritchie, A. W., Russell, J. M., Sanders, K., Thalmann, G. and James, N. D. (2012) Flexible trial design in practice - stopping arms for lack-of-benefit and adding research arms mid-trial in STAMPEDE: a multi-arm multi-stage randomized controlled trial. *Trials*, **13**, 1.

Tanniou, J., Van Der Tweel, I., Teerenstra, S. and Roes, K. C. (2016) Subgroup analyses in

confirmatory clinical trials: Time to be specific about their purposes. *BMC Medical Research Methodology*, **16**, 1–15.

Taves, D. R. (1974) Minimization: a new method of assigning patients to treatment and control groups. *Clinical Pharmacology and Therapeutics*, **15**, 443–453.

Thall, P. F. and Wathen, J. K. (2007) Practical bayesian adaptive randomization in clinical trials. *European Journal of Cancer*, **52**, 162–165.

Thompson, W. (1933) On the likelihood that one unknown probability exceeds another inview of the evidence of two samples. *Biometrika*, **25**, 285—294.

Titterington, D. M. (1983) On constrained balance randomization for clinical trials. *Biometrics*, **39**, 1083–1086.

Villar, S. S., Bowden, J. and Wason, J. (2015) Multi-armed Bandit Models for the Optimal Design of Clinical Trials: Benefits and Challenges. *Statistical Science*, **30**, 199–215.

Villar, S. S. and Rosenberger, W. F. (2018) Covariate-adjusted response-adaptive randomization for multi-arm clinical trials using a modified forward looking Gittins index rule. *Biometrics*, **74**, 49–57.

Wang, Y.-G. (1991) Gittins indices and constrained allocation in clinical trials. *Biometrika*, **78**, 101–111.

Wason, J. M., Abraham, J. E., Baird, R. D., Gournaris, I., Vallier, A. L., Brenton, J. D., Earl, H. M. and Mander, A. P. (2015) A Bayesian adaptive design for biomarker trials with linked treatments. *British Journal of Cancer*, **113**, 699–705.

Watkins, Christopher J. C. H.and Dayan, P. (1992) Q-learning. *Machine Learning*, **8**, 279–292.

Wei, L. J. and Durham, S. (1978) The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, **73**, 840–843.

Whitehead, J. (1992) *The Design and Analysis of Sequential Clinical Trials*. Ellis Horwood Ltd.

— (1993) Interim analyses and stopping rules in cancer clinical trials. *Br. J. Cancer*, **68**, 1179–1185.

Williamson, S. F., Jacko, P., Villar, S. S. and Jaki, T. (2017) A bayesian adaptive design for clinical trials in rare diseases. *Computational Statistics & Data Analysis*, **113**, 136 – 153.

Woodroofe, M. (1979) A One-Armed Bandit Problem With a Concomitant Variable. *Journal of the American Statistical Association*, **74**, 799–806.

Woods, D. C. and Atkinson, A. (2015) Designs for generalized linear models. In *Handbook of Design and Analysis of Experiments* (eds. A. Dean, M. Morris, J. Stufken and D. Bingham), chap. 13, 471–514. Florida: Chapman and Hall/CRC.

Woods, D. C., Lewis, S. M., Eccleston, J. A. and Russell, K. G. (2006) Designs for generalized linear models with several variables and model uncertainty. *Technometrics*, **48**, 284–292.

Woods, D. C., Overstall, A. M., Adamou, M. and Waite, T. W. (2017) Bayesian design of experiments for generalized linear models and dimensional analysis with industrial and scientific application. *Quality Engineering*, **29**, 91–103.

Wu, C. F. J. (1985) Efficient sequential designs with binary data. *Journal of the American Statistical Association*, **80**, 974–984.

Yang, Y. and Zhu, D. (2002) Randomized allocation with nonparametric estimation for a multi-armed bandit problem with covariates. *Annals of Statistics*, **30**, 100–121.

Yao, Q. and Wei, L. J. (1996) Play the winner for phase II/III clinical trials. *Statistics in medicine*, **15**, 2413–23; discussion 2455–8.

Yu, X., Chen, J. and Brant, R. (2016) Sequential design for binary dose-response experiments. *Journal of Statistical Planning and Inference*, **177**, 64–73.

Zhang, L. I., Hu, F., Cheung, S. H. and Chan, W. S. (2007) Asymptotic properties of covariate-adjusted response-adaptive designs. *Annals of Statistics*, **35**, 1166–1182.

Zhang, Y., Trippa, L. and Parmigiani, G. (2016) Optimal Bayesian adaptive trials when treatment efficacy depends on biomarkers. *Biometrics*, **72**, 414–421.

Zhang, Z., Chen, R., Soon, G. and Zhang, H. (2018) Treatment evaluation for a data-driven subgroup in adaptive enrichment designs of clinical trials. *Statistics in Medicine*, **37**, 1–11.

Zhu, H. (2015) Covariate-adjusted response adaptive designs incorporating covariates with and without treatment interactions. *Canadian Journal of Statistics*, **43**, 534–553.