Optimal Adaptive Designs for Dose Finding in Early Phase Clinical Trials

Muhammad Iftakhar Alam

School of Mathematical Sciences

Queen Mary, University of London

A thesis submitted for the degree of

Doctor of Philosophy

June 2015

To my late father

Acknowledgements

Foremost, I would like to thank my supervisors Dr Barbara Bogacka and Dr D. Stephen Coad. The thesis would not have seen light without their continuous support, valuable ideas and guidance. I am really grateful to them.

I would like to thank Queen Mary, University of London, for offering me a Doctoral Training Account Studentship. I would also like to thank the International Society for Clinical Biostatistics for selecting me for a Student Conference Award to present some of the work at its 34th Annual Conference in Munich, Germany, during 25-29 August 2013. Thanks should go to Queen Mary, University of London, for supporting me from the Postgraduate Research Fund to present a paper at the 35th Annual Meeting of the Society for Clinical Trials in Philadelphia, Pennsylvania, during 18-21 May 2014.

I have received support from many colleagues and friends. In particular, I would like to extend thanks to Lutfor, Shihan, Apon, Zakir and Mahbub for their lovely friendship and inspiration. I am also grateful to the Institute of Statistical Research and Training at the University of Dhaka in Bangladesh for approving my study leave to pursue doctoral research.

My gratitude goes to my parents, siblings and parents-in-law for their continuous support and encouragement during the entire period of study. Finally, special thanks to my wife Fariha for her invaluable patience, care and love.

Three research articles are planned to be published based on the thesis. The first

example in Chapter 6 will be submitted to *Statistical Methods in Medical Research*. The second example is aimed at *Clinical Trials: Journal of the Society for Clinical Trials*. The third example is likely to be submitted to the *Journal of Biopharmaceutical Statistics*. Of course, the supporting materials in each case will come from the previous chapters.

Declaration

I, Muhammad Iftakhar Alam, confirm that the research included within this thesis is my own work or that where it has been carried out in collaboration with, or supported by others, that this is duly acknowledged below and my contribution is indicated. Previously published material is also acknowledged below.

I attest that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge break any UK law, infringe any third party's copyright or other Intellectual Property Right, or contain any confidential material.

I accept that the college has the right to use plagiarism detection software to check the electronic version of the thesis.

I confirm that this thesis has not been previously submitted for the award of a degree by this or any other university.

Signature: Muhammad Iftakhar Alam

Date: 21/06/2015

Supervisors:

Dr Barbara Bogacka

Dr D. Stephen Coad

Abstract

A method of designing early clinical trials is developed for finding an optimum dose level of a new drug to be recommended for use in later phases. During the trial, the efficacious doses are allocated to the patients more often and those with a high probability of toxicity are less likely to be chosen. The method proposed is adaptive in the sense that the statistical models are updated after the data from each cohort of patients are collected and the dose level is adjusted at each stage based on the current data.

Two classes of designs are presented. Although both are for efficacy and toxicity responses, one of them also considers pharmacokinetic information. The dose optimisation criteria are based on the probability of success and on the determinant of the Fisher information matrix for estimation of the dose-response parameters. They can be constrained by both acceptable levels of the probability of toxicity and desirable levels of the area under the concentration curve or the maximum concentration.

The method presented is general and can be applied to various dose-response and pharmacokinetic models. To illustrate the methodology, it is applied to two different classes of models. In both cases, the pharmacokinetic model incorporates the population variability by making appropriate assumptions about the model parameters, while the dose responses are assumed to be either trinomial or bivariate binomial.

Various design properties of the method are examined by simulation studies. Efficiency measures and the sensitivity of the designs to the assumed prior parameter values are presented. All of the computations are conducted in *R*, where the *D*-

optimal sampling time points are obtained by using the package PFIM. The results show that the proposed adaptive method works well and could be appropriate as a seamless phase IB/IIA trial design.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Preliminary Concepts

Clinical trials, commonly classified into four phases, have become an integral part of drug development. Phase I is the first stage of testing in humans and designed to assess safety, tolerability and the pharmacokinetics (PK) of a drug. Phase II is designed to assess how well the drug works (pharmacodynamics, PD) and it also monitors safety in a large group of patients. Another goal of the early phases is to establish a dose level to recommend for further studies in later phases. Phase III assesses the effectiveness of the drug in comparison with the current standard treatments. Phase IV, also known as post-marketing surveillance, aims to detect any rare or long-term adverse effects over a large population.

1.1.1 Pharmacokinetics and Pharmacodynamics

Pharmacokinetics is generally defined as what the body does to the drug. It reflects the movement of the drug in the body, that is, how the drug enters into the body, how it is distributed throughout the body and how it leaves the body. It involves the study of the processes that affect the plasma concentration of drug in the body at any time after the administration of a dose (Rosenbaum, 2011). PK modelling helps in determining important PK parameters.

Pharmacodynamics is defined as what the drug does to the body. PD effects are usu-

ally classified as changes in biomarkers, surrogate endpoints and clinical endpoints. A biomarker reflects any pharmacological effect that has some link to the therapeutic benefit of the drug. A surrogate endpoint is a biomarker which is intended to substitute for a clinical endpoint. A clinical endpoint is a response variable measuring the direct benefit to a patient, that is, how a patient feels, functions or survives (Derendorf et al., 2000). For antihypertensive drugs, blood pressure is a biomarker and stroke is the clinical endpoint. For anticancer agents, tumour shrinkage and survival are the biomarker and clinical endpoint, respectively. PD modelling establishes the relationship between the dose and the resulting effect.

After administration of a dose, the PK mechanism transforms it into plasma concentration and through the systematic circulation of blood, it reaches the site of action and produces a response: see Figure 1.1. The extent of response depends on the concentration at the site of action. We cannot measure concentration at the site of action, but can measure the plasma, which reflects the concentration at the site. Concentration needs to be kept high enough to produce a desirable response, but low enough to avoid toxicity. Even if the same dose is given to a group of individuals, concentration profiles are very likely to be different since it depends on how a body functions. Therefore, for the efficient determination of the optimum dose, it is essential to monitor concentration.

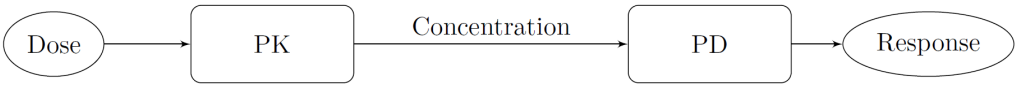

Figure 1.1: A schematic diagram showing how dose of a drug works.

Generally, concentration is modelled as a function of time for a given dose and the PD effect is often modelled as a function of dose. However, in the PK/PD approach, the PD response is modelled as a function of concentration. Concentration is inherently more informative than dose because unlike dose, which is only a nominal mass fixed by the clinicians, it gives biological information (Riviere, 2011). Such an approach establishes the dose-concentration-effect relationship and is capable of predicting the effect at any time after administering a dose. It also helps in estimat-

ing dose and dosing interval to achieve the effect of a desired level.

The PK/PD approach has been described in Hooker and Vicini (2005) and Davidian (2010). To illustrate the approach, we introduce an example that comprises a one-compartment PK model with bolus input and first-order elimination and the $E_{\max}$ model for the PD effect. The PK model for modelling the drug concentration is

$$y_1 = f_1(t_1; \boldsymbol{\theta}_1) + \epsilon_1 = \frac{x}{V} \exp\left(-\frac{Cl}{V} t_1\right) + \epsilon_1, \tag{1.1}$$

where x is the dose received by an individual, $\boldsymbol{\theta}_1 = (V, Cl)^T$ is the vector of PK parameters with V and Cl as the volume of distribution and clearance, respectively, t_1 denotes the sampling time for measuring concentration and ϵ_1 is the random error. More detailed explanations on the volume of distribution and clearance are available in Section 3.3.

The simple $E_{\max}$ model for modelling the PD effect is given as

$$y_2 = f_2(x; \boldsymbol{\theta}_2) + \epsilon_2 = E_0 + \frac{E_{\max}\ x}{C_{50} + x} + \epsilon_2, \tag{1.2}$$

where $\boldsymbol{\theta}_2 = (E_0, E_{\max}, C_{50})^T$ is the vector of PD parameters with E_0, $E_{\max}$ and C_{50} as the effect at baseline, the maximum effect and the dose needed to observe half of the maximum effect, respectively, and ϵ_2 is the random error. The parameter $E_{\max}$ measures the efficacy for a drug and C_{50} reflects a drug's potency.

The PK/PD approach expresses the PD effect as a function of the mean concentration. In our case, we can write

$$y_2 \quad = \quad E_0 + \frac{E_{\max}\ f_1(t_2; \boldsymbol{\theta}_1)}{C_{50} + f_1(t_2; \boldsymbol{\theta}_1)} + \epsilon_2, \tag{1.3}$$

where t_2 represents the sampling time for measuring the PD response. It is also assumed that ϵ_1 and ϵ_2 are independent.

There are two approaches to fit the above models: sequential fitting and simultaneous fitting. In sequential fitting, the PK model is fitted first to obtain the parameter estimates $\hat{\boldsymbol{\theta}}_1$. These parameters are then fixed and assumed known in fitting the model (1.3) to obtain $\hat{\boldsymbol{\theta}}_2$. This approach is based on the simplifying assumption that the two sets of parameters are independent. In the simultaneous fitting approach, the model (1.3) is fitted directly. It allows the assumption that the two sets of parameters are correlated (Hooker and Vicini, 2005).

The PK/PD approach requires the PD response to be a continuous random variable. However, the type of PD response we are interested in in this thesis is categorical. We want to see how the rate of success of a drug changes in the population of patients with the change of dose level. Therefore, the dose-response models that are going to be used are different from those in the PK/PD approach.

1.1.2 Dose-Response Relationship

The dose-response relationship describes the relation between a response and doses of a drug. The graphical presentation of such a relationship is known as a dose-response curve. More specifically, a dose-response curve considered in this thesis is a two-dimensional graph, where the x-axis represents dose and the y-axis represents the percentage of the population that exhibits the response. Depending on the response, which can be toxicity or desired effect, there are two types of important dose-response curves. A dose-efficacy curve is the one that describes the relationship between dose and some efficacy endpoint. The other one is the dose-toxicity curve that describes the relationship between dose and a toxicity endpoint (Chow and Liu, 2004). Although dose-response curves can assume any shape, there are many drugs for which the curves are S-shaped. Figure 1.2 shows such dose-response curves.

The dose-response relationship on toxicity is as equally important as that on efficacy. These relationships together help in identifying an appropriate dose. Development of these relationships is a central part in clinical trial studies to make the safe and

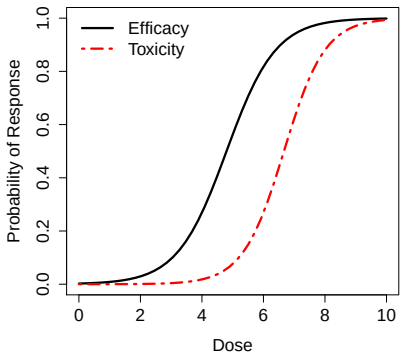

Figure 1.2: Typical dose-response curves.

effective use of a drug. Regression methods for modelling the probabilities of responses are used in developing such models. The main interest is to assess the efficacy and toxicity levels of a drug in a population of patients.

Although dose-response relationships mainly depend on dose, covariates are often considered to account for some known patient characteristics which may influence the dose effects. Generally, lower doses produce no desired effect and higher doses produce undesired side effects. To reduce the time for drug development and also to decrease the cost of the studies, accurate and early establishment of the relationships is essential.

1.1.3 Incorporating PK Information

A drug's pharmacokinetics are determined by the processes of absorption, distribution, metabolism and excretion (ADME). Also, concentrations at the site of action are determined by the ADME. It is the concentrations at the site of action which produce the PD response. The objectives in a early phase clinical trial are to quantify the ADME of the drug and also to identify an optimal dose for further studies. Quantifying these PK parameters is often straightforward, as it does not require the administration of an optimal dose (Piantadosi and Liu, 1998). However, obtaining the optimal dose is a difficult task. There are drugs for which the therapeutic range is narrow and as such careful dose escalation is essential since a small increase in dose may lead to a drastic change in outcomes.

Since the PD response depends on the concentration at the site of action, it may not be enough to consider only the administered dose in establishing the dose-response relationship. Rather, the resulting concentration from that dose has a vital role in it. One possible way to consider such pharmacokinetic data could be the formal inclusion of measures like area under the concentration curve or the maximum concentration in the dose-response model. One such method that puts the area under the concentration curve in the dose-response model is that of Piantadosi and Liu (1998). But that method requires many blood samples to be collected for measuring the concentration of a drug, as it uses the trapezoidal rule to obtain the area under the concentration curve. A larger number of blood samples gives more trapeziums to find an accurate estimate of the area.

Other possible ways of considering PK data in dose-finding studies could be constraining the area under the concentration curve or the maximum concentration during dose escalation. In this case, the measures are not incorporated in the dose-response model directly, rather they help the dose-finding procedure to identify an optimal dose. Such a PK-guided design will be more careful, so that patients do not receive doses which are too toxic or have no therapeutic benefit. This kind of approach is different from the standard PK/PD approach in the sense that unlike the PK/PD approach, we have a categorical PD response, and instead of replacing the dose by concentration, we take some PK parameters associated with the best dose.

1.2 Motivation for the Work

As indicated at the beginning, a phase I trial is the first step in applying a new drug to humans. The designs for phase I trials, particularly those in cancer, focus on the maximum tolerated dose (MTD). This MTD is based on the toxicity response and it ignores the efficacy response. Phase II designs aim to determine the efficacy level of an experimental drug assuming that a dose range has been established in phase

I.

For cytotoxic agents that are used in treating cancer, it is believed that the higher doses are more likely to be toxic, and also more likely to kill the cancer cells and therefore will be efficacious. As a consequence, cancer clinical trials assume that the dose-toxicity and dose-efficacy relationships are monotone non-decreasing functions of dose. Therefore, the MTD determined from the phase I trial will provide a dose with a desirable level of efficacy.

In recent years, targeted therapies have drawn attention for treating cancer patients. These are drugs or other substances that block the growth and spread of cancer by interfering with specific molecules that are involved in the growth, progression and spread of cancer. According to Cunningham et al. (2004), many clinical trials of targeted therapies are underway and many such therapies are already in use for curing the patients. For these targeted agents, the dose-efficacy curve may not increase with dose and therefore efficacy may occur at doses that lie in the middle of the range of possible doses. Therefore, conducting the phase II trials based on the MTD from phase I for those agents may not be an efficient way of running the trial. In such situations, it will be useful to consider the toxicity and efficacy simultaneously to come up with an optimum dose for further investigation in the next phase. Some examples that illustrate combined phase I/II trials are available in Gooley et al. (1994), O'Quigley et al. (2001), Braun (2002) and Zhang et al. (2006).

Zhang et al. (2006) consider an agent for targeted therapy which is theorised to boost haemoglobin levels (HbL) in patients for whom the HbL is below the standard range (14-18 g/dL for men; 12-16 g/dL for women). The response that is obtained following the administration of the agent can be categorised as follows: no response, success and toxicity. These successive categories mean that patients are under-stimulated and their HbL is still below the standard range; rightly stimulated and their HbL raised to the standard level; and over-stimulated and their HbL raised

above the standard level.

Also for drugs other than those for cancer, a dose will be acceptable only if it is efficacious and safe. Since it is quite difficult to find a dose which is both efficacious and non-toxic, efforts are necessary to make a trade-off between the two. The specific goal is to find a dose which has a high probability of efficacy and a low probability of toxicity.

Regardless of the method, phase I trials are small and as a consequence the dose-toxicity curves are not well estimated. They often determine the dose which is either unacceptably toxic or ineffective. Phase I/II trials can be larger and so the methods could potentially lead to more efficient dose selection.

The designs that we are going to consider are for combined phase IB/IIA trials. We also consider pharmacokinetic constraints during dose escalation in cases where it is important to avoid doses likely to cause toxicity, considered as possibly life-threatening. Since inter-patient variability in the plasma concentration is very likely, random effects PK models are being considered. The purpose of this thesis is to develop an efficient method for dose finding in early phase clinical trials.

1.3 Outline of Thesis

In Chapter 2, we review the literature on the commonly used designs in early phase clinical trials. The discussed designs are mainly used in phase I and II trials.

Chapter 3 contains the general population PK model and its Fisher information matrix. Some commonly used optimality criteria in designing experiments are also discussed. The chapter then introduces one-compartment PK models and the associated Fisher information matrices to be used in the simulation study. The general expressions for measuring the inter-patient variability in the area under the concentration curve and the maximum concentration following the administration of a

dose are presented. As special cases, analytical forms for these variances have been obtained for the PK models introduced.

Chapter 4 gives the dose-response models that will be used in the simulation study. The associated Fisher information matrices have been derived. Properties of the models are also discussed and some plausible dose-response scenarios for each of them are presented.

Chapters 5 and 6 are the ones with major contributions in dose-finding studies and they present new results. Chapter 5 introduces the general algorithm for the proposed dose-finding design. It also describes an up-and-down design to gather information prior to any estimation of the model parameters. Possible dose-optimisation criterion, constraints and stopping rules are also discussed in the chapter. To evaluate the quality of the designs, we introduce some performance measures.

Simulation studies using various dose-optimisation criteria and constraints are presented in Chapter 6. The numerical computations are implemented in R. The main purpose is to understand the behaviour of the designs numerically. Three examples are introduced with the dose-response and PK models introduced in the earlier chapters. Sensitivity analyses of the designs to the assumed values for the parameters are also presented. Finally, we discuss the major findings.

In Chapter 7, we draw conclusions of the work and discuss possible future research. It has been found from the simulation studies that the proposed dose-optimisation criteria and constraints can find the optimal dose accurately. The methods can limit toxic doses as the optimum dose by a considerable amount and assign most relevant doses to the cohorts during the trial. We have also seen that the efficiency of the design can be increased if it is possible to assume target values for the PK parameters like the area under the concentration curve and the maximum concentration.

Finally, we present the appendices. Appendix A gives a flow chart indicating the different steps of the proposed design and some supplementary material. Appendix B shows the R code that is used to simulate designs using various dose-optimisation criteria and constraints.

Chapter 2

Review of Early Clinical Trials

2.1 Introduction

This chapter gives a brief introduction to some commonly used designs in the early phases of clinical trials. Clinical trial designs are broadly classified as adaptive and non-adaptive. Here we present some adaptive early phase designs.

The goal of the early phases is to study the pharmacokinetics and pharmacodynamics of new drugs, and also to explore efficacy and toxicity profiles. For nontoxic agents, phase I trials are often conducted with healthy volunteers. But for toxic agents, such as those for cancer treatments, phase I trials are conducted among cancer patients at the last stage for whom standard treatments have failed. Since the benefits from such agents are believed to increase with dose, the highest possible dose is searched for during the development of an agent. However, toxicity also increases with dose. Therefore, the main challenge for those trials is to find a dose with a low chance of toxicity among the patients. This dose is usually referred to as the MTD.

Early phase clinical trials test drugs that were found promising in preclinical studies. For cytotoxic agents, usually the starting dose is one tenth of the LD_{10} in mice, that is, the dose lethal to 10% of mice, or one-third of the toxic dose low (TDL) in dogs or monkeys (Collins et al., 1986). The TDL is defined as the lowest dose that

produces drug-induced pathological alterations in haematological, chemical, clinical or morphological parameters and which, when doubled, produces no lethality (Prieur et al., 1973). In areas other than oncology, the approaches that are used to estimate the starting dose include: (1) the dose-by-factor approach; (2) the similar drug approach; (3) the pharmacokinetically-guided approach; and (4) the comparative approach. A detailed description of these methods is available in Reigner and Blesch (2002).

The dose-by-factor approach is based on the highest dose of the compound found to have no toxic effect in the most sensitive species tested in the preclinical toxicology studies. The maximum starting dose for a first-in-human study is the smallest of the following three doses: 1/10 of the highest no-effect dose in rodents, 1/6 of the highest no-effect dose in dogs and 1/3 of the highest no-effect dose in monkeys. This classical approach is used widely. However, it is often criticised as it ignores preclinical pharmacokinetic data.

The similar drug approach can be applied when human safety data are available for a drug similar to the one under investigation. The *similar drug* is usually of the same chemical class, with similar or related chemical structure. The method assumes that the ratio of the starting dose to the dose at which no adverse effect is observed is the same for *similar drug* and *investigational drug*, and from that relation it is possible to find the starting dose for the *investigational drug*. The dose thus obtained is again multiplied by a factor to accommodate the uncertainty about the safety. One can proceed with the method only if it is found that the assumption upon which it is based is valid.

The pharmacokinetically-guided approach uses systemic exposure rather than dose for the extrapolation from animal to human. The area under the concentration curve from the preclinical study for a dose at which no adverse effect is observed is multiplied by the predicted clearance for humans to obtain the starting dose. The

uncertainty about the prediction of human clearance is a shortcoming of the method. The practice of using this approach in pharmaceutical companies is increasing. The comparative approach utilises two or more methods to estimate a starting dose and then critically compares the results to arrive at the starting dose. This method can be criticised for being time consuming. However, obtaining similar results from several approaches is reassuring. This is not a very commonly used method.

Dose increment, dose assignment, cohort size and number of cohorts are the important components of dose-escalation schemes. Once the starting dose is determined, the subsequent dose levels need to be established. Sometimes they are determined by the modified Fibonacci sequence. A Fibonacci sequence is a sequence where each number is the sum of the previous two numbers in the sequence. For instance, $\{1, 1, 2, 3, 5, 8, 13, 21, \ldots\}$ is a Fibonacci sequence. Dose increments follow the percentages of increments in this sequence and are $\{100, 50, 67, 60, 63, 62, \ldots\}$. A slightly modified version of this sequence is used in practice where the increment decreases. However, simple dose levels like 5 mg/kg, 10 mg/kg, 15 mg/kg, 20 mg/kg are often used in many studies.

Dose assignment is the way in which new patients entering the trial are allocated to doses. It depends on the design chosen, as different designs have different allocation rules. A dose allocation design to be used in early phases should be such that it does not expose too many patients either to subtherapeutic or to toxic doses. Often, patients are treated in cohorts of size 3 or 6. Determining the number of cohorts that will be appropriate for a trial is also an interesting research problem.

By early phase, we mean phases I and II. Ratain et al. (1993) discussed the statistical and ethical issues that need to be addressed in the early phases of development of anticancer agents. Many phase I trials deal with the estimation of the best dose rather than testing a hypothesis about it. On the other hand, many standard phase II trials test a hypothesis in order to select the best dose out of a set of candidate

doses. It may seem contradictory that sometimes the patients in phase I trials experience less toxicity than those in phase II trials. But this may happen because many patients are often under-treated in phase I trials, whereas phase II trials treat many patients at a dose that can result in moderate to severe toxicity. The statistical issue in early phase trials is to locate the best dose in an efficient way. However, the ethical one concerns the minimisation of under- and over- treatment during a trial. A dose-finding design will be ideal if it reaches the best dose in an efficient and ethically appropriate way. In the following sections, we present some commonly used dose-allocation designs for early phases.

2.2 Phase I Designs

Although several improved statistical methods have been developed in recent years, many current studies still use a traditional 3+3 design (Le Tourneau et al., 2009). The 3+3 design, which we describe in more detail in the next subsection, is often used as the specific issues to be achieved in phase I trials are not stated clearly. According to O'Quigley et al. (1990), a phase I design should aim to: (1) minimise the number of under-treated patients and the number of over-treated patients; (2) minimise the number of patients needed to complete the study; and (3) rapidly escalate the dose in the absence of toxicity or rapidly de-escalate the dose in the presence of an unacceptable level of toxicity. It is possible in a 3+3 design to come to a conclusion by using only a few patients, but Reiner et al. (1999) and Lévy et al. (2001) showed that the probability of an incorrect recommendation for the MTD is very high for this design.

Phase I designs can be classified into two types: rule-based and model-based. Many rule-based designs only utilise information from the current cohort in allocating a dose to the next cohort. O'Quigley and Zohar (2006) called these memoryless designs, as the previous information is completely ignored. The opposite of these are the designs which carry information through the trial. Designs with memory are mostly model-based.

2.2.1 Rule-Based Designs

The essence of the rule-based designs is that they do not assume any parametric dose-response model, but they use instead pre-specified rules. Some of these designs are based on the up-and-down rule (Dixon and Mood, 1948), where escalation or de-escalation of dose depends on the occurrences of toxicity in the previous cohort. Commonly used rule-based designs are the 3+3 design, Storer's up-and-down designs, accelerated titration designs (Simon et al., 1997), pharmacologically-guided dose-escalation design and designs using isotonic regression.

3+3 Design

The 3+3 design is the most widely used design in clinical practice. Starting with a pre-specified number of doses $\{x^{(1)}, \ldots, x^{(d)}\}$, the design first assigns the dose $x^{(1)}$ to a cohort of three patients. Escalation to dose $x^{(2)}$ is carried out if none of the three patients experiences toxicity. The trial stops if at least two of the three patients have toxicities. The same dose $x^{(1)}$ is given to three additional patients if one of the initial three patients has a toxic response. Then, if only one of the six patients has toxicity, escalation to dose $x^{(2)}$ is made; otherwise, the trial stops. In such a design, the MTD is usually defined as the highest dose at which the observed toxicity rate is no more than 0.33. Some researchers claim that the MTD should be the dose at which 2 or fewer toxicities in six patients are observed. Therefore, it is recommended to check exactly six patients at the MTD, which may sometimes require a single de-escalation in the 3+3 design.

Simplicity of implementation and safety concerns made the 3+3 design very popular. However, the design is inefficient when the starting dose is very low and the dose increment is moderate. In such a case, the design requires an excessive number of steps to reach the desirable dose, which in turn means that many patients are treated at subtherapeutic doses and very few patients receive doses at or near the MTD. Also, the maximum probability of toxicity that the MTD can have is fixed once the definition is set. For instance, if we define the MTD as the dose at which

2 or fewer toxicities are observed in six patients, then the toxicity rate at that dose is less than or equal to 0.33.

Some modified versions of the design, such as 2+4, 3+3+3 and 3+1+1, are also available to accelerate the dose escalation (Storer, 2001). In the 2+4 design, an additional cohort of size 4 is added if one of the two individuals in the first cohort shows toxicity. The same stopping rule as the traditional 3+3 design is followed. In the 3+3+3 design, the same dose is applied to an additional cohort of size 3 if two individuals in the first two cohorts experience toxicity. The trial stops if three or more individuals in three cohorts show toxicity. The 3+1+1 design is a more aggressive design than 3+3 and is known as 'best-of-five' design in the literature. If one or even two of the individuals in the first cohort experience toxicity, the same dose is given to one more individual. If two individuals in the first four experience toxicity, the dose is administered to one more individual. The trial stops if three or more toxicities are observed in five individuals. Although the modified versions are aimed at accelerating the dose escalation, it is not clear in the literature whether they are completely better than the conservative 3+3 approach or even which is best out of all these modified versions.

Storer's Up-and-Down Designs

The 3+3 design and its modified versions that have been discussed only allow dose to be escalated upward. Therefore, as a precaution, the starting dose level is the lowest one and hence many patients are treated at the subtherapeutic doses. Another problem with the design is that, since many patients are treated at the low doses, it may take a long time to reach the MTD. To overcome these problems, Storer (1989) recommended three designs which allow both dose escalation and de-escalation, and do not require the lowest dose to be the starting dose. The proposed designs are as follows:

1. Cohort size is one at each dose. Escalate to the next higher dose if a nontoxic outcome is observed; otherwise, de-escalate to the next lower dose.

2. Similar to the design in 1, except that escalation is done only if two consecutive nontoxic responses are seen. De-escalate if a toxic response is seen.

3. Cohort size is three at each dose. Escalation is made if there is no toxic response, stay at the same dose level if one toxic response is seen and de-escalate if two or more toxic responses are seen.

Design 1 is not implemented as a single-stage design itself, but together with design 2 or 3 to make a two-stage design. All of these designs are implemented with a fixed sample size usually ranging from 12 to 36. At the end of a two-stage trial, a logistic regression model is fitted to the data and the MTD is determined from that model for a particular choice of target toxicity rate.

Pharmacologically-Guided Dose Escalation

Collins et al. (1990) proposed pharmacologically-guided dose escalation (PGDE). It needs the area under the concentration curve (AUC) in humans to be extrapolated from preclinical data and the AUC value at the LD_{10} in mice is usually used as the target. At the first stage of the design, dose escalation proceeds with one patient per dose level as long as the target AUC is not reached and typically at 100% dose increments between successive patients. When the target AUC is reached or if dose-limiting toxicity occurs, the design turns into stage 2 where dose escalation is carried out by following the traditional 3+3 design and the MTD is determined accordingly. Successive doses are increased by around 40% in the second stage.

The method has been found to produce good results for some cytotoxic agents, such as certain anthracyclines and platinum compounds. It has been found to be inappropriate for other classes of agents, such as antifolates where high inter-patient variability in pharmacokinetics exists (Berry et al., 2010). Some other practical issues, such as logistical difficulties in obtaining real-time pharmacokinetic results and problems in extrapolating preclinical pharmacokinetic data, impede the frequent use of PGDE. The method also suffers from the risk of exposing the next patient to a highly toxic dose if the AUC for the last patient was considerably lower due

to inter-patient variability in drug metabolism. One possible way to overcome this problem could be the consideration of inter-patient variability in the AUC.

Designs using Isotonic Regression

Leung and Wang (2001) introduced a design which uses the idea of isotonic regression to estimate the risk at each dose, so that toxicity is non-decreasing with dose. Generally speaking, if we have d dose levels $\{x^{(1)}, \ldots, x^{(d)}\}$ to be tested, the risk at $x^{(i)}$ $(1 \leq i \leq d)$ must satisfy the monotonic relationship with dose. For any dose $x^{(r)}$ below $x^{(i)}$ $(r \leq i)$ and any dose $x^{(s)}$ above $x^{(i)}$ $(s \geq i)$, the pooled estimate of risk can be expressed as

$$w_{r,i,s} = \frac{\sum_{j=r}^{s} \text{number of toxicities at } x^{(j)}}{\sum_{j=r}^{s} \text{number of patients tested at } x^{(j)}}. \tag{2.1}$$

The estimate of the risk of toxicity at $x^{(i)}$ is obtained by using the isotonic regression

$$\hat{q}_i = \min_{i \leq s \leq d} \max_{1 \leq r \leq i} w_{r,i,s}. \tag{2.2}$$

The idea is that $\hat{q}_i$ must be at least as large as any of $w_{1,i,s}, w_{2,i,s}, \ldots, w_{i,i,s}$ (or the maximum of these) for any s $(s \geq i)$. Similarly, $\hat{q}_i$ must be no larger than any of $w_{r,i,i}, w_{r,i,i+1}, \ldots, w_{r,i,d}$ (or the minimum of these) for any r $(r \leq i)$.

Starting with dose $x^{(i)}$, the algorithm for the design proceeds as follows:

1. Treat a cohort of c patients at dose $x^{(i)}$.

2. Evaluate the risk at different doses by using (2.2). Choose the dose for which $\hat{q}_i$ is closest to the target toxicity rate γ, where i is the level of the last dose used. If $\hat{q}_i < \gamma$, then escalate if $\gamma - \hat{q}_i \geq \hat{q}_{i+1} - \gamma$ for $i < d$; otherwise, continue at the same dose. If $\hat{q}_i \geq \gamma$, then de-escalate if $\gamma - \hat{q}_{i-1} < \hat{q}_i - \gamma$ for $i > 1$; otherwise, continue at the same dose.

3. Repeat steps 1-2 until some pre-specified stopping criterion is met.

Usually the target toxicity rate is set at $\gamma = 0.33$ and the cohort size is chosen to be 3. Although the method allows starting with any dose level between 1 and d, the safest option is to start with 1. The trial can be stopped based on two criteria: if the same dose has been assigned consecutively to 3 or 4 cohorts or if the trial reaches a sample of 24 patients. The MTD is the dose indicated for the next cohort when the trial stops. This design has the flexibility to choose any percentile as the target risk of toxicity and usually more than six patients are treated at the MTD, which lessens the variability of the estimate of the MTD.

The most attractive feature of the rule-based designs is that they are easy to implement and no specialised software is required for computation. But their operating characteristics are not very attractive. The designs often allocate doses based on the outcomes from the last cohort, rather than considering the cumulative data from all of the treated cohorts. Some of these designs are unable to establish a dose that meets any specific target toxicity. Despite all of these limitations, rule-based designs have been used in many clinical trials.

2.2.2 Model-Based Designs

Model-based designs are alternatives to rule-based designs and they assume a parametric model to establish the dose-response relationship. Such designs select a dose level that produces a target probability of toxicity using all of the accrued data in a trial. These designs are usually implemented under the Bayesian framework, as the sample size remains small at the early stages of a trial. The common model-based designs include the continual reassessment method, escalation with overdose control and others.

Continual Reassessment Method

The continual reassessment method (CRM) (O'Quigley et al., 1990) is a Bayesian model-based procedure for dose escalation. The design aims to reduce the number of patients at subtherapeutic doses and to obtain a more accurate estimate of the MTD.

Assume that d pre-specified doses $\{x^{(1)}, \ldots, x^{(d)}\}$ are available for an experimental drug. Let Y_k be the response from the kth patient ($k = 1, 2, \ldots, n$), which may be toxic or nontoxic with value 1 or 0, respectively. The method employs parametric models, such as the hyperbolic tangent model, logistic model or power model to characterise the dose-response relationship. For example, a one-parameter logistic model is given as

$$\psi(x_k, \beta) = \frac{\exp(3 + \beta x_k)}{1 + \exp(3 + \beta x_k)}, \tag{2.3}$$

where 3 is the assumed value for the intercept parameter (Ishizuka and Ohashi, 2001), β is the unknown slope parameter, x_k is the dose administered to the kth patient and ψ denotes the probability of toxicity. The advantage of using a one-parameter model is that it requires fewer patients to obtain precise estimate of the unknown parameter. However, such a model may not be flexible enough to represent the dose-response relationship accurately.

Let us assume that Θ is the parameter space for β. Initially, a prior distribution $g(\beta)$ for β is considered, which is updated sequentially. Denote the prior distribution for β for the kth patient by $f(\beta, S_k)$ where $S_k = \{y_1, y_2, \ldots, y_{k-1}\}$. This prior is the posterior distribution for β based on the outcomes from the first $k-1$ patients and is obtained as

$$f(\beta, S_k) = \frac{f(\beta, S_{k-1})\ L(\beta|x_{k-1}, y_{k-1})}{\int_\Theta f(\beta, S_{k-1})\ L(\beta|x_{k-1}, y_{k-1})\ d\beta} = \frac{g(\beta)\ \prod_{l=1}^{k-1} L(\beta|x_l, y_l)}{\int_\Theta g(\beta)\ \prod_{l=1}^{k-1} L(\beta|x_l, y_l)\ d\beta}, \tag{2.4}$$

where $L(\beta|x_l, y_l) = \{\psi(x_l, \beta)\}^{y_l}\ \{1 - \psi(x_l, \beta)\}^{1-y_l}$ is the likelihood function of β given the response y_l at x_l, the dose received by the lth patient. The mean of the posterior distribution of β, denoted by μ_k, is obtained as

$$\mu_k = \int_\Theta \beta\ f(\beta, S_k)\ d\beta. \tag{2.5}$$

With this mean, the risk of toxicity at each of the doses is updated by using the equation

$$\psi_{ik} = \psi(x^{(i)}, \mu_k), \quad i = 1, 2, \ldots, d. \tag{2.6}$$

For a pre-specified target probability of toxicity γ, the trial starts with the dose for which the prior probability of toxicity is around γ. That dose is chosen for the kth patient for which the absolute difference between the updated estimate of probability of toxicity and the target toxicity rate is minimum. The process continues until a fixed sample size n is achieved and the MTD is the dose that would be allocated to patient $n+1$ if he were in the trial.

To investigate how the CRM would perform if adopted in some completed trials, O'Quigley and Zohar (2006) reported some retrospective analyses. Their reanalysis of the studies in Giles et al. (2004), Gelmon et al. (2004), Bos et al. (2005), and Okamoto et al. (2006) clearly show the advantage of using the CRM over the 3+3 design. That is, it reaches the MTD more quickly and treats more patients at and close to the MTD.

There was considerable debate in the statistical literature about the original version of the CRM, as it starts with the initial MTD and also many patients are likely to be exposed to high toxicity because of skipping dose. Various modifications have been suggested to make it safer. Some of these include: (1) treating the first patient at the lowest dose level (Korn et al., 1994); (2) not allowing dose escalation for the next patient if a patient experiences toxicity (Faries, 1994); (3) increasing the dose by only one pre-specified level at a time (Goodman et al., 1995); and (4) treating several patients at the same dose level, especially for the higher dose levels (Goodman et al., 1995). Møller (1995) proposed a two-stage CRM in which the first stage involves an up-and-down method until the first toxicity is observed. The design then moves to the second stage and starts using the CRM. Also, it does not allow skipping of more than one dose level at a time during the second stage. Heyd and Carlin (1999) suggested stopping a trial if the estimated MTD achieves a pre-specified precision. The time-to-event continual reassessment method (TITE-CRM) is another modified version proposed by Cheung and Chappell (2000), which takes into account the time to toxicity for each patient in the dose-escalation procedure.

These modifications have been implemented in clinical practice.

O'Quigley and Shen (1996) proposed a new version of the CRM based on the classical likelihood approach of parameter estimation which is known as the continual reassessment maximum likelihood (CRML) method. There are some perceived difficulties by clinicians regarding the original version of the CRM. These include starting with the dose which is the best prior guess of the MTD, incorporating prior information regarding the model parameters and the numerical integration necessary in the implementation of the method. The CRML appears to deal with the above three difficulties. Since initially the new design uses an up-and-down design which starts with the lowest dose, it does not require the best prior guess of the MTD to start with. Also, as the design is based on the frequentist approach, it ignores any prior information regarding the parameters. The operating characteristics of the two methods are very similar. Although simulation studies show some minor differences in dose allocation during the trials, the final recommendations are almost the same. As the likelihood equation has no solution until a toxic outcome is observed, the CRML can be applied only after the occurrence of such an outcome. To overcome this, a suggestion is to initially use either a standard up-and-down procedure or the CRM until a toxic outcome is observed, after which dose allocation can be based on the CRML.

Escalation with Overdose Control

As the original version of the CRM received a lot of criticism for its potential for exposing patients to overly toxic doses, various suggestions have been made to improve it further as discussed earlier. In a further attempt, Babb et al. (1998) introduced an alternative method that directly reduces the chance of overdosing. The main argument is that allocating doses which are close to the MTD, as in the CRM, is not attractive from an ethical point of view, and, therefore, their design tries to allocate doses more cautiously.

This new method is known in the statistical literature as escalation with overdose control (EWOC) and it is designed to approach the MTD as fast as possible, subject to the constraint that the predicted proportion of patients who receive an overdose is equal to a specified value α, called the feasibility bound. The method is implemented by computing the posterior cumulative distribution function of the MTD after each assignment of a dose to a patient. At the kth stage of a trial, the posterior cumulative distribution function of the MTD is a function given by

$$\pi_k(x_{\text{MTD}}) = P\{\text{MTD} \leq x_{\text{MTD}} | S_k\},$$

where x_{MTD} is the dose expected to produce toxicity in a specified proportion γ of patients and S_k denotes the responses available from the previous patients. The EWOC method then selects the dose x_k for the kth patient such that

$$\pi_k(x_k) = \alpha.$$

That is, the method selects the dose for each patient so that the predicted probability that it exceeds the MTD is equal to α. Generally, a trial is continued until the maximum sample size n is reached. Upon completion of the trial, the MTD is estimated by minimising the posterior expected loss with respect to some choice of loss function. One should consider asymmetric loss functions since underestimation and overestimation have very different effects. The simulation study shown in the paper revealed that, relative to the CRM, the EWOC method overdosed a smaller proportion of patients, exhibited fewer toxicities and estimated the MTD with slightly lower average bias and marginally higher mean square error. Similarly, relative to designs based on up-and-down schemes, EWOC treated fewer patients at subtherapeutic and toxic doses, treated a higher proportion of patients at doses near the MTD and estimated the MTD with lower average bias and mean square error.

Although the original paper suggested α to be 0.25, the feasibility of varying α dur-

ing the trial has also been studied by different authors. These include Babb and Rogatko (2001, 2004) and Cheng et al. (2004). This approach is motivated by the fact that, in the early stages of a trial, the uncertainty about the MTD remains high and a small value of α can prevent the method from choosing doses higher than the MTD. The uncertainty decreases with the advancement of a trial and the likelihood of selecting a dose which is higher than the MTD becomes smaller. Chu et al. (2009) proposed a hybrid method in which EWOC begins with $\alpha = 0.1$ and then α gradually increases according to a fixed schedule up to $\alpha = 0.5$.

Tighiouart et al. (2005) addressed the issue of the choice of prior distributions for x_{MTD} and ρ_0, which are the MTD and the probability of toxicity at the starting dose, respectively. They extended the class of restrictive priors used in the original version of EWOC by relaxing some of the constraints placed on (ρ_0, x_{MTD}). Through simulation, they showed that a candidate joint prior distribution for (ρ_0, x_{MTD}) with negative a priori correlation between these two parameters could lead to a safer trial than the one which assumes independent priors.

Babb and Rogatko (2001) extended the idea of EWOC further to include covariates, so that patient-specific dose allocation is possible.

Cheung (2005) studied the coherence conditions of dose-finding studies in the context of phase I clinical trials. Many phase I designs are outcome-adaptive, as the selection of dose for the next patient depends on the accumulated observations. An escalation of dose for the next patient is said to be coherent when the outcome from the previous patient was not toxic. Similarly, a de-escalation is said to be coherent if the last outcome was a toxic one. A design with these conditions is called coherent. The feature limits the risk of exposing patients to highly toxic doses. As reported by Cheung (2005), most of the phase I designs like the 3+3, the CRM and EWOC have this attractive feature. Incoherence of these designs may happen due to ad hoc

modifications that are often made in the course of a trial.

2.3 Phase II Designs

In phase II studies, the focus moves from toxicity to efficacy. The investigators want to examine whether a drug has sufficient efficacy to be studied in the next phase. Phase II trials are mainly divided into two parts: IIA and IIB. The phase IIA trials are usually single-arm studies and devoted to assessing the efficacy of the experimental drug with the goal of identifying the best dose. The phase IIB trials are multi-arm studies and aim to compare the experimental drug with other standard drugs, so that large-scale comparison is possible in phase III for the most promising drug. Multi-stage designs are often useful in phase IIA trials, so that trials can stop early due to futility or efficacy.

A two-stage design is a special case of a multi-stage design. The essence of a two-stage design is that in the first stage a small group of patients are enrolled and enrolment of another group of patients in the second stage depends on the outcomes from the previous stage. The motivation behind a two-stage design is that we do not want to enrol a large group of patients if a drug is not found to be promising. A first of this type of two-stage design was proposed by Gehan (1961). The method received criticism since the sample size remains fixed in the first stage. Use of a fixed sample size ignores the possibility of stopping a trial early. Simon (1989) proposed two approaches to find a two-stage design. The first one is the minimax design that minimises the maximum trial sample size. The second is the optimal design that minimises the expected sample size under the null hypothesis that the true response probability is equal to a specified value.

Jung et al. (2001) proposed a graphical method to search for a design that is a compromise between the optimal and minimax designs. This approach helps in finding a design that has sample size close to that of the minimax design and expected sample size close to that of the optimal design. Jung et al. (2004) proposed a family

of two-stage designs that are admissible according to a Bayesian decision-theoretic criterion. It is based on an ethically justified loss function, which is a weighted average of the maximum sample size and the expected sample size under the null hypothesis. An admissible design is one that minimises the loss function for some chosen weights. The family, as special cases, includes the minimax, optimal and compromise designs. Mander et al. (2012) extended the methodology by incorporating an additional term in the loss function. More specifically, the expected sample size under the alternative hypothesis that the true probability of response is above the value specified in the null is also considered in the loss function.

Bryant and Day (1995) allowed the simultaneous monitoring of efficacy and toxicity in a two-stage design. Mander and Thompson (2010) suggested a two-stage design for cancer clinical trials that minimises the expected sample size under the alternative hypothesis. The Simon two-stage design allows early stopping for futility only. However, the above new design allows stopping for efficacy. Wason et al. (2011) showed that the consideration of a continuous endpoint in a two-stage cancer clinical trial can reduce the sample size. This gain is significant in the sense that it will reduce the development time of a drug.

On the other hand, phase IIB is carried out if a drug passes through phase IIA. The trials in this phase are smaller in comparison with the phase III trials. They are often randomised, multi-arm trials with the aim of identifying the optimal dose for extensive study in the next phase.

All of the phase I designs that we have described in the previous section utilise the toxicity data only and make the implicit assumption that higher efficacy rates are associated with higher doses. This assumption may not be true for all classes of drugs. Also, an independent study for establishing the efficacy of the drug through phase IIA trials can extend the time for development. To overcome all of these shortcomings, we need designs that can take care of both toxicity and efficacy data:

see Section 1.2 for more explanation. In the following section, we describe some such designs which combine the objectives of both phase I and phase II.

2.4 Designs using Efficacy and Toxicity as Endpoints

These designs are appropriate when we have no reason to assume that both toxicity and efficacy follow the same pattern of relationship with dose, and we also want to minimise the time and costs associated with the drug development. The main essence of these designs is that they find a dose for further studies which is both safe and efficacious.

Thall and Russell (1998) developed a dose-finding method that satisfies efficacy and also safety requirements. The dose-response outcome in this case is trinomial, accounting for both efficacy and toxicity. The outcome is categorised as neutral, efficacious or toxic. The dose-response relationship is modelled by assuming the proportional odds (PO) model (McCullagh and Nelder, 1989). The trial starts with the lowest dose assigned to a cohort of patients. At each step, it determines the set of acceptable doses based on minimum efficacy and maximum toxicity requirements. The acceptable dose, for which the efficacy criterion probability is largest, is allocated to the next cohort. The trial terminates if none of the doses are acceptable; otherwise, it continues until the maximum sample size is reached. The method treats a sufficient number of patients to estimate the rates of efficacy and toxicity at the selected dose with a given reliability. But it suffers from the limitation that in the settings where all of the doses have acceptable toxicity with higher efficacy at the higher doses, it does not escalate to the more desirable doses with high probability. So the method often fails to detect the best dose in the presence of a number of candidate doses.

Thall and Cook (2004) proposed another method based on the trade-offs between treatment efficacy and toxicity. The method uses the previous methodology for finding acceptable doses. However, it computes a desirability index for all of the

acceptable doses. The index depends on the marginal probabilities of efficacy and toxicity of the corresponding doses. As in the previous case, if no acceptable dose is available, the trial terminates and no dose is recommended for further evaluation. Otherwise, the dose with maximum desirability is assigned to the next cohort, subject to the condition that no untried dose may be skipped when escalating. The trial is continued until the maximum sample size is reached and the dose with maximum desirability is recommended for further studies. It provides a substantial improvement over the earlier version and also accommodates bivariate binary outcomes. The method is known as EffTox in the statistical literature.

Zhang et al. (2006) proposed another such idea considering trinomial responses. The design has a similar approach to that of Thall and Russell (1998), but it utilises a more flexible continuation ratio (CR) model. The design selects a dose based on a optimal dose selection criterion that is expressed as the difference between the probabilities of efficacy and toxicity multiplied by a scalar in the range between 0 and 1. However, the advantage of this scale parameter is not clear. There are dose-response scenarios where consideration of such an optimality criterion will lead to the recommendation of doses which are far away from the true optimal dose. The method is popularly known as TriCRM. The previous two approaches involve considerable effort to elicit priors. In that regard, TriCRM is a simple alternative.

Dragalin and Fedorov (2006) suggested an adaptive procedure considering efficacy and toxicity as endpoints. The modelling of these endpoints is based on either Gumbel bivariate binary logistic regression or the Cox bivariate binary model. They express a dose-finding problem in terms of a penalised D-optimality criterion. The design maximises the information under the control of a penalty function for treating patients at doses which are too low or too high.

Thall et al. (2008) presented a dose-finding procedure based on bivariate outcomes that incorporates patient covariates and dose-covariate interactions. This is an ex-

tension of the methodology in Thall and Cook (2004) to allow for covariate effects. More recently, Thall and Nguyen (2012) proposed a new approach based on elicited utilities of the possible dose-response outcomes.

Since the model-based designs use all of the available data, they are efficient in finding the best dose for further studies. They also avoid treating many patients at the subtherapeutic doses. Because most of the designs are based on a Bayesian framework, the success of the trials depends on the assumed prior distribution for the model parameters. Specialised software is required for identifying the best dose at each step of a trial. Since these designs have attractive operating characteristics, they are becoming more popular.

2.5 PK-Guided Designs

Although clinical researchers possess the opinion that PK information has an important role in clinical response (Govindarajulu, 1988), very little effort has been made so far to incorporate such information in dose-response studies. Piantadosi and Liu (1998) described a method for incorporating PK information as a covariate in a dose-response model suitable for binary toxicity responses. In particular, they put the AUC in the dose-response model. Their objective is to consider drug concentration rather than just dose administered. They want to prevent patients from receiving too high a dose of the experimental drug. The study demonstrates that the efficiency and accuracy of phase I clinical trials can be improved by incorporating such information. Although the implementation of the method requires accurate PK data, they have not used the theory of optimal design to collect the data. Moreover, estimating the AUC following the usual approaches will require many blood samples to be collected.

Whitehead et al. (2007) presented an approach for phase I trials based on simultaneous monitoring of pharmacokinetic and pharmacodynamic responses. Following a logarithmic transformation, a linear model is employed to relate dose to the pharma-

cokinetic endpoint and a quadratic model to relate the pharmacokinetic endpoint to the pharmacodynamic endpoint. A logistic model is used to relate the pharmacokinetic endpoint to the risk of an adverse event. The approach allows us to relate the amount of drug absorbed, reflected by the AUC, rather than the dose, to the risk of an adverse event. The doses which lead to an excessive plasma concentration or high risk of toxicity are avoided during each stage of the trial. The final dose at each stage is selected based on the predictive distribution of the pharmacodynamic endpoint. From the discussion, it appears that the method is appropriate for phase I trials only.

The methodology in Zhou et al. (2008) is for phase I trials in healthy volunteers. For each individual in the trial, it monitors two continuous pharmacokinetic measures AUC and $C_{\max}$, and a binary indicator variable for an undesirable event. The method uses no dose for which the posterior value of $P(\text{AUC} > L_{\text{AUC}})$, or $P(C_{\max} > L_C)$ or $P(\text{DLT})$ is greater than 0.2. The safety limits L_{AUC} and L_C are set prior to the start of a trial and are expected to be obtained from expert opinion in preclinical studies. The smallest dose that satisfies all three conditions is recommended at the end of each step of the trial.

Observe that the methods that incorporate PK measurements in dose-escalation are mostly for phase I trials. To our knowledge, no method is available in the literature that can explicitly take into account PK data along with efficacy and toxicity endpoints. Moreover, even in the methods for phase I trials, we have not seen the use of population optimum design for PK sampling. The use of optimum design can reduce the number of samples and can also provide precise estimates of the model parameters. This is very important, as often the parameters have physical meanings. Our design is different, since it is for seamless phase IB/IIA trials and it also considers population optimum design for PK sampling.

Chapter 3

Population PK Models and Design

3.1 Introduction

The concentration of a drug depends on absorption, distribution, metabolism and excretion. Since these mechanisms differ among patients, an approach to describe PK data should be such that it can capture the inter-patient variability. Moreover, it should allow us to generalise the findings to a population of patients. Population modelling, based on mixed effects models, is one such approach that helps in understanding how the individuals from a population differ from one another. Therefore, this chapter begins with an introduction to the population approach to modelling PK data.

Starting with the underlying idea, we introduce compartmental models that are commonly used in the analysis of PK data. These models are based on the solutions to differential equations representing the distribution of a drug in the body's compartments. Such models are non-linear in the parameters. Therefore, the population approach in this thesis will lead us to non-linear mixed effects models.

We introduce the general non-linear mixed effects model and derive the associated Fisher information matrix. To illustrate the general model, we consider two specific PK models and the analytical forms are obtained for the information matrices. We then discuss different optimality criteria in the context of design. The properties of

the parameters AUC and $C_{\max}$ are also studied in detail. The results will assist in the optimum dose selection in an adaptive clinical trial.

3.2 Population Approach

The population approach to PK modelling quantifies the effect of a drug in a population of patients. It allows the quantification of variability in plasma concentration over a patient population.

For a drug, the concentration profiles over individuals often have similar shapes but they may vary: see Figure 3.1. The data presented are from an experiment on the pharmacokinetics of the drug *indomethacin* (Kwan et al., 1976). The same dose was given to six individuals through bolus intravenous injection. The plasma concentrations of each of the individuals were then measured in mcg/ml 11 times during a period of 8 hours and 15 minutes. The measurement time points were the same for all six individuals. These data are also available in the *R* library *nlme* as the object *Indometh*. Figure 3.1 indicates that the concentration curves have a similar shape but they vary across the individuals. For instance, the peaks for individuals 1, 4 and 2 are close to each other, but they are lower than those for individuals 5, 6 and 3.

If we can assume the underlying mechanism is the same for all individuals, then we can use the same regression function but with individual parameter values for each subject. Thus, we can think of individual parameter values to be realisations of random variables. The mean parameters provide a profile for a typical individual. The population approach is based on non-linear mixed effects (NLME) models and have been widely used in the literature for many years (Sheiner et al., 1972; Yuh et al., 1994; Sheiner and Steimer, 2000).

Since sampling PK responses involves a cost, it is desirable to keep the number of samples and the number of patients as low as possible during an experiment. Also, a trial should be conducted in such a way that it ensures reliable estimates of the

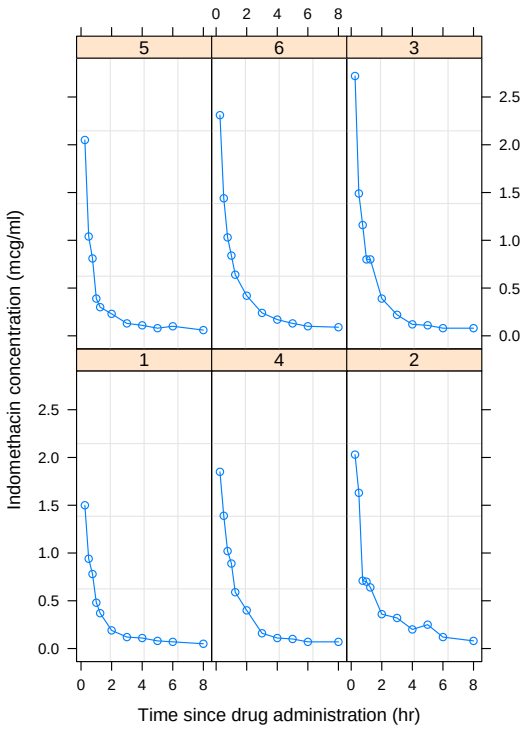

Figure 3.1: Concentration profiles of six individuals following the intravenous injection of *indomethacin.*

parameters. The theory of optimal design for non-linear mixed effects models helps to achieve this.

3.2.1 PK Compartmental Models

The recording of plasma concentration over time after the administration of a drug gives a concentration profile as in Figure 3.1. The concentration at any time depends on dose and the processes of absorption, distribution, metabolism and excretion. Therefore, the concentration of a drug at a given time point can be expressed through a function that incorporates dose and the rates of ADME. The models which serve this purpose are compartmental and are non-linear in the parameters.

In compartmental models, the body is thought to consist of several compartments. The central compartment is the one that remains in any model. A compartment is a homogeneous unit that is used to represent a group of tissues with similar rates of drug distribution. Usually one- to three-compartment models are used to quantify concentration-time profiles. If a drug is found to be distributed in the tissues of

a body very quickly, we may use a one-compartment model. If it appears that the drug is distributed at slower rates in different tissues, use of either a two- or a three-compartment model would be reasonable.

One-Compartment Model

In a one-compartment model, the body is regarded as a single compartment. The tissues in that compartment have a high rate of drug intake. It is assumed that the

Figure 3.2: A one-compartment model.

drug is instantaneously distributed throughout the body after the administration and also that it achieves equilibrium between the tissues. Figure 3.2 (Dubois et al., 2011) depicts the structure for a one-compartment model, where the central compartment has the volume of distribution V of the drug. The drug is cleared from that compartment at the rate k_e.

One-Compartment Model without Absorption

Assume that a drug is administered intravenously with bolus injection. It is then absorbed immediately in the central compartment. The drug is eliminated from the body following the first-order kinetics, where the speed of the elimination is proportional to the amount of drug left in the compartment. This can be expressed in terms of the equation

$$\frac{dX(t)}{dt} = -k_e X(t),$$

where $X(t)$ is the amount of drug in the central compartment at time t and k_e is the elimination rate. Initially, $X(0) = x$, where x is the dose that has been administered. To obtain the solution to this differential equation, we integrate both sides of the equation

$$\frac{dX(t)}{X(t)} = -k_e dt,$$

which leads to

$$X(t) = ce^{-k_e t},$$

where c is a constant. Using the initial condition at $t = 0$, the above equation gives $c = x$, and so

$$X(t) = xe^{-k_e t}.$$

Therefore, the concentration at time t can be written as

$$C(t) = \frac{X(t)}{V} = \frac{x}{V}e^{-k_e t} = \frac{x}{V}e^{-\frac{Cl}{V}t}, \tag{3.1}$$

where Cl is the clearance of the drug, V is the volume of distribution and $k_e = Cl/V$. A description of these parameters is given in Section 3.3.

One-Compartment Model with First-Order Absorption

Now we assume that the administered drug has an absorption phase following the first-order kinetics. That is, the speed of the absorption is proportional to the amount of drug that is yet to be absorbed. Then the processes of absorption and elimination can be described by the following set of differential equations:

$$\begin{aligned}\frac{dX_1(t)}{dt} &= k_a X_2(t) - k_e X_1(t),\\ \frac{dX_2(t)}{dt} &= -k_a X_2(t),\end{aligned}$$

with initial conditions $X_1(0) = 0$ and $X_2(0) = x$. Here, $X_1(t)$ is the amount of drug in the central compartment at time t, $X_2(t)$ is the amount of drug that is yet to be absorbed and k_a is the absorption rate. From the solution to this set of differential

equations given in Section A.1, it can be shown that the concentration at time t is

$$C(t) = \frac{X_1(t)}{V} = \frac{xk_a}{V(k_a - k_e)}(e^{-k_e t} - e^{-k_a t}). \tag{3.2}$$

Two-Compartment Model

For many drugs, the body cannot be assumed to be a single homogeneous unit, but it is sufficient to consider two compartments: a central one and a peripheral one. The tissues in the central compartment are highly perfused such as the heart, kidney, liver, lung and brain. The peripheral compartment comprises of less perfused tissues such as muscle and skin. So the tissues and plasma in the central compartment can

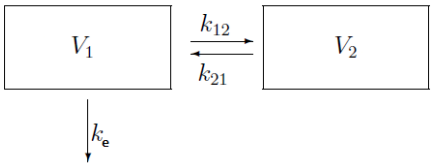

Figure 3.3: A two-compartment model.

absorb the drug rapidly, whereas those in the peripheral compartment absorb the drug at a slower rate. The central and peripheral compartments have V_1 and V_2 as the volumes of distribution, respectively. The rates of transfer from the central compartment to the peripheral compartment and back are k_{12} and k_{21}, respectively. The drug is eliminated from the central compartment at the rate k_e. Figure 3.3 represents such a model.

Two-Compartment Model with First-Order Absorption

The drug is absorbed and eliminated via the central compartment. The exchange between the central compartment and the peripheral compartment follows first-order kinetics. The overall process can be explained by the following set of differential equations:

$$\frac{dX_1(t)}{dt} = -(k_e + k_{12})X_1(t) + k_{21}X_2(t) + k_aX_3(t),$$
$$\frac{dX_2(t)}{dt} = k_{12}X_1(t) - k_{21}X_2(t),$$
$$\frac{dX_3(t)}{dt} = -k_aX_3(t),$$

with initial conditions $X_1(0) = 0$, $X_2(0) = 0$ and $X_3(0) = x$. Here, $X_1(t)$ and $X_2(t)$ are the amounts of drug in the central and peripheral compartments, respectively, and $X_3(t)$ is the amount of drug to be absorbed at time t. The concentration in the central compartment at a given time point after the administration of a dose can obtained by solving the above set of differential equations. The derivation will follow similar steps to the one shown in Section A.1.

Three-Compartment Model

This is an extension of the two-compartment model, where the drug is distributed at a very slow rate to certain tissues such as fat and bone. These tissues constitute the third compartment.

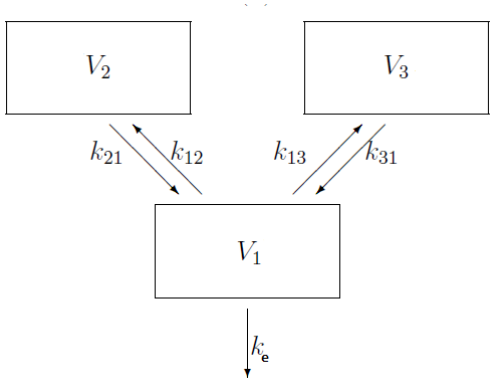

Figure 3.4: A three-compartment model.

The mechanism for such a three-compartment model is illustrated in Figure 3.4 (Dubois et al., 2011).

Three-Compartment Model with First-Order Absorption

For the three-compartment model when the drug is absorbed following the first-order absorption, the following set of differential equations express the structure of

the process:

$$\begin{aligned}
\frac{dX_1(t)}{dt} &= -(k_e + k_{12} + k_{13})X_1(t) + k_{21}X_2(t) + k_{31}X_3(t) + k_aX_4(t),\\
\frac{dX_2(t)}{dt} &= k_{12}X_1(t) - k_{21}X_2(t),\\
\frac{dX_3(t)}{dt} &= k_{13}X_1(t) - k_{31}X_3(t),\\
\frac{dX_4(t)}{dt} &= -k_aX_4(t),
\end{aligned}$$

with initial conditions $X_1(0) = 0$, $X_2(0) = 0$, $X_3(0) = 0$ and $X_4(0) = x$. Here, $X_1(t)$, $X_2(t)$ and $X_3(t)$ are the amounts of drug in the central and two peripheral compartments, respectively, and $X_4(t)$ is the amount of drug yet to be absorbed at time t. This set of differential equations can be solved to obtain the concentration in the central compartment at any time following the administration of a dose. The derivation is similar to the one shown in Section A.1 for the one-compartment model.

Since most of the kinetic functions are derived from differential equations, as shown above, they are non-linear in the parameters. These parameters often have physical interpretations. For example, those in (3.1) and (3.2) have their own meanings, as will be explained in Section 3.3.

From the discussion in Section 3.2, we know that variability in the concentration profiles is very likely to be present among patients after receiving the same dose. We always try to identify the sources of such variability in an attempt to model it accurately. The theory of mixed effects models allows us to do this. Since compartmental models, which are non-linear in the parameters, are to be used and as inter-patient variability in the concentration profiles exists, we plan to use non-linear mixed effects models for the purpose of modelling PK data.

3.3 Important PK Parameters

In this section, we list some commonly used PK parameters that will be used throughout the thesis. Some of these also appear in the functions that express

concentration as a function of time.

Volume of Distribution

The volume of distribution is defined as that volume of plasma in which the total amount of drug in the body would be required to be dissolved in order to reflect the drug concentration attained in plasma. It is denoted by V and quantifies the distribution of a medication between plasma and the rest of the body after dosing. It can be calculated by dividing the amount of drug in the body (X) by the concentration (C), that is,

$$V = \frac{X}{C}.$$

Clearance

The clearance is defined as the amount of plasma that is cleared of the drug per unit of time. It is denoted by Cl and its units are volume/time. Clearance can be expressed as

$$Cl = V \times k_e,$$

where k_e is the elimination rate.

Absorption Rate

The absorption rate, denoted by k_a, determines the time required for the administered drug to reach an effective plasma concentration. Therefore, the rate influences both the occurrence of the maximum concentration ($C_{\max}$) and the time to achieve it ($t_{\max}$).

Elimination Rate

The elimination rate describes the rate at which the drug is removed from the body and is denoted by k_e.

Area Under the Concentration Curve

The area under the concentration curve (AUC) reflects the overall amount of drug in the plasma after the administration of a dose. Its units are the product of concentration and time. For instance, they could be mg/l× hr. There are situations where a small increment in dose may produce a large increase in AUC (Graham and Workman, 1992). Also, drugs tend to produce toxicity at higher levels of AUC. Therefore, a drug can be administered safely by monitoring plasma concentration. The AUC can be regarded as one of the important PK parameters to guide in preventing patients being allocated toxic doses.

A non-parametric approach is often used to obtain the AUC. It employs the trapezoidal rule (Gabrielsson and Weiner, 2000) and requires a large number of samples per subject. An alternative is estimation based on a PK model. For many PK models, it is possible to obtain an analytical form for the AUC by integrating the model function over the design region. Use of optimal sampling time points will reduce the number of samples to be collected per subject, and so the parametric approach has an advantage over the non-parametric one. Hence, the parametric approach will be used here to find the AUC.

Maximum Concentration

The maximum concentration, denoted by $C_{\max}$, is the peak in the concentration that is achieved after the administration of a dose. The time at which this concentration is observed is denoted by $t_{\max}$. For a PK model, it is possible to find $t_{\max}$ and thus $C_{\max}$. A high value of $C_{\max}$ is likely to produce side effects and therefore monitoring this parameter is often crucial in clinical practice.

Figure 3.5 shows the concentration profiles for different individuals. Although they receive the same dose, the profiles are different. The shapes are similar but each individual has distinct values for the parameters. The AUC and $C_{\max}$ also vary over the individuals. To model these variabilities, an appropriate choice would be a

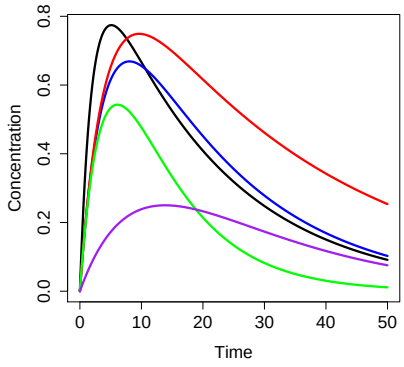

Figure 3.5: Concentration profiles for different individuals following the model function in (3.21).

non-linear mixed effects model.

3.4 Non-Linear Mixed Effects Model

We have a PK response to be measured on each individual i $(i = 1, \ldots, N)$. The response will be measured at n_i sampling times denoted by $\boldsymbol{\xi}_i = (t_{i1}, \ldots, t_{in_i})$. We describe the population design for the total number of observations $n = \sum_{i=1}^{N} n_i$ by

$$\boldsymbol{\Xi} = \{\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_N\}.$$

The responses can be represented through a known function $f : \mathbb{R}^+ \to \mathbb{R}^+$. Then, for the ith individual, we can write $\boldsymbol{f}(\boldsymbol{\xi}_i; \boldsymbol{\theta}_i) = (f(t_{i1}; \boldsymbol{\theta}_i), \ldots, f(t_{in_i}; \boldsymbol{\theta}_i))^T$, where $\boldsymbol{\theta}_i$ is the vector of all the individual parameters. The vector $\boldsymbol{\theta}_i$ has two components $\boldsymbol{\beta}$ and $\boldsymbol{b}_i$, where $\boldsymbol{\beta}$ is the p-vector of mean population parameters and $\boldsymbol{b}_i$ is the p-vector of random effects for the ith individual. These are related through a vector of functions $\boldsymbol{g}$, that is, $\boldsymbol{\theta}_i = \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i)$. The functions $\boldsymbol{g}$ are often chosen to be additive or exponential. We consider the simple case of $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$. It is assumed that $\boldsymbol{b}_i \sim N_p(\boldsymbol{0}, \boldsymbol{\Omega})$, with $\boldsymbol{\Omega}$ defined as a $p \times p$ diagonal matrix. We denote the kth diagonal element of $\boldsymbol{\Omega}$ by σ_k^2 $(k = 1, \ldots, p)$, which is the variance of the kth component of $\boldsymbol{b}_i$.

Let $\boldsymbol{y}_i = (y_{i1}, \ldots, y_{in_i})^T$ represent the vector of observations for the ith individual.

Then the statistical model is given by

$$\boldsymbol{y}_i = \boldsymbol{f}(\boldsymbol{\xi}_i; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i)) + \boldsymbol{\epsilon}_i, \tag{3.3}$$

where $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \ldots, \epsilon_{in_i})^T$ is the vector of random errors associated with the responses. We assume that $\boldsymbol{\epsilon}_i \sim N_{n_i}(\boldsymbol{0}, \sigma^2 \boldsymbol{I}_{n_i})$. The vector of all the population parameters to be estimated is represented by $\boldsymbol{\Psi} = (\boldsymbol{\beta}^T, \boldsymbol{\lambda}^T)^T$, where $\boldsymbol{\lambda} = (\sigma_1^2, \ldots, \sigma_p^2, \sigma^2)^T$ is the vector of variances.

3.4.1 Linearisation of the Model

Since our model is non-linear in the parameters, derivation of an analytical expression for the log-likelihood function is not possible. Various approximations to the information matrix have been proposed in the literature. A comparison of different methods is presented by Mielke (2012), who concludes that none of the methods is uniformly best. Lindstrom and Bates (1990) use a first-order Taylor series expansion of the model function f about the fixed parameters and the random effects at their estimates. At the design stage, there are no data to obtain the estimates, and so we use pre-specified values $\boldsymbol{\beta}^0$ for the fixed effects and the expectation of the random effects. That is, we approximate the log-likelihood using the first-order Taylor series expansion of the function $\boldsymbol{f}(\boldsymbol{\xi}_i; \boldsymbol{\theta}_i) = \boldsymbol{f}(\boldsymbol{\xi}_i; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))$ about $\boldsymbol{\phi}_i = (\boldsymbol{\beta}, \boldsymbol{b}_i)^T$ at $\boldsymbol{\phi}^0 = (\boldsymbol{\beta}^0, E(\boldsymbol{b}_i))^T = (\boldsymbol{\beta}^0, \boldsymbol{0})^T$.

Following the Taylor series expansion, we have

$$
\begin{aligned}
\boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i)) &\cong \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))\Big|_{\boldsymbol{\phi}^0} + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{\phi}_i}\right)^T\Big|_{\boldsymbol{\phi}^0}(\boldsymbol{\phi}_i-\boldsymbol{\phi}^0) \\
&= \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta}^0,\boldsymbol{0})) + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{\beta}}\right)^T\Big|_{\boldsymbol{\phi}^0}(\boldsymbol{\beta}-\boldsymbol{\beta}^0) \\
&\quad + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{b}_i}\right)^T\Big|_{\boldsymbol{\phi}^0}(\boldsymbol{b}_i-\boldsymbol{0}) \\
&= \boldsymbol{\mu_i} + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{\beta}}\right)^T\Big|_{\boldsymbol{\phi}^0}\boldsymbol{\beta} + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{b}_i}\right)^T\Big|_{\boldsymbol{\phi}^0}\boldsymbol{b}_i,
\end{aligned}
$$

where $\boldsymbol{\mu_i}$ is a $n_i \times 1$ vector of constants. With this approximation, the model in (3.3) becomes a linear mixed effects model of the form

$$
\boldsymbol{y_i} \;\cong\; \boldsymbol{\mu_i} + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{\beta}}\right)^T\Big|_{\boldsymbol{\phi}^0}\boldsymbol{\beta} + \left(\frac{\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial \boldsymbol{b}_i}\right)^T\Big|_{\boldsymbol{\phi}^0}\boldsymbol{b}_i + \boldsymbol{\epsilon}_i. \tag{3.4}
$$

With matrices $\boldsymbol{H}_i = (\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))/\partial\boldsymbol{\beta})^T\big|_{\boldsymbol{\phi}^0}$ and $\boldsymbol{L}_i = (\partial \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))/\partial\boldsymbol{b}_i)^T\big|_{\boldsymbol{\phi}^0}$, the model can be written as

$$
\boldsymbol{y_i} \;\cong\; \boldsymbol{\mu}_i + \boldsymbol{H}_i\boldsymbol{\beta} + \boldsymbol{L}_i\boldsymbol{b}_i + \boldsymbol{\epsilon}_i. \tag{3.5}
$$

Hence, we have

$$
E(\boldsymbol{y}_i) \cong \boldsymbol{E}_i = \boldsymbol{\mu}_i + \boldsymbol{H}_i\boldsymbol{\beta} \tag{3.6}
$$

and also, assuming that $\boldsymbol{b}_i$ and $\boldsymbol{\epsilon}_i$ are independent, we obtain

$$
\mathrm{Var}(\boldsymbol{y}_i) \cong \boldsymbol{V}_i = \boldsymbol{L}_i\boldsymbol{\Omega}\boldsymbol{L}_i^T + \sigma^2\boldsymbol{I}_{n_i}. \tag{3.7}
$$

Since the matrices $\boldsymbol{H}_i$ and $\boldsymbol{L}_i$ are evaluated at $\boldsymbol{\phi}^0$, they no longer depend on the parameters. However, the design solution will depend on the assumed prior values $\boldsymbol{\phi}^0$.

3.4.2 Fisher Information Matrix

The Fisher information matrix (FIM) for the ith individual and design $\boldsymbol{\xi}_i$ is given by

$$\boldsymbol{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i) = E\left[-\frac{\partial^2 l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i)}{\partial \boldsymbol{\Psi} \partial \boldsymbol{\Psi}^T}\right],$$

where $l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i)$ is the log-likelihood function for individual i.

Since $\boldsymbol{b}_i$ and $\boldsymbol{\epsilon}_i$ are assumed to be normal, the log-likelihood function is approximated by

$$l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i) \cong \log\left[(2\pi)^{-\frac{n_i}{2}} |\boldsymbol{V_i}|^{-\frac{1}{2}} \exp\{-\frac{1}{2}(\boldsymbol{y}_i - \boldsymbol{E}_i)^T \boldsymbol{V}_i^{-1}(\boldsymbol{y}_i - \boldsymbol{E}_i)\}\right].$$

That is,

$$-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i) \cong n_i \log(2\pi) + \log|\boldsymbol{V}_i| + (\boldsymbol{y}_i - \boldsymbol{E}_i)^T \boldsymbol{V}_i^{-1}(\boldsymbol{y}_i - \boldsymbol{E}_i). \tag{3.8}$$

Successive differentiation of (3.8) with respect to β_l and β_m gives

$$\frac{\partial(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \beta_l} \cong -2(\boldsymbol{y}_i - \boldsymbol{E}_i)^T \boldsymbol{V}_i^{-1}\left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_l}\right)$$

and

$$\frac{\partial^2(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \beta_m \partial \beta_l} \cong 2\left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_m}\right)^T \boldsymbol{V}_i^{-1}\left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_l}\right).$$

Therefore,

$$E\left(-\frac{\partial^2 l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i)}{\partial \beta_m \partial \beta_l}\right) \cong \left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_m}\right)^T \boldsymbol{V}_i^{-1}\left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_l}\right). \tag{3.9}$$

Differentiation of (3.8) with respect to λ_l gives

$$\frac{\partial(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \lambda_l} \cong \operatorname{tr}\left(\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\right) + (\boldsymbol{y}_i - \boldsymbol{E}_i)^T\left(-\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\right)(\boldsymbol{y}_i - \boldsymbol{E}_i), \tag{3.10}$$

since

$$\frac{\partial \log |\boldsymbol{V}_i|}{\partial \lambda_l} = \text{tr}\left(\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\right)$$

and

$$\frac{\partial \boldsymbol{V}_i^{-1}}{\partial \lambda_l} = -\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}.$$

Also,

$$\frac{\partial^2(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \beta_m \partial \lambda_l} \cong 2(\boldsymbol{y}_i - \boldsymbol{E}_i)^T\left(-\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\right)\left(-\frac{\partial \boldsymbol{E}_i}{\partial \beta_m}\right).$$

Therefore,

$$E\left(-\frac{\partial^2 l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i)}{\partial \beta_m \partial \lambda_l}\right) \cong \boldsymbol{0}. \tag{3.11}$$

Again, differentiation of (3.10) with respect to λ_m yields

$$\begin{aligned}\frac{\partial^2(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \lambda_m \partial \lambda_l} &\cong -\text{tr}\left(\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\right) \\ &+ (\boldsymbol{y}_i - \boldsymbol{E}_i)^T\left(\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1} + \boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\right)(\boldsymbol{y}_i - \boldsymbol{E}_i),\end{aligned}$$

so that

$$E\left(\frac{\partial^2(-2l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i))}{\partial \lambda_m \partial \lambda_l}\right) \cong -\text{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\right) + 2\ \text{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\right),$$

since, for a quadratic form, we have

$$E(\boldsymbol{X}^T\boldsymbol{A}\boldsymbol{X}) = \text{tr}(\boldsymbol{A}\boldsymbol{\Sigma}) + \boldsymbol{\mu}^T\boldsymbol{A}\boldsymbol{\mu},$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the mean and variance of $\boldsymbol{X}$, respectively.

Therefore,

$$E\left(-\frac{\partial^2 l_i(\boldsymbol{\Psi} \mid \boldsymbol{y}_i)}{\partial\lambda_m \partial\lambda_l}\right) \cong \frac{1}{2}\mathrm{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial\lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial\lambda_l}\boldsymbol{V}_i^{-1}\right). \tag{3.12}$$

Then the FIM for individual i can be approximated by the block diagonal matrix

$$\boldsymbol{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i) \cong \begin{bmatrix} \boldsymbol{A}_i & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{B}_i \end{bmatrix}, \tag{3.13}$$

where the elements of matrices $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ are

$$(\boldsymbol{A}_i)_{ml} = \left(\frac{\partial \boldsymbol{E}_i}{\partial\beta_m}\right)^T \boldsymbol{V}_i^{-1} \frac{\partial \boldsymbol{E}_i}{\partial\beta_l} \quad \text{for } m, l = 1, \ldots, p,$$

and

$$(\boldsymbol{B}_i)_{ml} = \frac{1}{2}\mathrm{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial\lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial\lambda_l}\boldsymbol{V}_i^{-1}\right) \quad \text{for } m, l = 1, \ldots, p+1.$$

The population FIM for the design $\boldsymbol{\Xi}$ is defined as the sum of N elementary Fisher information matrices, that is,

$$\boldsymbol{M}(\boldsymbol{\Psi}, \boldsymbol{\Xi}) = \sum_{i=1}^{N} \boldsymbol{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i).$$

For a single group of N individuals with identical designs, the population FIM is

$$\boldsymbol{M}(\boldsymbol{\Psi}, \boldsymbol{\Xi}) = N\boldsymbol{M}(\boldsymbol{\Psi}, \boldsymbol{\xi}). \tag{3.14}$$

Unlike in linear models, here the FIM depends on the parameters of the model. This is important as most of the criteria for design optimality are functions of the information matrix, as is explained in the next section.

Retout et al. (2001) use a first-order Taylor series expansion of the model function f about the expectation of the parameters to linearise the model. Although their linearised model is different from ours in (3.5), we have found that the FIMs in

both cases are the same. To derive the FIM in their approach, one needs to assume that $\boldsymbol{V}_i$ is independent of $\boldsymbol{\beta}$. But their expression for $\boldsymbol{V}_i$ clearly indicates that this will not be the case. However, in our approach, we do not need to make any such assumption. Both FIMs produce the same design. This is because, in the criteria for design optimality, only a function of the FIM is used and it will not matter whether the two linearised models are exactly the same.

3.5 Optimal Experimental Designs

Optimal design depends on the statistical model and it optimises a chosen design criterion. Compared to an optimal design, a non-optimal one requires a greater number of experimental runs to provide the same statistical efficiency. Use of optimal design thus reduces the cost of conducting an experiment. Utmost care should be taken in the choice of an optimality criterion so that it is consistent with the objective of an experiment.

3.5.1 Optimality Criteria

In this section, we discuss some widely used design criteria which are based on the books by Atkinson et al. (2007), Fedorov (1972), Berger and Wong (2009), and Fedorov and Hackl (1997). Out of all of the optimality criteria to be presented, our work only uses the D-criterion, since the objective is to minimise the variability in the parameter estimates.

D-Criterion

The uncertainty in a set of parameter estimators can be expressed in terms of the volume of a confidence ellipsoid. The precision of the estimators increases as the volume decreases. The D-optimality criterion minimises the volume of the confidence ellipsoid, which is a function of the determinant of the covariance matrix. More specifically, a design $\boldsymbol{\xi}_D^*$ is called D-optimal if it minimises $\Phi_D\{\boldsymbol{M}(\boldsymbol{\xi})\} = \log|\boldsymbol{M}^{-1}(\boldsymbol{\xi})|$, where $\boldsymbol{M}(\boldsymbol{\xi})$ is the FIM. That is,

$$\boldsymbol{\xi}_D^* = \arg\min_{\boldsymbol{\xi}} \log |\boldsymbol{M}^{-1}(\boldsymbol{\xi})|.$$

One important feature of the D-criterion is that the design obtained is invariant under linear transformations of the scale of the explanatory variables. This implies that, if the design region is changed, we can directly deduce the D-optimal design from the one previously constructed. This property may not hold for other optimality criteria.

A-Criterion

The A-optimality criterion minimises the sum or average of the variances of the parameter estimators. A design $\boldsymbol{\xi}_A^*$ is called A-optimal if it minimises $\Phi_A\{\boldsymbol{M}(\boldsymbol{\xi})\} = \text{trace}\{\boldsymbol{M}^{-1}(\boldsymbol{\xi})\}$. That is,

$$\boldsymbol{\xi}_A^* = \arg\min_{\boldsymbol{\xi}} \text{trace}\{\boldsymbol{M}^{-1}(\boldsymbol{\xi})\}.$$

This criterion suffers from a few drawbacks. Firstly, it is not invariant under linear transformations of the scale of the explanatory variables. Secondly, sometimes the variances of the parameter estimator may have very different magnitudes and therefore minimising the sum may mislead.

G-Criterion

Sometimes interest lies in predicting the response over the design region efficiently. A G-optimum design minimises the maximum standardised variance of the predicted response over the design region. A design $\boldsymbol{\xi}_G^*$ is called G-optimal if it minimises $\Phi_G\{\boldsymbol{M}(\boldsymbol{\xi})\} = \max_{t\in\mathcal{T}} \boldsymbol{\eta}^T(t)\boldsymbol{M}^{-1}(\boldsymbol{\xi})\boldsymbol{\eta}(t)$. That is,

$$\boldsymbol{\xi}_G^* = \arg\min_{\boldsymbol{\xi}} \max_{t\in\mathcal{T}} \{\boldsymbol{\eta}^T(t)\boldsymbol{M}^{-1}(\boldsymbol{\xi})\boldsymbol{\eta}(t)\},$$

where $\boldsymbol{\eta}(t) = (\partial f/\partial\theta_1, \ldots, \partial f/\partial\theta_p)^T$ is the vector of parameter sensitivities and $\mathcal{T}$ is the design region. This criterion has a connection to the D-optimality criterion.

Kiefer and Wolfowitz (1960) showed that the D- and G-optimality criteria are equivalent in the case of so-called continuous designs, which are probability measures on the discrete support of the design points. It is also known from their equivalence theorem that the standardised variance function for a continuous G-optimum design $\boldsymbol{\xi}^*$ is always less than or equal to the number of parameters p in the model, that is, $d(t, \boldsymbol{\xi}^*) \leq p$, with equality at the design points, where $d(t, \boldsymbol{\xi}^*) = \boldsymbol{\eta}^T(t)\boldsymbol{M}^{-1}(\boldsymbol{\xi}^*)\boldsymbol{\eta}(t)$ and $d(t^*, \boldsymbol{\xi}^*) = p$.

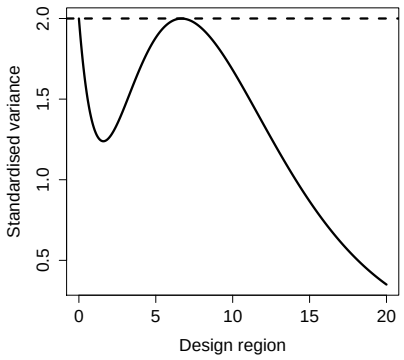

Figure 3.6: Standardised variance function plot for a continuous D-optimum design for the model in (3.15).

To give an illustration, let us consider the one-compartment PK model with bolus input and first-order elimination given by

$$y_i = \frac{x}{V}\exp\left(-\frac{Cl}{V}\ t_i\right) + \epsilon_i, \quad i = 1, \ldots, N, \tag{3.15}$$

where y_i is the concentration of a drug in the blood for the ith individual observed at time t_i, x is the dose received, V and Cl are the two parameters volume of distribution and clearance, and ϵ_i is the random error term. Figure 3.6 presents the standardised variance function plot for the evenly distributed D-optimal time points $\boldsymbol{\xi}^* = \{0, 6.627\}$, obtained for the values $\boldsymbol{\beta}^0 = (0.4, 0.06)^T$ of the parameters. The figure indicates that the design has a standardised variance function satisfying $d(t, \boldsymbol{\xi}^*) \leq 2$, where 2 is the number of parameters in the model, and also that equality holds at the design points. So the design is both G- and D-optimal. This property is quite often used to check the D-optimum designs obtained numerically and it is

used for constructing algorithms to calculate D-optimum designs. The equivalence theorem holds for continuous designs only. The D-optimum design that we use in the thesis is of discrete type and hence we cannot employ the theorem to judge the optimality of the design.

c-Criterion

When interest lies in a design to estimate a linear combination of the parameters, $\boldsymbol{c}^T\boldsymbol{\beta}$, with minimum variance, the criterion to be used is c. A design $\boldsymbol{\xi}_c^*$ is called c-optimal if it minimises $\Phi_c\{\boldsymbol{M}(\boldsymbol{\xi})\} = \boldsymbol{c}^T\boldsymbol{M}^{-1}(\boldsymbol{\xi})\boldsymbol{c}$. That is,

$$\boldsymbol{\xi}_c^* = \arg\min_{\boldsymbol{\xi}}\{\boldsymbol{c}^T\boldsymbol{M}^{-1}(\boldsymbol{\xi})\boldsymbol{c}\}.$$

This criterion can also be used to estimate a non-linear function of the parameters with minimum variance. For instance, assume that we have a non-linear function $g(\boldsymbol{\beta})$ to be estimated efficiently. The function can be linearised by a Taylor series expansion to give the components of vector $\boldsymbol{c}$ as

$$c_j(\boldsymbol{\beta}) = \frac{\partial g(\boldsymbol{\beta})}{\partial \beta_j}, \quad j = 1, \ldots, p,$$

evaluated at prior values $\boldsymbol{\beta}^0$.

3.5.2 Locally Optimal Design

In linear models, the information matrix does not depend on the model parameters. However, for non-linear models, the information matrix depends on the prior values of the model parameters. Therefore, the optimal design that we derive for a non-linear model is optimal locally only for a set of values of the parameters (Chernoff, 1953). Since the parameters are unknown and we want to estimate these through the experiment in an optimal way, the search for an optimal design starts with some prior values of the parameters.

Different approaches are available in the literature to overcome the dependence of

an optimum design on the unknown parameters. These include a sequential design, a Bayesian design and a maximin design (Atkinson et al., 2007). The sequential approach starts with some guess for the parameter values and as soon as the new estimates are available, they are replaced. The process continues until the parameter estimates are sufficiently precise or the experimental resources are exhausted. In the thesis, we employ the sequential approach for constructing designs. Since, in an adaptive dose-finding design, we update the estimates of the model parameters at each stage, they are likely to stabilise after several stages of the trial, as will the design points.

Optimisation Algorithms

In our work, we want to minimise the variability in the PK parameter estimates, and therefore the appropriate criterion for design optimality is D. The Cramér-Rao inequality tells us that the covariance matrix of the parameter estimators is greater than, and asymptotically approaches, the inverse of the FIM. Therefore, by minimising the inverse of the FIM through the D-criterion, we minimise the asymptotic lower bound for the variance of the estimated model parameters.

It has already been mentioned that the derived FIM for our non-linear mixed effects model has the same expression as in Retout, Duffull, and Mentré (2001). Bazzoli et al. (2010) developed *PFIM 3.2*, an *R* package, to evaluate and optimise designs in the context of population PK/PD experiments. It has a library of PK models with FIMs in Retout et al. (2001). Since both FIMs are the same, we can compute the D-optimum design for PK responses using *PFIM 3.2*.

The package applies two algorithms for design construction: the Fedorov-Wynn algorithm (Wynn, 1972) and the simplex algorithm (Nelder and Mead, 1965). The Fedorov-Wynn algorithm relies on the equivalence theorem of Kiefer and Wolfowitz (1960). It is an iterative algorithm that maximises the determinant of the FIM within a finite set of possible designs. It is important to note that maximisation of

the determinant of the FIM is equivalent to the minimisation of the determinant of the inverse of the FIM. Prior to the optimisation, one needs to provide the set of possible time points. It is therefore possible to avoid clinically unfeasible sampling times in this approach. Starting with an initial design, the algorithm then finds a design iteratively which satisfies the optimality criterion.

In many cases, feasible sampling times may not be known and therefore one can assume a continuous interval of time as the design region. The algorithm to be used in that case is the simplex algorithm, which is based on the method by Nelder and Mead (1965). The method uses the concept of a simplex, which is a generalised triangle in N dimensions. It iteratively generates a sequence of simplices to approximate an optimal point. An initial design is required to start the optimisation. From this initial design, initial vertices for the simplex algorithm are derived. At each iteration, the vertices of the simplex are ordered according to the objective function values. The worst vertex, where the function has the maximum value, is rejected and replaced with a new one. A new simplex is formed and the search is continued. The method thereby produces a sequence of simplices for which the function values at the vertices get smaller and smaller. The size of the simplex is reduced and the coordinates of the optimum point are found. In our work, we employ the simplex algorithm for the construction of designs, as the feasible sampling times are not known.

3.6 Parameter Estimation in NLME Models

A number of methods are available in the literature for fitting non-linear mixed effects models with ongoing debate as to which is the most accurate method. Sheiner and Beal (1980), Wolfinger and Lin (1997), and Lindstrom and Bates (1990) proposed methods based on linearisation. Pinheiro and Bates (1995), Vonesh (1996) and Wolfinger (1993) developed integral approximation methods. Kuhn and Lavielle (2005), Walker (1996) and Wang (2007) proposed methods that use the expectation-maximisation (EM) algorithm.

In the linearisation method, an approximate linear model is derived using a first-order Taylor series expansion of the model to obtain the analytical form for the likelihood function. A Laplace approximation, Gaussian quadrature or importance sampling are used in the integral approximation methods to obtain the marginal distribution of the response variable. The method then maximises the likelihood directly. Instead of a direct approximation to the marginal likelihood, the EM algorithm approximates the conditional expectation of the log-likelihood in the E-step and then maximises the expected log-likelihood to estimate the parameters in the M-step. Of these, the linearisation methods are the most popular due to their numerical simplicity.

The linearisation methods differ with respect to the expansion locus of the random effects. Beal and Sheiner (1982) suggested a method in which the likelihood function is based on a first-order Taylor series expansion of the model function about the mean of the random effects. Beal and Sheiner (1992) implemented this method in the software package *NONMEM*. Lindstrom and Bates (1990) developed a method in which the model function is expanded about the current estimate of the fixed parameters and their random effects. Pinheiro and Bates (2000) implemented the method in the package *nlme*, which is available in both the package *S-PLUS* and in *R*. This is the method which we use for estimation of the PK parameters in the examples presented in Chapter 6. It is worth mentioning that two different linearisation approaches are used in our work. These should not make a difference as one is for designing purposes and the other is related to parameter estimation.

3.7 PK Mixed Effects Model Examples

This section describes two models that we use for the concentrations for the purpose of simulations in Chapter 6. Although the route of administration of the dose is different, both models treat the whole body as a single compartment.

3.7.1 One-Compartment Model with Bolus Input and First-Order Elimination

The one-compartment PK mixed effects model with bolus input and first-order elimination is defined as

$$\begin{aligned} y_{il} &= f(t_{il};\boldsymbol{\theta}_i) + \epsilon_{il} \\ &= \frac{x}{V_i}\exp\left(-\frac{Cl_i}{V_i}\ t_{il}\right) + \epsilon_{il} \quad \text{for } i = 1,\ldots,N \text{ and } l = 1,\ldots,n_i, \end{aligned} \tag{3.16}$$

where y_{il} is the concentration of a drug in the blood for the ith individual observed at time t_{il}, x is the dose received by the individual, and $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i = (V_i, Cl_i)^T$ is the vector of parameters with $\boldsymbol{\beta} = (V, Cl)^T$ and $\boldsymbol{b}_i = (b_{V_i}, b_{Cl_i})^T$. We assume that $\boldsymbol{\epsilon}_i \sim N_{n_i}(\mathbf{0}, \sigma^2 \boldsymbol{I}_{n_i})$ and $\boldsymbol{b}_i \sim N_2(\mathbf{0}, \boldsymbol{\Omega})$, where $\boldsymbol{\Omega} = \text{diag}(\sigma_1^2, \sigma_2^2)$. We can write,

$$\boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{\theta}_i) = \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i)) = \begin{bmatrix} \frac{x}{V_i}\exp\left(-\frac{Cl_i}{V_i}\ t_{i1}\right) \\ \frac{x}{V_i}\exp\left(-\frac{Cl_i}{V_i}\ t_{i2}\right) \\ \vdots \\ \frac{x}{V_i}\exp\left(-\frac{Cl_i}{V_i}\ t_{in_i}\right) \end{bmatrix}. \tag{3.17}$$

To linearise the model, we expand it using a first-order Taylor series about $\boldsymbol{\phi}_i = (V, Cl, b_{V_i}, b_{Cl_i})^{\mathrm{T}}$ at $\boldsymbol{\phi}^0 = (V^0, Cl^0, 0, 0)^{\mathrm{T}}$. We obtain as an approximation the linear mixed effects model

$$\boldsymbol{y_i} \ \cong \ \boldsymbol{\mu}_i + \boldsymbol{H}_i\boldsymbol{\beta} + \boldsymbol{L}_i\boldsymbol{b}_i + \boldsymbol{\epsilon}_i,$$

where $\boldsymbol{\mu}_i$, $\boldsymbol{H}_i$ and $\boldsymbol{L}_i$ are as in Section 3.4.1. For this specific model, we have

$$
\begin{aligned}
\boldsymbol{H}_i &= \left[\begin{array}{cc} \dfrac{\partial f(t_{i1}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial V} & \dfrac{\partial f(t_{i1}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial Cl} \\ \vdots & \vdots \\ \dfrac{\partial f(t_{in_i}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial V} & \dfrac{\partial f(t_{in_i}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial Cl} \end{array} \right]_{\boldsymbol{\phi}^0} \\
&= \left[\begin{array}{cc} \dfrac{x}{(V^0)^2} \exp\left(-\dfrac{Cl^0}{V^0} t_{i1}\right) \left(\dfrac{Cl^0}{V^0} t_{i1} - 1\right) & -\dfrac{x\ t_{i1}}{(V^0)^2} \exp\left(-\dfrac{Cl^0}{V^0} t_{i1}\right) \\ \vdots & \vdots \\ \dfrac{x}{(V^0)^2} \exp\left(-\dfrac{Cl^0}{V^0} t_{in_i}\right) \left(\dfrac{Cl^0}{V^0} t_{in_i} - 1\right) & -\dfrac{x\ t_{in_i}}{(V^0)^2} \exp\left(-\dfrac{Cl^0}{V^0} t_{in_i}\right) \end{array} \right].
\end{aligned}
$$

Since $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$, we have $\boldsymbol{H}_i = \boldsymbol{L}_i$, where

$$
\boldsymbol{L}_i = \left[\begin{array}{cc} \dfrac{\partial f(t_{i1}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial b_{V_i}} & \dfrac{\partial f(t_{i1}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial b_{Cl_i}} \\ \vdots & \vdots \\ \dfrac{\partial f(t_{in_i}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial b_{V_i}} & \dfrac{\partial f(t_{in_i}; \boldsymbol{g}(\boldsymbol{\beta}, \boldsymbol{b}_i))}{\partial b_{Cl_i}} \end{array} \right]_{\boldsymbol{\phi}^0}.
$$

Furthermore, $E(\boldsymbol{y}_i) \cong \boldsymbol{E}_i = \boldsymbol{\mu}_i + \boldsymbol{H}_i\boldsymbol{\beta}$ and $\mathrm{Var}(\boldsymbol{y}_i) \cong \boldsymbol{V}_i = \boldsymbol{L}_i\boldsymbol{\Omega}\boldsymbol{L}_i^T + \sigma^2\boldsymbol{I}_{n_i}$. According to (3.13), the FIM for the ith individual can be approximated by the block diagonal matrix

$$
\boldsymbol{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i) \cong \left[\begin{array}{cc} \boldsymbol{A}_i & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{B}_i \end{array} \right], \tag{3.18}
$$

where

$$
(\boldsymbol{A}_i)_{ml} = \left(\frac{\partial \boldsymbol{E}_i}{\partial \beta_m}\right)^T \boldsymbol{V}_i^{-1} \frac{\partial \boldsymbol{E}_i}{\partial \beta_l} \quad \text{for } m, l = 1, 2 \text{ and } \beta_1 = V,\ \beta_2 = Cl,
$$

and

$$
(\boldsymbol{B}_i)_{ml} = \frac{1}{2}\mathrm{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial \lambda_m}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \lambda_l}\boldsymbol{V}_i^{-1}\right) \quad \text{for } m, l = 1, 2, 3 \text{ and } \boldsymbol{\lambda} = (\sigma_1^2, \sigma_2^2, \sigma^2)^T.
$$

Differentiation of $\boldsymbol{E}_i$ and $\boldsymbol{V}_i$ with respect to the components of $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}$, respectively, gives

$$\frac{\partial \boldsymbol{E}_i}{\partial V} = \boldsymbol{H}_i \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \frac{\partial \boldsymbol{E}_i}{\partial Cl} = \boldsymbol{H}_i \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

$$\frac{\partial \boldsymbol{V}_i}{\partial \sigma_1^2} = \boldsymbol{L}_i \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{L}_i^T, \qquad \frac{\partial \boldsymbol{V}_i}{\partial \sigma_2^2} = \boldsymbol{L}_i \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \boldsymbol{L}_i^T \qquad \text{and} \qquad \frac{\partial \boldsymbol{V}_i}{\partial \sigma^2} = \boldsymbol{I}_{n_i}.$$

Therefore,

$$(\boldsymbol{A}_i)_{12} = \begin{bmatrix} 1 & 0 \end{bmatrix} \boldsymbol{D}_i \begin{bmatrix} 0 \\ 1 \end{bmatrix} = (\boldsymbol{D}_i)_{12},$$

where $\boldsymbol{D}_i = \boldsymbol{H}_i^T \boldsymbol{V}_i^{-1} \boldsymbol{H}_i$.

Altogether, we have $\boldsymbol{A}_i = \boldsymbol{H}_i^T \boldsymbol{V}_i^{-1} \boldsymbol{H}_i$, since $\boldsymbol{e}_i^T \boldsymbol{X} \boldsymbol{e}_j = x_{ij}$, where $\boldsymbol{e}_i$ is a unit vector with 1 in the ith position and zeros elsewhere, and $\boldsymbol{X}$ is a matrix with elements x_{ij}. Continuing,

$$\begin{aligned}
(\boldsymbol{B}_i)_{12} &= \frac{1}{2}\mathrm{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial \sigma_1^2}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \sigma_2^2}\boldsymbol{V}_i^{-1}\right) \\
&= \frac{1}{2}\mathrm{tr}\left(\boldsymbol{L}_i \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{L}_i^T \boldsymbol{V}_i^{-1} \boldsymbol{L}_i \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \boldsymbol{L}_i^T \boldsymbol{V}_i^{-1}\right) \\
&= \frac{1}{2}\mathrm{tr}\left(\boldsymbol{W}_i \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{W}_i \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}\right),
\end{aligned}$$

where $\boldsymbol{W}_i = \boldsymbol{L}_i^T \boldsymbol{V}_i^{-1} \boldsymbol{L}_i$. Thus $(\boldsymbol{B}_i)_{12} = \frac{1}{2}(\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21}$. Similarly, $(\boldsymbol{B}_i)_{11} = \frac{1}{2}\{(\boldsymbol{W}_i)_{11}\}^2$, $(\boldsymbol{B}_i)_{22} = \frac{1}{2}\{(\boldsymbol{W}_i)_{22}\}^2$ and $(\boldsymbol{B}_i)_{21} = \frac{1}{2}(\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21}$. Also,

$$
\begin{aligned}
(\boldsymbol{B}_i)_{13} &= \frac{1}{2}\text{tr}\left(\frac{\partial \boldsymbol{V}_i}{\partial \sigma_1^2}\boldsymbol{V}_i^{-1}\frac{\partial \boldsymbol{V}_i}{\partial \sigma^2}\boldsymbol{V}_i^{-1}\right) \\
&= \frac{1}{2}\text{tr}\left(\boldsymbol{L}_i\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix}\boldsymbol{L}_i^T\boldsymbol{V}_i^{-1}\boldsymbol{I}_{n_i}\boldsymbol{V}_i^{-1}\right) \\
&= \frac{1}{2}\text{tr}\left(\boldsymbol{P}_i\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix}\right),
\end{aligned}
$$

where $\boldsymbol{P}_i = \boldsymbol{L}_i^T\boldsymbol{V}_i^{-2}\boldsymbol{L}_i$. Thus $(\boldsymbol{B}_i)_{13} = \frac{1}{2}(\boldsymbol{P}_i)_{11}$. Similarly, $(\boldsymbol{B}_i)_{23} = \frac{1}{2}(\boldsymbol{P}_i)_{22}$, $(\boldsymbol{B}_i)_{31} = \frac{1}{2}(\boldsymbol{P}_i)_{11}$ and $(\boldsymbol{B}_i)_{32} = \frac{1}{2}(\boldsymbol{P}_i)_{22}$. Furthermore, $(\boldsymbol{B}_i)_{33} = \frac{1}{2}\text{tr}\left(\boldsymbol{V}_i^{-2}\right)$. All of these expressions assist in obtaining (3.18), the approximate FIM for individual i as

$$
\boldsymbol{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i) \cong \begin{bmatrix}\boldsymbol{H}_i^T\boldsymbol{V}_i^{-1}\boldsymbol{H}_i & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{B}_i\end{bmatrix}, \tag{3.19}
$$

where

$$
\boldsymbol{B}_i = \frac{1}{2}\begin{bmatrix}\{(\boldsymbol{W}_i)_{11}\}^2 & (\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21} & (\boldsymbol{P}_i)_{11}\\ (\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21} & \{(\boldsymbol{W}_i)_{22}\}^2 & (\boldsymbol{P}_i)_{22}\\ (\boldsymbol{P}_i)_{11} & (\boldsymbol{P}_i)_{22} & \text{tr}\left(\boldsymbol{V}_i^{-2}\right)\end{bmatrix}.
$$

For a single group of N individuals with identical designs, the FIM simplifies to

$$
\boldsymbol{M}(\boldsymbol{\Psi}, \boldsymbol{\Xi}) = N\begin{bmatrix}\boldsymbol{H}_i^T\boldsymbol{V}_i^{-1}\boldsymbol{H}_i & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{B}_i\end{bmatrix}. \tag{3.20}
$$

The information matrix in (3.20) is used to obtain the D-optimum time points to measure the concentration of a drug in the blood for a cohort of N patients. We compute the D-optimum time points using the R package *PFIM 3.2* (Bazzoli et al., 2010). Figure 3.7 shows the optimum time points for our model for the true values of the parameters. We assume that the dose 0.5 mg/kg body weight has been administered to a cohort of size 3.

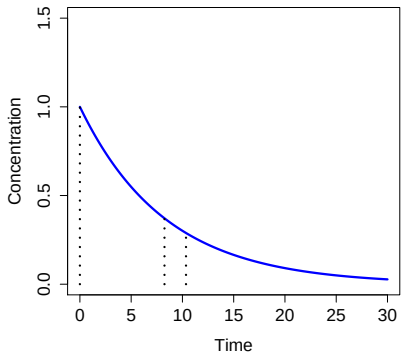

Figure 3.7: Location of optimum design points in the mean concentration profile for collecting blood samples. True parameter values are assumed as prior, that is, $\boldsymbol{\Psi^0} = (0.5, 0.06, 0.004, 0.00005, 0.000225)^T$.

The block diagonal form of the information matrix is useful in examining the sensitivity of the design to the prior parameter values $\boldsymbol{\beta}^0$. This is because

$$\begin{vmatrix} \boldsymbol{A}_i & 0 \\ 0 & \boldsymbol{B}_i \end{vmatrix} = |\boldsymbol{A}_i| \; |\boldsymbol{B}_i|,$$

which decreases the dimensions of the determinant. Keeping the prior values of the variances fixed and varying V and Cl, we can numerically obtain the values of the criterion for a given design for a range of V and Cl. The sensitivity factor in this case is the relative efficiency, defined as

$$\text{Efficiency} = \left(\frac{|\boldsymbol{M}(\tilde{\boldsymbol{\Psi}}, \boldsymbol{\xi}^*_{\text{true}})|}{|\boldsymbol{M}(\boldsymbol{\Psi}_{\text{true}}, \boldsymbol{\xi}^*_{\text{true}})|} \right)^{\frac{1}{5}},$$

where $\boldsymbol{\Psi}_{\text{true}} = (0.5, 0.06, 0.004, 0.00005, 0.000225)^T$, $\tilde{\boldsymbol{\Psi}} = (V, Cl, 0.004, 0.00005, 0.00$ $0225)^T$ and $\boldsymbol{\xi}^*_{\text{true}}$ is the D-optimum design obtained for $\boldsymbol{\Psi}_{\text{true}}$. Figure 3.8 shows the relative efficiencies of the designs for various choices of V and Cl that are within the three standard deviation of the means. It is clear that the design is not very sensitive to the parameter values, since the efficiency value is always very high, irrespective of the choice.

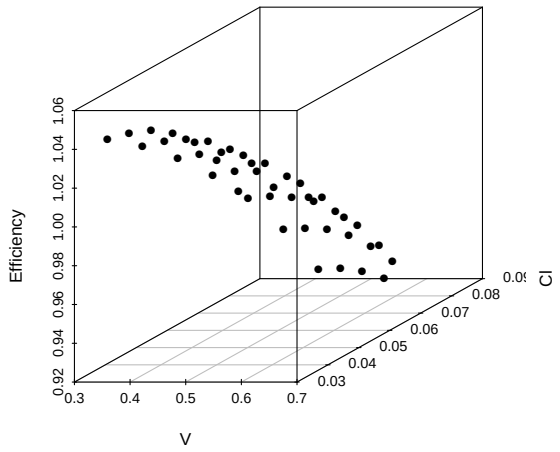

Figure 3.8: Sensitivity of the design to the assumed prior values of the parameters.

3.7.2 One-Compartment Model with First-Order Absorption

The one-compartment PK model with first-order absorption is defined as

$$\begin{aligned} y_{il} &= f(t_{il};\boldsymbol{\theta}_i) + \epsilon_{il} \\ &= \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}\left(e^{-k_{e_i}t_{il}} - e^{-k_{a_i}t_{il}}\right) + \epsilon_{il} \quad \text{for } i = 1,\ldots,N \text{ and } l = 1,\ldots,n_i, \end{aligned} \tag{3.21}$$

where y_{il} is the concentration of a drug in the blood for the ith individual observed at time t_{il}, x is the dose received by the individual, and $\boldsymbol{\theta}_i = \boldsymbol{\beta}+\boldsymbol{b}_i = (V_i, k_{e_i}, k_{a_i})^T$ is the vector of parameters with $\boldsymbol{\beta} = (V, k_e, k_a)^T$ and $\boldsymbol{b}_i = (b_{V_i}, b_{k_{e_i}}, b_{k_{a_i}})^T$. We assume that $\boldsymbol{\epsilon}_i \sim N_{n_i}(\boldsymbol{0}, \sigma^2\boldsymbol{I}_{n_i})$ and $\boldsymbol{b}_i \sim N_3(\boldsymbol{0}, \boldsymbol{\Omega})$, where $\boldsymbol{\Omega} = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$. We can write

$$\boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{\theta}_i) = \boldsymbol{f}(\boldsymbol{\xi}_i;\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i)) = \begin{bmatrix} \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}\left(e^{-k_{e_i}t_{i1}} - e^{-k_{a_i}t_{i1}}\right) \\ \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}\left(e^{-k_{e_i}t_{i2}} - e^{-k_{a_i}t_{i2}}\right) \\ \vdots \\ \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}\left(e^{-k_{e_i}t_{in_i}} - e^{-k_{a_i}t_{in_i}}\right) \end{bmatrix}. \tag{3.22}$$

To linearise the model, we expand it using a first-order Taylor series about $\boldsymbol{\phi}_i = (V,$

$k_e, k_a, b_{V_i}, b_{k_{e_i}}, b_{k_{a_i}})^T$ at $\boldsymbol{\phi}^0 = (V^0, k_e^0, k_a^0, 0, 0, 0)^T$, where V^0, k_e^0 and k_a^0 are some prior values of the population mean parameters. We obtain as an approximation the linear mixed effects model

$$\boldsymbol{y_i} \;\cong\; \boldsymbol{\mu}_i + \boldsymbol{H}_i\boldsymbol{\beta} + \boldsymbol{L}_i\boldsymbol{b}_i + \boldsymbol{\epsilon}_i,$$

where $\boldsymbol{\mu}_i$, $\boldsymbol{H}_i$ and $\boldsymbol{L}_i$ are as in Section 3.4.1. In particular, we have

$$\boldsymbol{H}_i = \left[\begin{array}{ccc} \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial V} & \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_e} & \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_a} \\ \vdots & \vdots & \vdots \\ \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial V} & \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_e} & \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_a} \end{array}\right]_{\boldsymbol{\phi}^0}.$$

The elements of $\boldsymbol{H}_i$ are given as

$$\left[\begin{array}{c} \dfrac{\partial f(t_{il};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial V} \\ \dfrac{\partial f(t_{il};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_e} \\ \dfrac{\partial f(t_{il};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial k_a} \end{array}\right]_{\boldsymbol{\phi}^0} = \left[\begin{array}{c} -\dfrac{xk_a^0}{(V^0)^2(k_a^0-k_e^0)}\left(e^{-k_e^0 t_{il}} - e^{-k_a^0 t_{il}}\right) \\ \dfrac{xk_a^0}{V^0(k_a^0-k_e^0)}\left(\dfrac{e^{-k_e^0 t_{il}} - e^{-k_a^0 t_{il}}}{k_a^0-k_e^0} - t_{il}e^{-k_e^0 t_{il}}\right) \\ \dfrac{x}{V^0(k_a^0-k_e^0)}\left\{-\dfrac{k_e^0\left(e^{-k_e^0 t_{il}} - e^{-k_a^0 t_{il}}\right)}{k_a^0-k_e^0} + k_a^0 t_{il}e^{-k_a^0 t_{il}}\right\} \end{array}\right].$$

Similar to Section 3.7.1, since $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$, here we also have $\boldsymbol{H}_i = \boldsymbol{L}_i$, where

$$\boldsymbol{L}_i = \left[\begin{array}{ccc} \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{V_i}} & \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{k_{e_i}}} & \dfrac{\partial f(t_{i1};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{k_{a_i}}} \\ \vdots & \vdots & \vdots \\ \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{V_i}} & \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{k_{e_i}}} & \dfrac{\partial f(t_{in_i};\boldsymbol{g}(\boldsymbol{\beta},\boldsymbol{b}_i))}{\partial b_{k_{a_i}}} \end{array}\right]_{\boldsymbol{\phi}^0}.$$

Following analogous derivations to those in Section 3.7.1, we obtain the approximate FIM for individual i as

$$\boldsymbol{M}_i(\boldsymbol{\Psi},\boldsymbol{\xi}_i) \cong \left[\begin{array}{cc} \boldsymbol{H}_i^T\boldsymbol{V}_i^{-1}\boldsymbol{H}_i & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{B}_i \end{array}\right], \tag{3.23}$$

where

$$\boldsymbol{B}_i = \frac{1}{2}\begin{bmatrix} \{(\boldsymbol{W}_i)_{11}\}^2 & (\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21} & (\boldsymbol{W}_i)_{13}(\boldsymbol{W}_i)_{31} & (\boldsymbol{P}_i)_{11} \\ (\boldsymbol{W}_i)_{12}(\boldsymbol{W}_i)_{21} & \{(\boldsymbol{W}_i)_{22}\}^2 & (\boldsymbol{W}_i)_{23}(\boldsymbol{W}_i)_{32} & (\boldsymbol{P}_i)_{22} \\ (\boldsymbol{W}_i)_{13}(\boldsymbol{W}_i)_{31} & (\boldsymbol{W}_i)_{23}(\boldsymbol{W}_i)_{32} & \{(\boldsymbol{W}_i)_{33}\}^2 & (\boldsymbol{P}_i)_{33} \\ (\boldsymbol{P}_i)_{11} & (\boldsymbol{P}_i)_{22} & (\boldsymbol{P}_i)_{33} & \mathrm{tr}\left(\boldsymbol{V}_i^{-2}\right) \end{bmatrix},$$

$(\boldsymbol{W}_i)_{lk}$ for $l, k = 1, 2, 3$ are elements of $\boldsymbol{W}_i = \boldsymbol{L}_i^T \boldsymbol{V}_i^{-1} \boldsymbol{L}_i$ and $(\boldsymbol{P}_i)_{ll}$ are diagonal elements of $\boldsymbol{P}_i = \boldsymbol{L}_i^T \boldsymbol{V}_i^{-2} \boldsymbol{L}_i$.

Since matrix $\boldsymbol{M}_i$ is 7×7, it is even more useful for numerical calculations to have a block diagonal structure to the matrix. The sensitivity analysis can be performed as in the previous case, but the graphical representation here would not be so clear.

3.8 Properties of the Derived PK Parameters

It is very common that, even if the same dose is given to a group of patients, the concentration profiles are different. This is due to the differences in biological factors among the patients. As the profiles vary, so do the AUC and $C_{\max}$. An obvious issue then is to quantify the variability in these parameters.

We employ a non-linear mixed effects model for the concentration data. The purpose is to describe the situation in a better way, since one can assume that each individual has distinct values for the model parameters. Because the AUC and $C_{\max}$ are derived from the model function, they rely on its parameters. This section describes way to find the variability in these PK parameters following the administration of a dose. Quantifying the variability is particulary important for two reasons. Firstly, it will tell us the extent of the variability from patient to patient, and, secondly, it may guide us in the selection of doses in a clinical trial design. Our purpose is to use it in the dose selection.

3.8.1 Area Under the Concentration Curve

Assume that the dose x is given to a cohort of patients and that $h(x, \boldsymbol{\theta_i})$ is the AUC for individual i in that group based on a PK model. Here, $\boldsymbol{\theta_i}$ is the vector of random PK parameters introduced in Section 3.4 and h is a differentiable function of the parameters. Since the parameters are random, individuals in the cohort will have different AUC values. We want to derive an expression so that the variability can be assessed. Since a PK model is non-linear in the parameters, the AUC will be too. Therefore, we linearise the function using a first-order Taylor series expansion of $h(x, \boldsymbol{\theta_i})$ about $\boldsymbol{\theta_i}$ at $E(\boldsymbol{\theta_i})$ to obtain

$$\begin{aligned} h(x, \boldsymbol{\theta_i}) &\cong h(x, \boldsymbol{\theta_i})\Big|_{E(\boldsymbol{\theta_i})} + \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} (\boldsymbol{\theta_i} - E(\boldsymbol{\theta_i})) \\ &= h(x, \boldsymbol{\beta}) + \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} \boldsymbol{b_i}, \end{aligned}$$

where $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$ and $E(\boldsymbol{\theta}_i) = \boldsymbol{\beta}$. Since $\boldsymbol{b}_i \sim N_p(\mathbf{0}, \boldsymbol{\Omega})$, $h(x, \boldsymbol{\theta}_i)$ will be approximately normally distributed as well.

Therefore,

$$E\{h(x, \boldsymbol{\theta_i})\} \cong h(x, \boldsymbol{\beta}) \tag{3.24}$$

and

$$\begin{aligned} \mathrm{Var}\{h(x, \boldsymbol{\theta_i})\} &\cong \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} \mathrm{Var}(\boldsymbol{b_i}) \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Bigg|_{E(\boldsymbol{\theta_i})} \\ &= \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} \boldsymbol{\Omega} \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Bigg|_{E(\boldsymbol{\theta_i})}. \end{aligned} \tag{3.25}$$

This derivation is based on the δ-method (Oehlert, 1992). The general expression in (3.25) can be used to find the variability in the AUC for a population based on a PK model after the administration of a dose. Also, we have $h(x, \boldsymbol{\theta_i}) \sim N[h(x, \boldsymbol{\beta}), \mathrm{Var}\{h(x, \boldsymbol{\theta_i})\}]$ approximately.

Example 1

For the one-compartment PK model with bolus input and first-order elimination in (3.16), the AUC for individual i over the range $[0, t_1]$ is defined as

$$\begin{aligned} h(x, \boldsymbol{\theta_i}) &= \int_0^{t_1} f(t; \boldsymbol{\theta_i})\, dt \\ &= \int_0^{t_1} \frac{x}{V_i} \exp\left(-\frac{Cl_i}{V_i}\, t\right)\, dt \\ &= \frac{x}{Cl_i}\left\{1 - \exp\left(-\frac{Cl_i}{V_i} t_1\right)\right\}. \end{aligned} \tag{3.26}$$

Assuming that $E(V_i) = V$ and $E(Cl_i) = Cl$, we obtain

$$E\{h(x, \boldsymbol{\theta_i})\} \cong h(x, \boldsymbol{\beta}) = \frac{x}{Cl}\left\{1 - \exp\left(-\frac{Cl}{V} t_1\right)\right\}. \tag{3.27}$$

Also,

$$\begin{aligned} \mathrm{Var}\{h(x, \boldsymbol{\theta_i})\} &\cong \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} \boldsymbol{\Omega} \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Bigg|_{E(\boldsymbol{\theta_i})} \\ &= \left(\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial V_i} \quad \frac{\partial h(x, \boldsymbol{\theta_i})}{\partial Cl_i} \right)_{E(\boldsymbol{\theta_i})} \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix} \begin{pmatrix} \dfrac{\partial h(x, \boldsymbol{\theta_i})}{\partial V_i} \\ \dfrac{\partial h(x, \boldsymbol{\theta_i})}{\partial Cl_i} \end{pmatrix}_{E(\boldsymbol{\theta_i})} \\ &= \left\{ \frac{\partial h(x, \boldsymbol{\theta_i})}{\partial V_i} \Big|_{E(\boldsymbol{\theta_i})} \right\}^2 \sigma_1^2 + \left\{ \frac{\partial h(x, \boldsymbol{\theta_i})}{\partial Cl_i} \Big|_{E(\boldsymbol{\theta_i})} \right\}^2 \sigma_2^2. \end{aligned} \tag{3.28}$$

Furthermore,

$$\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial V_i} \Big|_{E(\boldsymbol{\theta_i})} = -\frac{x t_1}{V^2} \exp\left(-\frac{Cl}{V} t_1\right)$$

and

$$\frac{\partial h(x, \boldsymbol{\theta_i})}{\partial Cl_i} \Big|_{E(\boldsymbol{\theta_i})} = \frac{x}{Cl} \exp\left(-\frac{Cl}{V} t_1\right) \left(\frac{1}{Cl} + \frac{t_1}{V}\right) - \frac{x}{Cl^2}.$$

Hence, we obtain the approximate variance of the AUC as

$$\begin{aligned}\mathrm{Var}\{h(x,\boldsymbol{\theta_i})\} \cong & \left\{-\frac{xt_1}{V^2}\exp\left(-\frac{Cl}{V}t_1\right)\right\}^2\sigma_1^2 \\ & + \left\{\frac{x}{Cl}\exp\left(-\frac{Cl}{V}t_1\right)\left(\frac{1}{Cl}+\frac{t_1}{V}\right)-\frac{x}{Cl^2}\right\}^2\sigma_2^2.\end{aligned}$$

The approximate AUC and its variance are used to impose a constraint in Section 5.5.1 in order to find optimal doses for the cohorts in a trial. The essence of the constraint is to be careful during dose escalation by taking into account the variability in the AUC among the patients in a cohort after receiving a dose.

Example 2

For the one-compartment PK model with first-order absorption that was introduced in (3.21), the AUC for individual i over the range $[0, t_1]$ is defined as

$$\begin{aligned}h(x,\boldsymbol{\theta_i}) &= \int_0^{t_1} f(t;\boldsymbol{\theta_i})\,dt \\ &= \int_0^{t_1} \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}(e^{-k_{e_i}t}-e^{-k_{a_i}t})\,dt \\ &= \frac{xk_{a_i}}{V_i(k_{a_i}-k_{e_i})}\left(\frac{1-e^{-k_{e_i}t_1}}{k_{e_i}}-\frac{1-e^{-k_{a_i}t_1})}{k_{a_i}}\right). \qquad (3.29)\end{aligned}$$

Assuming that $E(V_i) = V$, $E(k_{e_i}) = k_e$ and $E(k_{a_i}) = k_a$, we obtain

$$E\{h(x,\boldsymbol{\theta_i})\} \cong h(x,\boldsymbol{\beta}) = \frac{xk_a}{V(k_a-k_e)}\left(\frac{1-e^{-k_e t_1}}{k_e}-\frac{1-e^{-k_a t_1}}{k_a}\right). \qquad (3.30)$$

Also,

$$\begin{aligned}\text{Var}\{h(x,\boldsymbol{\theta_i})\} &\cong \left(\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T\Big|_{E(\boldsymbol{\theta_i})} \boldsymbol{\Omega} \left(\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)\Big|_{E(\boldsymbol{\theta_i})} \\ &= \left\{\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial V_i}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_1^2 + \left\{\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial k_{e_i}}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_2^2 \\ &\quad + \left\{\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial k_{a_i}}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_3^2. \end{aligned} \tag{3.31}$$

Putting $z_1 = 1 - e^{-k_e t_1}$ and $z_2 = 1 - e^{-k_a t_1}$, we obtain

$$\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial V_i}\Big|_{E(\boldsymbol{\theta_i})} = -\frac{xk_a}{V^2(k_a-k_e)}\left(\frac{z_1}{k_e} - \frac{z_2}{k_a}\right),$$

$$\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial k_{e_i}}\Big|_{E(\boldsymbol{\theta_i})} = \frac{xk_a}{V(k_a-k_e)^2}\left(\frac{z_1}{k_e} - \frac{z_2}{k_a}\right) + \frac{xk_a}{V(k_a-k_e)}\left(-\frac{z_1}{k_e^2} + \frac{t_1 e^{-k_e t_1}}{k_e}\right)$$

and

$$\frac{\partial h(x,\boldsymbol{\theta_i})}{\partial k_{a_i}}\Big|_{E(\boldsymbol{\theta_i})} = -\frac{xk_e}{V(k_a-k_e)^2}\left(\frac{z_1}{k_e} - \frac{z_2}{k_a}\right) + \frac{xk_a}{V(k_a-k_e)}\left(\frac{z_2}{k_a^2} - \frac{t_1 e^{-k_a t_1}}{k_a}\right).$$

These partial derivatives enable us to find the approximate variance in (3.31). The approximate mean and variance thus obtained can help in implementing the PK constraint in Section 5.5.1 when the model of interest is for one compartment with first-order absorption.

3.8.2 Maximum Concentration

Assume that $l(x,\boldsymbol{\theta_i})$ represents the maximum concentration for individual i for a given dose x, based on a PK model. Here, $\boldsymbol{\theta_i}$ is the vector of random PK parameters introduced in Section 3.4 and l is a differentiable function of those parameters. The main motivation is to derive an expression to assess the variability in $l(x,\boldsymbol{\theta_i})$. Since the PK models are non-linear in the parameters, the maximum concentration will be too. To obtain the inter-patient variability in $C_{\max}$, we linearise the function

using a first-order Taylor series expansion of $l(x, \boldsymbol{\theta_i})$ about $\boldsymbol{\theta_i}$ at $E(\boldsymbol{\theta_i})$ as

$$\begin{aligned} l(x, \boldsymbol{\theta_i}) &\cong l(x, \boldsymbol{\theta_i})\Big|_{E(\boldsymbol{\theta_i})} + \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Big|_{E(\boldsymbol{\theta_i})} (\boldsymbol{\theta_i} - E(\boldsymbol{\theta_i})) \\ &= l(x, \boldsymbol{\beta}) + \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Big|_{E(\boldsymbol{\theta_i})} \boldsymbol{b_i}, \end{aligned}$$

where $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$ and $E(\boldsymbol{\theta}_i) = \boldsymbol{\beta}$.

Therefore,

$$E\{l(x, \boldsymbol{\theta_i})\} \cong l(x, \boldsymbol{\beta}) \tag{3.32}$$

and

$$\begin{aligned} \mathrm{Var}\{l(x, \boldsymbol{\theta_i})\} &\cong \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Big|_{E(\boldsymbol{\theta_i})} \mathrm{Var}(\boldsymbol{b_i}) \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Big|_{E(\boldsymbol{\theta_i})} \\ &= \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Big|_{E(\boldsymbol{\theta_i})} \boldsymbol{\Omega} \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Big|_{E(\boldsymbol{\theta_i})}. \end{aligned} \tag{3.33}$$

The general expression in (3.33) can be used to find the variability in $C_{\max}$ for a population based on a PK model for a given dose x. Here, also, $l(x, \boldsymbol{\theta_i}) \sim N[l(x, \boldsymbol{\beta}), \mathrm{Var}\{l(x, \boldsymbol{\theta_i})\}]$ approximately.

Example 2

We are interested in the value of the concentration at a time point, denoted by $t_{\max}$, at which it is the largest at a given time interval. The $t_{\max}$ is obtained by differentiating the model function with respect to t and $C_{\max}$ is found by substituting that value for t into the model function. For a simple decay function, as in Example 1, it is always at $t = 0$. Here, we present a more interesting case of the one-compartment PK model with first-order absorption. The $t_{\max}$ is a solution of the equation

$$\frac{df(t;\boldsymbol{\theta_i})}{dt} = 0.$$

That is, solving

$$\frac{xk_{a_i}}{V_i(k_{a_i} - k_{e_i})}\left(-k_{e_i}\ e^{-k_{e_i}t} + k_{a_i}\ e^{-k_{a_i}t}\right) = 0$$

for t, we obtain

$$t_{\max_i} = \frac{\log(k_{a_i}) - \log(k_{e_i})}{k_{a_i} - k_{e_i}}.$$

Therefore, $C_{\max}$ for this model has the form

$$l(x, \boldsymbol{\theta_i}) = \frac{xk_{a_i}}{V_i(k_{a_i} - k_{e_i})}\left(e^{-k_{e_i}t_{\max_i}} - e^{-k_{a_i}t_{\max_i}}\right). \tag{3.34}$$

Assuming that $E(V_i) = V$, $E(k_{e_i}) = k_e$ and $E(k_{a_i}) = k_a$, we have

$$E\{l(x, \boldsymbol{\theta_i})\} \cong l(x, \boldsymbol{\beta}) = \frac{xk_a}{V(k_a - k_e)}\left(e^{-k_e t_{\max}} - e^{-k_a t_{\max}}\right). \tag{3.35}$$

Also,

$$\begin{aligned}
\mathrm{Var}\{l(x, \boldsymbol{\theta_i})\} &\cong \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right)^T \Bigg|_{E(\boldsymbol{\theta_i})} \boldsymbol{\Omega} \left(\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial \boldsymbol{\theta_i}}\right) \Bigg|_{E(\boldsymbol{\theta_i})} \\
&= \left\{\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial V_i}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_1^2 + \left\{\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial k_{e_i}}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_2^2 \\
&\quad + \left\{\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial k_{a_i}}\Big|_{E(\boldsymbol{\theta_i})}\right\}^2 \sigma_3^2.
\end{aligned} \tag{3.36}$$

After some simplification, it can be shown that

$$\frac{\partial l(x, \boldsymbol{\theta_i})}{\partial V_i}\Big|_{E(\boldsymbol{\theta_i})} = -\frac{xk_a}{V^2(k_a - k_e)}\left(e^{-k_e t_{\max}} - e^{-k_a t_{\max}}\right),$$

$$\left.\frac{\partial l(x,\boldsymbol{\theta_i})}{\partial k_{e_i}}\right|_{E(\boldsymbol{\theta_i})}$$
$$= \frac{xk_a}{V(k_a-k_e)^2}\left[e^{-k_e t_{\max}}(2-k_a t_{\max}) - e^{-k_a t_{\max}}\left\{1+\frac{k_a(1-k_e t_{\max})}{k_e}\right\}\right]$$

and

$$\left.\frac{\partial l(x,\boldsymbol{\theta_i})}{\partial k_{a_i}}\right|_{E(\boldsymbol{\theta_i})}$$
$$= \frac{x}{V(k_a-k_e)^2}\left[k_e e^{-k_e t_{\max}}\left\{-1+k_a(k_a t_{\max}-1)\right\} + e^{-k_a t_{\max}}\left\{k_e - k_a(k_e t_{\max}-1)\right\}\right].$$

These partial derivatives are plugged into (3.36) to obtain the approximate variance of $C_{\max}$ among the patients following the adminstration of a dose. The approximate mean and variance in (3.35) and (3.36) aid in implementing the constraint defined in Section 5.5.2.

It is clear from the above expressions that the mean and variance of the AUC and $C_{\max}$ depend on the unknown parameters $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}$ associated with the PK model. We use D-optimum design to collect blood samples to measure the concentrations. Since the AUC and $C_{\max}$, together with their variances, are non-linear functions of the model parameters, a natural question arises about the suitability of c-optimum design. We have three issues here: model parameter estimation, and estimation of the mean and variance of the derived PK parameters. It will be quite difficult, or even impossible, to find a design that will minimise the variability of each of these. Therefore, we rely on D-optimum design only to estimate precisely the model parameters. We believe that with such a set of estimates, it is also possible to obtain precise estimates for the derived PK parameters and their variances.

Chapter 4

Dose-Response Models

4.1 Introduction

This chapter introduces two dose-response models that will be used in the simulation studies in Chapter 6. We assume that the dose range is pre-specified from pre-clinical studies. The aim is to select the dose level which can be recommended for further investigation on a larger group of patients in a phase IIB trial, where it is compared with other standard treatments.

We consider adaptive designs in Chapter 5 in that context, where an analysis of the data is performed after treating each cohort of patients, and a decision on the dose to be allocated is made for the next cohort based on the updated knowledge of the responses. The first model assumes trinomial responses representing efficacy, toxicity and a neutral response. The second one assumes binary outcomes for efficacy and toxicity end points.

4.2 Trinomial Response

4.2.1 Model

We assume a trinomial response $\boldsymbol{Y} = (Y_0, Y_1, Y_2)^T$ for a given dose for each patient, where Y_i takes values 0 or 1 for $i = 0, 1, 2$ depending on the patient's response to

the drug. Here, $Y_0 = 1$ when neither toxicity nor efficacy occurs, $Y_1 = 1$ when efficacy occurs without toxicity and $Y_2 = 1$ when toxicity is the outcome, irrespective of the efficacy. Throughout the thesis, we call an efficacious but non-toxic response a success, since it is the outcome indicating some benefit from taking the drug. The probability of each of these outcomes depends on dose. It is commonly accepted that a drug's toxicity increases with dose. The probability of success also increases in many cases. However, it is possible for some drugs that it attains a plateau or increases and then decreases, as the dose is increased. For an experimental drug, let us assume that the probability of a neutral response decreases monotonically with dose and that the probability of toxicity increases monotonically with dose. The probability of success may be non-monotonic, increasing or decreasing. The corresponding cell probabilities are $\psi_0(x, \boldsymbol{\vartheta}), \psi_1(x, \boldsymbol{\vartheta})$ and $\psi_2(x, \boldsymbol{\vartheta})$, so that $\psi_0(x, \boldsymbol{\vartheta}) + \psi_1(x, \boldsymbol{\vartheta}) + \psi_2(x, \boldsymbol{\vartheta}) = 1$ for a given dose x, where $\boldsymbol{\vartheta}$ denotes the vector of dose-response parameters.

The continuation ratio model of McCullagh and Nelder (1989) is used to model the responses, and is given by

$$\log\left(\frac{\psi_1(x, \boldsymbol{\vartheta})}{\psi_0(x, \boldsymbol{\vartheta})}\right) = \vartheta_1 + \vartheta_2 x \tag{4.1}$$

and

$$\log\left(\frac{\psi_2(x, \boldsymbol{\vartheta})}{1 - \psi_2(x, \boldsymbol{\vartheta})}\right) = \vartheta_3 + \vartheta_4 x, \tag{4.2}$$

where $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)^T$ is the vector of parameters to be estimated. The parameter ϑ_1 represents the baseline log-relative probability, ϑ_2 reflects the contribution of dose in the log-relative probability of a success relative to the probability of a neutral outcome, ϑ_3 is the baseline log odds and ϑ_4 is the contribution of dose in the log odds when interest lies in the occurrence of a toxic outcome relative to a neutral outcome or a success.

From (4.1), we have

$$\psi_1(x, \boldsymbol{\vartheta}) = \psi_0(x, \boldsymbol{\vartheta})\ e^{\vartheta_1 + \vartheta_2 x},$$

and, from (4.2),

$$\begin{aligned} \psi_2(x, \boldsymbol{\vartheta}) &= \{\psi_0(x, \boldsymbol{\vartheta}) + \psi_1(x, \boldsymbol{\vartheta})\}\ e^{\vartheta_3 + \vartheta_4 x} \\ &= \psi_0(x, \boldsymbol{\vartheta})(1 + e^{\vartheta_1 + \vartheta_2 x}) e^{\vartheta_3 + \vartheta_4 x}. \end{aligned}$$

Since $\psi_0(x, \boldsymbol{\vartheta}) + \psi_1(x, \boldsymbol{\vartheta}) + \psi_2(x, \boldsymbol{\vartheta}) = 1$, we obtain

$$\psi_0(x, \boldsymbol{\vartheta}) = \frac{1}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})}.$$

Therefore,

$$\psi_1(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_1 + \vartheta_2 x}}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})}$$

and

$$\psi_2(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_3 + \vartheta_4 x}}{1 + e^{\vartheta_3 + \vartheta_4 x}}.$$

These three non-linear functions represent the probabilities for the trinomial responses. The selection of dose depends on the parameters in $\boldsymbol{\vartheta}$. Different values of $\boldsymbol{\vartheta}$ will lead to different dose-response curves, as shown in Figure 4.1.

4.2.2 Parameter Space

Defining the dose-response parameter space is an essential part in ensuring that the three non-linear functions exhibit the assumed behaviour.

Differentiation of $\psi_2(x, \boldsymbol{\vartheta})$ with respect to x gives

$$\frac{d\psi_2(x, \boldsymbol{\vartheta})}{dx} = \frac{\vartheta_4\ e^{-(\vartheta_3 + \vartheta_4 x)}}{\{1 + e^{-(\vartheta_3 + \vartheta_4 x)}\}^2},$$

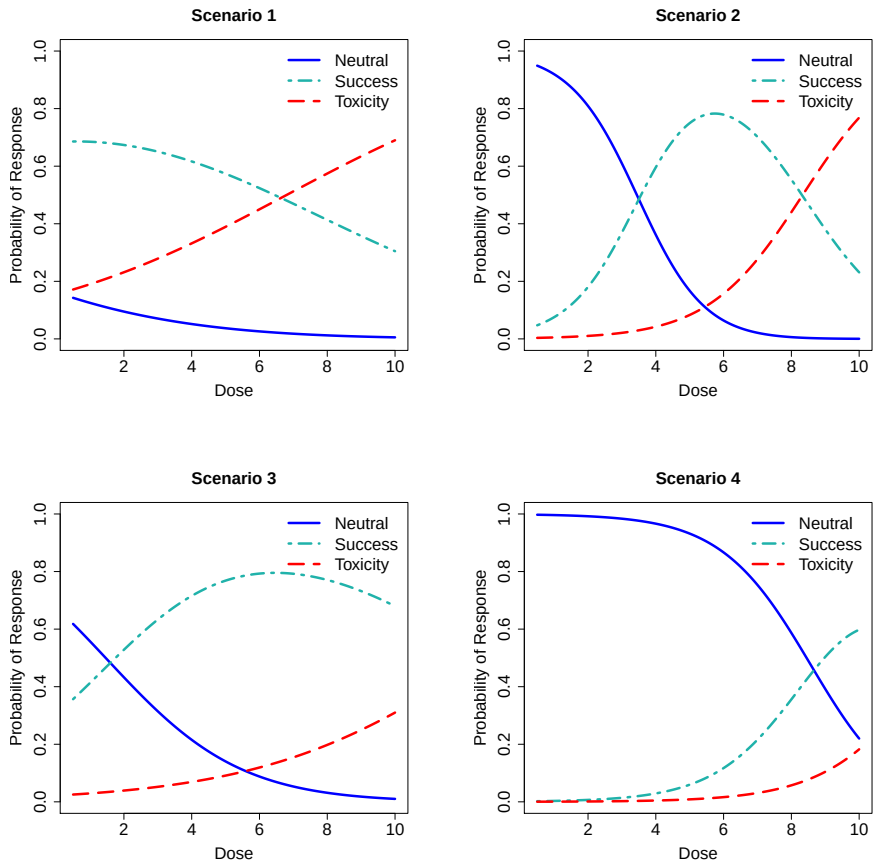

Figure 4.1: Dose-response scenarios for the continuation ratio model. The respective parameter values are: Scenario 1, $\boldsymbol{\vartheta} = (1.44, 0.26, -1.70, 0.25)^T$; Scenario 2, $\boldsymbol{\vartheta} = (-3.50, 1.00, -6.00, 0.72)^T$; Scenario 3, $\boldsymbol{\vartheta} = (-0.80, 0.50, -3.80, 0.30)^T$; and Scenario 4, $\boldsymbol{\vartheta} = (-6.50, 0.75, -8.00, 0.65)^T$.

which is positive only if $\vartheta_4 > 0$. Therefore, $\psi_2(x, \boldsymbol{\vartheta})$ is an increasing function of dose only if $\vartheta_4 > 0$.

Equation (4.1) indicates that ϑ_2 is the change in the log-relative probability due to a unit change in dose. Since we expect dose to have a positive impact on the transition from a neutral response to a success, we should have $\vartheta_2 > 0$.

When $x = 0$, we have

$$\log\left(\frac{\psi_1(0, \boldsymbol{\vartheta})}{\psi_0(0, \boldsymbol{\vartheta})}\right) = \vartheta_1, \quad \text{that is,} \quad \frac{\psi_1(0, \boldsymbol{\vartheta})}{\psi_0(0, \boldsymbol{\vartheta})} = e^{\vartheta_1}$$

and

$$\log\left(\frac{\psi_2(0,\boldsymbol{\vartheta})}{\psi_0(0,\boldsymbol{\vartheta})+\psi_1(0,\boldsymbol{\vartheta})}\right)=\vartheta_3,$$

from which it follows that

$$\frac{\psi_2(0,\boldsymbol{\vartheta})}{\psi_0(0,\boldsymbol{\vartheta})}=e^{\vartheta_3}(1+e^{\vartheta_1}).$$

We assume that, at very small doses, $\psi_1 \geq \psi_2$. Consequently,

$$\frac{\psi_1(0,\boldsymbol{\vartheta})}{\psi_0(0,\boldsymbol{\vartheta})} \geq \frac{\psi_2(0,\boldsymbol{\vartheta})}{\psi_0(0,\boldsymbol{\vartheta})},$$

which reduces to

$$\vartheta_1-\vartheta_3 \geq \log(1+e^{\vartheta_1})>0.$$

Therefore, we can conclude that $\vartheta_1 > \vartheta_3$.

Since ψ_2 represents the probability of toxicity, we expect it to be low at very small doses. This function reduces to $\psi_2(0,\boldsymbol{\vartheta}) = e^{\vartheta_3}/(1+e^{\vartheta_3})$ when $x=0$ and further reduces to $\frac{1}{2}$ when $\vartheta_3 = 0$. For any $\vartheta_3 > 0$, this probability is greater than $\frac{1}{2}$. Thus, the only choice to keep this probability low at the small doses is $\vartheta_3 < 0$.

So, the parameter space can be written as

$$\Theta=\{(\vartheta_1,\vartheta_2,\vartheta_3,\vartheta_4):\ \vartheta_1>\vartheta_3,\ \vartheta_3<0 \text{ and } \vartheta_2,\vartheta_4>0\}. \tag{4.3}$$

Figure 4.1 shows some plausible dose-response scenarios for different choices of the parameters in the continuation ratio model. These scenarios will be investigated in the simulation study in Chapter 6. It should be mentioned that different assumptions about the shapes of the functions ψ_0, ψ_1 and ψ_2 would give different parameter spaces Θ. Although the method of adaptive design that we propose in this thesis does not

depend on the choice of Θ, still it has to be specified.

4.2.3 Likelihood Function

Let us assume that we are at the kth stage in an adaptive clinical trial and that we want to conduct an interim analysis of the data. This means that k cohorts have been treated so far with selected doses from the dose range $\mathcal{X}$. Let $\boldsymbol{x}$ be a $k \times 1$ dose vector with components x_l and let $\boldsymbol{R}$ be a $k \times 3$ outcome matrix with $\tilde{\boldsymbol{R}}_l = (R_{l0}, R_{l1}, R_{l2})$ as the lth row $(l = 1, \ldots, k)$. It is important to note that $R_{l0} + R_{l1} + R_{l2} = c$, where c is the number of subjects in a cohort treated at dose x_l. The successive components of $\tilde{\boldsymbol{R}}_l$ are the counts of neutral, success and toxic responses for the lth cohort. Thus, for our dose-response model, the likelihood function is

$$L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) \propto \prod_{l=1}^{k} \{\psi_0(x_l, \boldsymbol{\vartheta})\}^{R_{l0}} \{\psi_1(x_l, \boldsymbol{\vartheta})\}^{R_{l1}} \{\psi_2(x_l, \boldsymbol{\vartheta})\}^{R_{l2}}.$$

The parameters in $\boldsymbol{\vartheta}$ can be estimated using either a Bayesian or frequentist approach. Since maximum likelihood estimation is unsuitable because of small sample sizes in the early stages of a trial, we employ a Bayesian approach.

The posterior estimates of the components of $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)^T$ at the kth stage are obtained as

$$\hat{\vartheta}_{ik} = \frac{\int_\Theta \vartheta_i \; p(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) d\boldsymbol{\vartheta}}{\int_\Theta p(\boldsymbol{\vartheta}) \; L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) \; d\boldsymbol{\vartheta}}, \quad i = 1, 2, 3, 4, \tag{4.4}$$

where $p(\boldsymbol{\vartheta})$ is the prior distribution of the parameters. We assume a uniform distribution on a restricted parameter space $\tilde{\Theta}$. Let $0 < \vartheta_2 < u_1, 0 < \vartheta_4 < u_2, l_1 < \vartheta_1 < l_2$ and $l_3 < \vartheta_3 < l_4$. Then the condition $\vartheta_1 > \vartheta_3$ leads to $l_3 < \vartheta_3 < \vartheta_1 < l_2$. We define $\tilde{\Theta}$ as

$$\tilde{\Theta} = \{\boldsymbol{\vartheta} : \; l_3 < \vartheta_3 < \vartheta_1 < l_2, \quad 0 < \vartheta_2 < u_1, \; 0 < \vartheta_4 < u_2\}.$$

As the probability density function for a uniform distribution is constant, we let $p(\boldsymbol{\vartheta}) = A$ and can write

$$\int_{\tilde{\Theta}} A \; d\vartheta_3 \; d\vartheta_1 \; d\vartheta_4 \; d\vartheta_2 = 1.$$

It follows that

$$A \int_0^{u_1} \int_0^{u_2} \int_{l_3}^{l_2} \int_{l_3}^{\vartheta_1} \; d\vartheta_3 \; d\vartheta_1 \; d\vartheta_4 \; d\vartheta_2 = 1,$$

which yields $A = 2/\{u_1 u_2 (l_2 - l_3)^2\}$. Therefore,

$$p(\boldsymbol{\vartheta}) = \frac{2}{u_1 u_2 (l_2 - l_3)^2}, \quad \boldsymbol{\vartheta} \in \tilde{\Theta}. \tag{4.5}$$

This function will be used in our example of simulating clinical trials in Section 6.3.

4.2.4 Fisher Information Matrix

The likelihood function for a single cohort at dose x is

$$L(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}}) = \frac{c!}{R_0! R_1! R_2!} (\psi_0)^{R_0} \; (1 - \psi_0 - \psi_2)^{c - R_0 - R_2} \; (\psi_2)^{R_2},$$

and so the log-likelihood function can be written as

$$l(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}}) = \text{constant} + R_0 \; \log(\psi_0) + (c - R_0 - R_2) \; \log(1 - \psi_0 - \psi_2) + R_2 \; \log(\psi_2).$$

The information matrix associated with $\boldsymbol{\vartheta}$ can be obtained as

$$I(x, \boldsymbol{\vartheta}) = \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}} \right)^T \; I(x, \boldsymbol{\psi}) \; \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}} \right), \tag{4.6}$$

where

$$I(x, \boldsymbol{\psi}) = E \left[- \frac{\partial^2 l(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}})}{\partial \boldsymbol{\psi} \; \partial \boldsymbol{\psi}^T} \right]$$

and $\boldsymbol{\psi} = (\psi_0, \psi_2)$. The second-order partial derivatives of the components of

$l(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}})$ are

$$\frac{\partial^2 \log(\psi_0)}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T} = \begin{bmatrix} -\frac{1}{\psi_0^2} & 0 \\ 0 & 0 \end{bmatrix},$$

$$\frac{\partial^2 \log(1-\psi_0-\psi_2)}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T} = -\frac{1}{(1-\psi_0-\psi_2)^2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$$

and

$$\frac{\partial^2 \log(\psi_2)}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T} = \begin{bmatrix} 0 & 0 \\ 0 & -\frac{1}{\psi_2^2} \end{bmatrix}.$$

Therefore,

$$\begin{aligned} I(x, \boldsymbol{\psi}) &= E\left[-\frac{\partial^2 l(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}})}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T}\right] \\ &= E(R_0) \begin{bmatrix} \frac{1}{\psi_0^2} & 0 \\ 0 & 0 \end{bmatrix} + E(c - R_0 - R_2) \frac{1}{(1-\psi_0-\psi_2)^2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \\ &\quad + E(R_2) \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{\psi_2^2} \end{bmatrix}. \end{aligned} \tag{4.7}$$

Using a property of the trinomial distribution, we have $E(R_0) = c\psi_0$, $E(c - R_0 - R_2) = c(1 - \psi_0 - \psi_2)$ and $E(R_2) = c\psi_2$. Thus, (4.7) reduces to

$$\begin{aligned} I(x, \boldsymbol{\psi}) &= c\left(\begin{bmatrix} \frac{1}{\psi_0} & 0 \\ 0 & \frac{1}{\psi_2} \end{bmatrix} + \frac{1}{(1-\psi_0-\psi_2)} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}\right) \\ &= \frac{c}{(1-\psi_0-\psi_2)} \begin{bmatrix} \frac{1-\psi_2}{\psi_0} & 1 \\ 1 & \frac{1-\psi_0}{\psi_2} \end{bmatrix}. \end{aligned} \tag{4.8}$$

Furthermore,

$$\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}} = \begin{bmatrix} \dfrac{\partial \psi_0}{\partial \vartheta_1} & \dfrac{\partial \psi_0}{\partial \vartheta_2} & \dfrac{\partial \psi_0}{\partial \vartheta_3} & \dfrac{\partial \psi_0}{\partial \vartheta_4} \\ \dfrac{\partial \psi_2}{\partial \vartheta_1} & \dfrac{\partial \psi_2}{\partial \vartheta_2} & \dfrac{\partial \psi_2}{\partial \vartheta_3} & \dfrac{\partial \psi_2}{\partial \vartheta_4} \end{bmatrix},$$

where the derivatives are

$$\frac{\partial \psi_0}{\partial \vartheta_1} = -\frac{e^{\vartheta_1+\vartheta_2 x}}{(1+e^{\vartheta_1+\vartheta_2 x})^2\ (1+e^{\vartheta_3+\vartheta_4 x})} = -\frac{\psi_1}{1+e^{\vartheta_1+\vartheta_2 x}},$$

$$\frac{\partial \psi_0}{\partial \vartheta_2} = -\frac{x e^{\vartheta_1+\vartheta_2 x}}{(1+e^{\vartheta_1+\vartheta_2 x})^2\ (1+e^{\vartheta_3+\vartheta_4 x})} = x\frac{\partial \psi_0}{\partial \vartheta_1},$$

$$\frac{\partial \psi_0}{\partial \vartheta_3} = -\frac{e^{\vartheta_3+\vartheta_4 x}}{(1+e^{\vartheta_1+\vartheta_2 x})\ (1+e^{\vartheta_3+\vartheta_4 x})^2} = -\psi_0\psi_2,$$

$$\frac{\partial \psi_0}{\partial \vartheta_4} = -\frac{x e^{\vartheta_3+\vartheta_4 x}}{(1+e^{\vartheta_1+\vartheta_2 x})\ (1+e^{\vartheta_3+\vartheta_4 x})^2} = x\frac{\partial \psi_0}{\partial \vartheta_3},$$

$$\frac{\partial \psi_2}{\partial \vartheta_1} = 0, \qquad \frac{\partial \psi_2}{\partial \vartheta_2} = 0,$$

$$\frac{\partial \psi_2}{\partial \vartheta_3} = \frac{e^{\vartheta_3+\vartheta_4 x}}{1+e^{\vartheta_3+\vartheta_4 x}} - \left(\frac{e^{\vartheta_3+\vartheta_4 x}}{1+e^{\vartheta_3+\vartheta_4 x}}\right)^2 = \psi_2(1-\psi_2)$$

and

$$\frac{\partial \psi_2}{\partial \vartheta_4} = x\left\{\frac{e^{\vartheta_3+\vartheta_4 x}}{1+e^{\vartheta_3+\vartheta_4 x}} - \left(\frac{e^{\vartheta_3+\vartheta_4 x}}{1+e^{\vartheta_3+\vartheta_4 x}}\right)^2\right\} = x\frac{\partial \psi_2}{\partial \vartheta_3}.$$

That is, we have the matrix of derivatives

$$\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}} = \begin{bmatrix} -\dfrac{\psi_1}{1+e^{\vartheta_1+\vartheta_2 x}} & x\dfrac{\partial \psi_0}{\partial \vartheta_1} & -\psi_0\psi_2 & x\dfrac{\partial \psi_0}{\partial \vartheta_3} \\ 0 & 0 & \psi_2(1-\psi_2) & x\dfrac{\partial \psi_2}{\partial \vartheta_3} \end{bmatrix}.$$

Applying the results to (4.6), after some simplification, we obtain the FIM for $\boldsymbol{\vartheta}$ as

$$I(x,\boldsymbol{\vartheta}) = c \begin{bmatrix} \dfrac{\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & \dfrac{x\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & 0 & 0 \\ \dfrac{x\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & \dfrac{x^2\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & 0 & 0 \\ 0 & 0 & \psi_2(1-\psi_2) & x\psi_2(1-\psi_2) \\ 0 & 0 & x\psi_2(1-\psi_2) & x^2\psi_2(1-\psi_2) \end{bmatrix}. \tag{4.9}$$

This matrix is of rank 2. It is block diagonal and has non-zero submatrices of rank 1. So we need at least two different doses, say x_1 and x_2, to obtain a combined nonsingular information matrix $I(x_1,\boldsymbol{\vartheta}) + I(x_2,\boldsymbol{\vartheta})$. This is important when considering D-optimum dose selection, where the optimality criterion is defined as the determinant of the FIM. We present such criteria in Sections 5.3.2 and 5.3.3. In Section 6.6, we apply this form of the FIM to determine the optimum dose in an adaptive clinical trial.

4.3 Bivariate Binary Response

4.3.1 Model

Here, we assume that efficacy and toxicity are two 0/1 binary variables and denote these by Y and Z, respectively. This will result in four possible (y,z) outcomes (0,0), (0,1), (1,0) and (1,1). The Cox model (Cox, 1970) treats each of the four possible outcomes as a separate response category. The corresponding cell probabilities are $\psi_{00}(x,\boldsymbol{\vartheta})$, $\psi_{01}(x,\boldsymbol{\vartheta})$, $\psi_{10}(x,\boldsymbol{\vartheta})$ and $\psi_{11}(x,\boldsymbol{\vartheta})$, so that $\psi_{00}(x,\boldsymbol{\vartheta})+\psi_{01}(x,\boldsymbol{\vartheta})+$

$\psi_{10}(x, \boldsymbol{\vartheta}) + \psi_{11}(x, \boldsymbol{\vartheta}) = 1.$

The probabilities of the four outcomes are given by

$$\psi_{00}(x, \boldsymbol{\vartheta}) = \frac{1}{1 + e^{\vartheta_1 + \vartheta_2 x} + e^{\vartheta_3 + \vartheta_4 x} + e^{\vartheta_5 + \vartheta_6 x}},$$

$$\psi_{01}(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_1 + \vartheta_2 x}}{1 + e^{\vartheta_1 + \vartheta_2 x} + e^{\vartheta_3 + \vartheta_4 x} + e^{\vartheta_5 + \vartheta_6 x}},$$

$$\psi_{10}(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_3 + \vartheta_4 x}}{1 + e^{\vartheta_1 + \vartheta_2 x} + e^{\vartheta_3 + \vartheta_4 x} + e^{\vartheta_5 + \vartheta_6 x}}$$

and

$$\psi_{11}(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_5 + \vartheta_6 x}}{1 + e^{\vartheta_1 + \vartheta_2 x} + e^{\vartheta_3 + \vartheta_4 x} + e^{\vartheta_5 + \vartheta_6 x}},$$

where $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4, \vartheta_5, \vartheta_6)^T$ is the vector of parameters.

Table 4.1: Relation between trinomial and bivariate binary responses.

		Binary responses	
		Efficacy	Toxicity
Trinomial responses	Neutral	No	No
	Success	Yes	No
	Toxic	No	Yes
		Yes	Yes

In this set up, we call the outcome (1,0) a success, since it is the outcome that produces efficacy without exposing a patient to toxicity. The trinomial response model that was introduced in Section 4.2.1 considers (0,1) and (1,1) as a single category, *toxic*, irrespective of the efficacy. Thus, it ignores the simultaneous occurrence of efficacy and toxicity as an individual category, as shown in Table 4.1. This is important when toxicity is a serious concern and compromising it may lead to serious health issues.

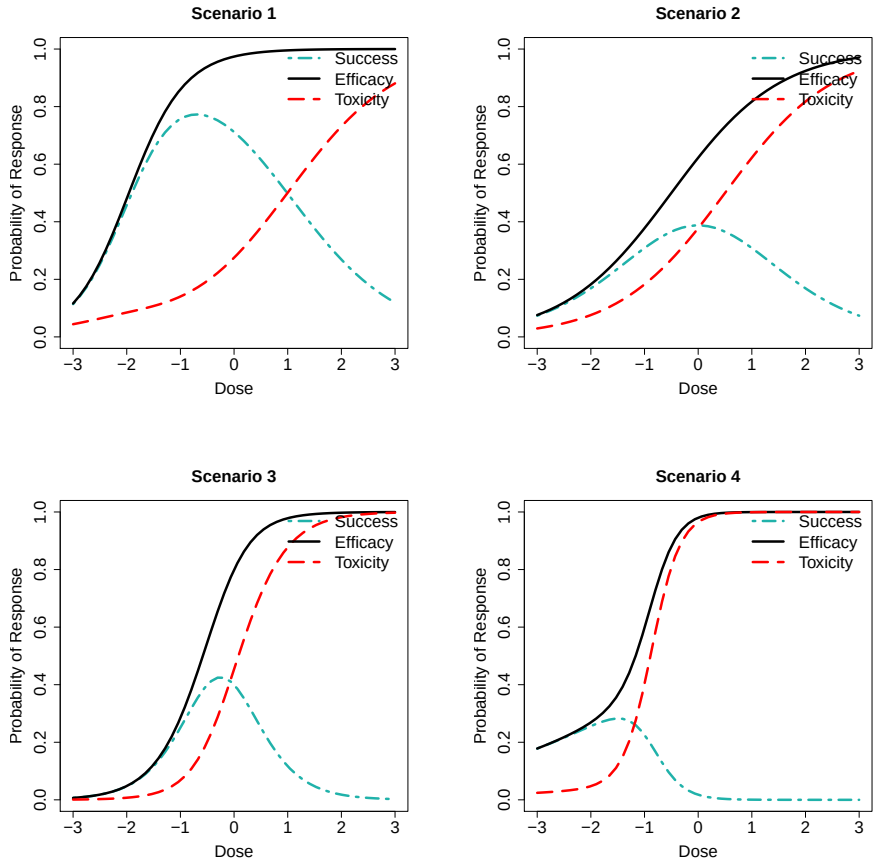

Figure 4.2: Dose-response scenarios for the Cox model. The respective parameter values are $\boldsymbol{\vartheta} = (0.0, 1.0, 4.0, 2.0, 3.0, 3.0)^T$, $\boldsymbol{\vartheta} = (-0.5, 1.0, 0.5, 1.0, 0.0, 2.0)^T$, $\boldsymbol{\vartheta} = (-1.0, 2.0, 1.0, 2.0, 1.0, 4.0)^T$ and $\boldsymbol{\vartheta} = (-2.0, 0.5, 0.0, 0.5, 4.0, 4.0)^T$.

As before, we are interested in designing the trial so that the recommended dose maximises the probability of success. However, this model gives more flexibility to consider other outcomes.

Again, we want to avoid doses with too high a risk of toxicity. Hence, we consider the total risk of toxicity. Also, we may be interested in the probability of efficacy, whatever the toxicity outcome. In Figure 4.2, we present three curves for various scenarios: the probability of success ((1,0) outcome), the probability of toxicity ((0,1) or (1,1) outcomes) and the probability of efficacy ((1,0) or (1,1) outcomes). We are primarily interested in the first one. However, the toxicity curve is used to control the risk of undesired side effects. The efficacy curve is included for additional information.

For the simulation study in Section 6.4, we consider some plausible dose-response scenarios, as shown in Figure 4.2, which are taken from Dragalin and Fedorov (2006). The doses in these scenarios are on a log scale. The curve which is labelled as *success* is simply obtained by plotting $\psi_{10}(x, \boldsymbol{\vartheta})$ against x. The other two curves are the marginal probabilities of efficacy and toxicity, obtained as $\psi_{1.}(x, \boldsymbol{\vartheta}) = \psi_{10}(x, \boldsymbol{\vartheta}) + \psi_{11}(x, \boldsymbol{\vartheta})$ and $\psi_{.1}(x, \boldsymbol{\vartheta}) = \psi_{01}(x, \boldsymbol{\vartheta}) + \psi_{11}(x, \boldsymbol{\vartheta})$, respectively.

4.3.2 Properties

This section introduces some properties of the Cox model. Although these are not explicitly used in our work, they give some insight into the nature of the model. In this model, the marginal probabilities for Y and Z are neither logistic nor necessarily monotonic in dose. Rather, it is the conditional probabilities of the responses that are logistic in dose. For example,

$$P(Y = 1|Z = 0, x) = \frac{e^{\vartheta_3 + \vartheta_4 x}}{1 + e^{\vartheta_3 + \vartheta_4 x}}.$$

According to Murtaugh (1989), the Cox bivariate binary model has the following properties:

1. If $\vartheta_1 + \vartheta_3 = \vartheta_5$ and $\vartheta_2 + \vartheta_4 = \vartheta_6$, then the marginal probabilities for Y and Z are logistic in dose.

2. If the marginal dependence of Y and Z on x is logistic and $\vartheta_2 \neq 0$ or $\vartheta_4 \neq 0$, then $\vartheta_1 + \vartheta_3 = \vartheta_5$ and $\vartheta_2 + \vartheta_4 = \vartheta_6$.

3. The variables Y and Z are independent if and only if $\vartheta_1 + \vartheta_3 = \vartheta_5$ and $\vartheta_2 + \vartheta_4 = \vartheta_6$.

4. The correlation between efficacy and toxicity has the form

$$\text{Corr}(Y, Z|x) = \frac{e^{z_3} - e^{z_1 + z_2}}{\sqrt{(e^{z_2} + e^{z_3})(1 + e^{z_1})(e^{z_1} + e^{z_3})(1 + e^{z_2})}},$$

   where $z_1 = \vartheta_1 + \vartheta_2 x$, $z_2 = \vartheta_3 + \vartheta_4 x$ and $z_3 = \vartheta_5 + \vartheta_6 x$.

In the continuation ratio model case, we derived the parameter space under some assumptions about the shapes of the functions. Although we have tried a similar approach here, we could not obtain any solution, since the Cox model is more complex with more functions and parameters than the previous one. Therefore, we use the unrestricted parameter space

$$\Theta = \{(\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4, \vartheta_5, \vartheta_6) : \ -\infty < \vartheta_i < \infty \text{ for } i = 1, \ldots, 6\}. \tag{4.10}$$

4.3.3 Likelihood Function

With the advancement of an adaptive clinical trial, we update the likelihood function, and hence obtain up-to-date estimates of the model parameters. Assume that different doses from the dose range $\mathcal{X}$ have been assigned to the successive cohorts up to the kth stage of a trial. Denote the k-dimensional column vector of the assigned doses by $\boldsymbol{x}$. The responses of the cohorts to the doses are represented through a $k \times 4$ matrix $\boldsymbol{R}$. More specifically, each row $\tilde{\boldsymbol{R}}_l = (R_{l0}, R_{l1}, R_{l2}, R_{l3})$ consists of the outcomes that result upon receiving dose x_l $(l = 1, \ldots, k)$. The successive elements in $\tilde{\boldsymbol{R}}_l$ are the counts of (0,0), (0,1), (1,0) and (1,1) responses for the lth cohort. Moreover, $R_{l0} + R_{l1} + R_{l2} + R_{l3} = c$, where c is the number of subjects in a cohort. Thus, the likelihood function for our model is

$$L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) \propto \prod_{l=1}^{k} \{\psi_{11}(x_l, \boldsymbol{\vartheta})\}^{R_{l3}} \{\psi_{10}(x_l, \boldsymbol{\vartheta})\}^{R_{l2}} \{\psi_{01}(x_l, \boldsymbol{\vartheta})\}^{R_{l1}} \{\psi_{00}(x_l, \boldsymbol{\vartheta})\}^{R_{l0}}.$$

Since at an early stage in a trial the sample size is small, the Bayesian approach is employed to estimate the parameters. We obtain the posterior estimates of the components of $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4, \vartheta_5, \vartheta_6)^T$ at the kth stage as

$$\hat{\vartheta}_{ik} = \frac{\int_\Theta \vartheta_i \ p(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) d\boldsymbol{\vartheta}}{\int_\Theta p(\boldsymbol{\vartheta}) \ L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) \ d\boldsymbol{\vartheta}}, \quad i = 1, 2, 3, 4, 5, 6, \tag{4.11}$$

where $p(\boldsymbol{\vartheta})$ is the prior distribution of the parameters. Let us assume that $l_{i1} < \vartheta_i < l_{i2}$ for $i = 1, \ldots, 6$ and that the prior joint probability density function of the parameters is uniform. Define $\tilde{\Theta} = \{\boldsymbol{\vartheta} : \ l_{i1} < \vartheta_i < l_{i2} \text{ for } i = 1, \ldots, 6\}$. Then we

have

$$p(\boldsymbol{\vartheta}) = \frac{1}{(l_{12}-l_{11})(l_{22}-l_{21})(l_{32}-l_{31})(l_{42}-l_{41})(l_{52}-l_{51})(l_{62}-l_{61})}, \quad \boldsymbol{\vartheta} \in \tilde{\Theta}. \tag{4.12}$$

This function will be used in our example of simulating clinical trials in Section 6.4.

4.3.4 Fisher Information Matrix

The likelihood function for a single cohort at dose x is

$$L(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}}) = \frac{c!}{R_0!R_1!R_2!R_3!}(\psi_{11})^{R_3}(\psi_{10})^{R_2}(\psi_{01})^{R_1}(1-\psi_{11}-\psi_{10}-\psi_{01})^{(c-R_3-R_2-R_1)},$$

and so the log-likelihood function is

$$\begin{aligned} l(\boldsymbol{\vartheta}|x, \tilde{\boldsymbol{R}}) &= \text{constant} + R_3\log(\psi_{11}) + R_2\log(\psi_{10}) + R_1\log(\psi_{01}) \\ &\quad +(c-R_3-R_2-R_1)\log(1-\psi_{11}-\psi_{10}-\psi_{01}). \end{aligned}$$

As before, the information matrix associated with $\boldsymbol{\vartheta}$ can be obtained using the formula (4.6) for $\boldsymbol{\psi} = (\psi_{11}, \psi_{10}, \psi_{01})$. Analogous derivations as for the previous model, presented below, lead to the form of the FIM which can further be used in optimal dose selection. The partial derivatives are

$$\frac{\partial^2 \log(\psi_{11})}{\partial\boldsymbol{\psi}\,\partial\boldsymbol{\psi}^T} = \begin{bmatrix} -\frac{1}{\psi_{11}^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \frac{\partial^2 \log(\psi_{10})}{\partial\boldsymbol{\psi}\,\partial\boldsymbol{\psi}^T} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -\frac{1}{\psi_{10}^2} & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

$$\frac{\partial^2 \log(\psi_{01})}{\partial\boldsymbol{\psi}\,\partial\boldsymbol{\psi}^T} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\psi_{01}^2} \end{bmatrix}$$

and

$$\frac{\partial^2 \log(1-\psi_{11}-\psi_{10}-\psi_{01})}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T} = -\frac{1}{(1-\psi_{11}-\psi_{10}-\psi_{01})^2}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

Therefore,

$$\begin{aligned} I(x,\boldsymbol{\psi}) &= E\left[-\frac{\partial^2 l(\boldsymbol{\vartheta}|x,\tilde{\boldsymbol{R}})}{\partial \boldsymbol{\psi}\ \partial \boldsymbol{\psi}^T}\right] \\ &= E(R_3)\begin{bmatrix} \frac{1}{\psi_{11}^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + E(R_2)\begin{bmatrix} 0 & 0 & 0 \\ 0 & \frac{1}{\psi_{10}^2} & 0 \\ 0 & 0 & 0 \end{bmatrix} + E(R_1)\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{1}{\psi_{01}^2} \end{bmatrix} \\ &\quad + E(c-R_3-R_2-R_1)\frac{1}{(1-\psi_{11}-\psi_{10}-\psi_{01})^2}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}. \end{aligned} \tag{4.13}$$

Since $E(R_3) = c\psi_{11}$, $E(R_2) = c\psi_{10}$, $E(R_1) = c\psi_{01}$ and $E(c - R_3 - R_2 - R_1) = c(1-\psi_{11}-\psi_{10}-\psi_{01})$, after some simplification, (4.13) reduces to

$$\begin{aligned} &I(x,\boldsymbol{\psi}) \\ &= c\left(\begin{bmatrix} \frac{1}{\psi_{11}} & 0 & 0 \\ 0 & \frac{1}{\psi_{10}} & 0 \\ 0 & 0 & \frac{1}{\psi_{01}} \end{bmatrix} + \frac{1}{(1-\psi_{11}-\psi_{10}-\psi_{01})}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}\right) \\ &= \frac{c}{(1-\psi_{11}-\psi_{10}-\psi_{01})}\begin{bmatrix} \frac{1-\psi_{10}-\psi_{01}}{\psi_{11}} & 1 & 1 \\ 1 & \frac{1-\psi_{11}-\psi_{01}}{\psi_{10}} & 1 \\ 1 & 1 & \frac{1-\psi_{11}-\psi_{10}}{\psi_{01}} \end{bmatrix}. \end{aligned} \tag{4.14}$$

It can also be shown that

$$\left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}}\right)^T = \begin{bmatrix} -\psi_{11}\psi_{01} & -\psi_{10}\psi_{01} & \psi_{01} - \psi_{01}^2 \\ -x\psi_{11}\psi_{01} & -x\psi_{10}\psi_{01} & x(\psi_{01} - \psi_{01}^2) \\ -\psi_{11}\psi_{10} & \psi_{10} - \psi_{10}^2 & -\psi_{01}\psi_{10} \\ -x\psi_{11}\psi_{10} & x(\psi_{10} - \psi_{10}^2) & -x\psi_{01}\psi_{10} \\ \psi_{11} - \psi_{11}^2 & -\psi_{10}\psi_{11} & -\psi_{01}\psi_{11} \\ x(\psi_{11} - \psi_{11}^2) & -x\psi_{10}\psi_{11} & -x\psi_{01}\psi_{11} \end{bmatrix}.$$

Furthermore, we have

$$I(x, \boldsymbol{\psi}) \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}}\right) = c \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & x \\ 0 & 0 & 1 & x & 0 & 0 \\ 1 & x & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Finally, we obtain the FIM for a cohort of patients as

$$I(x, \boldsymbol{\vartheta}) = \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}}\right)^T I(x, \boldsymbol{\psi}) \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{\vartheta}}\right)$$

$$= c \begin{bmatrix} \psi_{01} - \psi_{01}^2 & x(\psi_{01} - \psi_{01}^2) & -\psi_{01}\psi_{10} & -x\psi_{01}\psi_{10} & -\psi_{01}\psi_{11} & -x\psi_{01}\psi_{11} \\ x(\psi_{01} - \psi_{01}^2) & x^2(\psi_{01} - \psi_{01}^2) & -x\psi_{01}\psi_{10} & -x^2\psi_{01}\psi_{10} & -x\psi_{01}\psi_{11} & -x^2\psi_{01}\psi_{11} \\ -\psi_{10}\psi_{01} & -x\psi_{10}\psi_{01} & \psi_{10} - \psi_{10}^2 & x(\psi_{10} - \psi_{10}^2) & -\psi_{10}\psi_{11} & -x\psi_{10}\psi_{11} \\ -x\psi_{10}\psi_{01} & -x^2\psi_{10}\psi_{01} & x(\psi_{10} - \psi_{10}^2) & x^2(\psi_{10} - \psi_{10}^2) & -x\psi_{10}\psi_{11} & -x^2\psi_{10}\psi_{11} \\ -\psi_{11}\psi_{01} & -x\psi_{11}\psi_{01} & -\psi_{11}\psi_{10} & -x\psi_{11}\psi_{10} & \psi_{11} - \psi_{11}^2 & x(\psi_{11} - \psi_{11}^2) \\ -x\psi_{11}\psi_{01} & -x^2\psi_{11}\psi_{01} & -x\psi_{11}\psi_{10} & -x^2\psi_{11}\psi_{10} & x(\psi_{11} - \psi_{11}^2) & x^2(\psi_{11} - \psi_{11}^2) \end{bmatrix}.$$

The information matrix $I(x, \boldsymbol{\vartheta})$ has rank 3, since it has three linearly independent rows or columns. Therefore, we require at least three different doses to be assigned to the cohorts to obtain a nonsingular information matrix. This is particularly important as both D and the combined criterion are defined in terms of the determinant of the FIM.

Chapter 5

Adaptive Designs

5.1 Introduction

This chapter introduces approaches for dose finding in phase IB/IIA clinical trials. In one of the approaches, along with dose-response outcomes, we also incorporate PK information in dose escalation. The aim is to develop an efficient dose-finding method, so that it exposes not too many patients to either subtherapeutic or toxic doses and also it can identify the best dose for further study in the next phase.

We begin with the general algorithm in Section 5.2. Assume that the patients enter a clinical trial sequentially and cohorts of the same size are treated with a dose level determined from the updated information. The choice of dose level for each cohort is model based and satisfies an optimisation criterion. There are various possible criteria, such as the maximum tolerated dose, the biologically optimum dose or the D-optimum dose. The maximum tolerated dose is the dose level for which the probability of toxicity attains a maximum permissible value. This criterion is often used in oncology trials, as it is usually assumed that both the efficacy and toxicity probabilities increase with dose level. However, in cases where we can observe efficacy with no toxicity, the outcome which we call a success, it makes sense to consider a criterion which allows for the highest chance of such an outcome. Alternatively, one can consider a criterion which in principle should lead to the best dose-response model prediction and so the best indication of the efficacious dose

level.

5.2 General Algorithm

The proposed method is model based. Hence, we assume that a dose-response model and a PK model, when appropriate, are known apart from the parameters. The choice of these models for a specific drug is generally elicited from the experts in the area. Pre-clinical studies of the candidate drug or previous studies of similar drugs can provide guidance. We assume that the set of d ordered doses $\mathcal{X} = \{x^{(1)}, \ldots, x^{(d)}\}$ of a candidate drug is available for experimentation. The goal is to find the best of these doses, which we call the *optimum dose* (OD), for further study in the next phase. Below we present the main steps of the adaptive design, where k represents the stage in a trial. We set $k = 1$ initially.

Step 1: Treat cohort k with the current best dose.

Step 2: Observe the PK responses at the locally D-optimal sampling time points, when appropriate.

Step 3: Observe the dose-response outcomes.

Step 4: Estimate the model parameters and update the models.

Step 5: Select the best dose for the next cohort based on the chosen dose optimisation criterion and constraints.

Step 6: Stop if the stopping rule is met, otherwise set $k = k + 1$ and repeat Steps 1-5.

Step 7: Carry out a complete analysis of the data to recommend a dose for further study.

At any stage of the trial, if it is found that none of the available doses satisfy the constraints, then we assign the lowest dose to the next cohort. The above adaptive procedure requires estimation of the dose-response model parameters, and also the PK parameters in the case when PK constraints are used. With a small cohort size, we need to gather data from a few cohorts before we estimate any parameters. Hence, we start the trial with an up-and-down procedure, which is run for the

first four cohorts before we start the fully adaptive parametric algorithm described above. This is similar to the procedure of Ivanova (2006), where a dose is increased, stays at its present level or is decreased depending on the responses of the most recent cohort. Here, however, we take into account responses from all cohorts up to the most recent one (Bogacka et al., 2014). The basis of the method lies in the toxic outcomes, which can stop the trial if toxicity is above the acceptable level.

Assume that we are at the kth stage in a trial and that patients in the successive k cohorts have received doses from a pre-specified sequence of doses $\mathcal{X}$, the same dose within a cohort. Let us denote the proportion of toxic responses up to cohort k by $\hat{p}_k$, that is,

$$\hat{p}_k = \frac{1}{kc} \sum_{i=1}^{k} R_{\text{tox}_\text{i}},$$

where c is the cohort size and R_{tox_i} is the number of toxic responses for the ith cohort upon receiving a dose. The algorithm starts with the lowest dose from $\mathcal{X}$. Then, for given thresholds p_L, p_M and p_U, we increase, stay at the same dose level, decrease or stop the trial, depending on the value of $\hat{p}_k$. In our study, we set $p_L = \gamma/3$, $p_M = 2\gamma/3$ and $p_U = \gamma$, where γ is the maximum acceptable level for the probability of toxicity. More specifically, the algorithm has the following structure to follow:

$$\hat{p}_k \begin{cases} \leq p_L & \text{increase the dose to the next level if not at the highest} \\ & \text{level, otherwise stay at the highest level,} \\ \in (p_L, p_M) & \text{stay at the current dose,} \\ \in [p_M, p_U) & \text{decrease the dose by one level if not at the lowest} \\ & \text{level, otherwise stay at the lowest level,} \\ \geq p_U & \text{stop the trial.} \end{cases}$$

Although we choose the above steps to be 1/3, it can be changed if some indications of the chances of toxicity suggest that other thresholds may be more appropriate. The main underlying idea of this stage of the trial is to gather information on the

responses, and thus on the parameters without exposing many patients to a risk of toxicity.

Once the trial is finished with the up-and-down stage, it moves to the adaptive parametric algorithm. As we have seen in Chapter 4, since the FIMs for a cohort are singular, to facilitate the computation of the D-optimum design for the continuation ratio model, we ensure that at least two of the four doses are different in the up-and-down stage. Assume that a trial stops after revisiting the above two algorithms K times. Then we will have K cohorts treated at the doses $x_1, \ldots, x_K$, selected from $\mathcal{X}$ in the different stages. A flow chart indicating different steps of the proposed design is given in Section A.2.

Since the PK models are non-linear in the parameters, the search for the optimal time points depends on the parameter values. As they are unknown initially, we begin with some guesses, which may be quite inaccurate. It is therefore important to check the sensitivity of the design to appreciate the impact of such a choice of parameter values. However, as the trial proceeds, more accurate estimates are calculated and hence we obtain more reliable optimal sampling time points.

The dose-selection criteria at each stage and also the stopping rule are important decision functions. They will depend on the objective of the trial.

5.3 Criteria for Dose Optimisation

At each stage in the trial, we select that dose for the next cohort for which the desired criterion is maximised, subject to the condition that a number of constraints are satisfied. Often, it is not advisable to skip dose levels when they are increased for application in the next cohort. Therefore, in our method, we introduce an option of constraining the increase by any number of dose levels. However, we impose no such constraint on the levels when they are decreased. For example, the next best dose could be five levels higher then the previous one, but with a constraint

of not skipping more than one dose level, we apply the one two levels higher than the previous one. The following sections describe some possible dose-optimisation criteria that can be used in a trial. We use these criteria in the simulation studies in Chapter 6.

5.3.1 Maximisation of Probability of Success

This criterion is based on the probability of success. Recall that a success is defined as that outcome which produces a desired level of efficacy without severe toxicity.

Assume that we are at the kth stage in a trial and based on the current data, we have the estimates $\hat{\boldsymbol{\vartheta}}_k$ of the dose-response parameters. Then we select the dose x_{k+1} for the next cohort of patients so that

$$x_{k+1} = \arg\max_{x \in \mathcal{X}} \psi_S(x, \hat{\boldsymbol{\vartheta}}_k), \tag{5.1}$$

where ψ_S is the probability of success at a given dose. In the trinomial dose-response model introduced in Chapter 4, ψ_1 is taken as ψ_S. For bivariate binary outcomes, it is ψ_{10}.

The above criterion is defined so that it allocates the most efficacious doses to the cohorts during a trial. The issue of allocating the most efficacious doses during a trial is very important from an ethical point of view. The criterion, by definition, does not take into account the efficiency of parameter estimation. As mentioned in Chapter 4, the estimates $\hat{\boldsymbol{\vartheta}}_k$ can be obtained by using either a Bayesian or frequentist approach. Some authors (Pronzato, 2000; Fedorov et al., 2011) report that the above algorithm may converge to a sub-optimal dose. These studies are based on the frequentist approach, the least squares or the maximum likelihood estimates of the parameters.

We are using Bayesian parameter estimation, and, to our knowledge, the convergence properties are not known. The above algorithm may converge to a sub-optimal dose

due to the possibility of poor estimates of the parameters. This possibility is embedded in the criterion (5.1), which chooses the best dose for the efficacy purpose rather than for precise estimation of the parameters. It may happen that, when the responses are efficacious for the same dose given to a number of consecutive cohorts, the algorithm will not explore other doses very much and consequently will choose that dose as the optimum one. Hence, for a small number of cohorts, we observe that on some occasions the algorithm chooses a sub-optimal dose. However, in the majority of cases, the algorithm stops at the true optimum dose, as shown in Figures 6.2-6.5 and 6.15-6.18.

Fedorov et al. (2011) showed that the convergence property holds for the penalised D-optimum design in Section 5.3.3 when a large number of patients, such as 400, are enrolled in the trial. For small numbers of patients, all known methods can lead to a sub-optimal dose. The criterion in (5.1) has been found to outperform the penalised D-optimum design in (5.8) for three of the presented scenarios: see Tables 6.13-6.16. This provides evidence of where a design with the convergence property may end up with a sub-optimal dose. Employing a very large number of patients is difficult, as most of the early phase designs aim for a small number like 30-60. Consequently, in a practical situation, there will always be a risk of selecting a sub-optimal dose whatever the convergence properties of the applied method.

5.3.2 Maximisation of Determinant of FIM

This approach allocates those doses to the cohorts which contribute most to the efficient estimation of the dose-response parameters. Assume that we are at the kth stage in a trial. So, the doses allocated to the cohorts are $\boldsymbol{\xi}_k = \{x_1, x_2, \ldots, x_k\}$. Also, based on the data so far we have the estimates $\hat{\boldsymbol{\vartheta}}_{\boldsymbol{k}}$. Let us define

$$M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \frac{k}{k+1} M(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) + \frac{1}{k+1} I(x, \hat{\boldsymbol{\vartheta}}_k), \tag{5.2}$$

where $M(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \sum_{i=1}^{k} I(x_i, \hat{\boldsymbol{\vartheta}}_k)$ and $I(x, \hat{\boldsymbol{\vartheta}}_k)$ is the Fisher information matrix for a cohort which received the dose x. For both of our dose-response models, the

expressions for $I(x, \boldsymbol{\vartheta})$ are shown in Chapter 4. Atkinson et al. (2014) discuss the construction of designs based on augmented FIMs. Bogacka et al. (2014) use the same approach for the construction of design.

Then we select the dose x_{k+1} for the next cohort of patients such that

$$x_{k+1} = \arg\max_{x \in \mathcal{X}} \Phi_D\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}, \tag{5.3}$$

where $\Phi_{\mathrm{D}}\{M\} = |M|$. It is well known that a D-optimum design tends to assign doses from the extremes of $\mathcal{X}$. Therefore, patients are likely to receive non-efficacious and toxic doses during a trial.

5.3.3 Combined Criterion

Clinicians may be interested in achieving several objectives, such as efficient estimation of the model parameters and allocation of the most efficacious doses to the cohorts during a clinical trial. The combined criterion in (5.6) and the penalised version in (5.9) balance these two objectives. The criterion is defined so that one can obtain the design which ensures either efficient parameter estimation or efficacious dose allocation or a combination of both.

The combined criterion that we want to utilise for dose selection is a linear combination of the determinant of the Fisher information matrix for the dose-response model and the probability of success. At each stage of the trial, we will be selecting that dose for which the criterion is maximised. The main idea is to find a dose that will be appropriate from the efficacy point of view and will also lead to the efficient estimation of the dose-response parameters.

To implement the method, we initially determine the doses that maximise the determinant of the FIM and the probability of success

$$x_{k+1}^{D} = \arg\max_{x \in \mathcal{X}} \Phi_D\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}$$

and

$$x_{k+1}^{\psi_S} = \arg\max_{x \in \mathcal{X}} \psi_S(x, \hat{\boldsymbol{\vartheta}}_k),$$

where $M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)$ is defined in (5.2). Since the determinant and the probability of efficacy may have quite different magnitudes, we scale them at the dose x as

$$E_D(x) = \frac{\Phi_D\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}}{\Phi_D\{M(x_{k+1}^D|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}} \tag{5.4}$$

and

$$E_{\psi_S}(x) = \frac{\psi_S(x, \hat{\boldsymbol{\vartheta}}_k)}{\psi_S(x_{k+1}^{\psi_S}, \hat{\boldsymbol{\vartheta}}_k)}. \tag{5.5}$$

The combined criterion then selects the dose x_{k+1} for the next cohort of patients so that

$$x_{k+1} = \arg\max_{x \in \mathcal{X}} \left\{ a E_D(x) + (1-a) E_{\psi_S}(x) \right\}, \tag{5.6}$$

where a is some weight such that $0 \leq a \leq 1$.

It is clear that, when $a = 1$, the combined criterion is simply the D-criterion. Similarly, for $a = 0$, we have dose selection based on the probability of success only. This design is expected to allocate the most efficacious doses to the cohorts compared to the D-optimum design and also to recommend the best dose for further study.

Penalty Function

In the D-optimal design for the dose-response model, defined in (5.3), the doses allocated to the cohorts are often found at the extremes of the design region. Therefore, such a design suffers from the limitation that patients in a trial may be exposed to non-efficacious or highly toxic doses. To reduce the possibility of exposing patients to such doses, we introduce a penalty function in the search for D-optimum doses

and consider a penalised D-optimum design (Dragalin and Fedorov, 2006).

We denote the penalty function for an observation taken at the design point x by $\varphi(x, \boldsymbol{\vartheta})$. Following Dragalin and Fedorov (2006), a simple choice for that function would take the form

$$\varphi(x, \boldsymbol{\vartheta}) = \{\psi_S(x, \boldsymbol{\vartheta})\}^{-C_S}\{1 - \psi_T(x, \boldsymbol{\vartheta})\}^{-C_T}, \tag{5.7}$$

where ψ_S and ψ_T are the probabilities of success and toxicity at a given dose, respectively, C_S and C_T are control parameters used to construct an appropriate penalty function at the low and high dose levels. The function is defined so that the lower the probability of success and/or the higher the probability of toxicity at the assigned dose, the higher the penalty for the observation taken.

For a given x and $\boldsymbol{\vartheta}$, the larger the values of C_S and C_T, the higher the value of the penalty function will be, and thus the design will avoid allocating doses with a low probability of efficacy or a high probability of toxicity to the patients. Assume that $C_S = C_T = C = 0$, which in turn means no penalty. Then this will lead to the D-optimum design. The larger the value of C is, the further it is from the D-optimum design and the less efficient the parameter estimation.

Penalised D-Criterion

The penalised D-criterion takes into account a penalty function, as shown in (5.7). In this version of the D-optimum design, we select the dose as

$$x_{k+1} = \arg\max_{x \in \mathcal{X}} \Phi_{\mathrm{PD}}\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}, \tag{5.8}$$

where $\Phi_{\mathrm{PD}}\{M\} = |M/\boldsymbol{\varphi}|$.

The matrix $M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)$ is defined in (5.2) and also

$$\boldsymbol{\varphi}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \frac{k}{k+1}\boldsymbol{\varphi}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) + \frac{1}{k+1}\varphi(x, \hat{\boldsymbol{\vartheta}}_k),$$

where

$$\boldsymbol{\varphi}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \sum_{i=1}^{k} \varphi(x_i, \hat{\boldsymbol{\vartheta}}_k).$$

Penalised Combined Criterion

The penalised combined criterion is similar to the one in (5.6). The only difference is that, instead of the D-criterion, we are considering here the penalised D-criterion. As before, we need to determine the doses based on the maximisation of the determinant of the FIM and the probability of success. The penalised D-optimal dose at stage k of the adaptive procedure is obtained as

$$x_{k+1}^{PD} = \arg\max_{x \in \mathcal{X}} \Phi_{PD}\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}.$$

We scale the determinant value at the dose x as

$$E_{PD}(x) = \frac{\Phi_{PD}\{M(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}}{\Phi_{PD}\{M(x_{k+1}^{PD}|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}}.$$

Then the penalised combined criterion selects the dose x_{k+1} such that

$$x_{k+1} = \arg\max_{x \in \mathcal{X}}\{aE_{PD}(x) + (1-a)E_{\psi_S}(x)\}, \tag{5.9}$$

where $E_{\psi_S}(x)$ is obtained as in (5.5). If $a = 1$, then (5.9) simply becomes the penalised D-criterion. For $a = 0$, the dose selection is based on the probability of success only. When $a = 0.5$, it gives equal weight to both.

Although the penalised D-criterion introduces a penalty function to improve the quality of treatments during dose escalation, we have found the quality not to be improved as expected. Therefore, further effort has been taken through the above

combined criterion. Although the probability of success appears twice in the criterion, through $E_{PD}(x)$ and $E_{\psi_S}(x)$, the roles they play are different. In the penalised D-criterion, the probability of success is used to scale the FIM so that doses with low efficacy or high toxicity result in small values for the determinant of the FIM. The E_{PD} part of the criterion guides us to choose the dose that will provide maximum information regarding parameter estimation taking into account the efficacy levels of all of the available doses. The maximisation of the probability of success on the other hand tends to choose a dose from the efficacy view point, without caring for parameter estimation. A combination of them, as presented above, is expected to facilitate a balance between the two approaches.

A design is presented in Section 6.6 that utilises the penalised combined criterion. The gains over the penalised D-criterion are quite evident from the results in Tables 6.13-6.16. All of the performance measures are found to be improved there. Most importantly, we notice an appreciable improvement in the quality of treatment allocation, reflected through the sampling efficiency measure, SE. The quality of optimum dose selection for the next phase, presented through DE, is also found to be improved.

5.4 Constraint on Probability of Toxicity

Let γ be the maximum acceptable level for the probability of toxicity. Then, for the next cohort of patients, we select the dose x_{k+1} for which the chosen dose-optimisation criterion is maximised subject to the constraint

$$\psi_T(x_{k+1}, \hat{\boldsymbol{\vartheta}}_k) \leq \gamma, \tag{5.10}$$

where ψ_T is the probability of toxicity evaluated at the current estimates of the parameters. In the trinomial dose-response case, ψ_T is simply ψ_2. But, in the case of bivariate binary outcomes, it is the marginal probability of toxicity obtained as $\psi_{\cdot 1} = \psi_{01} + \psi_{11}$.

The main purpose of introducing this constraint is to select a subset of doses which are safe, and then to administer one of them to the next cohort based on a dose-optimisation criterion. The constraint restricts us from choosing a dose for which the estimated probability of toxicity is above the acceptable level.

5.5 PK-Constrained Dose Optimisation

Here, we focus on the constrained dose-optimisation algorithm based on pharmacokinetic parameters, such as the AUC and $C_{\max}$. Too low a concentration provides no effect and similarly too high a concentration leads to a toxic outcome. It is possible to avoid excessive drug concentration by putting restrictions on the AUC or $C_{\max}$. Therefore, from the safety point of view, restricting these parameters will be worthwhile. Although one can restrict the choice of a dose by the probability of toxicity constraint, it does not take into account population variability in PK measures. However, the additional PK constraint incorporates this variability.

5.5.1 Area Under the Concentration Curve

We have explained how the concentration of a drug contributes towards the effects in Section 1.1.1. The area under the concentration curve is an important PK parameter and it depends on the concentration: see Section 3.3. The outcomes in dose-finding studies are often dichotomised, and, as a result, we lose some information. In particular, an outcome which is not toxic may be just below the cut-off point or a toxic outcome may be far above the cut-off point. Similarly, the efficacy is dichotomised. Usually these issues are not considered in the dose-escalation methods. However, they can be considered implicitly by taking into account continuous measures like the AUC in the dose-finding methods.

The proposed constraint is related to the total concentration of the drug in the body so that the curative purpose is likely to be achieved, and this is expressed by the area under the concentration curve over time. We select a dose x_{k+1} so that the

chosen criterion is maximised subject to the condition

$$\frac{h(x_{k+1}, \hat{\boldsymbol{\beta}}_k) - \text{AUC}^0}{\widehat{\text{SD}}\{h(x_k, \boldsymbol{\theta}_i)\}} \leq \delta(x_k, \hat{\boldsymbol{\vartheta}}_k). \tag{5.11}$$

The vector of estimates of the population mean PK parameters is $\hat{\boldsymbol{\beta}}_k$ and $h(x, \hat{\boldsymbol{\beta}}_k)$ is the estimate of the approximate population mean AUC at stage k. The estimate of the approximate standard deviation of the AUC is denoted by $\widehat{\text{SD}}\{h(x_k, \boldsymbol{\theta}_i)\}$. This notation is explained in Section 3.8.1. Also, AUC^0 is a value for the AUC that is considered to be desirable and $\delta(x_k, \hat{\boldsymbol{\vartheta}}_k) = 1/\psi_S(x_k, \hat{\boldsymbol{\vartheta}}_k)$. A desirable AUC is one for which the curative purpose is likely to be achieved, allowing some acceptable level of toxicity. The choice of such a value will require expert opinion. Previous studies of similar drugs or pre-clinical studies can help in this context.

The left-hand side of (5.11) represents a relative difference between $h(x_{k+1}, \hat{\boldsymbol{\beta}}_k)$ and AUC^0. We constrain the choice of x_{k+1} so that, for large values of the estimated probability of success, this difference is small. This 'forces' convergence of the dose to the one giving the required exposure to the drug. On the other hand, when the estimated probability of success is small, the constraint is weak, allowing for a wider choice for the next dose level. The PK constraint (5.11) is dynamic, that is, the value of $\delta(x_k, \hat{\boldsymbol{\vartheta}}_k)$ changes during the trial according to the current estimate of the probability of success. This gives some flexibility to the algorithm. A fixed δ might lead to choosing a sub-optimal dose, and, in any case, it would be difficult to decide on its value.

Hence, instead of using a fixed δ, we are expressing it in terms of ψ_S, whose value is updated after each cohort. Further justification for choosing δ as the reciprocal of ψ_S is that, for a dose at which the probability of success is high, we can assume that the optimum dose is in the neighbourhood of the current dose and therefore would like to restrict the search through the small value of δ. On the other hand, for a dose with a small probability of success, we want the algorithm to be able to search a wide dose region. The δ can be regarded as the tolerance, since its values have

an impact on the selection of the optimum dose. The relationship between δ and ψ_S is intended to have small tolerance towards the end of a trial, so that it enables the choice of a dose which has a mean AUC close to the target. If the tolerance is high, the margin to the right of the target will be high, and there is a possibility to choose a dose which is excessively toxic.

It follows from the constraint that $h(x_{k+1}, \hat{\boldsymbol{\beta}}_k) \leq \text{AUC}^0 + \delta(x_k, \hat{\boldsymbol{\vartheta}}_k)\ \widehat{\text{SD}}\{h(x_k, \boldsymbol{\theta}_i)\}$. If $\psi_S(x_k, \hat{\boldsymbol{\vartheta}}_k)$ attains the maximum possible value 1, then $\delta(x_k, \hat{\boldsymbol{\vartheta}}_k)$ will have the value 1 and consequently we will choose a dose with a mean AUC within one standard deviation of the target, which accommodates the population variability. However, in the majority of cases, δ will have a larger value than 1. Therefore, we will usually be selecting a dose with a mean AUC within more than one standard deviation of the target value. This constraint, as well as the next, is introduced as an additional precaution against allocation of too toxic doses. They work differently to (5.10), which is the constraint on the estimated probability of toxicity. The AUC and $C_{\max}$ not only take into account the population variability, but also directly constrain the pharmacokinetic parameters responsible for the drug's action.

To the best of our knowledge, such constraints have not been considered before in the adaptive design set-up. We have tried other functions of ψ_S for a choice of δ, such as $\delta = 1/2\psi_S$ and $\delta = 1/(1+\psi_S)$, but they were too conservative in identifying the optimum dose.

5.5.2 Maximum Concentration

The quantity $C_{\max}$ is related to the side effects that may result from excessive drug concentration. Therefore, it is required to be maintained at a level to ensure safety. We introduce another dynamic constraint for $C_{\max}$ to avoid unacceptable adverse events. The criterion tells us to select the dose x_{k+1} based on a dose-optimisation

criterion so that

$$\frac{l(x_{k+1}, \hat{\boldsymbol{\beta}}_k) - C^0_{\max}}{\widehat{\mathrm{SD}}\{l(x_k, \boldsymbol{\theta}_i)\}} \leq \delta(x_k, \hat{\boldsymbol{\vartheta}}_k). \tag{5.12}$$

The vector of estimates of the population mean PK parameters is $\hat{\boldsymbol{\beta}}_k$, $l(x, \hat{\boldsymbol{\beta}}_k)$ is the estimate of the approximate population mean $C_{\max}$ at stage k and $\widehat{\mathrm{SD}}\{l(x_k, \boldsymbol{\theta}_i)\}$ is the estimate of the approximate standard deviation of the maximum concentration: see Section 3.8.2. Also, $C^0_{\max}$ is a value for the $C_{\max}$ that is considered to be desirable, and, as before, we have $\delta(x_k, \hat{\boldsymbol{\vartheta}}_k) = 1/\psi_S(x_k, \hat{\boldsymbol{\vartheta}}_k)$.

It is important to note that the constraints alone are not used for dose finding, but rather together with a dose-optimisation criterion, they help in dose finding. The choice of optimisation criterion and constraints depends on the purpose to be achieved.

5.6 Stopping Rules

A variety of stopping rules are possible and the choice of one solely depends on the purpose of a trial. In many of the designs, a trial is run for a fixed number of cohorts. O'Quigley and Reiner (1998) stop the trial early on the basis that continuing it would not lead to a change in the dose recommendation. The essence of the method is that, before recruiting all of the patients, if we can predict with high probability what the final recommendation will be, the trial can stop. This early stopping rule is based on precise probabilistic calculations and not straightforward to implement. Heyd and Carlin (1999) recommend a stopping rule based on an approximate confidence interval for the probability of toxicity at the recommended dose. A trial is stopped if this confidence interval contains some pre-specified range of target toxicities. However, the approach suffers from the limitation that reaching a precise interval would require a reasonable number of patients to be recruited in the trial. Also, finding the range for target toxicities is a major concern. A simple rule is proposed by O'Quigley (2002). It stop the trial if the dose recommended to

the next patient has already been allocated m times, where m is some number fixed at the beginning of the trial. Zhang et al. (2006) terminate a trial after treating at least n_1 patients, provided at least n_0 are treated at the recommended dose, or a maximum number of n_2 subjects are treated. Dragalin and Fedorov (2006) stop a trial after reaching a fixed sample size.

A simple rule is employed in our examples in Chapter 6. We stop a trial when the same dose is repeated for r cohorts or when the trial reaches the maximum number of m cohorts, whichever comes first. The idea behind early stopping is the saving of resources if it is found that the same dose is being selected repeatedly. The assumption is that no further improvement is possible for the current trial. Otherwise, it will run for the maximum number of cohorts available. For early stopped trials, the OD is defined as the dose that has been repeated r times. However, for the trials that reach the maximum number of cohorts m, we carry out a complete analysis of the data and define the OD as the dose that would be allocated to cohort $m+1$ if that cohort were in the trial.

5.7 Evaluation of the Designs

To compare the produced adaptive designs based on different criteria and constraints, we introduce some performance measures. The measures are based on the simulation results. By a simulation, we refer to the completion of the up-and-down procedure and the adaptive algorithm to recommend a dose. We have two dose-response models and different dose-response scenarios for each of them, which were introduced in Chapter 4. For each of the scenarios, we have the true values of the dose-response parameters $\boldsymbol{\vartheta}$. The following sections define the performance measures.

5.7.1 Distribution of Dose Allocation

To find the distribution of the allocated doses to the cohorts over all simulations, we define

$$\bar{\xi} = \left\{ \begin{array}{ccc} x^{(1)} & , \ldots , & x^{(d)} \\ w^{(1)} & , \ldots , & w^{(d)} \end{array} \right\}, \tag{5.13}$$

where $w^{(j)}$ is the proportion of times that dose $x^{(j)}$ was allocated to the cohorts over all simulated trials, $j = 1, \ldots, d$, and d is the number of available dose levels. So

$$w^{(j)} = \frac{1}{n_{\text{total}}} \sum_{i=1}^{U} n_{ji}, \tag{5.14}$$

where $n_{\text{total}} = \sum_{i=1}^{U} n_{(i)}$ and $n_{(i)}$ denotes the number of cohorts used in the ith simulated trial. Thus, n_{total} is the total number of cohorts used in all simulations and n_{ji} is the number given dose $x^{(j)}$ in the ith simulation.

5.7.2 Distribution of Optimum Dose

The distribution of the OD is presented as

$$\tilde{\xi} = \left\{ \begin{array}{ccc} x^{(1)} & , \ldots , & x^{(d)} \\ \tilde{w}^{(1)} & , \ldots , & \tilde{w}^{(d)} \end{array} \right\}, \tag{5.15}$$

where $\tilde{w}^{(j)}$ denotes the proportion of times that dose $x^{(j)}$ was recommended for the next phase. It is obtained as

$$\tilde{w}^{(j)} = \frac{q_j}{U}, \tag{5.16}$$

where q_j is the number of times that dose $x^{(j)}$ was recommended for the next phase in U trials.

5.7.3 Decision Efficiency

Since, at the end of each trial a dose is recommended for the next phase, we introduce an efficiency criterion which is called the *decision efficiency (DE)*. Denote the ratio of ψ_S at a recommended dose x for the next phase to ψ_S at the dose x_{OD} by ρ. The latter dose is the true OD obtained from a scenario assumed for the simulation study. Therefore, we have

$$\rho(x, \boldsymbol{\vartheta}_{\text{true}}) = \frac{\psi_S(x, \boldsymbol{\vartheta}_{\text{true}})}{\psi_S(x_{\text{OD}}, \boldsymbol{\vartheta}_{\text{true}})}. \tag{5.17}$$

The ratio assesses how good the recommended dose level is. To measure the global efficiency of choosing the OD, we calculate the weighted sum

$$DE = \sum_{j=1}^{d} \tilde{w}^{(j)} \ \rho(x^{(j)}, \boldsymbol{\vartheta}_{\text{true}}) \ I_A(x^{(j)}), \tag{5.18}$$

where $I_A(x^{(j)})$ is an indicator function defined as

$$I_A(x^{(j)}) = \begin{cases} 1 & \text{if } x^{(j)} \notin A, \\ 0 & \text{if } x^{(j)} \in A, \end{cases} \tag{5.19}$$

and A is the set of doses for which the probability of toxicity is above the acceptable level. If x_{OD} is recommended in each of the simulated trials, then the weight is 1 corresponding to the OD and 0 for all other dose levels, and hence $DE = 1$. Similarly, if all of the recommended doses are from the toxic region, then the indicator function takes the value 0 for each of them, and hence $DE = 0$. In general, $0 \leq DE \leq 1$.

5.7.4 Sampling Efficiency

We define the *sampling efficiency (SE)* based on the information on the allocated doses to the cohorts over the simulations. With the $w^{(j)}$ defined in (5.14), we have

$$SE = \sum_{j=1}^{d} w^{(j)} \ \rho(x^{(j)}, \boldsymbol{\vartheta}_{\text{true}}) \ I_A(x^{(j)}). \tag{5.20}$$

This criterion assesses the quality of treatment received by the patients during the trials. If the cohorts in the trials only receive the OD, then the SE will have the value 1. In that case, there will be no experimental variation in the dose level. In general, we can expect the SE to be appreciably less than the DE, unless the variation in ψ_S over $\mathcal{X}$ is small. We have $0 < SE \leq 1$.

Hardwick et al. (2003) consider similar efficiency measures. However, they do not penalise for doses which have a probability of toxicity above the acceptable level. We have found that, in some dose-response scenarios, ρ may be greater than 1 for a few doses which have a probability of toxicity above the acceptable level. As a consequence, a design choosing these doses more frequently will have larger values for the efficiency measures than one that does not. This should not be the case, since these are the doses which a good design should not choose because of their toxicity. Therefore, we have presented measures that penalise for toxic doses. Our measures are capable of judging a design selecting low efficacious or high toxic doses as a poor one.

Chapter 6

Simulation Studies

6.1 Introduction

This chapter aims to study the operating characteristics of the designs obtained using the various dose-optimisation criteria and constraints introduced in the previous chapter. We present three examples utilising the PK and dose-response models introduced in Chapters 3 and 4. The first example comprises a one-compartment PK model with bolus input and first-order elimination and the continuation ratio model for the probabilities of the dose-response outcomes. The second one uses one-compartment PK model with first-order absorption and the Cox bivariate binary model as the dose-response model. The third example only uses dose-response data and the continuation ratio model is employed.

Section 6.2 gives a detailed description of the software used. For both dose-response models, in Chapter 4, we presented some plausible scenarios. Simulation studies are based on those scenarios to explore the behaviour of the designs. Section 6.3 contains a detailed description of the first example. The second example is presented in Section 6.4. Since the presented designs depend on various parameters, we also conduct sensitivity analyses to assess the robustness of the designs, the results of which are given in Section 6.5. The simulation results for the third example are presented in Section 6.6. Finally, there is a discussion in Section 6.7.

6.2 Software Used

All of the computations involved in the simulation studies are conducted using code written in *R* (R Core Team, 2014), which is available in Appendix B. The computer program follows the adaptive procedure which is shown schematically in the flow chart in Appendix A. The program has the following major parts: running the up-and-down design; obtaining the *D*-optimal time points for measuring concentration; the generation of the concentrations and dose-response outcomes; PK and dose-response parameter estimation; dose selection for the next cohort; checking the stopping rules; and identifying the OD.

The *D*-optimal time points for PK sampling are obtained by using *PFIM 3.2* (Bazzoli et al., 2010), an *R* package to evaluate and optimise designs in the context of population PK/PD experiments. The package *PFIM 3.2* includes two folders: PFIM 3.2 and examples. Some illustrations of the package using various models are available in the folder examples. The folder PFIM 3.2 contains three principal files: PFIM3.2.r, model.r and Stdin.r. These are the main program file, the model file and the input file, respectively. They need to be put in a working directory. Then we specify the working directory and the program directory in the file PFIM3.2.r. The package has a library of PK and PD models. In addition to these, users can define their own model. Since our models are available in the library, we can specify them in the model file. The input file requires specification of the optimisation method and essential parameter values to obtain the optimal time points. We have kept the model file general, so that the current estimates from the trial can be fed into it to find the locally optimal time points.

To obtain the maximum likelihood estimates of the PK parameters, we employ the *R* procedure *nlme* (Pinheiro and Bates, 2000). Bayesian estimates of the dose-response parameters are obtained by numerical integration using *cubature*, an *R* package (Johnson and Narasimhan, 2013). The combined criterion for dose optimisation requires the penalised *D*-optimal dose to be found, which is achieved by

some further code written in R. Here, we search for the dose, out of a set of discrete doses, for which the determinant of the FIM is maximum.

The program is structured in such a way that one has to first specify the dose range, acceptable level for the probability of toxicity, cohort size, dose-response and PK models, true parameter values, prior values, steps for the up-and-down design, dose-optimisation criterion and constraints, target values for the AUC and $C_{\max}$, early stopping rule and maximum number of cohorts to be employed. At the end of each simulation, it records the OD, the allocated doses to the cohorts and the parameter estimates. The simulations were implemented on a DELL PC with an Intel Core 2 Duo processor running at 3.00 GHz and RAM 4.00 GB. The processing time for 1,000 simulations depends on the choice of the models. It takes 6-8 hours when we only consider dose-response data. For both dose-response and PK data, it takes 10-15 hours.

6.3 Example 1

We assume a one-compartment model with bolus input and first-order elimination, introduced in Section 3.7.1, for modelling the concentrations of the experimental drug collected from the patients. The dose-response outcomes are assumed to be trinomial and the continuation ratio model, introduced in Section 4.2.1, is used to model the outcomes. Based on the updated information, we select a new dose for the next cohort so that the estimated probability of success is maximum, subject to a set of constraints. The trial stops according to the stopping rules described in the previous chapter.

6.3.1 Simulation Settings

Choice of Design Parameters

We assume that the set of doses of an experimental drug is $\mathcal{X} = \{0.5, 1.0, \ldots, 10.0\}$. Four hypothetical dose-response scenarios are considered: see Figure 4.1. Scenario

1 has a monotonically decreasing efficacy curve with dose 0.5 as the OD. Scenarios 2 and 3 depict non-monotonic efficacy curves with respective ODs of 5.5 and 6.5. A monotonically increasing efficacy curve is considered in Scenario 4 with 10 as the OD. Two kinds of responses need to be generated for the simulation study: the concentration of the drug in the blood and the trinomial dose-response outcomes. The values of the PK parameters for the simulation study are $V = 0.5$, $Cl = 0.06$, $\sigma_1^2 = 0.004$, $\sigma_2^2 = 0.00005$ and $\sigma^2 = 0.000225$. The parameter values are chosen such that the coefficient of variation is around 12%.

Each trial starts with the lowest dose of 0.5 mg/kg body weight. The acceptable toxicity level γ is taken to be 0.2. Doses for the first four cohorts in each trial are allocated according to the up-and-down design in Section 5.2. The value of AUC^0 is taken to be the AUC obtained from (3.27) for the true OD in the scenario and the true mean PK parameters. Although we consider the same γ for each scenario, we have different values of AUC^0. The doses which satisfy the safety level $\gamma = 0.2$ are 1 mg/kg, 6 mg/kg, 8 mg/kg and 10 mg/kg for the four scenarios, respectively. For each trial, we set the maximum number of cohorts to be $m = 20$. To allow the trials to stop early when it is found that no further improvement in dose selection is possible, we set $r = 6$.

Generation of PK and Dose-Response Outcomes

A vector of random effects $\boldsymbol{b}_i$ for individual i is generated from the normal distribution $N_2(\mathbf{0}, \boldsymbol{\Omega})$, where $\boldsymbol{\Omega} = \text{diag}(0.004, 0.00005)$. The PK parameters for that individual are then obtained as $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \boldsymbol{b}_i$, where $\boldsymbol{\beta} = (0.5, 0.06)^T$. The next step is to find the individual concentrations at the D-optimal time points. The design region for the sampling times is $\mathcal{T} = [0, t_1]$ hours, where $t_1 = 30$. To decide what is the best number of sampling times, we use the relative efficiency defined in (6.1). Although we consider the prior $\boldsymbol{\Psi}^0$ in each case, the designs have a different number of design points. It has been found that the efficiency of a 3-point design relative to a 2-point one is double. The efficiency of a 4-point design relative to a 3-point one

falls to 1.15. As we increase the number of design points, such efficiencies become closer to 1: see Figure A.1. This means that the gain is substantial if we consider 3 design points rather than 2. The increase from 3 design points to 4, or from 4 to 5 is negligible. The model has two so-called parameter sensitivities, which are the partial derivatives of the model function with respect to the parameters V and Cl. It is known from Section 3.7 that the FIM depends on the parameter sensitivities. For one of them, the maximum lies at the beginning of the design region, while, for the other, it is towards the end. Therefore, if the number of design points is increased, a more uniform distribution of the points over the region can be expected, as shown in Figure A.1. Furthermore, collecting many blood samples is often not possible. Therefore, three optimal time points are considered for the individuals in each cohort of size $c = 3$, and so $n_i = 3$ for all i. We use the R package *PFIM 3.2* (Bazzoli et al., 2010) to find these optimal time points. The random errors of the observations at the optimal sampling time points are then generated from $N_3(\mathbf{0}, \sigma^2 \boldsymbol{I}_3)$, where $\sigma^2 = 0.000225$, and added to the previously generated individual concentrations to produce the simulated PK responses for individual i. The same scheme is followed to simulate the responses for all individuals in each cohort. Figure 6.1 shows the generated concentrations for a cohort of three patients who received the dose 0.5 mg.

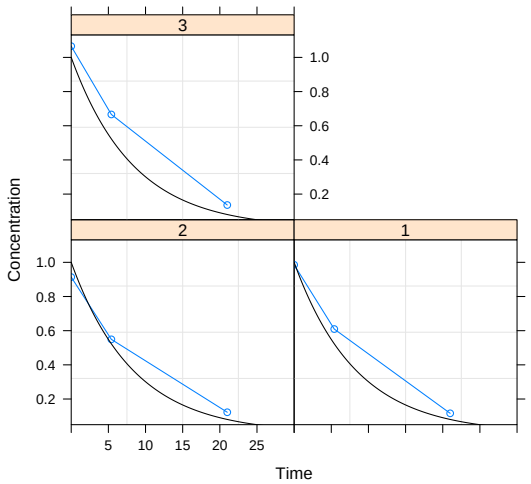

Figure 6.1: Simulated concentrations following the administration of the lowest dose to a cohort. The black curve indicates the true mean PK profile.

Under each of the scenarios in Figure 4.1, we have specific probabilities at each dose to generate the trinomial dose-response outcomes, that is, we also assume some true values for the model parameters.

Priors

We chose the prior values $\mathbf{\Psi^0} = (V^0, Cl^0, (\sigma_1^2)^0, (\sigma_2^2)^0, (\sigma^2)^0)^T = (0.1, 0.005, 0.0007, 0.0000006, 0.000004)^T$ to obtain the optimal time points for the first four cohorts in the up-and-down stage of the trial. The prior values are quite far away from the true values $\mathbf{\Psi} = (0.5, 0.06, 0.004, 0.00005, 0.000225)^T$ of the parameters, which would normally be unknown and could be wrongly assumed at the beginning of the trial. For the fifth cohort onwards, the current maximum likelihood estimates are used.

We use a joint uniform prior distribution for $\boldsymbol{\vartheta}$, given in (4.5). The parameter space $\tilde{\Theta}$ is chosen for each scenario so that the true values of the parameters lie in the middle of the corresponding intervals. For instance, since Scenario 1 has the true parameters $\boldsymbol{\vartheta} = (1.44, 0.26, -1.70, 0.25)^T$, $\tilde{\Theta}$ has $0 < \vartheta_1 < 2.88$, $0 < \vartheta_2 < 0.52$, $-3.40 < \vartheta_3 < 0$ and $0 < \vartheta_4 < 0.50$. More specifically, $u_1 = 0.52$, $u_2 = 0.50$, $l_1 = 0$, $l_2 = 2.88$, $l_3 = -3.40$ and $l_4 = 0$ in (4.5). The same approach is followed for the other scenarios. The chosen priors allow for any extreme scenario, as shown in Figure A.2. Since the parameters ϑ_2 and ϑ_4 have the value 0 in each case, the graphs in the left panel are identical. In evaluating the integrals in (4.4) with the uniform distribution specified in (4.5), the prior distributions in the numerator and denominator cancel out, as they are constants.

Model Fitting

Once we have data on the concentrations and the dose-response outcomes, we can update the fitted model. We obtain the maximum likelihood estimates of the PK parameters using the *R* procedure *nlme* (Pinheiro and Bates, 2000). The posterior estimates of the dose-response parameters in (4.4) are obtained by numerical integration using *cubature*, an *R* package (Johnson and Narasimhan, 2013). The package

carries out the adaptive multidimensional integration over hypercubes. One needs to specify the tolerance limit and the maximum number of function evaluations desired. The smaller the tolerance limit or the larger the maximum number of function evaluations, the more accurate an estimate is. Since the computational time is also an important consideration, we set these to be 0.001 and 5,000, respectively.

Dose Selection for the Next Cohort

Once we have the updated PK and dose-response parameter estimates, we select the dose for the next cohort based on the dose-optimisation criterion in (5.1), subject to two different sets of constraints in two separate runs of the simulations, that is, subject to (5.10) and to (5.10) together with (5.11). As discussed in Section 5.3, we do not allow the design to skip more than one dose level at a time when it is increased.

Checking the Stopping Rules and the OD Selection

We continue the process of allocating doses to the cohorts until the stopping rules in Section 5.6 are satisfied. Once a trial reaches m cohorts, we carry out a complete analysis to find the OD. At the end of each trial, we record the doses allocated to the cohorts, the PK and dose-response parameter estimates, whether the trial stopped early or not and the OD selected. Each of the four scenarios is investigated through 1,000 simulated trials. The following section summarises the results.

6.3.2 Numerical Results

We compare the operating characteristics of the PK-guided design, incorporating the AUC constraint, with the one that does not take into account the PK information. Dose selection in the latter design is based on the estimated probability of success and the toxicity condition only, defined in (5.1) and (5.10). Although the method is similar to the one presented by Zhang et al. (2006), their design criterion is the maximisation of the difference between the estimated probabilities of success and toxicity, subject to a toxicity constraint. The simulation results for our two designs

are presented in Table 6.1 and Figures 6.2-6.5. In these figures, the first row shows the summaries when the additional PK constraint is employed. The summaries in the second row are based on the toxicity constraint only. The bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at the allocated doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

Table 6.1 clearly shows the advantages of incorporating the additional constraint on the AUC profile. The gain in percentage of the doses correctly recommended for further studies depends on the scenario, but, in all cases considered, the PK-guided designs are uniformly better.

Table 6.1: Percentage of best doses recommended for further studies (%BD), percentage of doses recommended as optimum, but carrying the probability of toxicity above the maximum allowed threshold (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD).

Scenario	Best Doses	%BD		%TD		%AD	
		PK	No PK	PK	No PK	PK	No PK
1	0.5	99.0	52.4	0.6	32.7	65.3	31.8
2	5.5 and 6.0	80.2	66.2	0.9	9.5	41.1	33.5
3	5.5-7.5	91.7	85.7	0.0	2.5	52.4	49.2
4	10.0	47.9	46.3	0.0	0.0	17.8	17.5

As seen in Table 6.1, as well as in the left panels of Figure 6.2, the largest benefit is shown in Scenario 1, where the best dose is the first one and small doses have a high probability of toxicity. Scenario 1 is an example of a dose range which is not well defined, as the smallest dose level gives the highest probability of success and all other doses give toxicity rates too high with a lower chance of success. Such a situation may be rather rare in real applications, though it is not unlikely to happen. Hence, it is important to know how the methodology works in such an extreme scenario. The PK-guided design avoids doses with a high chance of toxicity, while the other design has not prevented this from happening. Despite the toxicity probability increasing sharply with dose levels in the upper dose range, there are still

high doses allocated to patients when only the toxicity condition restricts the design.

The dose range is well chosen in Scenario 2, where the best dose levels are in the middle of the range. Although the rate of success is very similar for two middle dose levels, it drops off rather fast on both sides of the middle dose range. The PK-guided approach in Scenario 2 selects 5.5 as the OD in 39.7% of the trials. It selects dose 6.0 in 40.5% of the trials. This happens as the true probabilities of success at these doses are quite close. These two figures together make 80.2% of what we call in Table 6.1 "best doses". The corresponding figure for the other approach is 66.2%. Again, we observe that the PK-guided design avoids recommending doses with a high probability of toxicity, and, moreover, such doses are used much less in the simulated trials in this case: see Figure 6.3.

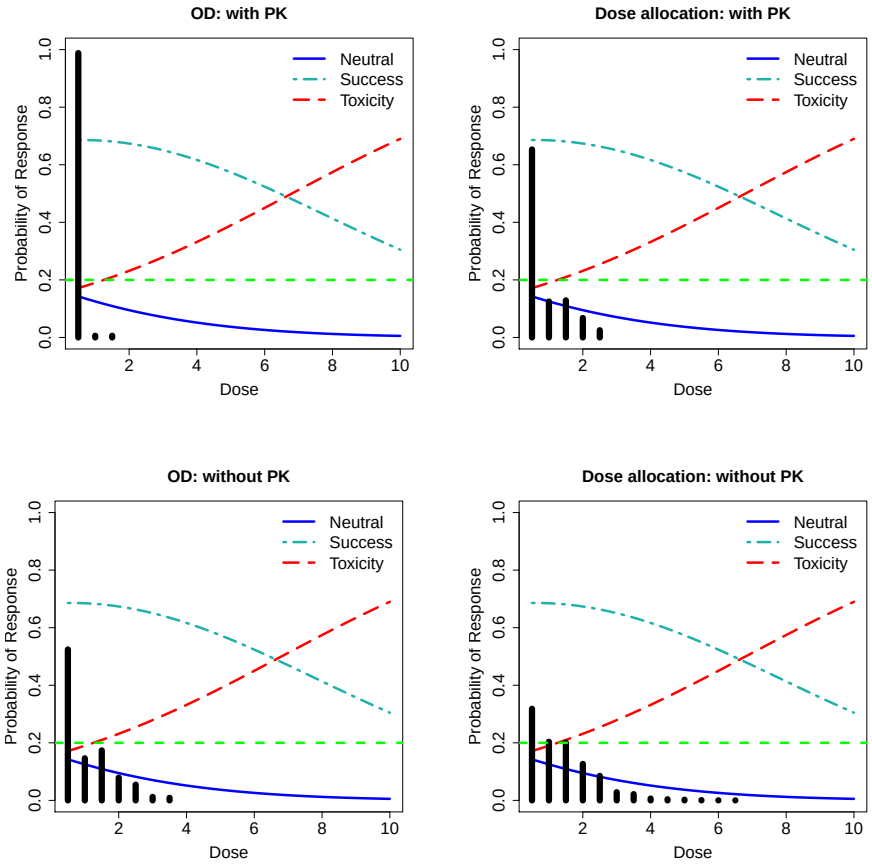

Figure 6.2: Scenario 1 with the OD at 0.5.

The new approach in Scenario 3 identifies exactly 6.5 as the OD in 17.8% of the trials. Because of the flat shape of the success curve, this scenario has a number of doses

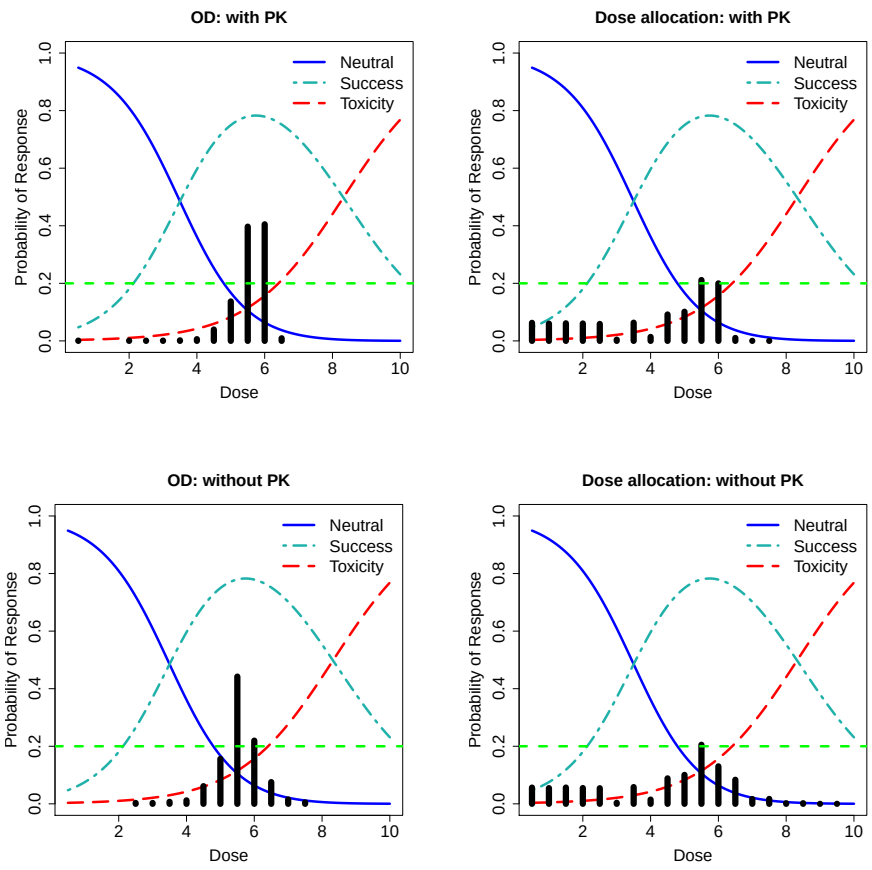

Figure 6.3: Scenario 2 with the OD at 5.5.

with probabilities of success quite close to that for the OD. The doses are 5.5, 6.0, 6.5, 7.0 and 7.5, and these "best doses" are selected in 91.7% of the trials. Although the other approach selects these doses in 85.7% of the trials, it recommends doses above the toxicity probability threshold in 2.5% of cases. Furthermore, from Figure 6.4, the allocation of doses in the trials is again more ethical in the PK-guided design.

There is little difference between the two designs in Scenario 4, as shown in Table 6.1 and also in Figure 6.5. This is the case where both the probability of success and the probability of toxicity increase with dose, where only higher doses have a better chance of having an effect and all doses are below the toxicity threshold. This scenario illustrates a very cautiously chosen dose range. As a consequence, this leads to slow learning in the trial and requires the collection of a lot of information before a recommendation can be made.

This is also shown in Figure 6.8, where, for this scenario, all available cohorts were

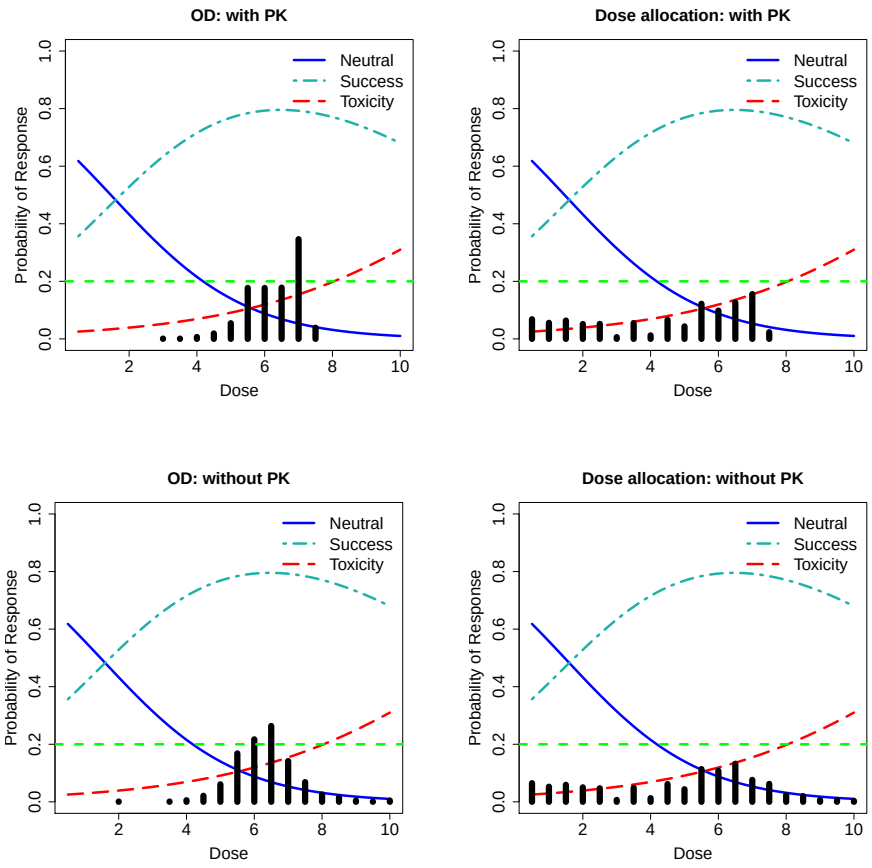

Figure 6.4: Scenario 3 with the OD at 6.5.

used almost all of the time. This is in contrast to Scenario 1, where, especially in the PK-guided design, the learning process was fast and there were much smaller numbers of cohorts required in the trials. The other two scenarios use slightly lower numbers of cohorts than the maximum, since the OD lies in the middle of the dose range.

Now we try to understand how the efficiency of the design and also the optimal sampling time points change in successive stages of a trial. We define the relative D-efficiency of a design $\boldsymbol{\xi}_k^*$ to $\boldsymbol{\xi}_{\text{true}}^*$ as

$$\text{Relative Efficiency} = \left(\frac{|\boldsymbol{M}(\hat{\boldsymbol{\Psi}}_k, \boldsymbol{\xi}_k^*)|}{|\boldsymbol{M}(\boldsymbol{\Psi}_{\text{true}}, \boldsymbol{\xi}_{\text{true}}^*)|} \right)^{\frac{1}{p}}, \tag{6.1}$$

where $\boldsymbol{\xi}_k^*$ is the optimum design obtained at the kth stage of a trial using the current estimates of the parameters $\hat{\boldsymbol{\Psi}}_k$, $\boldsymbol{\xi}_{\text{true}}^*$ is the optimum design obtained for the

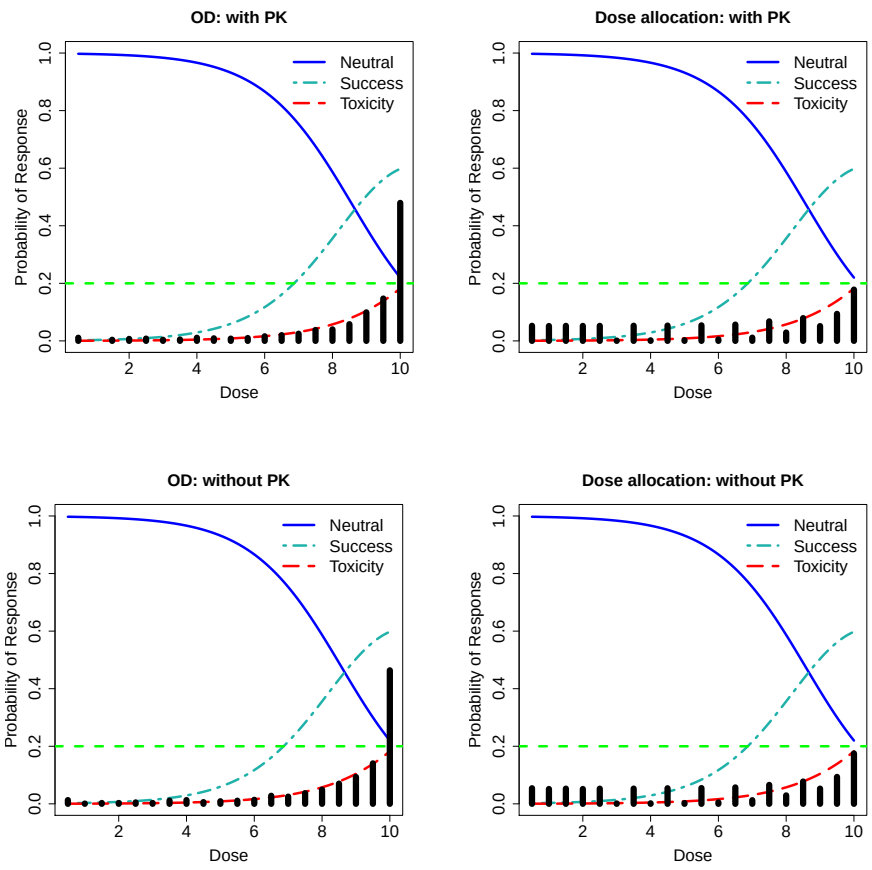

Figure 6.5: Scenario 4 with the OD at 10.0.

true values of the parameters and p is the number of parameters in the model. In both cases, dose remains fixed and it is the dose administered to the cohort at the kth stage. Since both designs depend on the parameter values, it is possible, with completely different parameter values to the true ones, to have a design for which the relative efficiency is very high. A larger determinant of the information matrix means a smaller general variance of the estimators. But, in our case, that will mean that the variance is underestimated. Hence, we want the numerator to be close to the denominator in (6.1). This, in turn, means that we want to have an optimum design which is obtained for values around the true values of the parameters. This can be achieved in a trial as we update the parameter estimates at each stage. After a sufficient number of stages, the estimates will be stable and the design points will be too.

Figure 6.6 shows the relative D-efficiency of the designs computed at each of the stages using (6.1) for Scenario 2. Though not presented, we have found underestimated variances for the initial four stages. Recall that we use the up-and-down

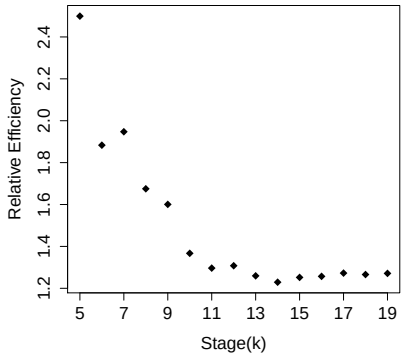

Figure 6.6: Relative D-efficiency in a randomly selected trial from Scenario 2.

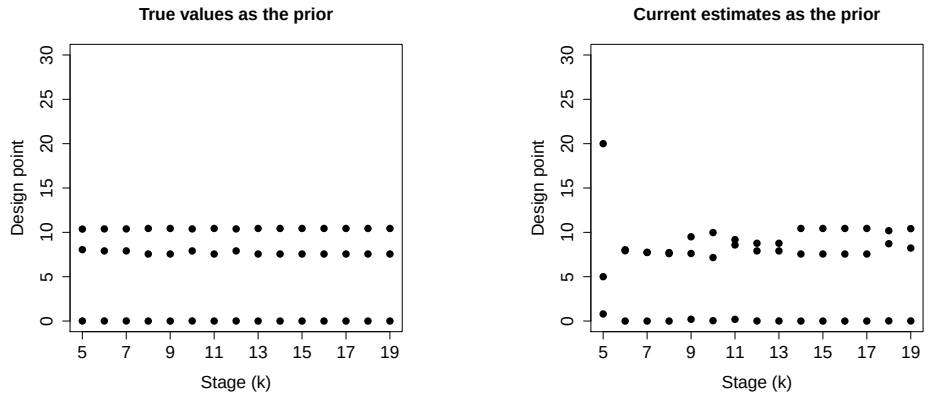

Figure 6.7: Optimal design points in a trial. The left one shows the points for the true values of the PK parameters and the one on the right gives the points using the current estimates.

design for the first four cohorts and that blood samples for these cohorts are collected at the optimal time points which are based on an initial guess about the parameter values. From the fifth cohort onwards, we use the current estimates obtained from the trial data. Here, we observe a decreasing trend in the efficiency. Since the estimates stabilise as the trial proceeds, there is not much change at the later stages and also the efficiency approaches to one.

The design points are displayed in Figure 6.7. With the true values of the parameters, the design points are very similar at the various doses received by the cohorts in the trial. This is because the shape of the concentration profile does not change much with dose. But, of course, the amount of concentration changes rapidly with dose. There is good agreement between the time points at the later stages with the

corresponding true ones. This suggests that, after a reasonable number of steps, the optimal time points stabilise.

It has already been mentioned that we stop a trial early if the same dose is repeated for six of the cohorts and that we call the associated dose the OD. We have found that, as the location of the OD moves from left to right in the dose region of a scenario, more cohorts are needed to stop early. Most of the early stopped trials identify the OD accurately. Figure 6.8 compares the average numbers of cohorts used by the two approaches. It is observed that the PK-guided approach utilises fewer cohorts in each scenario.

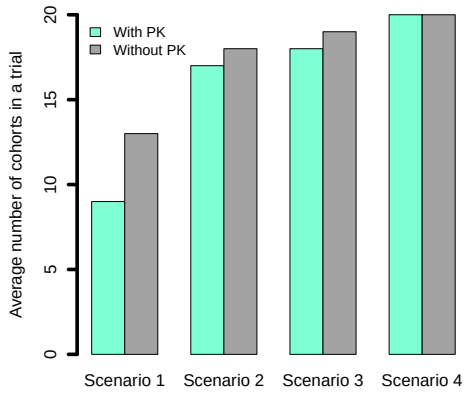

Figure 6.8: Average numbers of cohorts used in the four scenarios by the two dose-allocation methods.

Figures 6.9 and A.3 indicate a small bias and variance for the PK parameter estimates for all of the scenarios. Since the design employs the D-criterion for measuring the PK responses, it is ensuring accuracy and efficiency in parameter estimation. Figures 6.10, A.4 and A.5 show that the dose-response parameter estimates obtained from the two approaches are similar. Obviously, they are not as good as the PK estimates. This is due to the fact that the information on the trinomial dose-response model is not gathered in a way that would be optimum for parameter estimation. Here, we focussed on the criterion which would provide a good dose for further studies in an ethical trial, which is particularly important in classes of drugs where toxicity can be very serious.

The decision and sampling efficiencies of the designs, obtained using (5.18) and (5.20), are presented in Table 6.2. Both measures are larger for the PK-guided design than for the other design. These again reflect the fact that the former design is capable of identifying the OD more accurately and also allocating the most efficacious doses to the cohorts in a trial.

We have checked the sensitivity of the optimal design to the prior values of the parameters in Section 3.7.1. It has been found that the design is not sensitive to the values. We now check the sensitivity of our dose-finding design to the target value of the AUC.

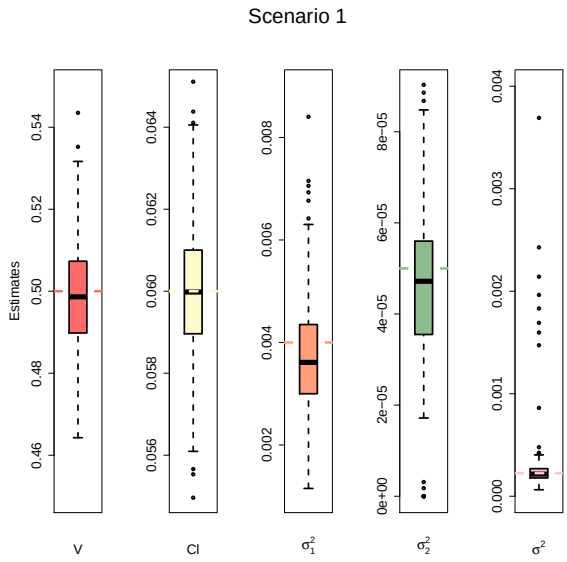

Figure 6.9: Boxplots of the PK parameter estimates obtained from the simulations for Scenario 1. The horizontal dashed lines indicate the true parameter values.

Table 6.2: Decision and sampling efficiencies of the designs.

Scenario	DE		SE	
	PK	No PK	PK	No PK
1	0.994	0.670	0.778	0.522
2	0.975	0.882	0.694	0.608
3	0.991	0.964	0.859	0.830
4	0.806	0.796	0.453	0.449

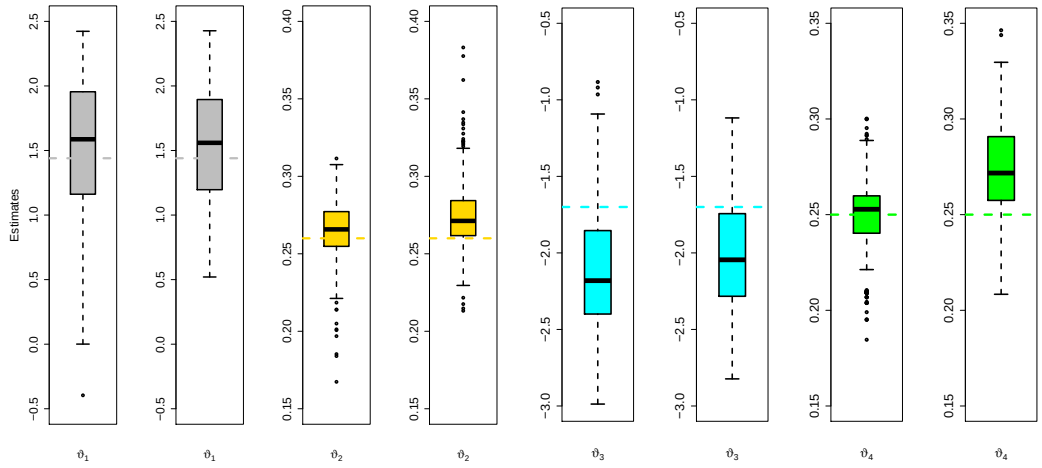

Figure 6.10: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenario 1. The horizontal dashed lines indicate the true parameter values. For each parameter, the left boxplot corresponds to the design which takes into account the AUC and the right boxplot to the one which ignores it.

Sensitivity to Target AUC

The PK-guided design depends on the target value of the AUC. In our simulations so far, we have considered it as the one at the true OD. To assess the sensitivity of the design to the target value, we set it at doses other than the true OD. Scenario 2 is studied for this purpose. This scenario has 5.5 and 6.0 as the best doses. Table 6.3 gives a summary of the results. The notation in the table is defined as follows: percentage of best doses recommended for further studies (%BD), percentage of doses recommended as optimum, but carrying the probability of toxicity above the maximum allowed threshold (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD). Also, % of 3-BD is the percentage of the three best doses. DE and SE are the decision and sampling efficiencies, respectively.

Table 6.3: Sensitivity of the design to the assumed target for AUC in Scenario 2.

Dose	%BD	% of 3-BD	%TD	%AD	DE	SE
4.5	0.5	88.0	0.0	2.0	0.948	0.647
5.0	72.7	96.0	0.1	33.6	0.983	0.696
5.5	80.2	94.0	0.9	41.1	0.973	0.694
6	72.0	93.1	3.5	34.2	0.950	0.697
6.5	62.9	84.8	8.5	34.0	0.892	0.633
No PK	66.2	81.7	9.5	33.5	0.882	0.608

By the three best doses, we mean the doses 5.0, 5.5 and 6.0, since the probabilities of success at these are close to each other. The figures in the table indicate that the design is sensitive to the choice of target value for the AUC. %BD is smaller as the target is further away from that at the true OD. If we choose a target below that at the true OD, the design will avoid recommending a toxic dose as the OD and will also not allocate toxic doses to the cohorts. But that will have a negative impact on the correct identification of the optimum dose. The distributions of optimum dose selection and dose allocation are presented in Figures 6.11 and 6.12.

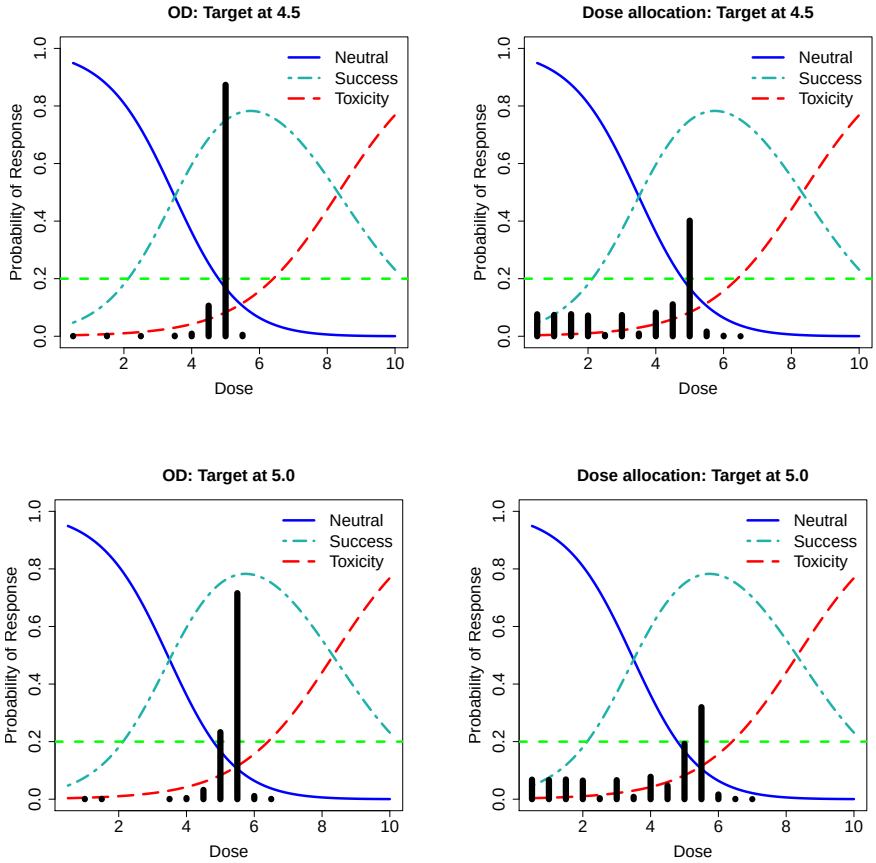

Figure 6.11: Optimum dose selection and dose allocation when the target AUC is taken at the doses below the true optimum dose.

Sensitivity to Dose-Skipping Constraint

Section 5.3 discusses a restriction of not skipping more than one dose level at a time when the level is increased. This is to make the design more cautious to avoid any unacceptable toxicity from a high dose following a relatively small dose. In both

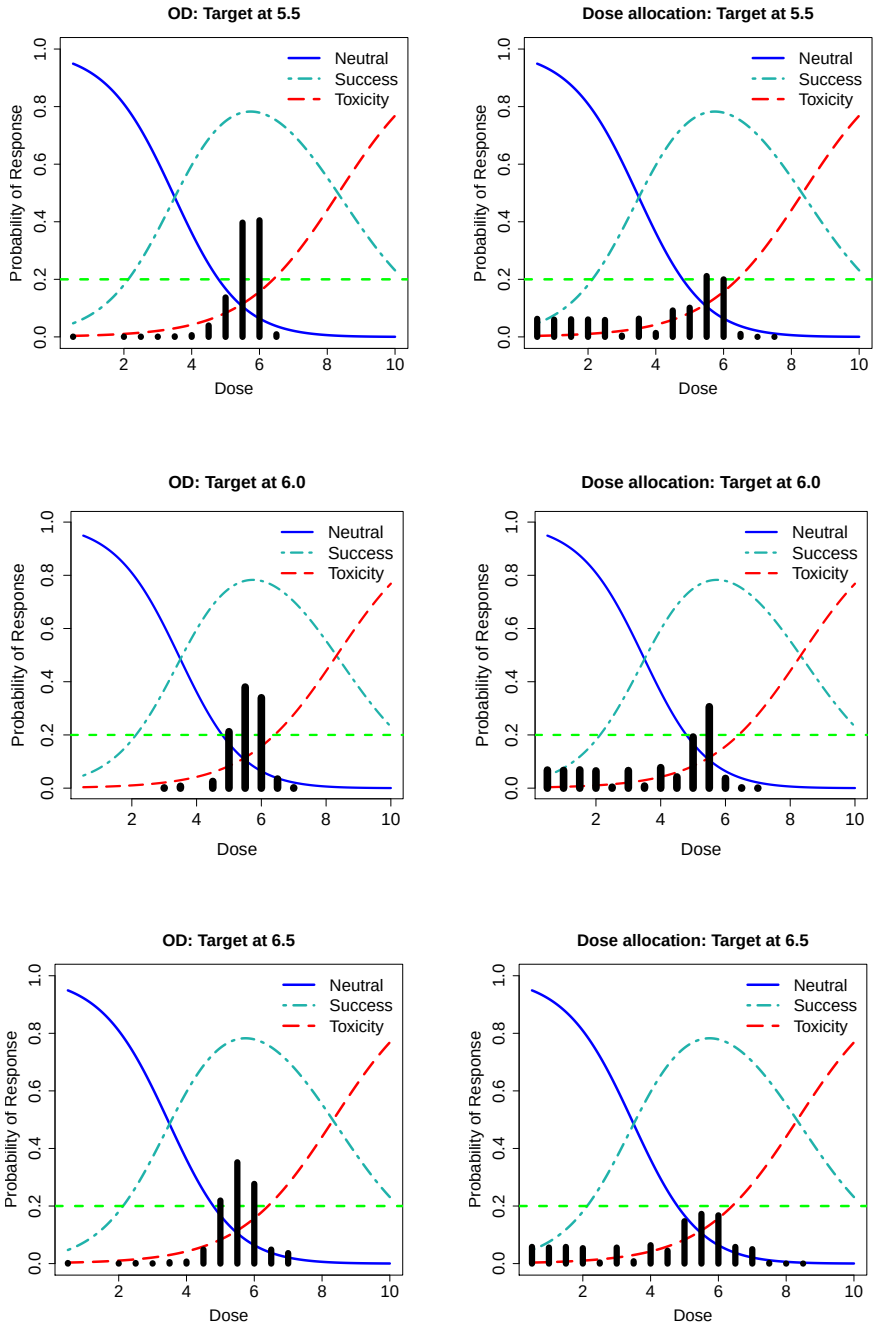

Figure 6.12: Optimum dose selection and dose allocation when the target AUC is taken at the true optimum dose and also at the doses above it.

PK-guided and other designs, we have employed this skipping constraint. To see the impact, we compare them with those that do not employ any such constraint. Again, we study Scenario 2. Tables 6.4 and 6.5 summarise the results from 1,000 simulated trials. It is found that there is no appreciable difference between the figures under the two different situations, except for dose allocation. That might be due to the reason that, when there is no constraint to avoid dose skipping, the cohorts after the up-and-down phase receive doses which are close to the most efficacious ones.

Therefore, patients are likely to be treated more often at the optimum dose at an earlier stage of the trial than in the case of a slow-dose increase.

Figure 6.13 shows the dose allocation to successive cohorts for four randomly chosen trials for both cases. As we can see, there may be situations when there is a sharp increase in the dose level to the upper end of the dose range, which has a very high

Table 6.4: Sensitivity of the PK-guided design to the dose-skipping constraint.

Constraint	%BD	%TD	%AD	DE	SE
Yes	80.2	0.9	41.1	0.975	0.694
No	79.0	0.7	47.3	0.979	0.692

Table 6.5: Sensitivity of the design to the dose-skipping constraint in the absence of PK information.

Constraint	%BD	%TD	%AD	DE	SE
Yes	66.2	9.5	33.5	0.882	0.608
No	67.3	9.5	40.0	0.886	0.582

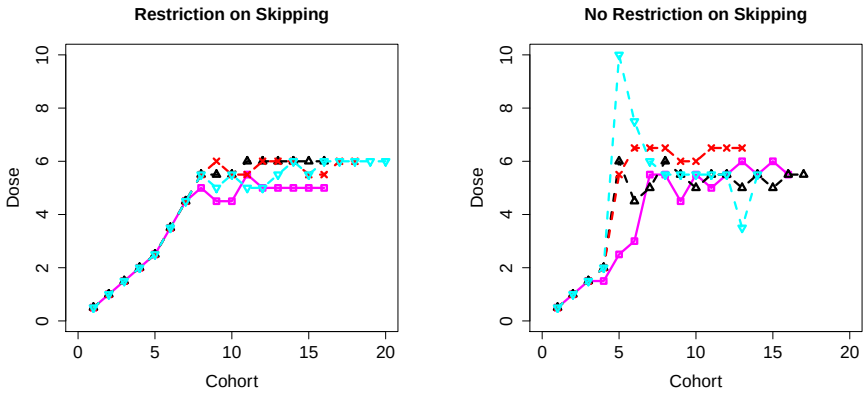

Figure 6.13: Dose allocation to successive cohorts in four randomly chosen trials for the PK-guided design for Scenario 2.

probability of toxicity. This kind of event is avoided by the dose-skipping restriction. In some cases, where toxicity is of less concern, a clinician may choose to skip more than one dose level. This is one of the input parameters in the computer program.

To summarise the first example, we have seen that the efficiency of a dose-finding

design can be improved by the use of pharmacokinetic information. Such a design can restrict the recommendation of toxic doses for further studies and also allow allocation of more efficacious doses to the cohorts. We have found the design to be sensitive to the choice of the target AUC and therefore care is needed in setting it.

6.4 Example 2

In this example, we assume that the drug is administered orally and that the whole body is a single compartment. Assume that the dose-response outcomes are bivariate binary. The Cox model for bivariate binary responses is used to model the dose-response outcomes. This is accompanied by the one-compartment PK model with first-order absorption. The details of these models are given in Sections 4.3 and 3.7.2, respectively. Here, we present two adaptive designs: one considers PK data and the other does not. The dose-optimisation criterion in both cases is the same, which is the maximisation of the estimated probability of success, discussed in Section 5.3.1. The toxicity constraint in Section 5.4 is also common to both designs. However, the PK-guided design considers an additional constraint, introduced in Section 5.5.2, that is, the constraint on the maximum concentration. The aim is to study the operating characteristics of these designs and also to compare them.

6.4.1 Simulation Settings

Choice of Design Parameters

Assume that our experimental drug has 11 available doses on the log scale between -3 and 3 with a grid width of 0.6. Therefore, the set of doses is $\mathcal{X} = \{-3.0, -2.4, \ldots, 3.0\}$. We investigate four plausible dose-response scenarios, as shown in Figure 4.2, and taken from Dragalin and Fedorov (2006). We assume that the acceptable level for the probability of toxicity is $\gamma = 0.33$. The assumed optimal dose that maximises the probability of success in the first three scenarios is -0.6 and in Scenario 4, is -1.8. These scenarios have doses 0, -0.6, -0.6 and -1.2 for which the probability of toxicity is no larger than γ. The PK parameter values

assumed for the simulation study are $V = 2.0$, $k_e = 0.05$, $k_a = 0.50$, $\sigma_1^2 = 0.16$, $\sigma_2^2 = 0.0001$, $\sigma_3^2 = 0.01$ and $\sigma^2 = 0.00005$. The parameters are chosen in such a way that the coefficient of variation is 20%.

Starting with the lowest dose -3.0, the up-and-down design in Section 5.2 continues for the first four cohorts to gather information prior to any parameter estimation. The target value of $C_{\max}$ is taken to be the $C_{\max}$ at the true OD in the scenario, obtained from (3.35) for the true mean PK parameters.

Generation of PK and Dose-Response Outcomes

Following the administration of a dose to a cohort of patients, we observe the concentration of the drug in the blood samples and the dose-response outcomes. The design that considers PK data also requires the population D-optimal time points to collect blood samples to measure the concentrations. The purpose is to ensure efficiency in the estimation of model parameters, even if we collect only a few samples from each patient. As before, the time points are obtained using the R package *PFIM 3.2* (Bazzoli et al., 2010). The PK sampling region is assumed to be $\mathcal{T} = [0, 50]$ hours.

To decide on the optimal number of sampling points, we have checked the eligibility of either 3, 4, 5 or 6 points. A 2-point design is not possible, as the model has more parameters, and therefore has been avoided as a candidate. As shown in Figure A.6, the efficiency of a 4-point design relative to a 3-point one is much higher. If we increase the number of design points, the efficiency at each point relative to the previous point decreases. The PK model has three parameter sensitivities for its parameters V, k_e and k_a. We have checked that two of the sensitivities have an extremum at the beginning, while the other has an minimum further to the right of the design region. Some graphical investigation of the concentration profiles with random parameters, chosen within three standard deviations of the means, also showed the variability to be high in those areas. Therefore, consideration of many design points will lead to the choice of design points from those neighbourhoods. Since

patients are usually unwilling to give many samples and also not much is gained by considering many design points, we decide to collect blood samples at the four optimal time points for the individuals in each cohort of size $c = 3$. More specifically, $n_i = 4$ for all i in (3.21).

To obtain the concentrations at these time points, we generate a vector of random effects $\boldsymbol{b}_i$ for individual i from the normal distribution $N_3(\mathbf{0}, \boldsymbol{\Omega})$, where $\boldsymbol{\Omega} = \text{diag}(0.16, 0.0001, 0.01)$. The random PK parameters are then obtained as $\boldsymbol{\theta}_i = \boldsymbol{\beta}+\boldsymbol{b}_i$, where $\boldsymbol{\beta} = (2.0, 0.05, 0.50)^T$. We generate the random errors from $N_4(\mathbf{0}, \sigma^2\boldsymbol{I}_4)$, where $\sigma^2 = 0.00005$. These are added to the generated concentrations to produce the simulated PK responses for an individual. The same principle is adopted to generate responses for the remaining individuals in a cohort. Such generated concentrations for a cohort are presented in Figure 6.14.

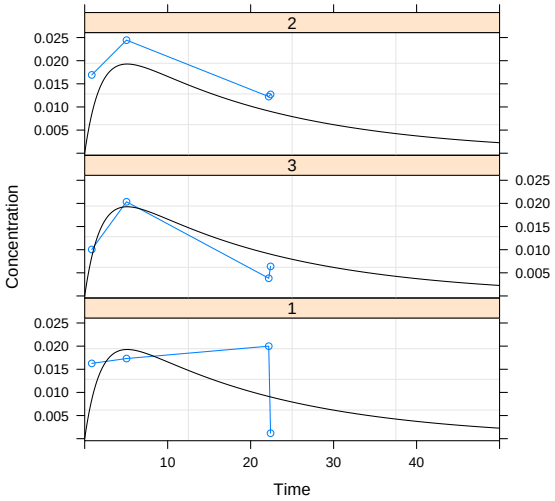

Figure 6.14: Simulated concentrations at the locally D-optimum time points following the administration of the lowest dose to a cohort. The true mean PK profile is indicated by the black curve.

The dose-response scenarios to be used in the simulation study are presented in Figure 4.2. We have the true parameter value $\boldsymbol{\vartheta}$ for each of these scenarios. The true probabilities of the four possible dose-response outcomes are then available corresponding to the doses in $\mathcal{X}$. Thus, the dose-response outcomes are generated

from a multinomial distribution with four categories having these probabilities.

Priors

To obtain the optimal time points for the first four cohorts in the up-and-down stage of a trial, we chose the values $\boldsymbol{\Psi}^0 = (V^0, k_e^0, k_a^0, (\sigma_1^2)^0, (\sigma_2^2)^0, (\sigma_3^2)^0, (\sigma^2)^0)^T = (3.20, 0.08, 0.80, 0.04, 0.000025, 0.0025, 0.0000125)^T$ for the PK parameters. The priors are chosen such that the mean PK parameters are three standard deviations above the true values, and the variance components and error variance are one quarter of the true values. The motivation behind choosing such priors is to make them as vague as possible, since in reality these would normally be unknown and could be wrongly assumed at the beginning of the trial. For any cohort after the up-and-down stage, the current maximum likelihood estimates of the parameters are used.

A joint uniform prior distribution for $\boldsymbol{\vartheta}$, given in (4.12), is used for Bayesian estimation of the dose-response parameters. The parameter space $\tilde{\Theta}$ is chosen for each scenario so that the margin is 3 on either side of the true values. As Scenario 1 has the true parameters $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4, \vartheta_5, \vartheta_6)^T = (0.0, 1.0, 4.0, 2.0, 3.0, 3.0)^T$, $\tilde{\Theta}$ takes the values $-3 < \vartheta_1 < 3$, $-2 < \vartheta_2 < 4$, $1 < \vartheta_3 < 7$, $-1 < \vartheta_4 < 5$, $0 < \vartheta_5 < 6$ and $0 < \vartheta_6 < 6$. The other scenarios follow the same approach. As we utilise the uniform prior distribution, the distributions cancel out in evaluating the integrals in (4.11).

Model Fitting

Upon receiving the PK and dose-response data, we update the estimates of the model parameters. A similar approach to that described for Example 1 is used for fitting the models.

Dose Selection for the Next Cohort

Dose selection for the first four cohorts is based on the up-and-down design in Section 5.2. After the up-and-down stage, we estimate the PK and dose-response parameters

based on the data to hand and proceed to select the dose for the next cohort. The dose-optimisation criterion to be used is the maximisation of the estimated probability of success. Assume that we are at the kth stage of the trial, and, based on the current data, we have the estimates $\hat{\boldsymbol{\Psi}}_k$ and $\hat{\boldsymbol{\vartheta}}_k$ of the PK and dose-response parameters, respectively. Then we select the dose x_{k+1} for the next cohort of patients based on the dose-optimisation criterion in Section 5.3.1, subject to two different sets of constraints in two separate runs of the simulations, that is, subject to (5.10), and to (5.10) and (5.12).

Checking the Stopping Rules and the OD Selection

The same approach as was applied in the previous example is used for stopping a trial and in identifying the optimum dose. Like before, we have $r = 6$ and $m = 20$ in this example. Each scenario is investigated through 1,000 simulated trials.

6.4.2 Numerical Results

The performance of the PK-guided design, which constrains $C_{\max}$, is compared with the one which does not use such a constraint. Other than this constraint, the designs are the same. Tables 6.6-6.7 and Figures 6.15-6.18 summarise the simulation results for the four scenarios. As before, in these figures, the bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at these doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

Table 6.6 shows the results for the designs in terms of three performance measures: percentage of best doses recommended for further studies (%BD), percentage of doses recommended as optimum, but carrying the probability of toxicity above the maximum allowed threshold (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD). The higher the values of %BD and %AD, the better the design is. Similarly, we would expect %TD to be as small as possible.

Except for Scenario 4, the advantage of the PK-guided design is clearly illustrated by the figures in the table. However, in Scenario 4, there is a slight improvement in avoiding the doses with a high chance of toxicity.

Table 6.6: Percentage of best doses recommended for further studies (%BD), percentage of doses recommended as optimum, but carrying the probability of toxicity above the maximum allowed threshold (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD).

Scenario	Best Doses	%BD		%TD		%AD	
		PK	No PK	PK	No PK	PK	No PK
1	-0.6	62.9	54.0	2.0	5.6	37.5	32.6
2	-0.6	63.1	46.1	1.5	21.7	35.6	26.9
3	-0.6	94.8	83.1	0.0	11.5	52.4	43.9
4	-1.8 and -1.2	80.7	88.2	0.1	1.1	62.2	66.6

Figure 6.15 shows the distributions of optimum dose selection and dose allocation for Scenario 1. The left panel shows the OD selection and the right panel shows the dose allocation for the two different adaptive designs. The figures are obtained by following Sections 5.7.1 and 5.7.2. The scenario has -0.6 as the true OD. The PK-guided approach selects this as the OD in 62.9% of the trials compared to 54.0% by the other approach. The designs also select -1.2 as the OD in an appreciable number of trials. This is due to the fact that the probabilities of success at these doses are quite close. These two doses are selected in 90.2% of the trials, while the figure is 82.3% for the other approach. The PK-guided approach lessens the selection of highly toxic doses as the OD. It also allocates more patients to the most efficacious doses.

The PK approach in Scenario 2 identifies -0.6 as the OD in 63.1% of the trials. Since its probability of success is close, it selects -1.2 in 26.7% of the trials. These two doses are recommended in 89.8% of the trials. The design is very careful in recommending toxic doses as the OD. It also treats a good proportion of cohorts with the best dose throughout the trials. The corresponding figure for the doses -1.2 and -0.6 in the other approach is 68.1%.

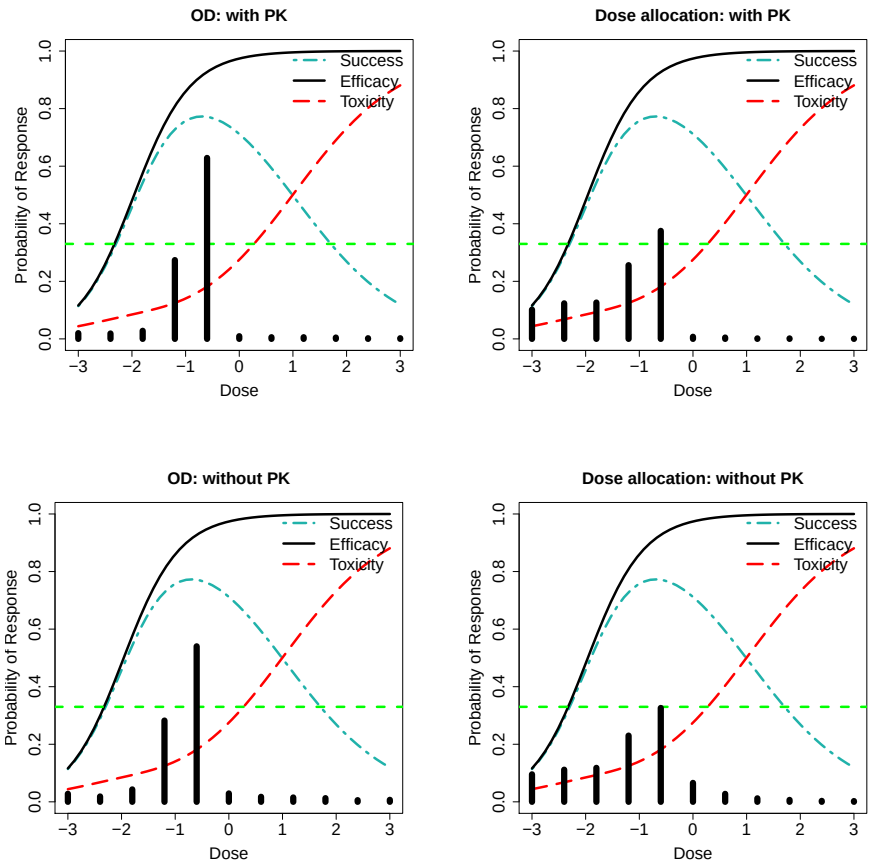

Figure 6.15: Scenario 1 with the OD at -0.6.

The dose -0.6 has been chosen in 94.8% of the trials in Scenario 3 by the PK-guided approach. The other approach selects it in 83.1% of the trials. More than half of the cohorts have been treated at the best dose in the simulated trials. The PK approach avoids selecting toxic doses as the OD in a good percentage of trials.

In Scenario 4, the true OD is -1.8. But the doses -1.2 and -1.8 have probabilities of success which are quite close. These two doses are selected in 80.7% of the trials. The other approach selects these in 88.2% of the trials. This is the scenario where the PK approach is not performing more satisfactorily than the other approach. Since we put the PK constraint on the true OD, the design fails to select -1.2, which is next to -1.8.

As we move through the scenarios, the steepness of the toxicity curve increases. The differences in the results for different scenarios can be related to this. We see that gradually more patients are saved from toxic doses by the PK-guided approach.

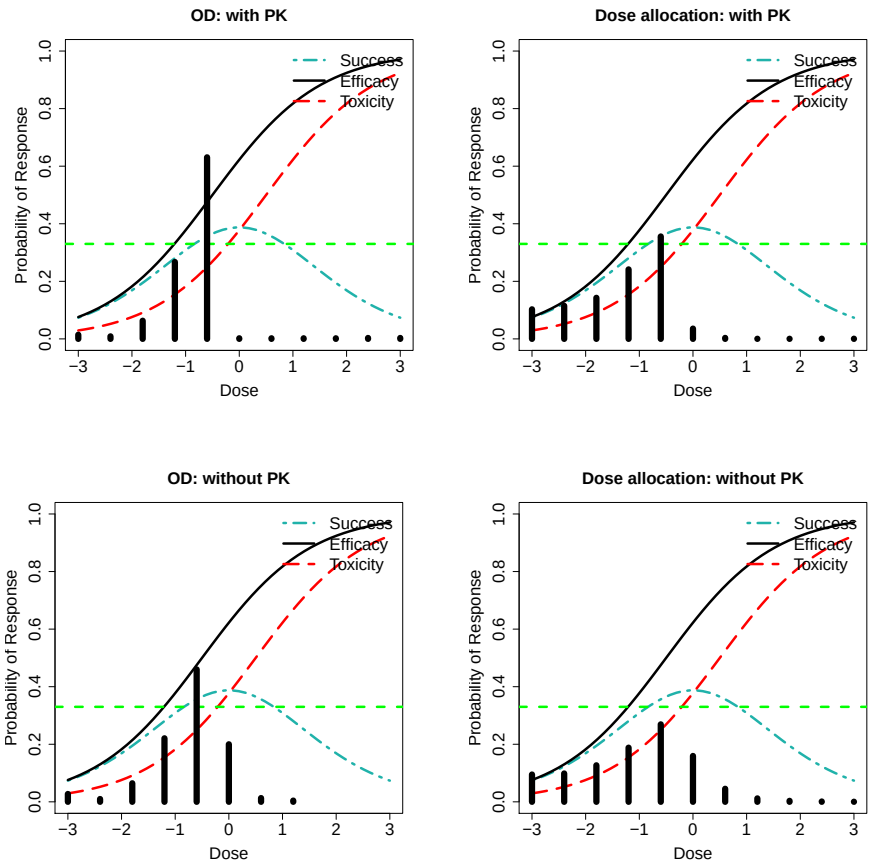

Figure 6.16: Scenario 2 with the OD at -0.6.

Figure 6.19 gives the relative efficiencies at successive stages of a trial in Scenario 1 using (6.1). Although there are fluctuations at the beginning, the efficiency measure stabilises at the later stages. This, in turn, means that the parameter estimates are becoming close to the true values towards the end of the trial. Also, the associated design points stabilise as the trial proceeds: see Figure 6.20.

Figure 6.21 shows the average numbers of cohorts used in a trial by the dose-finding approaches. The difference in the numbers of cohorts is not significant, since the figures are quite similar for the two approaches. The most noticeable feature is that, although we allow a trial to use a maximum number of 20 cohorts, the average in all scenarios is far less. This indicates that most of the trials stop before reaching the maximum sample size.

Table 6.7 presents the decision and sampling efficiencies for the scenarios, obtained

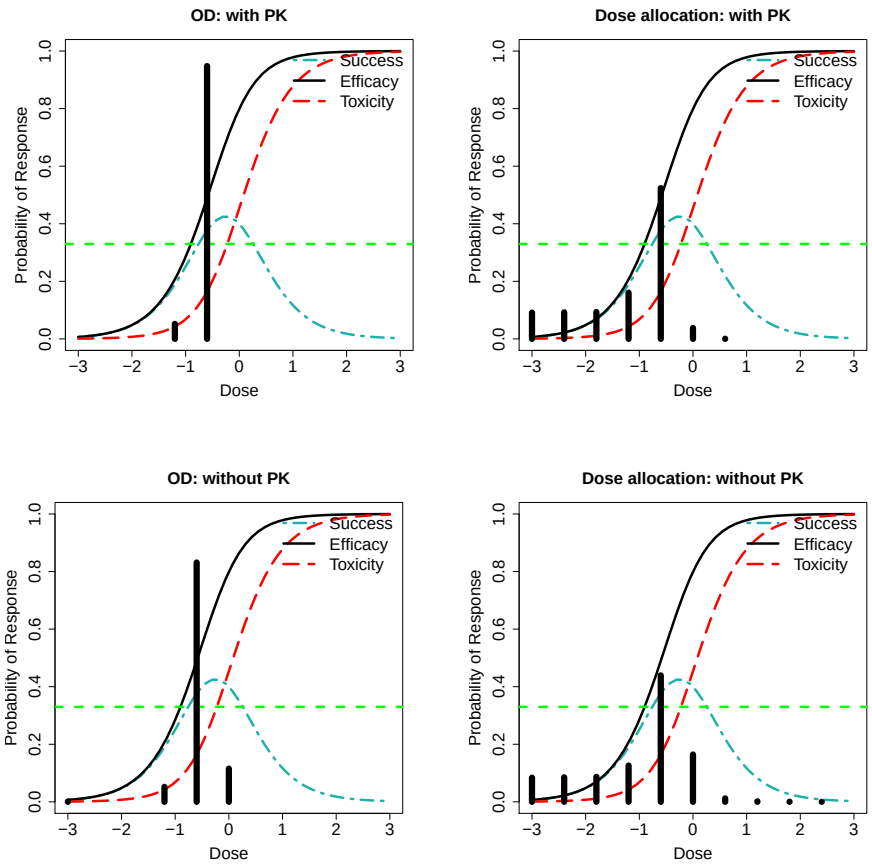

Figure 6.17: Scenario 3 with the OD at -0.6.

using (5.18) and (5.20). The decision efficiencies are usually higher in the PK-guided design than in the design which does not constrain $C_{\max}$. The largest difference between the DEs is seen in Scenario 2. Since the PK-guided design does not recommend toxic doses as the OD many times compared to the other design, it gives the maximum gain in decision efficiency. The DEs for the two designs are close in Scenario 4. This is because the optimum dose selection is more accurate in the other approach, and also the difference in %TD between the two designs is small.

The sampling efficiency for the PK-guided design is also higher than that of the other design for all of the scenarios. However, in Scenario 4, they are quite close. This is the case since the distribution of dose allocation in the two designs is very similar: see Figure 6.18.

Figure 6.22 summarises the distribution of the PK parameter estimates obtained from simulations for Scenario 1. We see some outliers in the estimates for each of

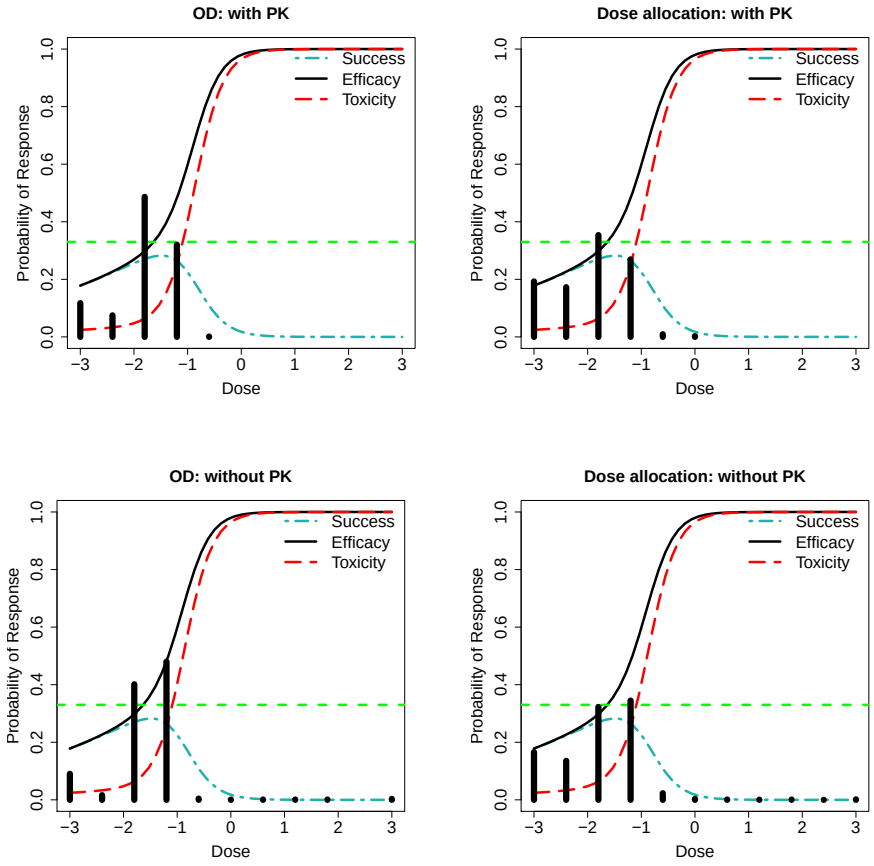

Figure 6.18: Scenario 4 with the OD at -1.8.

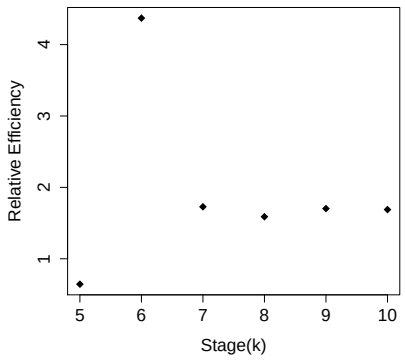

Figure 6.19: Relative efficiencies in a randomly selected trial from Scenario 1.

the parameters. Apart from these, the bias and variance of the estimates are small. The distributions for the other scenarios are shown in Figures A.7-A.8. We find a similar behaviour of the estimates as in the first scenario. Since the approach utilises D-optimal time points for collecting concentrations, it ensures accuracy and efficiency in parameter estimation. Figure 6.23 displays the distribution of the dose-response parameter estimates for Scenario 1. We have considered wide uniform

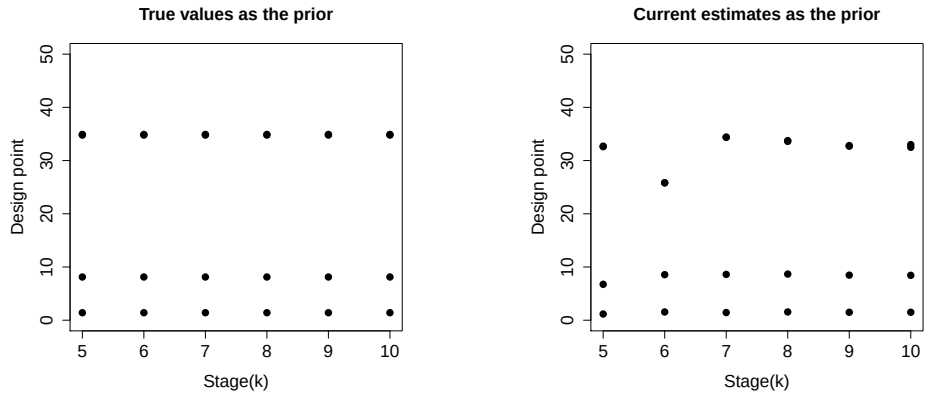

Figure 6.20: Optimal design points in a trial. The left one shows the points for the true values of the PK parameters and the one on the right gives the points using the current estimates.

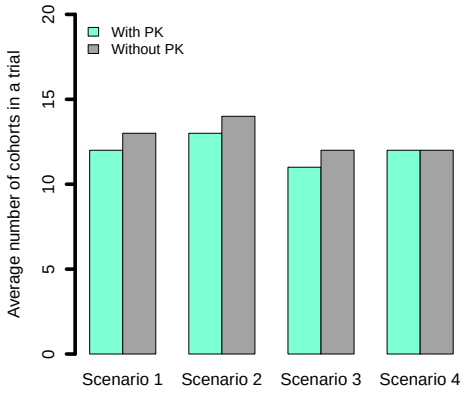

Figure 6.21: Average numbers of cohorts used by the two dose-allocation methods.

priors for the scenarios to allow extreme cases that might possibly occur.

Table 6.7: Decision and sampling efficiencies of the designs.

Scenario	DE		SE	
	PK	No PK	PK	No PK
1	0.925	0.875	0.774	0.743
2	0.882	0.679	0.685	0.540
3	0.974	0.858	0.629	0.525
4	0.939	0.944	0.887	0.882

As a consequence, we have found outliers in the estimates of the parameters. A similar pattern is seen for the other scenarios presented in Figures A.9-A.10. These estimates are not as good as the PK estimates. The reason is that we have allocated doses to the cohorts so that the estimated probability of success is maximum, rather

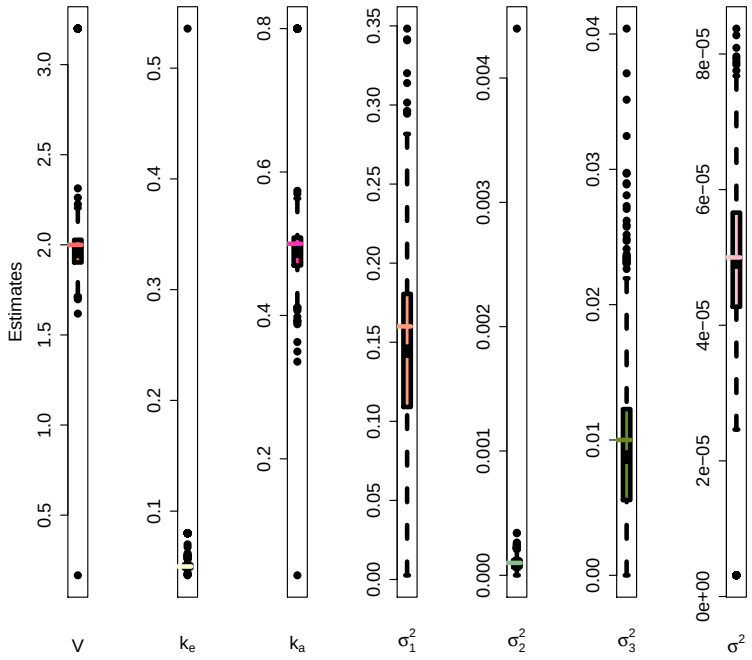

Figure 6.22: Boxplots of the PK parameter estimates obtained from the simulations for Scenario 1. The horizontal dashed lines indicate the true parameter values.

than allocating doses for which the efficiency in parameter estimation is maintained. Moreover, the dose-response parameter estimates obtained from the PK-guided approach are similar to those from the other approach. These indicate that consideration of PK data is not improving the estimates of the dose-response parameters.

The percentages of the trials that stopped in the up-and-down stage are 15, 10, 3 and 8, respectively, for the four scenarios. The figures are almost identical for the two approaches. Scenario 3 has the lowest percentage, since the probabilities of toxicity at the early doses are very low compared with the other scenarios.

We have learnt from this example that a dose-finding design can be made more efficient if we can allocate doses to the cohorts entering a trial around a target value for $C_{\max}$. Of course, the gain in efficiency depends on the underlying scenario. The gain is maximum in a scenario with a steep toxicity curve. Along with dose-response data, both examples have utilised pharmacokinetic information. They are similar in both cases apart from the different PK constraint. The intention is to keep the

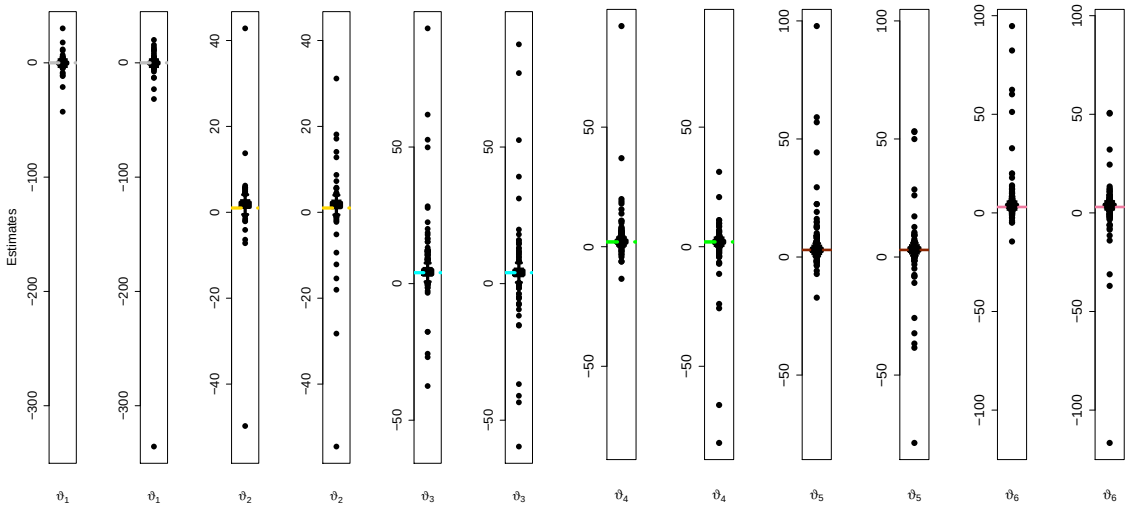

Figure 6.23: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenario 1. The horizontal dashed lines indicate the true parameter values. For each parameter, the left boxplot corresponds to the design which takes into account $C_{\max}$ and the right boxplot to the one which ignores it.

method a flexible one, so that one can constrain either of the PK parameters AUC and $C_{\max}$, depending on the availability of the information and the interest of the researchers.

6.5 Sensitivity Analysis

In this section, we present the sensitivity analysis of the proposed design to the different parameters via simulation studies. We restrict ourselves to Scenario 1 in Example 2, since checking all of the scenarios involves considerable computational time. However, we would expect similar results for the other scenarios. There are a number of parameters in the design that one can vary. However, we are particularly interested in the most crucial ones. These include the priors for the PK and dose-response parameters, and also the target value for the maximum concentration. Each case is investigated through 1,000 simulated trials. All other parameters are kept fixed during the simulations to see the effect of the parameter being varied.

6.5.1 Priors for Dose-Response Parameters

The prior distribution plays an important role in Bayesian estimation of parameters. Here, we express the uncertainty about the parameters in terms of a probability distribution and combine this with the new data to form the posterior distribution. Then we find the posterior means of the parameters, which we call the Bayesian estimates. It is always a good idea to assess the sensitivity of the estimates to the assumed prior values. To ensure the prior has less effect and the data dominate, we use non-informative priors in both examples presented in Sections 6.3 and 6.4. More specifically, we use the uniform distribution over some sensible range of values as the prior. The ranges are chosen in such a way that they encompass any possible extreme case.

In the simulation studies in Section 6.4, we assumed priors with a margin of 3 on either side of the true values. These priors are wide enough. To see what happens in the presence of narrower priors, we plan to employ priors which have margins of 2 and 1. The true parameter values in Scenario 1 are $\boldsymbol{\vartheta} = (0.0, 1.0, 4.0, 2.0, 3.0, 3.0)^T$. The margin 3 allows the parameter space $\tilde{\Theta}$ to have the values $-3 < \vartheta_1 < 3$, $-2 < \vartheta_2 < 4$, $1 < \vartheta_3 < 7$, $-1 < \vartheta_4 < 5$, $0 < \vartheta_5 < 6$ and $0 < \vartheta_6 < 6$. Similarly, for margin 2, the parameter space consists of the values $-2 < \vartheta_1 < 2$, $-1 < \vartheta_2 < 3$, $2 < \vartheta_3 < 6$, $0 < \vartheta_4 < 4$, $1 < \vartheta_5 < 5$ and $1 < \vartheta_6 < 5$. The margin 1 gives the values $-1 < \vartheta_1 < 1$, $0 < \vartheta_2 < 2$, $3 < \vartheta_3 < 5$, $1 < \vartheta_4 < 3$, $2 < \vartheta_5 < 4$ and $2 < \vartheta_6 < 4$. The possible scenarios that might occur at the boundaries of these priors are shown in Figure 6.24. The graphs in the right panel are fairly similar and also similar to the original scenario. We have found that the graphs keep a similar shape for any margin above the true values of the parameters. Although the graphs in the left panel are different in shape, if we increase the margin further, they converge to a shape where the marginal probability of toxicity decreases as dose increases. This contradicts the usual form of toxicity curves. Also, in all of the current left-panel graphs, toxicity decreases initially and then increases with dose. Moreover, since there are six parameters, any change in one of the parameters will lead to a change

in the shape of the curves. With all of these considerations, the set of priors that we are going to use can be regarded as vague enough.

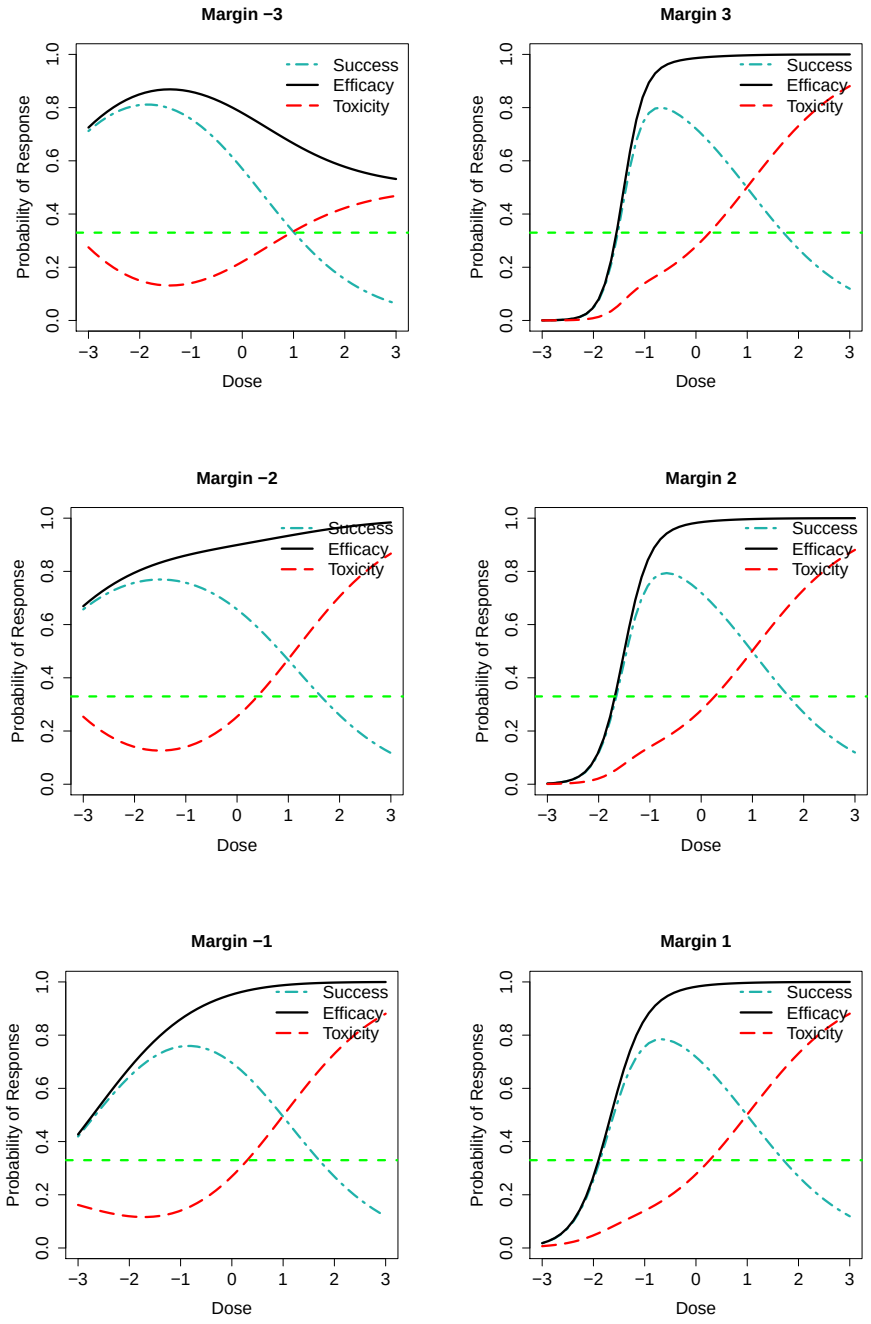

Figure 6.24: Possible dose-response curves for margins of 3, 2 and 1 on either side of the true parameter values. The left panel represents the cases where the parameter values are below the true values and those in the right panel represent the cases where the parameter values are above the true values.

Figure 6.25 shows the distributions of optimum dose selection and dose allocation for the PK-guided design for various prior distributions. Table 6.8 summarises the results. All of the performance measures improve as the margin decreases.

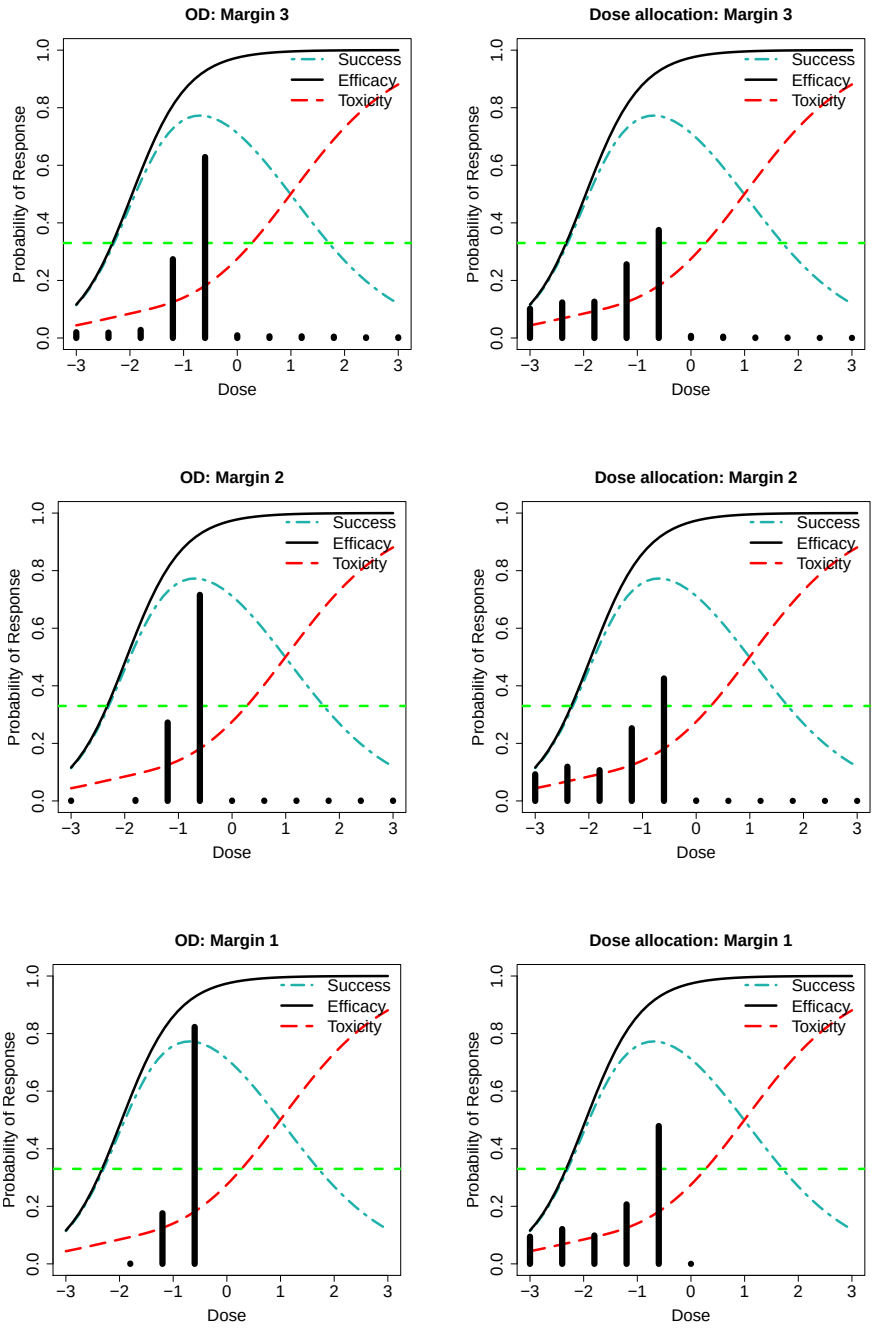

Figure 6.25: Optimum dose selection and dose allocation under different prior values for Scenario 1 in Example 2.

Table 6.8: Sensitivity of the design to the assumed priors for the dose-response parameters in Scenario 1 that takes into account PK information.

Margin	%BD	%TD	%AD	DE	SE
3	62.9	2.0	37.5	0.925	0.774
2	71.6	0.6	42.5	0.977	0.799
1	82.3	0.0	47.9	0.989	0.803

Table 6.9: Sensitivity of the design to the assumed priors for the dose-response parameters in Scenario 1 that ignores PK information.

Margin	%BD	%TD	%AD	DE	SE
3	54.0	5.6	32.6	0.875	0.743
2	71.3	0.2	41.3	0.956	0.737
1	81.7	0.0	49.0	0.971	0.780

The correctness in the optimum dose selection (%BD) increases as the ranges of the parameter values decrease. Fewer trials recommend a dose for further studies which carries the probability of toxicity above the acceptable level. The percentage of such trials completely disappears for a margin of 1. Also, the allocation of the best doses to the cohorts over the trials increases quite significantly for the dense priors. The decision and sampling efficiencies much improve as we choose narrower ranges. As %BD and %AD increase and %TD decreases for such a choice of priors, it is obvious that DE and SE will improve. The trend in improvements is similar in Table 6.9, which summaries the results for the design without the PK constraint. Though not presented, we have found the similar distributions for the optimum dose selection and dose allocation to those presented. On the whole, the results show that the efficiency of the proposed adaptive design can be improved by the choice of more appropriate priors for the dose-response parameters.

6.5.2 Priors for PK Parameters

Since our PK model is non-linear in the parameters, the Fisher information matrix depends on them. Therefore, in the search for the D-optimal time points, we need to assume some prior values for the parameters: see Section 3.5.2. We start with some best guess about the parameters. However, once the estimates are available from the trial, we use them to find the optimal time points. Also, the maximum likelihood estimation procedure for the PK parameters requires some initial values. To see whether the assumed prior values have an impact on the design, we conduct some simulation studies.

Six sets of prior values are investigated in the simulation studies. Each set consists

of the mean PK parameters which are some multiple of the standard deviation away from the true values. The main idea is to make them as far away as possible to see the impact on the design when we have complete ignorance about the parameters. The first set assumes that the mean PK parameters $\boldsymbol{\beta}$ are 3 standard deviations below the true ones. The variance components and error variance are assumed to

Table 6.10: PK parameter values for the sensitivity analysis.

Set	V	k_e	k_a	σ_1^2	σ_2^2	σ_3^2	σ^2
1	0.80	0.02	0.20	0.040	0.000025	0.0025	0.0000125
2	1.20	0.03	0.30	0.053	0.000033	0.0033	0.0000166
3	1.60	0.04	0.40	0.080	0.000050	0.0050	0.0000250
4	2.40	0.06	0.60	0.080	0.000050	0.0050	0.0000250
5	2.80	0.07	0.70	0.053	0.000033	0.0033	0.0000166
6	3.20	0.08	0.80	0.040	0.000025	0.0025	0.0000125

be one quarter of the true values. The second set contains the mean parameters which are 2 standard deviations below the true values, whereas the variance components and error variance are one third of the true values. In a similar way, we have the third set of prior values. In the fourth set, the PK parameters are one standard deviation above the true values, with the variance components and error variance one half of the true values. Similarly, we construct the fifth and sixth sets of priors.

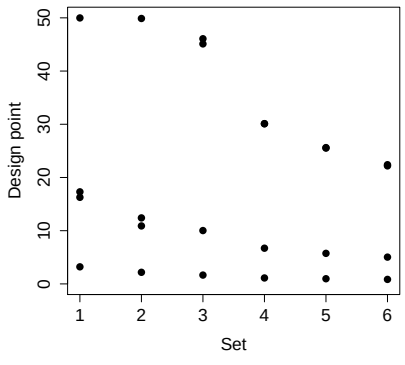

Figure 6.26: Design points obtained at various sets of prior values following the administration of the lowest dose to a cohort.

Table 6.10 gives the six sets of values that we obtain. The D-optimal time points obtained for the various prior values are shown in Figure 6.26. Some of the points corresponding to different sets are almost identical and therefore we see three points

instead of four. It seems that the time points change with the different priors. Therefore, emphasis should be given to the priors. We do this by replacing the prior values with the current knowledge in advance of a trial.

Table 6.11: Sensitivity of the design to the assumed priors for the PK parameters in Scenario 1.

Set	%BD	%TD	%AD	DE	SE
1	65.5	2.1	39.8	0.932	0.779
2	65.5	2.4	39.5	0.928	0.776
3	67.1	0.8	39.9	0.951	0.781
4	66.1	1.1	39.5	0.941	0.775
5	65.1	1.1	38.5	0.952	0.786
6	62.9	2.0	37.5	0.925	0.774

Table 6.11 summarises the simulation results obtained under different sets of priors. It is evident from the figures that the design is less sensitive to the assumed prior values for the PK parameters. The small variability that is seen in the numerical results is mainly due to randomness. Set 6 is the one that has been used in the original simulation study for Example 2. Although the priors have an impact on optimal design, they do not have an impact on the adaptive design. This might be due to the fact that the learning through gathering information in the first few cohorts is substantial and it stabilises the estimates.

6.5.3 Target Maximum Concentration

The implementation of the design presented in Example 2 requires a target value for the maximum concentration. We set that target at the true optimal dose. That is, we obtain the target $C_{\max}$ using (3.35) for the true values of the OD and the mean PK parameters $\boldsymbol{\beta}$. However, the target may be misspecified, and, therefore, this section is devoted to assessing the impact of such misspecification on the design. Other than taking it at the true OD, we will set this target at the doses which are one and two levels away. Moreover, we check it for Scenario 1 only. Checks for the other scenarios are possible using the same idea.

The true optimal dose in Scenario 1 is -0.6. The other doses that we are interested in are -1.8, -1.2, 0.0 and 0.6. As discussed earlier, we set the target $C_{\max}$ at these doses in separate runs of the trials.

Table 6.12: Sensitivity of the design to the assumed target for $C_{\max}$ in Scenario 1.

Dose	%BD	% of 3-BD	%TD	%AD	DE	SE
-1.8	0.8	10.5	2.0	3.0	0.672	0.627
-1.2	0.8	85.4	2.0	3.8	0.885	0.722
-0.6	62.9	91.2	2.0	37.5	0.925	0.774
0.0	59.4	91.6	0.6	34.8	0.935	0.790
0.6	55.3	85.9	4.0	33.4	0.889	0.756
No PK	54.0	85.1	5.6	32.6	0.875	0.743

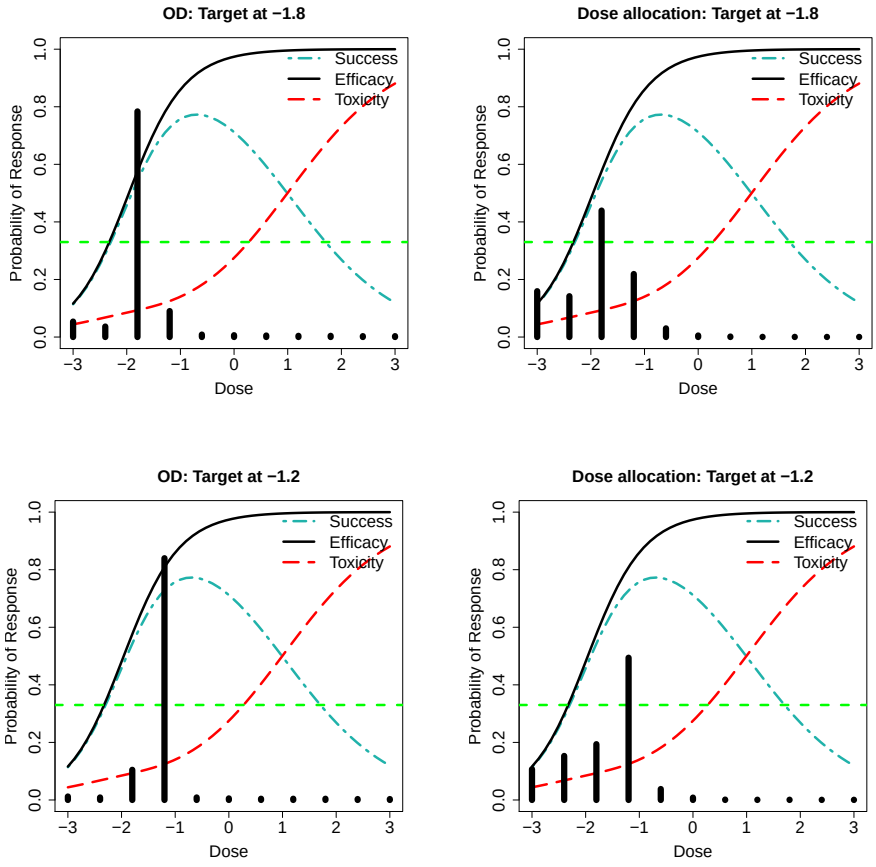

Figure 6.27: Optimum dose selection and dose allocation when the target $C_{\max}$ is taken at the doses below the true optimum dose.

Table 6.12 contains a summary of the simulation results obtained from 1,000 trials. Figures 6.27-6.28 present the distributions of the recommended doses and allocated doses to the cohorts under different choices of target values. When targets are taken at the doses below the true optimum dose, all performance indicators give poor val-

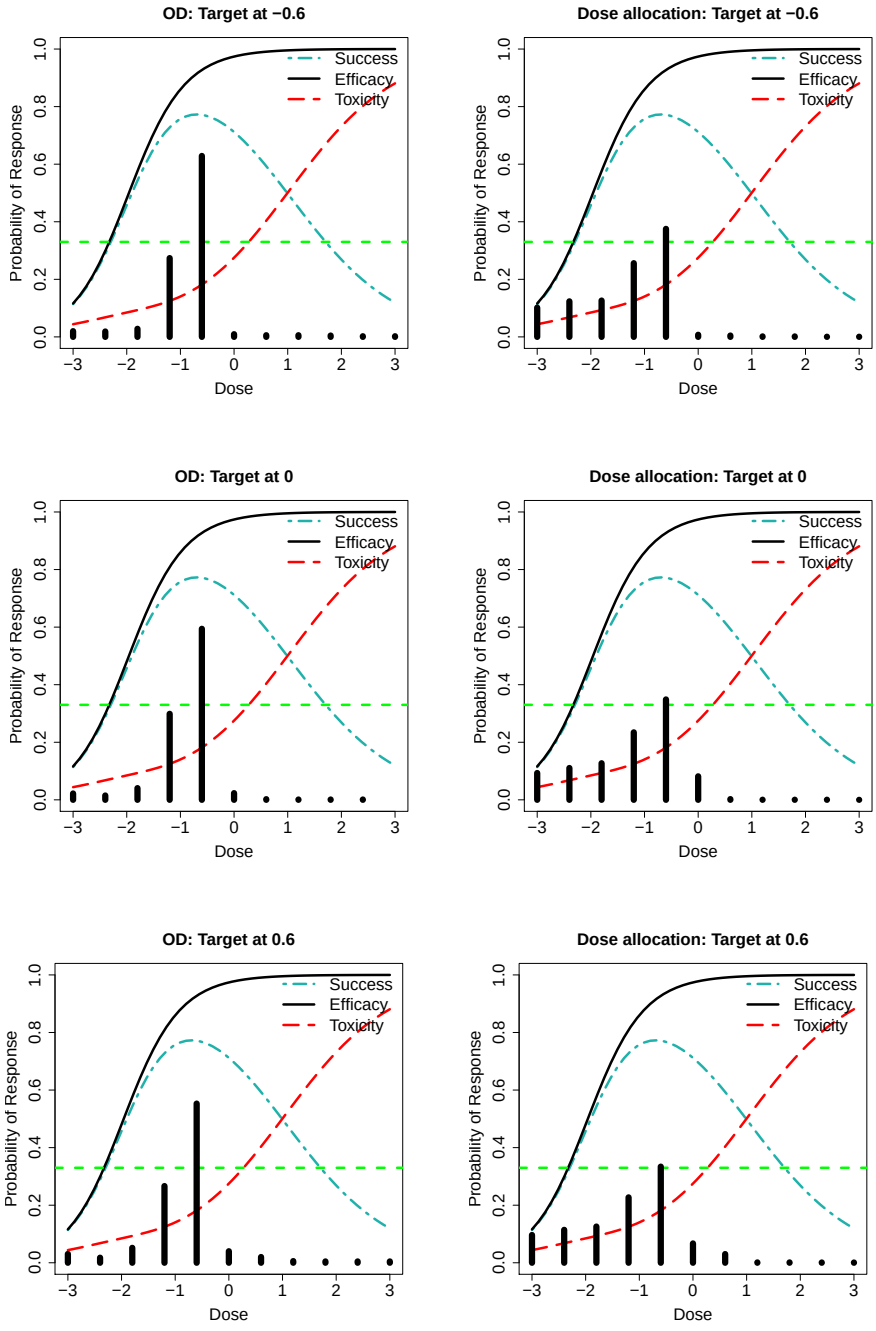

Figure 6.28: Optimum dose selection and dose allocation when the target $C_{\max}$ is taken at the true optimum dose and the doses above it.

ues. In such cases, the design identifies the optimum dose very poorly. Although the decision efficiency for -1.8 is poor, the measure provides a satisfactory value for -1.2. This is because of the fact that the probabilities of a successful outcome at doses -1.2 and -0.6 are quite close.

The scenario has three doses -1.2, -0.6 and 0.0 for which the probabilities of success are 0.73, 0.77 and 0.71, respectively. Therefore, we create a column % of 3-BD

in Table 6.12 showing how frequently these were recommended for further studies. As the target moves to the doses above the true OD, correct identification of the optimum dose is less likely. Also, fewer patients are allocated to the true OD. % of 3-BD gives a better picture of the recommended doses. Although we have a very small figure for %BD corresponding to -1.2, % of 3-BD gives a very high figure. This means that many of the trials recommended doses -1.2 and -0.6. There is a similar behaviour for dose allocation. As a result, we have fairly large values for DE and SE.

The DE corresponding to dose 0.0 is higher than that at dose -0.6, since the above three doses are selected more often than for the original design. Similarly, for dose 0.0, SE is slightly higher and that is evident from the relevant distribution of dose allocation in Figure 6.28. It is important to note that the performance indicators are more sensitive when the design considers the target $C_{\max}$ at a dose below the true OD. However, it is less sensitive when the dose is above the true OD. All of these results demonstrate the need for care in setting the target value for the maximum concentration, although there is some margin for misspecification.

6.6 Example 3

This example tries to explore the behaviour of a design which uses the combined criterion in Section 5.3.3 for dose optimisation. As mentioned earlier, the criterion is intended to serve two purposes: allocation of the most efficacious doses to the cohorts and improvement in the estimation of the dose-response parameters. We conduct simulation studies, details of which are given in the following section, to investigate whether the intended objectives are met.

6.6.1 Simulation Settings

Here, we use the continuation ratio dose-response model introduced in Section 4.2. The dose-response scenarios in Figure 4.1 are investigated. The same scenarios have been used as in Example 1 and more details of them are available in Section 6.3.1. We assume that the acceptable level for the probability of toxicity is $\gamma = 0.2$.

The design follows the general algorithm described in Section 5.2. Each trial assigns the lowest available dose of 0.5 mg/kg body weight to the first cohort of patients. Escalation or de-escalation of doses to the first four cohorts is based on the up-and-down design presented in Section 5.2. Since the Fisher information matrix in (4.9) is singular, to have a non-zero value for the determinant of $M(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)$ in Section 5.3.2 immediately after the up-and-down stage, we need to ensure that at least two of the doses allocated to the first four cohorts are different.

The design does not use any PK constraint and therefore only requires dose-response outcomes to be simulated. These are drawn from a trinomial distribution with the true probabilities corresponding to a dose obtained for the respective scenario. Once the trial passes the up-and-down stage, the model-based procedure starts. This requires estimation of the dose-response parameters and assignment of the best dose to the next cohort. The Bayesian estimates of the dose-response parameters are obtained at each steps after the up-and-down stage, using the same set of priors as in Example 1.

The penalised combined criterion in (5.9) is utilised to assign doses at each step of a trial. As the dose-optimisation criterion to be used already penalises observations for low efficacy or high toxicity or both, not additional constraint is imposed. The rules presented in Section 5.6 are used to stop a trial. Here, the optimum dose is defined as the dose that has been repeated r times. However, for the trials that utilise the maximum number of cohorts m, the optimum dose is defined as the one for which the estimated probability of success is maximum, subject to the constraint that the estimated probability of toxicity at that dose is no more than γ. Utilisation of the maximum number of cohorts means that the trial came through a reasonable number of steps. Consequently, we can expect the estimates of the parameters to be reliable, as the determinant of the FIM is in the optimisation criterion. Therefore, we can use these estimates to predict the dose-response curves more accurately to

obtain the dose which is best in terms of the probabilities of success and toxicity. We set $r = 6$ and $m = 20$ in our simulations.

The penalty function to be used in the simulation study was introduced in (5.7). It has the control parameters C_S and C_T. To motivate on the values for these parameters, we conduct an investigation described below.

Optimum Values for C_S and C_T

Figure 4.1 shows the dose-response scenarios that we are going to consider in this example. In Scenario 1, the true optimal dose lies at the beginning. Since the probability of success is low and the probability of toxicity is high at the higher doses, the penalty resulting from high C_S and C_T values will restrict the dose selection from the end of the dose region: see Figure 6.29. In Scenario 2, success is low at the beginning and toxicity is high at the higher doses. Therefore, penalising doses from both ends will keep the dose selection in the middle, where the true optimal dose lies. The success curve in Scenario 3 is of a flat type. Even at the early doses, success is quite high, and toxicity at the higher doses is not as high as in the previous two scenarios. As a result, the penalty for taking observations from both ends is smaller than that for the previous scenario. However, it still keeps the dose selection in the middle of the dose region. The true optimal dose in Scenario 4 lies at the end of the dose region. The probabilities of success and toxicity remain very low for more than half of the available doses. Also, we do not have any dose for which the probability of toxicity is very high. Penalising observations for low success and high toxicity gives very high penalty values for most of the early doses in this scenario and thereby will keep the dose selection from the upper end of the dose region. To summarise, for all of our scenarios, we plan to penalise for taking a dose which is either low officious or highly toxic or both.

Having decided on the non-zero values for the control parameters, we assume that they take the same value, that is, $C_S = C_T = C$. Since a non-zero value of C will

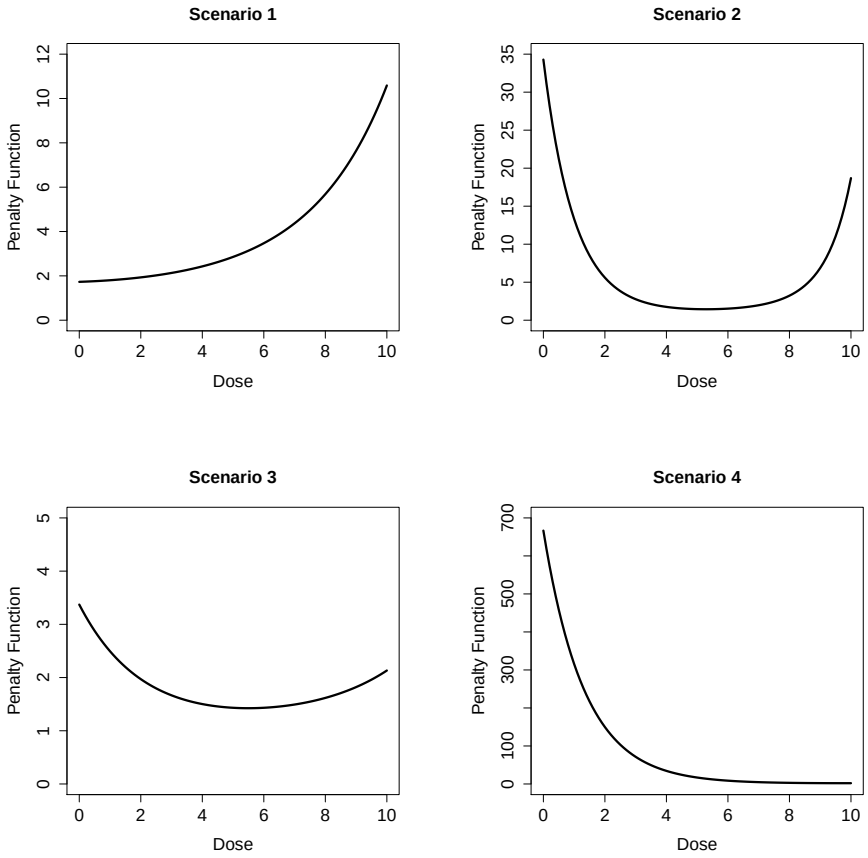

Figure 6.29: Penalty function for the four scenarios assuming $C_S = C_T = 1$.

keep patients away from low efficacious or highly toxic doses, we need to find its optimum value. We have tried various values of C. Figures 6.30-6.32 depict some features of the resulting designs to guide the choice for C.

In each case, the results are based on 200 simulations of the respective scenario. In Figure 6.30, we have the percentages indicating how often toxic doses are allocated to the patients. It also shows the percentages of toxic doses selected as the OD. The percentages for correct identification of the OD are displayed too. It reveals that, as the value of C increases, fewer patients are treated with toxic doses. In all of the scenarios, there is a fall in the percentage of toxic doses as the OD when the value of C increases from 0 to 1. It increases and then decreases for the other values of C. However, it levels off for Scenario 3 at $C = 1$. As C increases, the percentage for correct identification of the OD usually increases initially. It then levels off in Scenario 3 and gradually decreases in Scenario 1. The corresponding

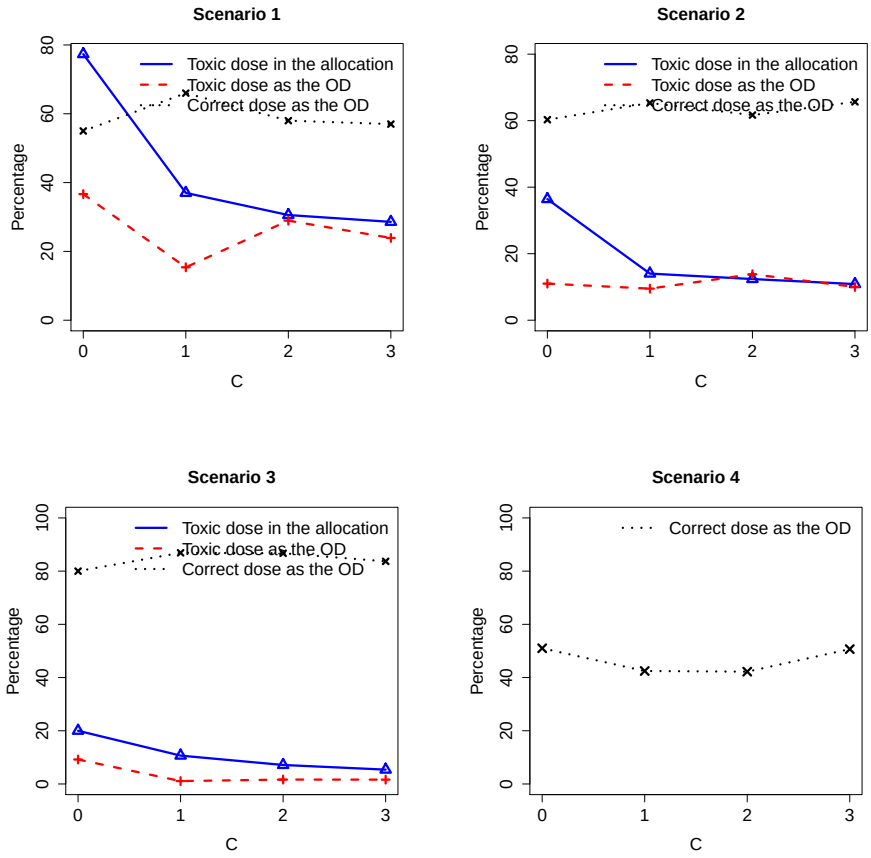

Figure 6.30: Percentage of cohorts treated at the toxic doses during the trials, percentage of toxic doses recommended as the optimum dose and percentage of trials with the correct OD selection. Different values for the control parameters C_S and C_T have been used assuming $C_S = C_T = C$.

change in Scenario 2 is very slight. Since Scenario 4 is one where we do not have any dose for which the toxicity level is higher than the acceptable level, although we cannot draw the first two curves related to toxic doses, we can show the percentage for correct identification of the OD. All of these results suggest that $C = 1$ could be a reasonable choice.

Figures 6.31-6.32 present the bias and mean square error for different scenarios. As C increases, one would expect the penalised D-optimum design to be further away from the D-optimum design. Consequently, the bias and mean square error would also increase. This was found to be true on some occasions, but not always. Small number of simulations could be a possible reason. Therefore, together with the previous results, we decide to use $C = 1$ in this example.

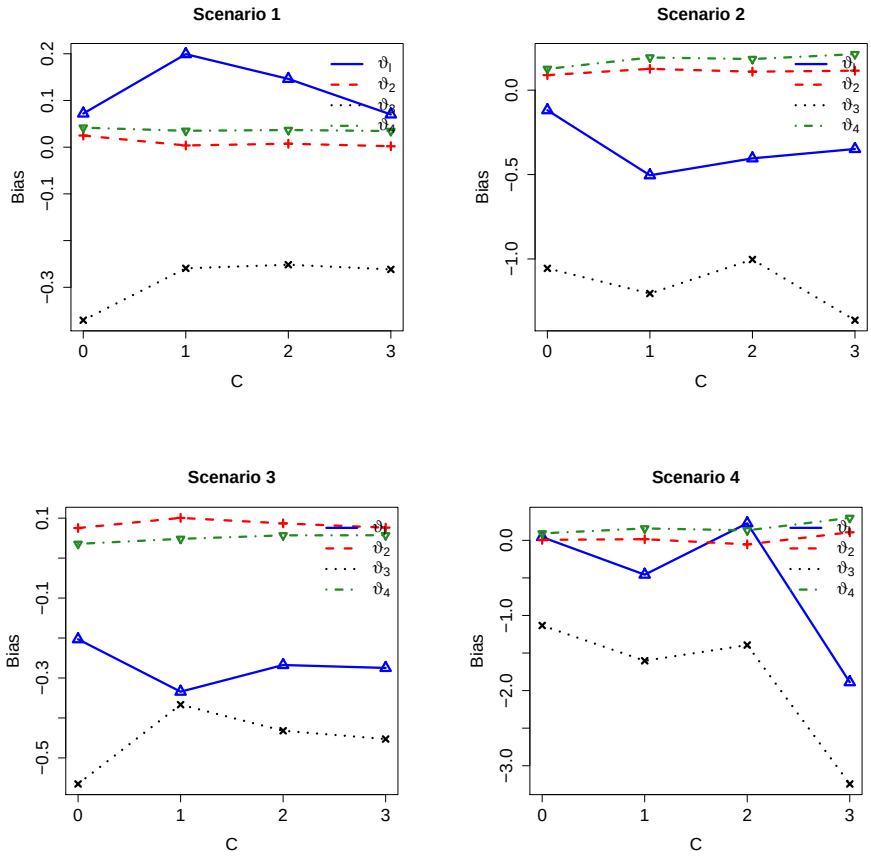

Figure 6.31: Biases of the parameter estimates for different choices of control parameters C_S and C_T assuming $C_S = C_T = C$.

6.6.2 Numerical Results

With the control parameter values set to 1, we run the penalised combined criterion for each of the scenarios. One thousand simulated trials are generated in each case for various values of the weight a. It is important to note that, when $a = 0$, the dose selection is based on the probability of success only with no PK constraint. On the other hand, the criterion reduces to the penalised D for $a = 1$: see (5.9). Tables 6.13-6.16 illustrate the simulation results.

Table 6.13: Combined criterion for Scenario 1.

a	%BD	%TD	%AD	DE	SE
0.0	53.5	31.8	31.9	0.681	0.516
0.2	64.0	27.8	35.2	0.722	0.515
0.4	74.0	21.9	37.8	0.780	0.478
0.6	86.0	12.6	42.0	0.873	0.502
0.8	94.3	4.7	43.8	0.952	0.513
1.0	81.1	12.3	32.0	0.877	0.377

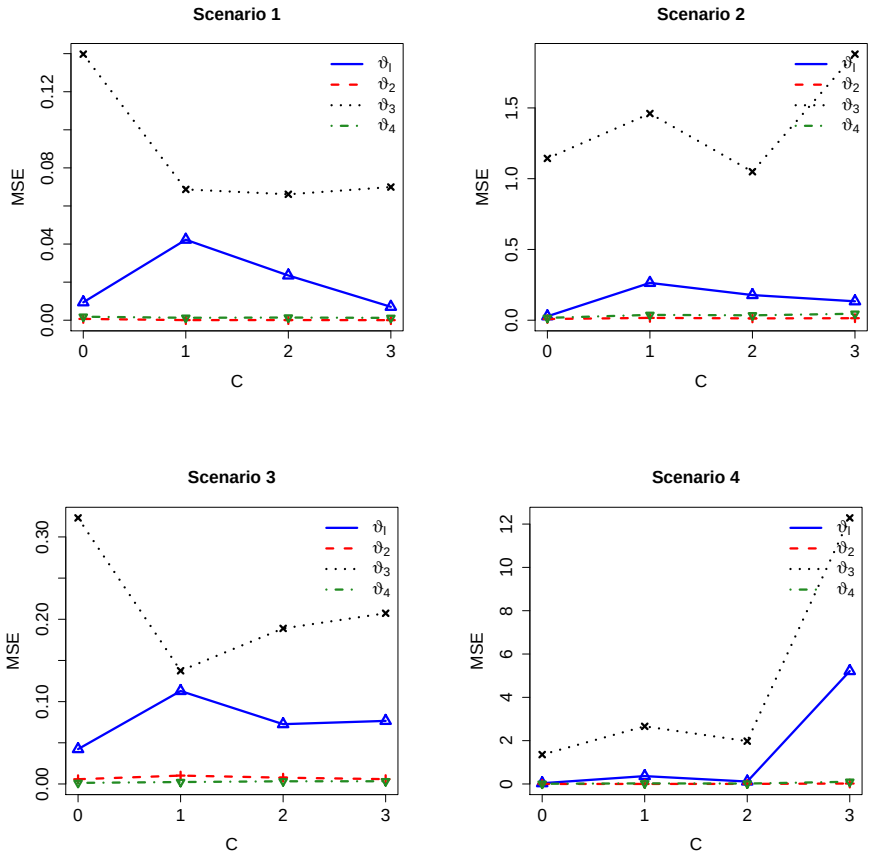

Figure 6.32: Mean square errors of the parameter estimates for different choices of control parameters C_S and C_T assuming $C_S = C_T = C$.

Table 6.13 shows the results for Scenario 1. The penalised D-criterion gives better results than maximisation of the probability of success. It is seen for all measures apart from SE. This is expected, as the criterion based on the maximisation of the probability of success only is defined so that more patients are treated during the trial with efficacious doses more often. The penalised D-optimum criterion does not take care of this. The design performs the best overall when $a = 0.8$. Here, the design can identify the true optimum dose 0.5 most accurately. Only a few trials gave toxic doses for further study in the next phase. About 44% of the cohorts are treated at the OD throughout the trials. On the whole, the combined criterion enhances the performance for Scenario 1.

The results for Scenario 2 for different values of a are compared in Table 6.14. The best result in terms of all the performance indicators is attained at the weight 0. However, there is little variation in the performance for the values of a between 0

and 0.6. The only noticeable difference is observed for $a = 1$. For this choice of weight, the design identifies the OD less accurately compared with the other weights. Fewer cohorts are treated with the best dose as the weight increases and there is also a sharp drop when the weight is 1. The efficiency measures DE and SE are both considerably smaller in this case.

We obtain similar results for the weights ranging between 0 and 0.8 in Scenario 3. As in the previous scenario, we observe a decreasing trend in the dose allocation, with a sharp drop at weight 1. All of the performance values indicate that the penalised D is performing poorly compared to the other cases. However, it might be worth combining the criteria in this case, with a in [0.2,0.8].

Table 6.14: Combined criterion for Scenario 2.

a	%BD	%TD	%AD	DE	SE
0.0	68.0	7.0	36.2	0.913	0.638
0.2	66.0	7.8	35.6	0.903	0.629
0.4	66.0	7.8	35.3	0.901	0.622
0.6	64.2	7.4	34.0	0.905	0.616
0.8	57.6	8.3	30.4	0.888	0.590
1.0	49.7	8.1	16.2	0.843	0.493

Table 6.15: Combined criterion for Scenario 3.

a	%BD	%TD	%AD	DE	SE
0.0	84.5	2.2	49.5	0.968	0.842
0.2	86.9	1.7	48.9	0.973	0.838
0.4	84.6	2.2	47.6	0.969	0.824
0.6	85.4	2.6	45.8	0.964	0.801
0.8	85.6	1.1	42.0	0.971	0.771
1.0	62.1	3.8	22.6	0.899	0.652

The optimum value of a in Scenario 4 is 0.4, since both DE and SE attain their largest values. However, looking at the figures for the best dose recommendation and allocation, it can be claimed that the design is performing similarly for the weights between 0.4 and 0.8. That is, higher values of a gives the design the best performance. But, as in Scenarios 2 and 3, the penalised D-optimum design is not performing well in this scenario.

When $a = 0$, for the first three scenarios, the performance values are consistent with the ones obtained for the case 'No PK' in Tables 6.1-6.2. Although the combined criterion does not consider the toxicity constraint in Section 5.4 during dose allocation, the figures are very similar. However, for $a = 0$ in Scenario 4, the results are better than those that we obtained earlier. Since the optimum dose is the last dose in the dose region, the toxicity constraint often restricts us from choosing it for the cohorts. This can be regarded as one possible reason for having better results than before.

Table 6.16: Combined criterion for Scenario 4.

a	%BD	%TD	%AD	DE	SE
0.0	69.5	0	27.0	0.899	0.485
0.2	69.5	0	27.3	0.888	0.484
0.4	73.4	0	28.2	0.914	0.488
0.6	72.0	0	28.2	0.893	0.481
0.8	75.0	0	29.6	0.901	0.481
1.0	47.3	0	26.4	0.784	0.460

It is worth mentioning that SE is large when the weight is 0 and small when the weight is 1 for Scenarios 1-3. This is because maximisation of probability of success allocates most of the cohorts to the most efficacious doses compared to the D-criterion. Even though we penalise the D-criterion for low efficacious or high toxic doses, the improvement in SE is not that noticeable. In Scenario 4, the SE values are close to each other, apart from the one corresponding to $a = 1$. This is a scenario where we do not have any dose for which the probability of toxicity exceeds the acceptable level, but it has many doses for which the probability of success is very low. Therefore, the differences among the SEs are not very substantial for this scenario.

We have looked at three approaches for dose finding in our examples. In one approach, the intention is to allocate doses to the patients that are best in terms of current knowledge. This kind of approach is known as the "best intention" approach

in the literature. The other approach targets the most effective gathering of information and is based on the theory of optimal design. The third approach tries to make a trade-off between the two. Best intention designs are ethically attractive, as they take care of the patients, but, unlike the one based on optimal design, it has limitations in terms of convergence. Since here we have small numbers of cohorts, the convergence property of the D-optimum design is not that important. However, combining the two approaches, on the whole, may improve the performance of the adaptive design.

6.7 Discussion of the Results

This chapter has made detailed simulation studies to learn about the adaptive designs. Two classes of designs are presented: one considers efficacy and toxicity end points and the other, along with these outcomes, also considers pharmacokinetic information. Three examples are introduced in this context and the major results are discussed below.

The design presented in Section 6.3 is conceptually similar to that of Zhang et al. (2006), but their design does not incorporate pharmacokinetic information. They find a dose that maximises the difference between the estimated success probability and λ times the estimated toxicity probability, given that the estimated toxicity probability is smaller than a pre-specified level, where $0 \leq \lambda \leq 1$. The value of λ can be varied to include toxicity in the criterion and no general recommendation is made about its value. But, in many real scenarios, such a difference with a non-zero λ may lead to doses which are not optimum. Scenario 3 in Example 1 is such a case, where the difference at some sub-optimal doses is the same as that at the optimal dose. Therefore, we decide to use $\lambda = 0$. The decision also helps to avoid double dependence on the probability of toxicity.

In this example, along with dose-response outcomes, we have considered an important PK measure, AUC, and its inter-patient variability in the dose escalation.

The main purpose of this study was to investigate the role of PK measures in dose finding, and, by means of detailed comparisons, we showed that utilising the PK information can be very beneficial.

The simulation results from four different dose-response scenarios indicate that the incorporation of such measures can improve the accuracy of dose-finding studies. It is also shown that the method is capable of limiting overdosing by a considerable amount depending on the location of the OD. The proposed PK-guided approach can therefore be used as a reliable dose-finding procedure in situations where more careful escalation is essential to avoid toxicity. However, we have found the design to be sensitive to the target AUC, but less sensitive to the dose-skipping constraint.

Section 6.4 presents the second example with a different set of PK and dose-response models. The dose-optimisation criterion and the toxicity constraint used are the same as in Example 1. However, this time the PK constraint is on the maximum concentration. The reason for using the maximum concentration is that often clinicians are more interested in it than the AUC because of the simplicity in comparing it with a threshold to decide on the nature of the dose-response outcome. The inter-patient variability in $C_{\max}$ is also taken into consideration in the dose escalation.

Results from the simulation study of four plausible dose-response scenarios show a gain in the efficiency of dose finding due to considering additional PK information. We can identify the optimum dose accurately and can allocate most relevant doses to the cohorts in a trial by using the proposed design. It also restricts us from recommending a toxic dose as the optimum dose for further study in the next phase.

We achieve similar gains in both examples when PK information is used. It is important to mention that, in both cases, we obtained a small bias and mean square error for the PK parameter estimates. This happens as the D-criterion has been employed to obtain the time points to collect blood samples to measure the concen-

tration of the drug. The bias and mean square error of the dose-response parameter estimates obtained using the two approaches are very similar. Therefore, consideration of the additional pharmacokinetic information is not improving the quality of the dose-response parameter estimates, but rather it is helping us to find the best dose.

To assess whether the observed differences between the PK-guided design and the other design are appreciable, we present confidence intervals for the various measures in Tables A.1 and A.2. Non-overlapping intervals are found in most cases, which confirms our claim regarding the PK-guided design. Because of time constraints, the number of simulations was restricted to 1,000 in both examples. We expect that a larger number of simulations would not change the estimated values very much, and, therefore, we would obtain similar results for optimum dose selection and dose allocation. However, as the standard error of the proportions would decrease with the increased number of simulations, shorter confidence intervals for the measures would be obtained. That is, the significance of our results would then be even more pronounced.

Sensitivity analyses of the designs to the different parameters are presented in Section 6.5. Since the computations are very time consuming, we restrict ourselves to Scenario 1 in Example 2. However, investigation for the other examples and scenarios will be straightforward.

Section 6.5.1 shows the results for the assumed priors of the dose-response parameters. It is seen that, as the priors become narrower, the efficiency of the method increases. That is, the method can identify the optimum dose more accurately and does not recommend toxic doses as the optimum dose very often. Also, more patients are allocated to the best doses throughout the trials. This means that the more one knows about the scenario, the better the chance of a successful trial. A summary of the results is also presented for the design which does not consider the

constraint on $C_{\max}$ in Table 6.9. It is observed that both methods produce similar results as the priors become more dense. Therefore, the additional constraint on the PK parameter is useful in the presence of a vague prior. If a reliable prior is available, the method itself is able to produce reliable results, even if we do not consider pharmacokinetic information.

The priors for the PK parameters are mainly used to initiate the finding of D-optimal time points and the maximum likelihood estimates of the parameters. Once estimates from the trial data are available, they are used. The sensitivity of the design to the assumed priors for the PK parameters is shown in Section 6.5.2 for six sets of values. The results are very similar to each other. Therefore, we can conclude that the design is not especially sensitive to the priors. The whole idea was to see what happens in the presence of vague priors. If the initial values for the parameters are far away from the true ones, these may lead the PK sampling times to be different initially, but that will not have a great impact on the optimum dose finding for the next phase.

Section 6.5.3 contains the simulation results to check the sensitivity of the design to the assumed value of $C_{\max}$. The design is much more sensitive when the target is taken at a dose below the true optimum dose. The degree of sensitivity is less when the target is at a dose above the true optimum dose. Also, the results are still better when we do not have any PK constraint. On the whole, the design is sensitive to the target value of the maximum concentration. Therefore, to implement the proposed PK-guided design, one needs to carefully select the target value, since the future of the trial depends highly on it.

We have explained in Section 1.1.1 how the response following a dose depends on the concentration of the drug. Although concentration is a function of the dose, for a fixed value of the latter, due to population variability, we can obtain different values of the former for different patients. The proposed PK constraints in (5.11)

and (5.12) take into account the concentration along with the inter-patient variability. Since a population D-optimum design is used to collect blood samples, we achieve precise estimates of the parameters. The constraints are defined in terms of the population mean parameters and their standard errors, as shown in Section 3.8. Although the constraints are working like a brake on dose escalation, they are adapted during the trial and work in a flexible way by modifying the strength of restriction. Other constraints such as the toxicity limit (5.10) or the dose-skipping constraint work differently. They are not flexible and are purely for safety reasons. For instance, the constraint of not skipping more than one dose level at a time is used in both PK-guided and the other designs: see Section 5.3. The impact was reported in Section 6.3.2. The results in Tables 6.4 and 6.5 tell us that, other than the difference in dose allocation, the dose-skipping constraint does not have a noticeable effect on optimum dose selection. So, the gain that we claim in the presented designs is largely due to the PK constraint. The PK constraint and the constraint on dose skipping work differently. The dose-skipping constraint aims to minimise jumps over multiple doses. Although it slows down the escalation procedure to make it safer, this may still lead to unacceptably toxic doses being applied. The PK constraints are intended to limit such occurrences, as well as to choose the best dose more accurately.

Although PK-guided designs are found to be sensitive to target values, we have noticed that they still outperform the other designs in the cases where the target is at the dose above the true OD: see Tables 6.3 and 6.12. As the PK-guided designs are sensitive, we recommend using them in situations where such targets can be reasonably well assessed. With such targets, people would expect improved performance of the design, and we have shown it numerically in our work.

PK information is commonly collected in early clinical trials and it is often analysed for the purpose of dose selection. Here, we propose a systematic method with a two-fold aim: first, to obtain the best dose level for further study in the next phase with the minimum chance of toxic responses during the trial, and, second, to obtain

the most efficient estimates of the population PK parameters. The second goal is achieved by using the population D-optimum design for the sampling times. A PK-guided trial will make sense only if the PK information is accurate and we assure this by the choice of design.

The results obtained in Section 6.6 indicate that the proposed combined criterion is a promising one in dose-finding studies. It has been found that, for Scenarios 1 and 4, the combined criterion produces the best results for large values of the weight. Dose-optimisation criteria like maximisation of the probability of success and the penalised D-criterion are not that attractive for these scenarios. For Scenarios 2 and 3, there is not much difference between the results for the combined criterion and those for the criterion that maximises the probability of success. However, as we expect to obtain better estimates of the dose-response function, it would be good to use the combined criterion in these scenarios too. We have also found the combined criterion to always outperform the penalised D-criterion. Except in Scenario 1, maximisation of the probability of success is found to produce better results than the penalised D-criterion.

Chapter 7

Conclusions and Future Work

The development of a new drug product is a lengthy and costly process. Since a lengthy development time is not very well accepted, continuous efforts have been made by researchers to shorten it. As a consequence, adaptive dose-finding designs are available in the literature to minimise the drug development time and costs. The main essence of an adaptive design is that the dose level to be used for the next patient depends on the outcomes obtained from the previous patients who have already received different doses from the dose region.

Although many adaptive designs are available, we believe that there is still scope for further improvements. The whole purpose of our research has been to develop a flexible model-based design to find the optimum dose more often without exposing many patients to either subtherapeutic or toxic doses. Furthermore, if recommended, by applying the population D-optimum design for blood sampling, the method allows us to efficiently estimate the population PK parameters, and, also to use the fitted PK models to enhance the dose-finding studies.

7.1 Conclusions

The expressions derived in Chapter 3 for measuring the inter-patient variability in the area under the concentration curve and the maximum concentration are general. The methodology can be adapted for any underlying pharmacokinetic model.

People researching in this area often use a non-parametric approach to find these quantities. But that requires collecting too many blood samples to measure the concentration of the drug. Even if the AUC and $C_{\max}$ are obtained from such data, there is no straightforward method to assess the inter-patient variability. In that way, the proposed method is very flexible. The use of the D-optimum design will ensure the efficiency in parameter estimation with fewer blood samples. This will lead efficient estimation of the area under the concentration curve and the maximum concentration, and also their inter-patient variabilities.

Chapter 4 has the continuation ratio and Cox models, which we use to model the probabilities of the dose-response outcomes. It has been shown that the associated Fisher information matrix in each case is singular. Although other work mentions the singularity, there are no explicit results for the information matrices. The analytical forms will reduce the time for any optimisation problem involving these matrices. Also, knowing the rank of the FIM tells us how many different dose levels are necessary to ensure non-singularity of the matrix.

The essence of the methodology presented in Chapter 5 is that, following the assignment of the current best dose to a cohort of patients, the dose-response outcomes are observed. The algorithm runs following an up-and-down design for the initial four cohorts. The up-and-down stage need not be exactly for four cohorts. Since the ranks for the continuation ratio and Cox models are 2 and 3, respectively, to implement the combined criterion, we have to run the up-and-down design so that at least 2 or 3 different doses are assigned to the cohorts. Therefore, we have set it as four to keep it general. Once the trial is past the up-and-down stage, the dose is selected for each cohort based on a chosen dose-optimisation criterion and constraints evaluated at the current estimates of the model parameters. If we plan to incorporate pharmacokinetic information, we also need to measure the concentration of a drug in the blood at the sequentially obtained locally D-optimal time points after each assignment of dose. The algorithm, dose-optimisation criterion

and constraints presented in this chapter are general. The criterion that maximises the estimated probability of success is ethical, since the current knowledge about success is employed in dose escalation. The penalised D-criterion gathers information so that the variability in the parameter estimates is minimised. The combined criterion provides a bridge between these two. All of the presented methodology can be applied to any dose-response and pharmacokinetic models.

To study the design properties under different dose-optimisation criteria and constraints, we have introduced three different examples in this thesis. The first example is a one-compartment PK mixed-effects model with bolus input and first-order elimination and a continuation ratio dose-response model. In this case, the dose-optimisation criterion is the maximisation of the estimated probability of success with constraints on toxicity and the area under the concentration curve. The second example assumes the Cox model for binary efficacy and toxicity responses. This is accompanied by a one-compartment PK mixed-effects model with first-order absorption. The dose-optimisation criterion employed is the maximisation of the estimated probability of success with constraints on toxicity and the maximum concentration. The last example utilises the penalised combined criterion for finding doses for each successive cohort in a trial. This example is based on the continuation ratio model for trinomial dose-response outcomes. Unlike the other examples, here we apply the criterion without any additional constraint, since it utilises a penalty function to address the issue of low efficacy or high toxicity.

We have developed code to conduct all of the computations in *R*. The code is written so that it can be applied to other dose-response models. But, of course, that will require us to define any such model in *R*. It is important to note that the part of the code that deals with the D-optimal time points is based on the package *PFIM 3.2*. Since the package can work with other PK models from the literature, it will not be difficult to extend the methodology to those models.

Simulation studies for various plausible dose-response scenarios in the different examples show that the designs are able to identify the optimum dose accurately. They limit toxic doses as the optimum dose by a considerable amount and assign the most relevant doses to the cohorts during a trial. We have also seen that the efficiency of a dose-finding design can be increased if it is possible to assume a target value for the AUC or $C_{\max}$. Finally, we can conclude that the presented designs are efficient and ethical, and can be used as reliable dose-finding procedures.

Implementation of the PK-guided approach requires a reliable value for the target AUC and $C_{\max}$. Since we did not have one for our examples, for the simulations, we set the target value as the one obtained for the true PK parameters and the true OD. In real trials for new drugs, it may be elicited from the experiences of the clinicians. Previous studies of similar drugs or extrapolation from pre-clinical studies of the same drug could be some possible options as well. If no information on the target values is available, we cannot implement the methods. In that case, the combined criterion will be the best option.

Although we have listed existing designs in Chapter 2 for combined phase I/II trials, none of them utilise pharmacokinetic information. Our motivation was to see if any improvement is possible due to the use of this kind of data. Also, we wanted to combine the objectives of best intention and optimum designs into a single framework. The purpose is to ensure that patients are allocated at the most efficacious doses and also we obtain precise estimates of the parameters.

7.2 Future Work

The current research effort is dedicated to developing a flexible dose-finding design for early phase clinical trials. In the future, we want to focus on the following issues, which are not covered in the thesis.

The penalised combined criterion has been used for the continuation ratio dose-

response model. Although we believe that it will perform in a similar way, it would be interesting to apply it to the Cox model. This will require us to develop code for computing the penalised D-optimum design. Since we have the analytical form for the information matrix, this will be straightforward.

We want to develop an R package to have functions to run the presented methods in the thesis. These functions will enable us to find a dose for the next cohort of patients and to generate simulated trials under a specified dose-response relationship. That is, given data from a trial, the package will guide us to administer doses to the successive cohorts. Also, it will have the feature to summarise the behaviour of dose selection for a specific scenario through simulations.

In many practical situations, the probability of response may not be limited to dose only, but may also depend on the covariates of the individual. If such a dependence is overlooked in the modelling, we may end up with unreliable estimates of the model parameters. For accurate prediction of the dose-response relationship in such a situation, it will be worthwhile to consider the covariates. Therefore, another possible direction would be the extension of the presented methodology to include covariates.

Appendix A

Flow Chart for the Design and Supplementary Material

A.1 Solutions to the Differential Equations: One-Compartment PK Model with First-Order Absorption

From Section 3.2.1, we have

$$\frac{dX_1(t)}{dt} = k_a X_2(t) - k_e X_1(t)$$

and

$$\frac{dX_2(t)}{dt} = -k_a X_2(t).$$

For notational simplicity, we write

$$X_1' = k_a X_2 - k_e X_1$$

and

$$X_2' = -k_a X_2,$$

which in matrix form can be expressed as $\boldsymbol{X'} = \boldsymbol{AX}$, where $\boldsymbol{X'} = (X_1', X_2')^T$, $\boldsymbol{X} = (X_1, X_2)^T$ and

$$\boldsymbol{A} = \begin{bmatrix} -k_e & k_a \\ 0 & -k_a \end{bmatrix}.$$

We obtain the eigenvalues $\lambda = -k_a$ and $\lambda = -k_e$ from the characteristic equation $|\boldsymbol{A} - \lambda \boldsymbol{I}_2| = 0$. It can be shown that the corresponding eigenvectors are $(k_a/(k_a - k_e), -1)^T$ and $(k_a/(k_a - k_e), 0)^T$. Therefore, the general solution to the differential equations is

$$\begin{pmatrix} X_1(t) \\ X_2(t) \end{pmatrix} = c_1 \begin{pmatrix} \dfrac{k_a}{k_a - k_e} \\ -1 \end{pmatrix} e^{-k_a t} + c_2 \begin{pmatrix} \dfrac{k_a}{k_a - k_e} \\ 0 \end{pmatrix} e^{-k_e t}.$$

That is, $X_1(t) = c_1 k_a/(k_a - k_e)e^{-k_a t} + c_2 k_a/(k_a - k_e)e^{-k_e t}$ and $X_2(t) = -c_1 e^{-k_a t}$. Using the initial conditions, we have $c_1 = -x$ and $c_2 = x$. Thus, we find that

$$X_1(t) = \frac{x k_a}{k_a - k_e}(e^{-k_e t} - e^{-k_a t})$$

and $X_2(t) = x e^{-k_a t}$.

A.2 Structure of the Proposed Design

The flow chart below shows the different steps of the general algorithm presented in Section 5.2.

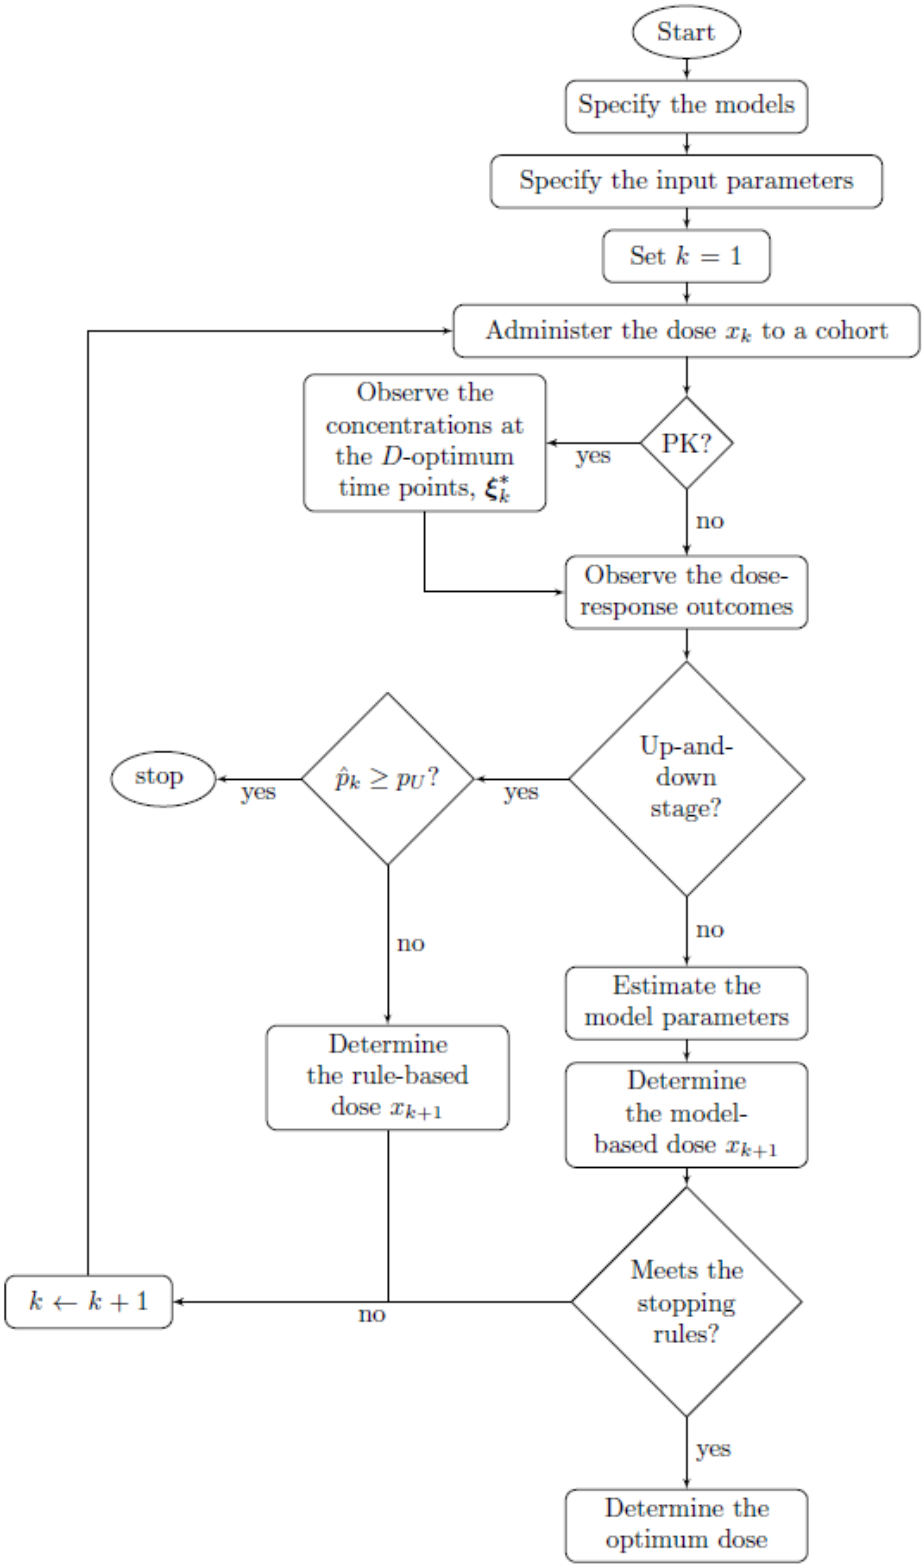

A.3 Efficiency versus Design Points: Example 1

The graphs below are for deciding the number of blood samples to be considered per patient in Section 6.3.1.

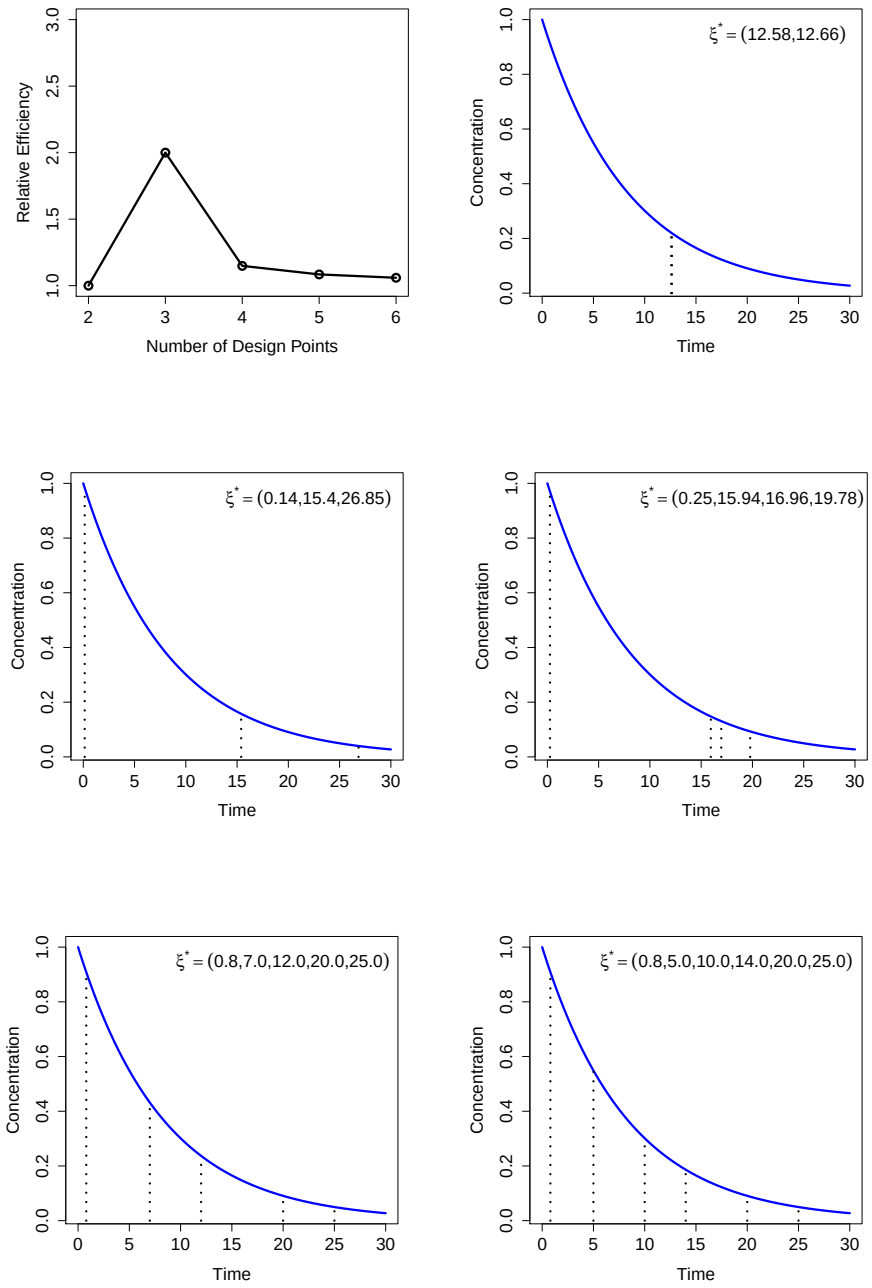

Figure A.1: Rationale for setting the number of design points in the one-compartment PK model with bolus input and first-order elimination. The locally D-optimum design points are obtained using the initial prior values $\boldsymbol{\Psi}^0$ assuming that the lowest dose is given to a cohort of 3 patients.

A.4 Dose-Response Scenarios at the Prior Ends

The graphs below are the scenarios that can be obtained at the prior ends.

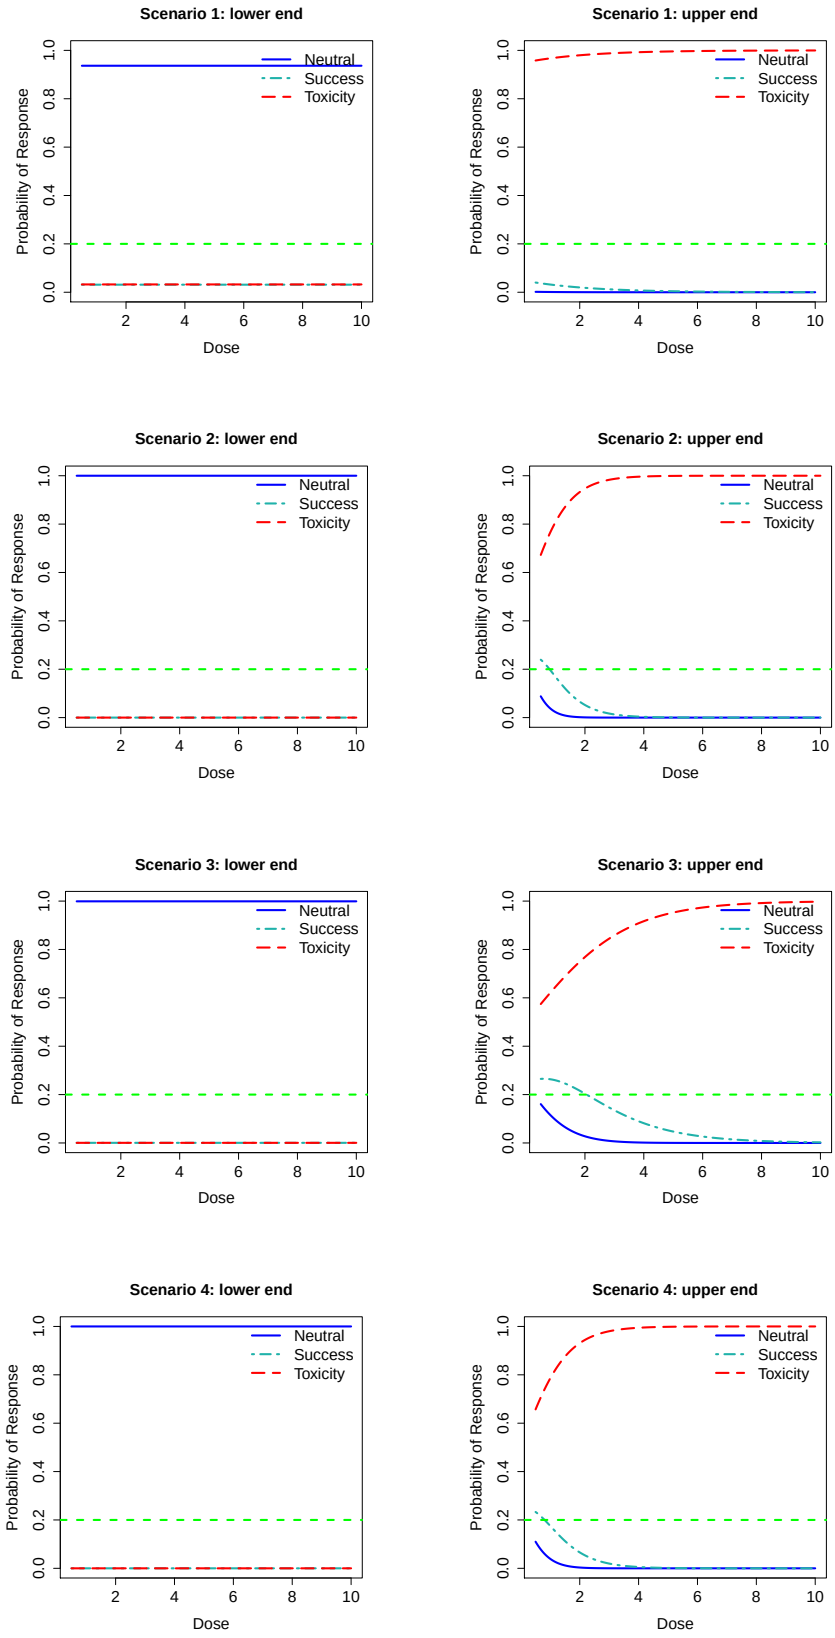

Figure A.2: Dose-response curves at the lower and upper ends of the priors used in the simulation study. These are for the continuation ratio model in Example 1.

A.5 Boxplots of PK Parameter Estimates Obtained in Example 1

The boxplots below are for the PK parameter estimates in Example 1 in Section 6.3.

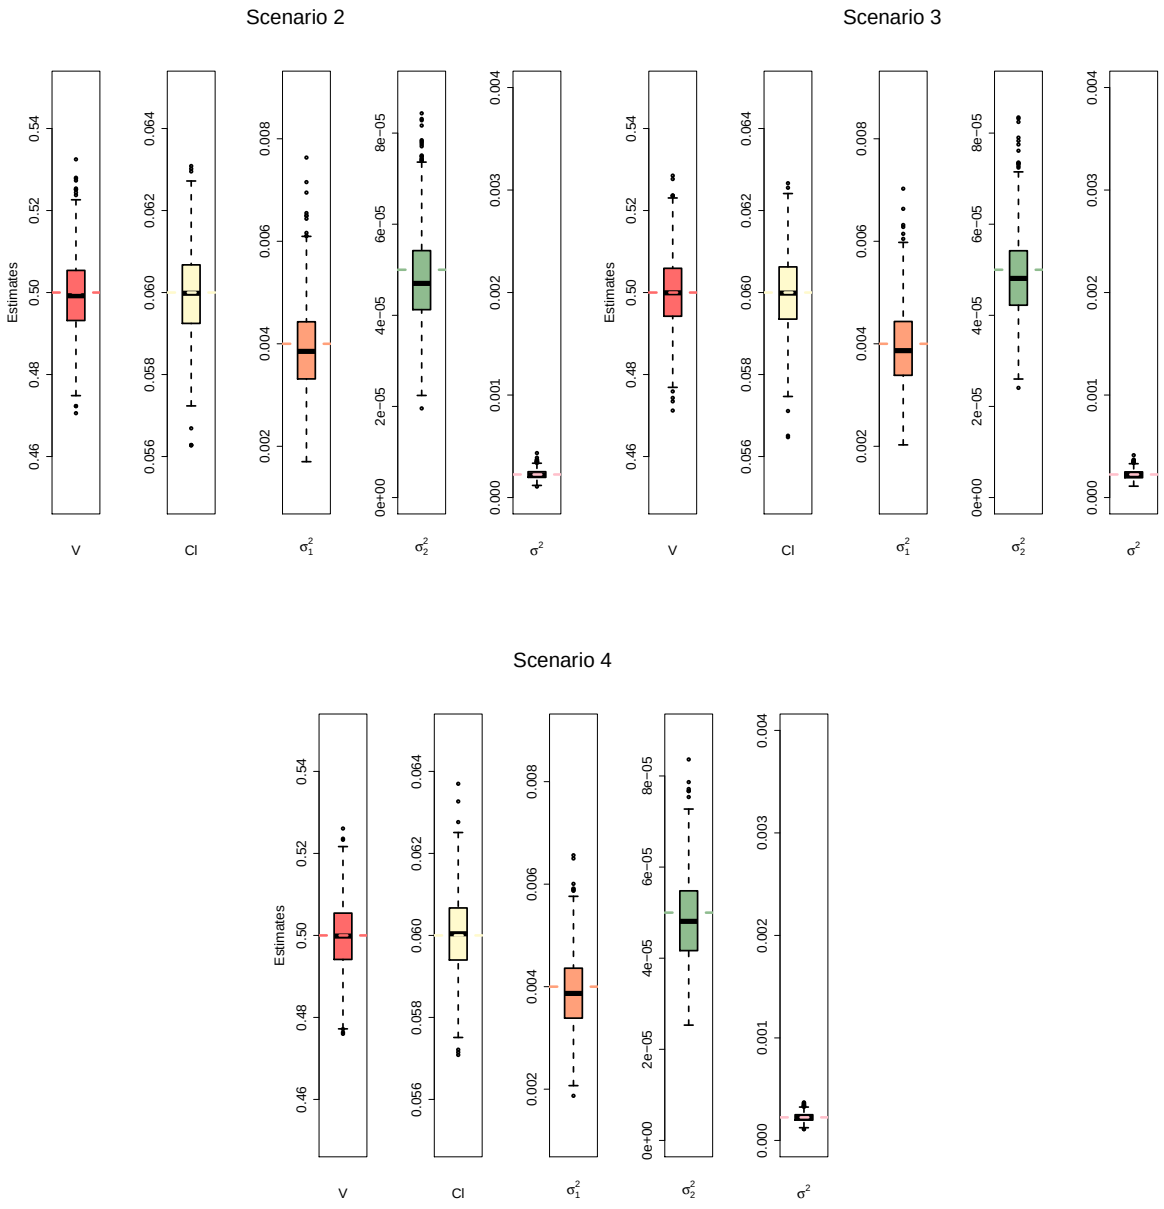

Figure A.3: Boxplots of the PK parameter estimates obtained from the simulations. The horizontal dashed lines indicate the true parameter values.

A.6 Boxplots of Dose-Response Parameter Estimates Obtained in Example 1

The boxplots below are for the dose-response parameter estimates in Example 1 in Section 6.3.

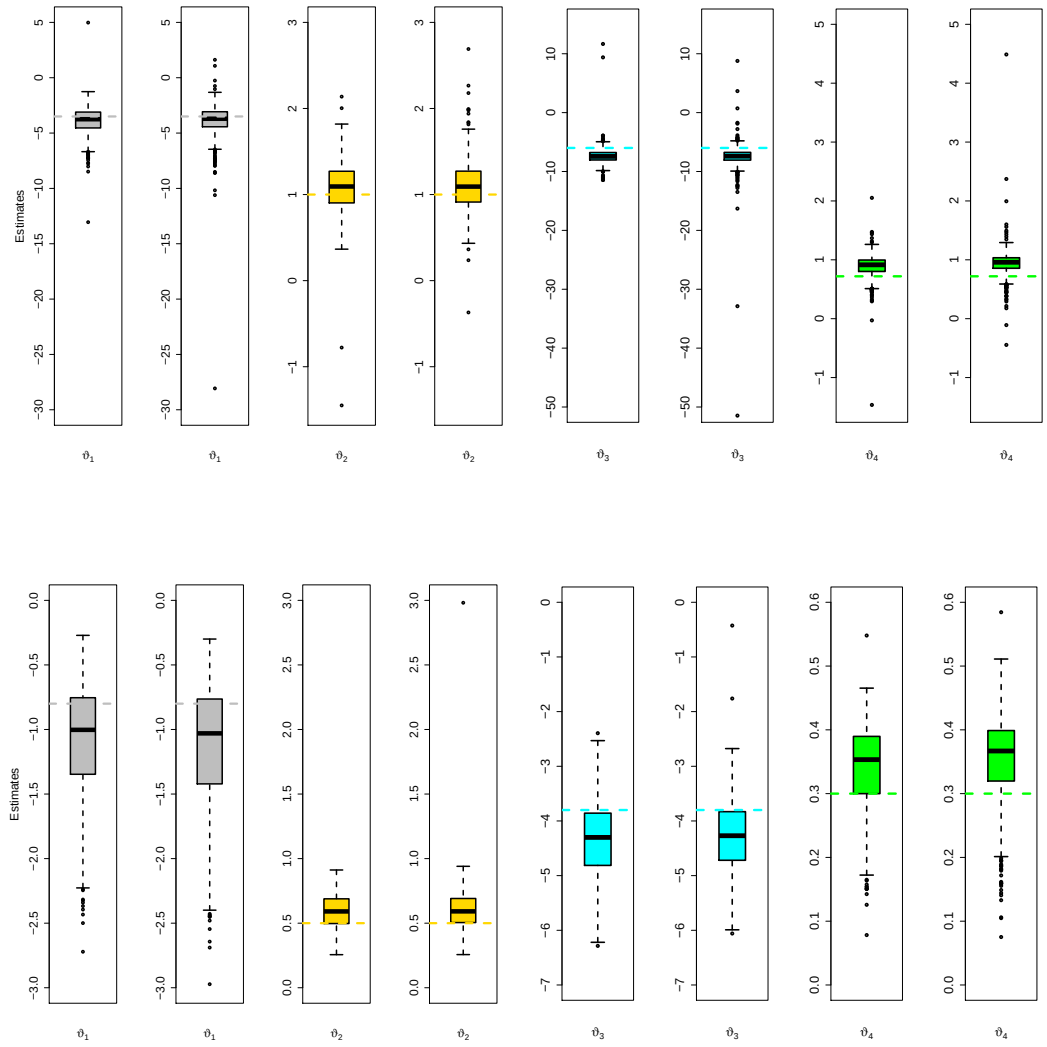

Figure A.4: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenarios 2 and 3. The horizontal dashed lines indicate the true parameter values. For each parameter, the left boxplot corresponds to the design which takes into account the AUC and the right boxplot to the one which ignores it.

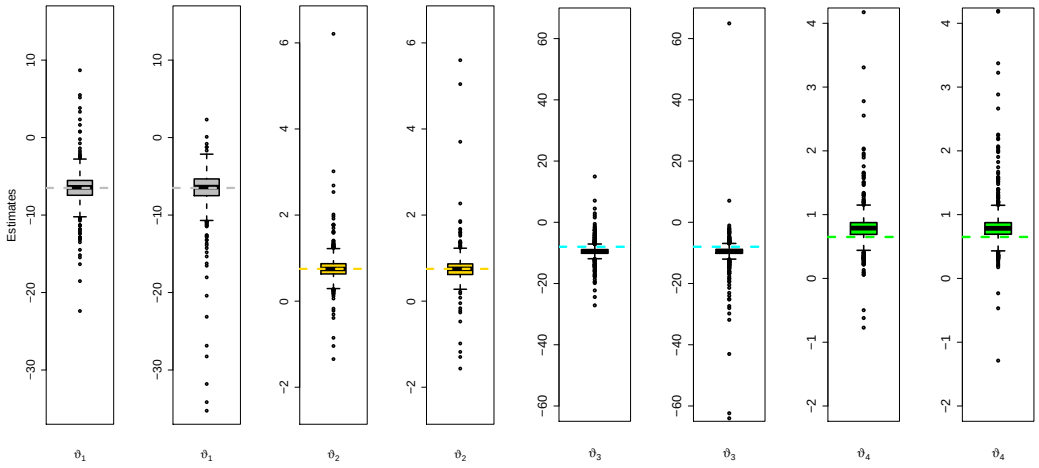

Figure A.5: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenario 4. The horizontal dashed lines indicate the true parameter values. For each parameter, the left boxplot corresponds to the design which takes into account the AUC and the right boxplot to the one which ignores it.

A.7 Efficiency versus Design Points: Example 2

The graphs below are for deciding the number of blood samples to be considered per patient in Section 6.4.1.

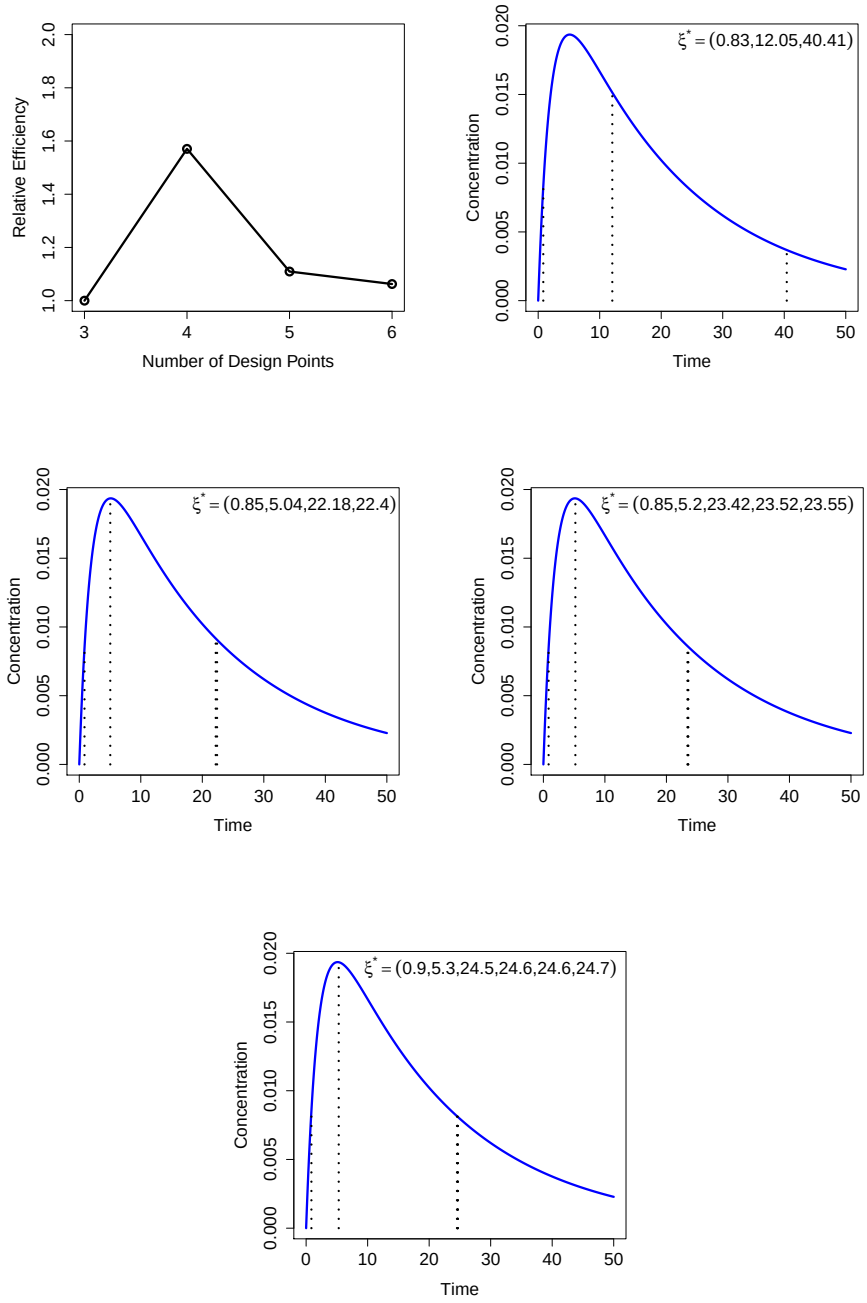

Figure A.6: Rationale for setting the number of design points in the one-compartment PK model with first-order absorption. The locally D-optimum design points are obtained using the initial prior values $\mathbf{\Psi}^0$ assuming that the lowest dose is given to a cohort of 3 patients.

A.8 Boxplots of PK Parameter Estimates Obtained in Example 2

The boxplots below are for the PK parameter estimates in Example 2 in Section 6.4.

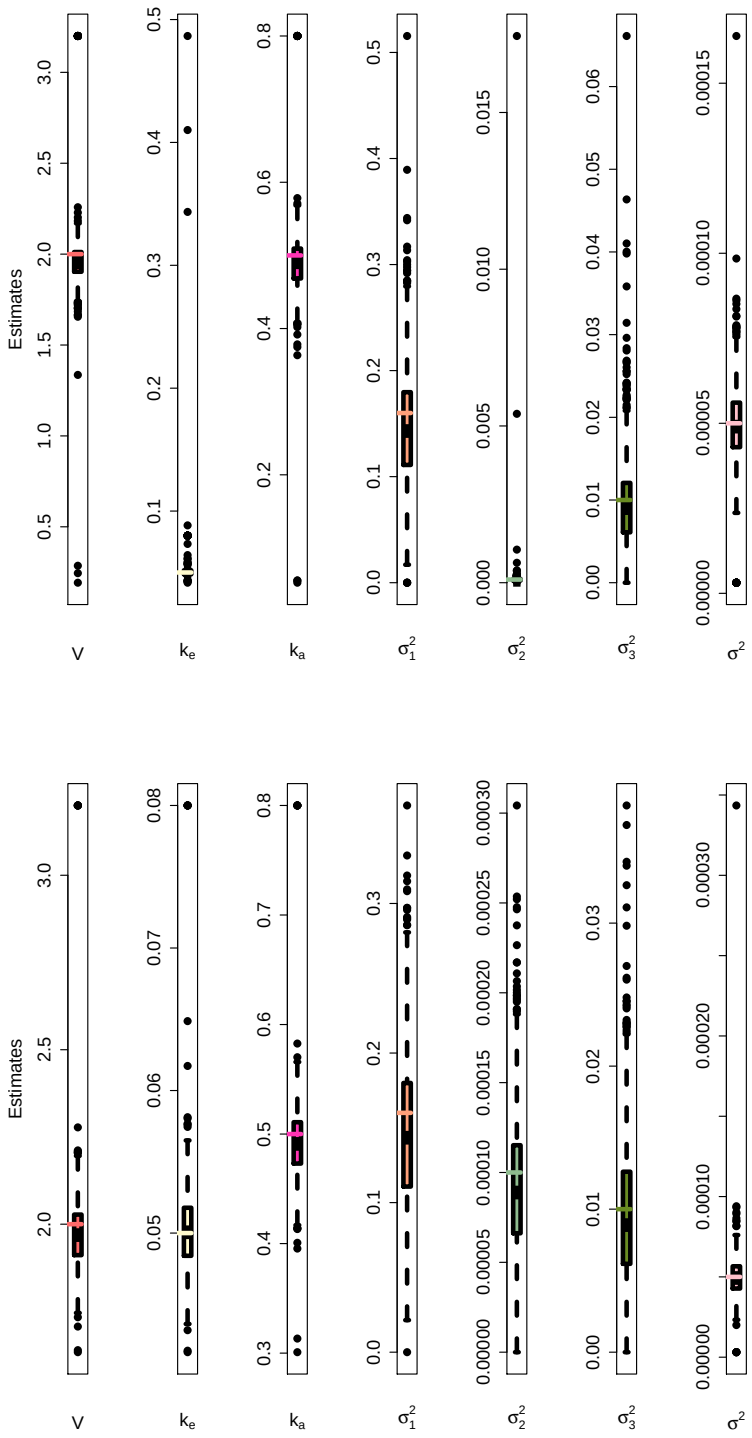

Figure A.7: Boxplots of the PK parameter estimates obtained from the simulations for Scenarios 2 and 3. The horizontal dashed lines indicate the true parameter values.

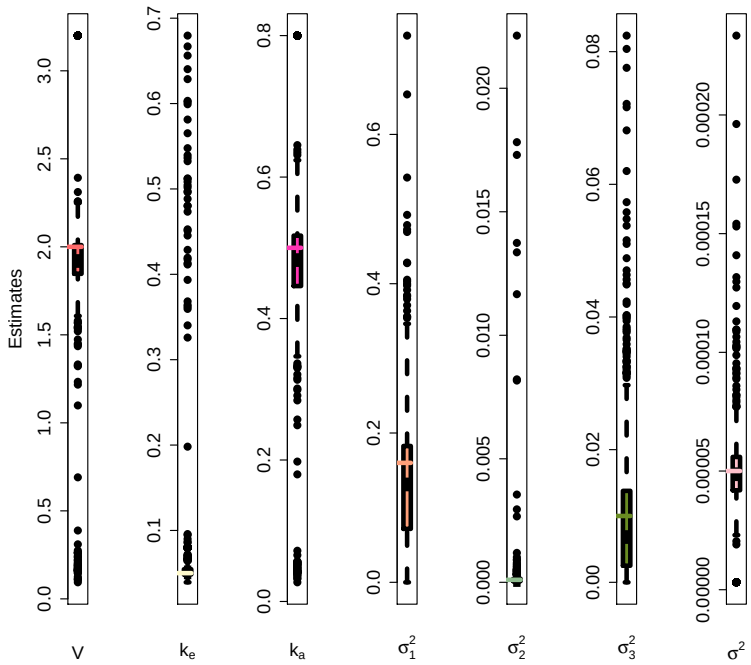

Figure A.8: Boxplots of the PK parameter estimates obtained from the simulations for Scenario 4. The horizontal dashed lines indicate the true parameter values.

A.9 Boxplots of Dose-Response Parameter Estimates Obtained in Example 2

The boxplots below are for the dose-response parameter estimates in Example 2 in Section 6.4.

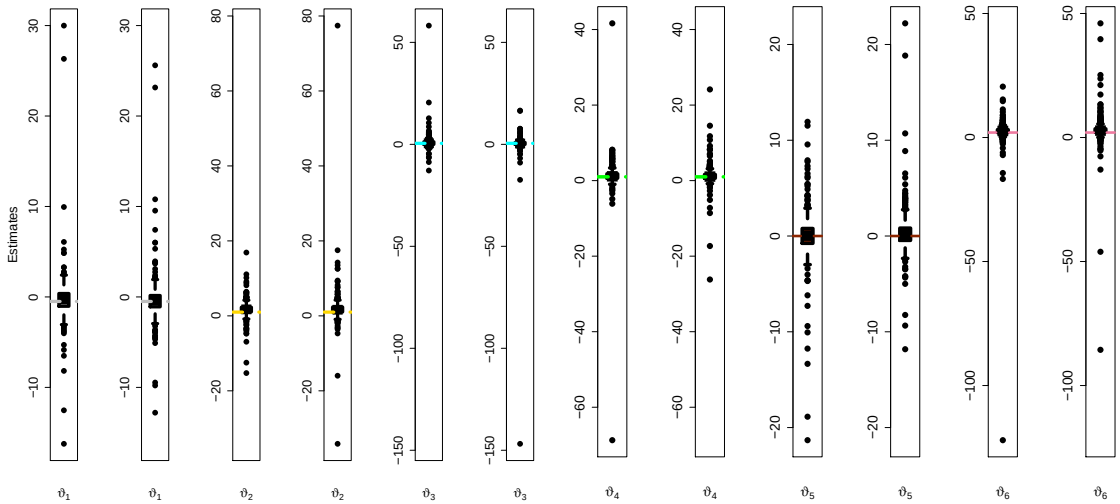

Figure A.9: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenario 2. The horizontal dashed lines indicate the true parameter values.

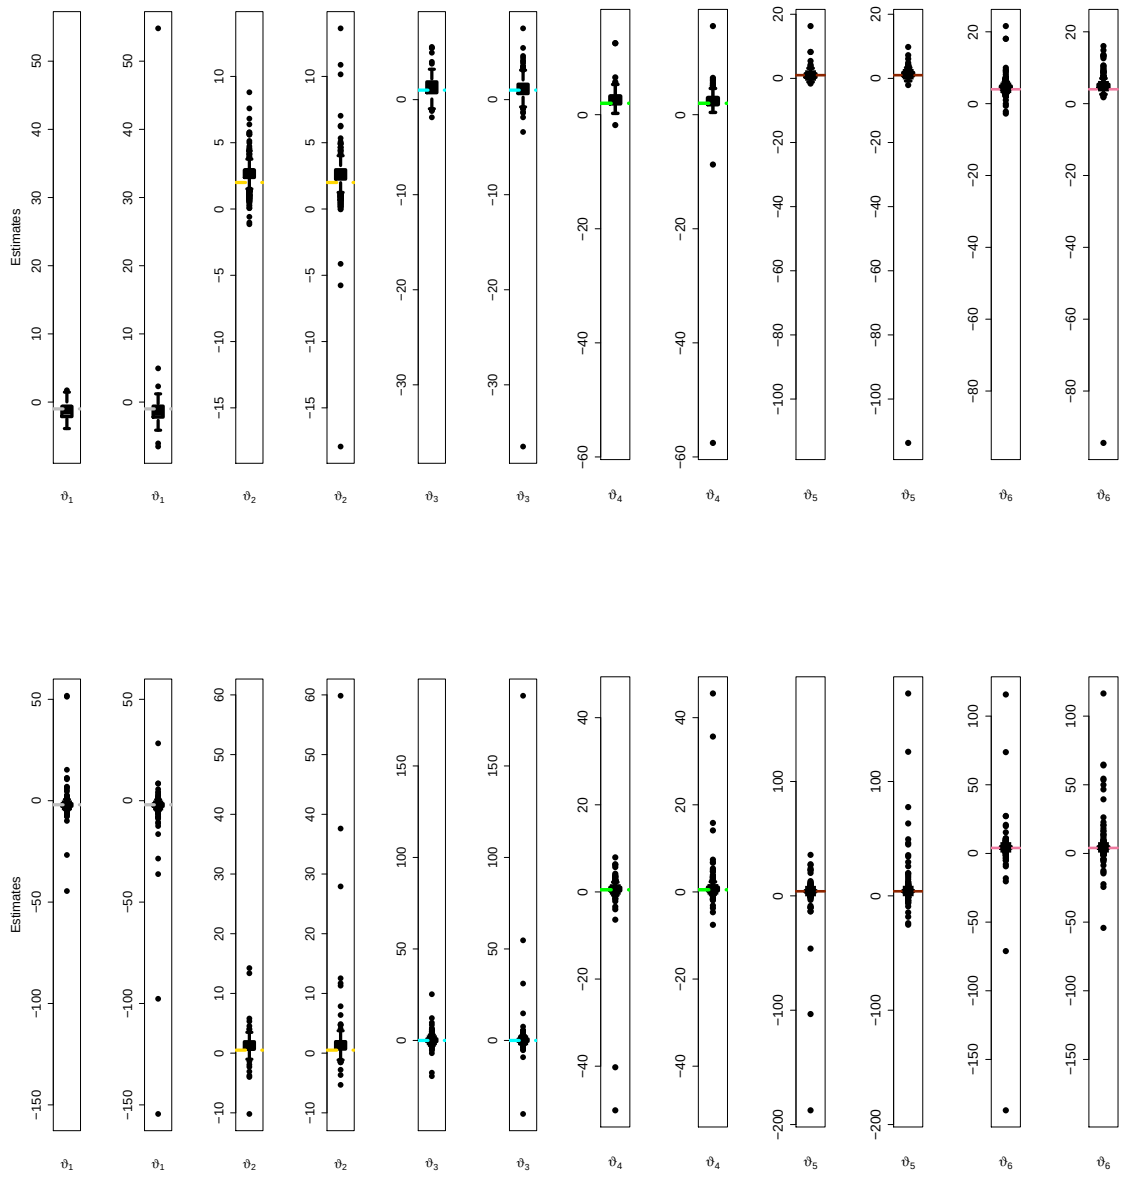

Figure A.10: Boxplots of the dose-response parameter estimates obtained from the simulations for Scenarios 3 and 4. The horizontal dashed lines indicate the true parameter values. For each parameter, the left boxplot corresponds to the design which takes into account the $C_{\max}$ and the right boxplot to the one which ignores it.

A.10 Confidence Intervals for Dose Selections

The tables below present confidence intervals for each of the performance measures in Examples 1 and 2. They correspond to Tables 6.1 and 6.6.

Table A.1: The 95% confidence intervals for the measures in Example 1.

Scenario	Best Doses	%BD		%TD		%AD	
		PK	No PK	PK	No PK	PK	No PK
1	0.5	(98.38, 99.61)	(49.30, 55.49)	(0.12, 1.08)	(29.79, 35.61)	(62.34, 68.25)	(28.91, 34.68)
2	5.5 and 6.0	(77.70, 82.67)	(63.26, 69.13)	(0.31, 1.49)	(7.68, 11.32)	(38.05, 44.14)	(30.57, 36.42)
3	5.5-7.5	(90.00, 93.40)	(83.53, 87.86)	(0.0, 0.0)	(1.53, 3.47)	(49.30, 55.49)	(46.10, 52.29)
4	10.0	(44.80, 50.99)	(43.20, 49.39)	(0.0, 0.0)	(0.0, 0.0)	(15.42, 20.17)	(15.14, 19.85)

Table A.2: The 95% confidence intervals for the measures in Example 2.

Scenario	Best Doses	%BD		%TD		%AD	
		PK	No PK	PK	No PK	PK	No PK
1	-0.6	(59.90, 65.90)	(50.90, 57.08)	(1.13, 2.87)	(4.17, 7.02)	(34.49, 40.50)	(29.69, 35.50)
2	-0.6	(60.11, 66.09)	(43.01, 49.19)	(0.75, 2.25)	(19.14, 24.25)	(32.63, 38.56)	(24.15, 29.64)
3	-0.6	(93.42, 96.17)	(80.78, 85.42)	(0.0, 0.0)	(9.52, 13.48)	(49.30, 55.49)	(40.82, 46.97)
4	-1.8 and -1.2	(78.25, 83.14)	(86.20, 90.19)	(0.0[a], 0.30)	(0.45, 1.75)	(59.19, 65.21)	(63.67, 69.52)

[a]Found to be -0.09 and therefore rounded to 0.0 as a negative value is not plausible.

Appendix B

R Code

B.1 *R* Program

Below is the R code for simulating designs under various dose-optimisation criteria and constraints introduced in Chapter 5.

```
#####################################################
################ Calling PFIM #######################
#####################################################

source("E:\\PKmodels\\elimination\\optimisation\\PFIM3.2.r")

#main programme

source("E:\\PKmodels\\elimination\\optimisation\\model.r")

#first-order elimination model (model 1)

source("E:\\PKmodels\\absorption\\optimisation\\PFIM3.2.r")

#main programme

#source("E:\\PKmodels\\absorption\\optimisation\\model.r")

#first-order absorption model (model 2)

#if working directory changes, directory

#in PFIM3.2.r needs to be changed

#####################################
###### Libraries to be used ########
#####################################

library(mvtnorm)

library(nlme)

library(deSolve)

library(lattice)

library(cubature)
```

```
library(compiler)


#####################################################
####### Functions to be used in the program #######
#####################################################

source("E:\\Programme functions\\functions.R")

#####################################################
##############  Design parameters #################
#####################################################

lower.dose<-0.5

upper.dose<-10

increment<-0.5

dose.scale<-"not log"

#first eg.

#lower.dose<--3

#upper.dose<-3

#increment<-0.6

#dose.scale<-"log"

#second eg.

true.dose<-seq(lower.dose,upper.dose,increment)

#sequence of doses

cohort.size<-3

trial.cohort<-20 #number of cohorts in a trial

stop.freq<-6

#a trial stops early when the same dose is repeated r times (r=6)

dose.skip<-2

#to ensure a dose not more than 2 level higher

up.down<-4

#no. of cohorts in the up-and-down procedure

#########################################################
##############   Methods and criteria   #################
#########################################################

scenario<-1

nsim<-1000

#number of simulations to run

######## Include PK?  ############

method<-"PK"

#method<-"not PK"

############ Dose-response model ############
```

```
doseresponse.model<-"CR"
#doseresponse.model<-"Cox"
########### PK model ###########
PK.model<-1
#bolus input and first-order elimination
#PK.model<-2
#first-order absorption
########  Dose-optimisation criteria   ##########
criterion<-"maximisation of prob. of success"
#criterion<-"combined"
######### PK constraint #############
# PK.constraint<-"none"
PK.constraint<-"auc"
#PK.constraint<-"cmax"
###### Acceptable level for probability of toxicity #######
target.toxicity<-0.20
#target.toxicity<-0.33
#### Control parameters and weight for the combined criterion ####
#CS<-1
#CT<-1
#a<-0.6
#weight in the combined criterion
#a=1 means penalised D-criterion and a=0 means
#max. of prob. of success
######## Thresholds in up-and-down design ########
pL<-target.toxicity/3
pM<-(2*target.toxicity)/3
pU<-target.toxicity
## Tolerance limit and max. function evaluation ##
tol.limit<-1e-3
#tolerance limit for the integration
maxEval.limit<-5000
#maximum number of function evaluations needed
##### True PK parameters for simulation #####
if(method=="PK") {
if(PK.model==1) {
```

```
lower.sampling<-0  #sampling region for PK
upper.sampling<-30
#sampling time for PK (in hours)
fixed<-c(0.5,0.06) #fixed PK parameters: (v,cl)
var.comp<-c(0.004,0.00005) #variance-components of the parameters
error.variance<-0.000225 #error variance
                }
if(PK.model==2) {
lower.sampling<-0
upper.sampling<-50
#sampling time for PK (in hours)
#theta=(v,ke,ka)
fixed<-c(2.0,0.05,0.50)
var.comp<-c(0.1600, 0.0001, 0.0100) # CV is exactly 20%
error.variance<-0.00005 #since my obs. are also tiny (STD is 0.007)
                }
                 }
#########################################################
### True dose-response parameters for simulation  ###
#########################################################
############## CR model ##############
if(scenario==1 & doseresponse.model=="CR") {
theta1.true<-1.44
theta2.true<-0.26
theta3.true<--1.70
theta4.true<-0.25
lower.lim<-c(-3.4,0,-3.4,0)
upper.lim<-c(2.88,0.52,2.88,0.50)

                                          }
if(scenario==2 & doseresponse.model=="CR") {
theta1.true<--3.5
theta2.true<-1.0
theta3.true<--6
theta4.true<-0.72
lower.lim<-c(-12,0,-12,0)
upper.lim<-c(0,2,0,1.44)
```

```
                                                }
if(scenario==3 & doseresponse.model=="CR") {
theta1.true<--0.80
theta2.true<-0.50
theta3.true<--3.80
theta4.true<-0.30
lower.lim<-c(-7.6,0,-7.6,0)
upper.lim<-c(0,1,0,0.60)

                                                }
if(scenario==4 & doseresponse.model=="CR") {
theta1.true<--6.50
theta2.true<-0.75
theta3.true<--8.00
theta4.true<-0.65
lower.lim<-c(-12,0,-12,0)
upper.lim<-c(0,1.5,0,1.30)

                                                }
################## Cox model ##################
if(scenario==1 & doseresponse.model=="Cox") {
theta1.true<-0
theta2.true<-1
theta3.true<-4
theta4.true<-2
theta5.true<-3
theta6.true<-3
lower.lim<-c(-3,-2,1,-1,0,0)
upper.lim<-c(3,4,7,5,6,6)
#margin: 3
                                                }
if(scenario==2 & doseresponse.model=="Cox") {
theta1.true<--0.5
theta2.true<-1.0
theta3.true<-0.5
theta4.true<-1.0
```

```
theta5.true<-0.0
theta6.true<-2.0
lower.lim<-c(-3.5,-2,-2.5,-2,-3,-1)
upper.lim<-c(2.5,4,3.5,4,3,5)
#margin: 3
                                                }
if(scenario==3 & doseresponse.model=="Cox") {
theta1.true<--1.0
theta2.true<-2.0
theta3.true<-1.0
theta4.true<-2.0
theta5.true<-1.0
theta6.true<-4.0
lower.lim<-c(-4,-1,-2,-1,-2,1)
upper.lim<-c(2,5,4,5,4,7)
#margin: 3
                                                }
if(scenario==4 & doseresponse.model=="Cox") {
theta1.true<--2.0
theta2.true<-0.5
theta3.true<-0.0
theta4.true<-0.5
theta5.true<-4.0
theta6.true<-4.0
lower.lim<-c(-5,-2.5,-3,-2.5,1,1)
upper.lim<-c(1,3.5,3,3.5,7,7)
#margin: 3
                                                }
#####  Initial values for a trial to start with  #####
starting.dose.level<-1
if(method=="PK") {
if(PK.model==1) {
ini.pkpara<-c(0.1,0.005) #PK parameters
ini.varcomp<-c(0.0007,0.0000006) #variance-components
ini.error.std<-0.002  #error std.
```

```
ini.time<-c(1,5,20) #time points. sampling region [0,30]
                 }
if(PK.model==2) {
#PK parameters : (v,ke,ka)
ini.pkpara<-c(3.20, 0.08, 0.80) #fixed+3*sqrt(var.comp)
ini.varcomp<-c(4.0e-02, 2.5e-05, 2.5e-03) #var.comp/4
ini.error.std<-0.003535534  #error std. #sqrt(error.variance/4)
ini.time<-c(1,15,30,40)   #time points. sampling region [0,50]
                 }
                   }
##### True probabilities of dose-response  #####
##### outcomes under a specific scenario   #####
if(doseresponse.model=="CR") {
true.neu<-psi0(theta1.true,theta2.true,theta3.true,theta4.true,
               true.dose)
true.effi<-psi1(theta1.true,theta2.true,theta3.true,theta4.true,
                true.dose)
true.toxic<-psi2(theta3.true,theta4.true,true.dose)
                              }
if(doseresponse.model=="Cox") {
true.psi00<-psi00(theta1.true,theta2.true,theta3.true,theta4.true,
                  theta5.true,theta6.true, true.dose)
true.psi01<-psi01(theta1.true,theta2.true,theta3.true,theta4.true,
                  theta5.true,theta6.true,true.dose)
true.effi<-psi10(theta1.true,theta2.true,theta3.true,theta4.true,
                 theta5.true,theta6.true,true.dose)
#true prob. of success
true.psi11<-psi11(theta1.true,theta2.true,theta3.true,theta4.true,
                  theta5.true,theta6.true,true.dose)
#true.effi<-psi1.(theta1.true,theta2.true,theta3.true,theta4.true,
                  theta5.true,theta6.true,true.dose)
#marginal prob. of efficacy. no need
true.toxic<-psi.1(theta1.true,theta2.true,theta3.true,theta4.true,
                  theta5.true,theta6.true,true.dose)
#marginal prob. of toxicity
                              }
```

```
##### Identifying the true optimal dose #####
index.safe<-which(true.toxic<=target.toxicity)
#index of safe doses
candi.true.effi<-true.effi[index.safe]
true.efficacy.od<-max(candi.true.effi)
#prob.of success at the true OD
true.od.level<-which(true.effi==true.efficacy.od &
                     true.toxic<=target.toxicity)
true.od<-true.dose[true.od.level]
#true OD
true.toxicity.od<-true.toxic[true.od.level]
#prob. of toxicity at the true OD
##### Identifying the target AUCs #####
if(PK.constraint=="auc") {
if(PK.model==1) {
auc.true<-auc1(fixed[1],fixed[2],upper.sampling,true.dose)
#first-order elimination
                }
if(PK.model==2) {
auc.true<-auc2(fixed[1],fixed[2],fixed[3],upper.sampling,true.dose)
#first-order absorption
                }
auc.lim<-auc.true[true.od.level]
#value for target AUC
                        }
######## Identifying the target Cmax ########
if(PK.constraint=="cmax") {
if(PK.model==2) {
tmax.true<-(log(fixed[3])-log(fixed[2]))/(fixed[3]-fixed[2])
cmax.true<-cmax2(fixed[1],fixed[2],fixed[3],tmax.true,true.dose)
#first-order absorption
                }
cmax.lim<-cmax.true[true.od.level]
#value for target Cmax
                         }
```

```
##########################################################
##########   Declaring variables to store the   ##########
##########   estimates over the simulations     ##########
##########################################################

pk1.hat=c()

pk2.hat=c()

pk3.hat=c()

varcomp1.hat=c()

varcomp2.hat=c()

varcomp3.hat=c()

stddev.error.hat=c()

#to store PK parameter estimates

od=c()

#to store optimal doses

theta1.hat=c()

theta2.hat=c()

theta3.hat=c()

theta4.hat=c()

theta5.hat=c()

theta6.hat=c()

#to store dose-response parameter estimates

count.stops<-0

#counts trial with early stopping (achieving r)

count.stops.updown<-0 #counts trial stopped at up-and-down stage

updown.dose=c() #dose at which a trial stops in up-and-down stage

updown.cohort=c() #cohorts used in up-and-down stage

od.complete=c() #from the complete analysis

stopped.at=c() #no. of cohorts used in a trial when it stops early

stopped.dose=c() #optimum dose selected in such a trial

##########################################################
##############    Start of the Program     ##############
##########################################################

for(l in 1:nsim) {

########### Values to start with ###########

h<-starting.dose.level  #starting dose level

d<-1  #to create an index for the subjects

if(method=="PK") {

pk.esti<-ini.pkpara
```

```
var.comp.esti<-ini.varcomp
error.esti<-ini.error.std
# initial values for PK parameters
                  }
next.dose<-true.dose[h] #dose for a trial to start with
previ.level<-h #to be used in the skip calculation
stage<-1 #stage of a trial
######### End of values to start with ###########
if(doseresponse.model=="CR") {
alloc.first<-true.dose[h]
tri.res<-rmultinom(1,size=cohort.size,prob=c(true.neu[h],
                  true.effi[h],true.toxic[h]))
#generating dose-response outcomes for the starting dose
r0.1<-tri.res[1,1] #neutral response
r1.1<-tri.res[2,1] #efficacious response
r2.1<-tri.res[3,1] #toxic response
alloc.dose=c(alloc.first) #allocated doses to cohorts in a trial
r0=c(r0.1) #neutral outcomes over cohorts in a trial
r1=c(r1.1) #efficacious outcomes over cohorts in a trial
r2=c(r2.1) #toxic outcomes over cohorts in a trial
                              }
if(doseresponse.model=="Cox") {
alloc.first<-true.dose[h]
dose.res<-rmultinom(1,size=cohort.size,prob=c(true.psi00[h],
                   true.psi01[h],true.effi[h],true.psi11[h]))
#dose-response outcomes for the starting dose
r0.1<-dose.res[1,1] #(0,0)
r1.1<-dose.res[2,1] #(0,1)
r2.1<-dose.res[3,1] #(1,0)
r3.1<-dose.res[4,1] #(1,1)
alloc.dose=c(alloc.first) #allocated doses to cohorts in a trial
r0=c(r0.1) #(0,0) outcomes over cohorts in a trial
r1=c(r1.1) #(0,1) outcomes over cohorts in a trial
r2=c(r2.1) #(1,0) outcomes over cohorts in a trial
r3=c(r3.1) #(1,1) outcomes over cohorts in a trial
                              }
```

```
stop.iter<-0
for(i in 1:trial.cohort) {
#################################################################
#############    Generation of PK responses    ##############
#################################################################
if(method=="PK") {
con.cohort=c()
#concentrations obtained from a cohort
#after receiving a specific dose
if(PK.model==1) {
source("E:\\PKmodels\\elimination\\optimisation\\stdin.r")
#input file
res<-tryCatch(PFIM(),error=function(e)NULL)
if(!is.null(res)) {
time<-round(c(res$prot.opti[[1]][[1]][[1]][1],res$prot.opti[[1]]
        [[1]][[2]][1],res$prot.opti[[1]][[1]][[3]][1]),digits=2)
#extracting the optimal time points
                  } else {
stop.iter<-stop.iter+1
                          }
                 }
       #end of (PK.model==1)
if(PK.model==2) {
source("E:\\PKmodels\\absorption\\optimisation\\stdin.r")
res<-tryCatch(PFIM(),error=function(e)NULL)
if(!is.null(res)) {
time<-round(c(res$prot.opti[[1]][[1]][[1]][1],res$prot.opti[[1]]
[[1]][[2]][1],res$prot.opti[[1]][[1]][[3]][1],res$prot.opti[[1]]
[[1]][[4]][1]),digits=2)
                   } else {
stop.iter<-stop.iter+1
                           }
                 }
        #end of (PK.model==2)
k<-1
for(a in 1:cohort.size) {
```

```
mean.b<-rep(0,length(beta))
varcov.b<-diag(var.comp,nrow=length(mean.b),ncol=length(mean.b))
b<-rmvnorm(1,mean.b,varcov.b) #random effects for an individual
theta<-fixed+b  #PK parameters for an individual
x<-true.dose[h]
if(dose.scale=="log") {
x<-exp(x)
                      }
Theta1<-theta[1,1]
Theta2<-theta[1,2]
f=c()
for(g in 1:length(time)) {
if(PK.model==1) {
f[g]<-f1(x,Theta1,Theta2,time[g])
                } else {
Theta3<-theta[1,3]
f[g]<-f2(x,Theta1,Theta2,Theta3,time[g])
                       }
                          }
mean.error<-rep(0,length(time))
varcov.error<-diag(error.variance,nrow=length(mean.error),
ncol=length(mean.error))
epsi<-rmvnorm(1,mean.error,varcov.error)
y<-f+epsi   #concentration for an individual
s<-1
for(u in k:(k+length(time)-1)) {
con.cohort[u]<-y[1,s]
s<-s+1
                                }
k<-k+length(time)
                       }
        #end of the loop with index 'a'
subject<-c(rep(d,length(time)),rep(d+1,length(time)),
           rep(d+2,length(time)))
dose<-rep(x,cohort.size*length(time))
t<-rep(time,cohort.size)
```

```
data.cohort=data.frame(subject=subject,dose=dose,t=t,
                       conc=con.cohort)
if(d==1) {
data.d<-data.cohort
         } else {
data.d<-rbind(data.d,data.cohort)
               }
d<-d+cohort.size
                  }
       #end of (method=="PK")
################################################
##  Dose selection based on up-and-down design  ##
################################################
if(stage<up.down) {
if((stage==(up.down-1))&(all(alloc.dose[1]==alloc.dose))) {
h<-h+1
                                                          } else {
if(doseresponse.model=="CR") {
prop.toxic<-(sum(r2)/(cohort.size*stage))
                              }
if(doseresponse.model=="Cox") {
prop.toxic<-(sum(r1+r3)/(cohort.size*stage))
                               }
if(prop.toxic<=pL) {
if(h==length(true.dose)) {
h<-h                     } else {
h<-h+1                          }
                   }
else if((prop.toxic>pL)&(prop.toxic<pM)) {
h<-h
                                        }
else if((prop.toxic>=pM)&(prop.toxic<pU)) {
if(h==1) {
h<-h     } else {
h<-h-1
               }
```

```
                                          } else {
count.stops.updown<-count.stops.updown+1
updown.dose<-c(updown.dose,true.dose[h])
updown.cohort<-c(updown.cohort,stage)
break #breaks the trial to the end
                                                  }
                                                                    }
                  } else {
     ##  start of stage>up.down: model-based approach  ##
if(method=="PK") {
#applies when PK information is considered
#########################################################
#############   Fitting the PK  model   ###############
#########################################################
if(PK.model==1) {
grouped.data<-groupedData(formula=conc~t|subject,data=data.d)
model.d<-tryCatch(nlme(conc~f1(dose,Theta1,Theta2,t),
fixed=Theta1+Theta2~1,data=grouped.data,
random=Theta1+Theta2~1,start=list(fixed=pk.esti))
,error=function(e)NULL)
if(!is.null(model.d)) {
summ<-summary(model.d)
pk.esti<-c(summ$coefficients$fixed[[1]],
           summ$coefficients$fixed[[2]])
var.comp.esti<-c(as.numeric(VarCorr(summ)[1,1]),
                 as.numeric(VarCorr(summ)[2,1]))
error.esti<-summ$sigma } else {
stop.iter<-stop.iter+1
                               }

                 }
        #end of (PK.model==1)
if(PK.model==2) {
grouped.data<-groupedData(formula=conc~t|subject,data=data.d)
model.d<-tryCatch(nlme(conc~f2(dose,Theta1,Theta2,Theta3,t),
fixed=Theta1+Theta2+Theta3~1,data=grouped.data,
random=Theta1+Theta2+Theta3~1,start=list(fixed=pk.esti))
```

```
,error=function(e)NULL)
if(!is.null(model.d)) {
summ<-summary(model.d)
pk.esti<-c(summ$coefficients$fixed[[1]]
,summ$coefficients$fixed[[2]],summ$coefficients$fixed[[3]])
var.comp.esti<-c(as.numeric(VarCorr(summ)[1,1]),
as.numeric(VarCorr(summ)[2,1]),as.numeric(VarCorr(summ)[3,1]))
error.esti<-summ$sigma } else {
stop.iter<-stop.iter+1
                              }

                 }
        #end of (PK.model==2)
                 }
       #end of(method=="PK")
#################################################################
##############  Fitting the dose-response model  ################
#################################################################
if(doseresponse.model=="CR") {
denomit.integ<-adaptIntegrate(lf1.cond,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
#integrates the likelihood function
theta1.integ<-adaptIntegrate(theta1.f1,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta2.integ<-adaptIntegrate(theta2.f1,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta3.integ<-adaptIntegrate(theta3.f1,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta4.integ<-adaptIntegrate(theta4.f1,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta1<-(theta1.integ$integral/denomit.integ$integral)
theta2<-(theta2.integ$integral/denomit.integ$integral)
theta3<-(theta3.integ$integral/denomit.integ$integral)
theta4<-(theta4.integ$integral/denomit.integ$integral)
#posterior estimates
                              }
if(doseresponse.model=="Cox") {
denomit.integ<-adaptIntegrate(lf2.c,lowerLimit=lower.lim,
```

```
       upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
#integrates the likelihood function
theta1.integ<-adaptIntegrate(theta1.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta2.integ<-adaptIntegrate(theta2.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta3.integ<-adaptIntegrate(theta3.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta4.integ<-adaptIntegrate(theta4.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta5.integ<-adaptIntegrate(theta5.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta6.integ<-adaptIntegrate(theta6.f2,lowerLimit=lower.lim,
        upperLimit=upper.lim,tol=tol.limit,maxEval=maxEval.limit)
theta1<-(theta1.integ$integral/denomit.integ$integral)
theta2<-(theta2.integ$integral/denomit.integ$integral)
theta3<-(theta3.integ$integral/denomit.integ$integral)
theta4<-(theta4.integ$integral/denomit.integ$integral)
theta5<-(theta5.integ$integral/denomit.integ$integral)
theta6<-(theta6.integ$integral/denomit.integ$integral)
#posterior estimates
                              }
############################################################
###########  Dose selection for the next cohort  ###########
############################################################
if(doseresponse.model=="CR") {
psi1.hat<-psi1(theta1,theta2,theta3,theta4,true.dose)
#estimates of prob. of success at different doses
psi2.hat<-psi2(theta3,theta4,true.dose)
#estimates of prob. of toxicity at different doses
                             }
if(doseresponse.model=="Cox") {
psi1.hat<-psi10(theta1,theta2,theta3,theta4,theta5,theta6,true.dose)
#estimates of prob. of success at different doses
psi2.hat<-psi.1(theta1,theta2,theta3,theta4,theta5,theta6,true.dose)
#estimates of prob. of toxicity at different doses
```

```
                                }
#########################################################
#### Criterion: maximisation of prob. of success #####
#########################################################
if (criterion=="maximisation of prob. of success" ) {
if(method=="not PK") {
index<-which(psi2.hat<=target.toxicity)
#identifying indexes of the doses for which the condition is met
                     } else {
                    #when (method=="PK")
if(PK.model==1) {
if(PK.constraint=="auc") {
auc.hat<-auc1(pk.esti[1],pk.esti[2],upper.sampling,true.dose)
#AUC estimates at various doses
c.v.f1.esti<-c.v.f1(pk.esti[1],pk.esti[2],upper.sampling,next.dose)
c.cl.f1.esti<-c.cl.f1(pk.esti[1],pk.esti[2],upper.sampling,next.dose)
c.var.esti<-(((c.v.f1.esti^2)*var.comp.esti[1])+((c.cl.f1.esti^2)*
               var.comp.esti[2]))
#variance of AUC
                         }
                 }
if(PK.model==2) {
if(PK.constraint=="auc") {
auc.hat<-auc2(pk.esti[1],pk.esti[2],pk.esti[3],upper.sampling,
              true.dose)
c.v.f2.esti<-c.v.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                    upper.sampling,next.dose)
c.ke.f2.esti<-c.ke.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                      upper.sampling,next.dose)
c.ka.f2.esti<-c.ka.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                      upper.sampling,next.dose)
c.var.esti<-(((c.v.f2.esti^2)*var.comp.esti[1])+((c.ke.f2.esti^2)*
            var.comp.esti[2])+((c.ka.f2.esti^2)*var.comp.esti[3]))
#variance of AUC
                          } else {
                 #when (PK.constraint=="cmax")
tmax.hat<-(log(pk.esti[3])-log(pk.esti[2]))/(pk.esti[3]-pk.esti[2])
```

```
cmax.hat<-cmax2(pk.esti[1],pk.esti[2],pk.esti[3],tmax.hat,true.dose)
cmax.v.f2.esti<-cmax.v.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                          tmax.hat,next.dose)
cmax.ke.f2.esti<-cmax.ke.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                            tmax.hat,next.dose)
cmax.ka.f2.esti<-cmax.ka.f2(pk.esti[1],pk.esti[2],pk.esti[3],
                            tmax.hat,next.dose)
cmax.var.esti<-(((cmax.v.f2.esti^2)*var.comp.esti[1])+
                ((cmax.ke.f2.esti^2)*var.comp.esti[2])+
                ((cmax.ka.f2.esti^2)*var.comp.esti[3]))
#variance of Cmax
                                    }
                }
if(PK.constraint=="auc") {
c.std.esti<-sqrt(c.var.esti)
rel.c<-(auc.hat-auc.lim)/c.std.esti

                          } else {
cmax.std.esti<-sqrt(cmax.var.esti)
rel.c<-(cmax.hat-cmax.lim)/cmax.std.esti

                                  }
delta<-(1/psi1.hat[previ.level])
index<-which(psi2.hat<=target.toxicity & rel.c<=delta)
#indexes of the doses for which both conditions are met
                              }
                 #end of (method=="PK")
if(length(index)==0) {
h<-1
                       } else {
candi.psi1.hat<-psi1.hat[index]
#extracting efficacy estimates for these doses
h<-which(psi1.hat==max(candi.psi1.hat))
#the dose level with maximum efficacy estimate
                              }
                                                  } else {
```

```
                        #end of first dose-optimsation criterion

############################################################
#############  Criterion: combined criterion  ##############
############################################################

m.mat<-matrix(c(0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0),

               nrow=4,ncol=4,byrow="TRUE")

#initial 'M' matrix

for(i in 1:length(alloc.dose)){

x<-alloc.dose[i]

I11.x<-I11(theta1,theta2,theta3,theta4,x)

I12.x<-I12(theta1,theta2,theta3,theta4,x)

I22.x<-I22(theta1,theta2,theta3,theta4,x)

I33.x<-I33(theta3,theta4,x)

I34.x<-I34(theta3,theta4,x)

I44.x<-I44(theta3,theta4,x)

I<-matrix(c(I11.x,I12.x,I13,I14,I12.x,I22.x,I23,I24,I31,I32,I33.x,

            I34.x,I41,I42,I34.x,I44.x),nrow=4,ncol=4,byrow=TRUE)

#FIM for an individual

m.mat<-m.mat+(cohort.size*I)

#contructs the matrix 'M'

                              }

penalty.upto<-penalty(theta1,theta2,theta3,theta4,alloc.dose)

#penalty at the adminstered doses

penalty.sum<-sum(penalty.upto)

#sum of penalties at the administered doses

#formation of the objective function

combined.cri<-function(x) {

I11.x<-I11(theta1,theta2,theta3,theta4,x)

I12.x<-I12(theta1,theta2,theta3,theta4,x)

I22.x<-I22(theta1,theta2,theta3,theta4,x)

I33.x<-I33(theta3,theta4,x)

I34.x<-I34(theta3,theta4,x)

I44.x<-I44(theta3,theta4,x)

I<-matrix(c(I11.x,I12.x,I13,I14,I12.x,I22.x,I23,I24,I31,I32,I33.x,

            I34.x,I41,I42,I34.x,I44.x),nrow=4,ncol=4,byrow=TRUE)

#FIM for an individual

penalty.x<-penalty(theta1,theta2,theta3,theta4,x)
```

```
det(((stage/(stage+1))*m.mat+(1/(stage+1))*cohort.size*I)/
    ((stage/(stage+1))*penalty.sum+(1/(stage+1))*penalty.x))
# penalised D-criterion is here
                          }
combined.cri.values=c()
x.dose<-lower.dose
for (i in 1:length(true.dose)) {
combined.cri.values[i]<-combined.cri(x.dose)
x.dose<-x.dose+increment
                               }
scaled.d<-combined.cri.values/max(combined.cri.values)
scaled.psi1<-psi1.hat/max(psi1.hat)
#standardising the values
cri.values<-a*scaled.d+(1-a)*scaled.psi1
#combining two criterion
if(stage<trial.cohort) {
index<-which(cri.values==max(cri.values))
if(length(index)==0) {
h<-1
                     } else {
h<-index
                            }
                       } else {
#when ("stage==trial.chort"), the last estimates
#are used to find the OD
index<-which(psi2.hat<=target.toxicity)
#identifying indexes of the doses for which the condition is met
if(length(index)==0){
h<-1
                    } else {
candi.psi1.hat<-psi1.hat[index]
#extracting the estimates of the prob. of success at these doses
h<-which(psi1.hat==max(candi.psi1.hat))
#dose level with the maximum prob. of success
                           }
```

```
                              }
                                                              }
                                    #end of the combined criterion
############################################################
###########  Restriction on dose escalation  ###############
############################################################
if(stage<=trial.cohort) {
#with <=, the restriction is on the last stage also
if((h-previ.level)>dose.skip){
h<-previ.level+dose.skip
#to avoid restriction on the last stage
                              } else {
h<-h
                                    }
                         }
                           }
        ####### end of stage>up.down  #########
next.dose<-true.dose[h]
#selected dose for the next cohort
previ.level<-h
#current dose level is going to be the previous
#dose level in the next stage
stage<-stage+1
#counts stage for a trial
if(doseresponse.model=="CR") {
tri.res<-rmultinom(1,size=cohort.size,prob=c(true.neu[h],
                  true.effi[h],true.toxic[h]))
#dose-response outcomes for a cohort which received the new dose
r0<-c(r0,tri.res[1,1])
r1<-c(r1,tri.res[2,1])
r2<-c(r2,tri.res[3,1])
alloc.dose<-c(alloc.dose,true.dose[h])
                               }
if(doseresponse.model=="Cox") {
dose.res<-rmultinom(1,size=cohort.size,prob=c(true.psi00[h],
                   true.psi01[h],true.effi[h],true.psi11[h]))
#dose-response outcomes for a cohort which receives the new dose
```

```
r0<-c(r0,dose.res[1,1])
r1<-c(r1,dose.res[2,1])
r2<-c(r2,dose.res[3,1])
r3<-c(r3,dose.res[4,1])
alloc.dose<-c(alloc.dose,true.dose[h])
                              }
max.freq<-max(table(alloc.dose))
#determines the maximum frequency that a dose occurred with
if(max.freq==stop.freq & stage<=(trial.cohort-1)) {
count.stops<-count.stops+1
stopped.at<-c(stopped.at,length(alloc.dose))
stopped.dose<-c(stopped.dose,true.dose[h])
break
#breaks the trial and moves to store the estimates
                                                  }
                          }
             #end of the loop with index 'i'
#################################################################
#######  Storing the estimates only when the trial is  #######
#######  above the up-and-down phase                   #######
#################################################################
if(stage>up.down) {
if(method=="PK") {
pk1.hat<-c(pk1.hat,pk.esti[1])
pk2.hat<-c(pk2.hat,pk.esti[2])
pk3.hat<-c(pk3.hat,pk.esti[3])
varcomp1.hat<-c(varcomp1.hat,var.comp.esti[1])
varcomp2.hat<-c(varcomp2.hat,var.comp.esti[2])
varcomp3.hat<-c(varcomp3.hat,var.comp.esti[3])
stddev.error.hat<-c(stddev.error.hat,error.esti)
                  }
       #end of (method=="PK")
if(stage==trial.cohort+1) { #recommended dose for the next phase
od.complete<-c(od.complete,true.dose[h])
#OD from complete analysis
alloc.sim<-capture.output(alloc.dose[-length(alloc.dose)])
#removing the recommended dose for the next
```

```
#phase from the list of allocated doses
cat(alloc.sim,file="E:/Office/alloc
    dose/alloc.txt",sep="\n",append=TRUE)
#saving the doses allocated to cohorts to an external file
                          } else {
alloc.sim<-capture.output(alloc.dose)
cat(alloc.sim,file="E:/Office/alloc
    dose/alloc.txt",sep="\n",append=TRUE)
                                  }
od<-c(od,true.dose[h])
#records all the ODs: early and complete
theta1.hat<-c(theta1.hat,theta1)
theta2.hat<-c(theta2.hat,theta2)
theta3.hat<-c(theta3.hat,theta3)
theta4.hat<-c(theta4.hat,theta4)

#theta5.hat<-c(theta5.hat,theta5)
#theta6.hat<-c(theta6.hat,theta6)
                   }
           #end of stage>up.down
                  }
    #end of the loop with index 'l'

##################  End of the program  ###################
```

B.2 Functions in *R*

Below are the functions used in the program.

```
#####################################################
######## One-compartment PK model with bolus #######
########  input and first-order elimination  #######
#####################################################

################ Model ###################

f1<-function(x,v,cl,t) {

(x/v)*exp(-(cl/v)*t)

                       }

#'x' represents dose

############ AUC and its variance ###############

auc1<-function(v,cl,t1,x) {

(x/cl)*(1-exp(-(cl/v)*t1))

                          }

#t1 is the upper sampling time

c.v.f1<-function(v,cl,t1,x) {

((x*t1)/(v^2))*(-exp(-(cl/v)*t1))

                            }

#derivative w.r.t. v

c.cl.f1<-function(v,cl,t1,x) {

(x/cl)*(exp(-(cl/v)*t1))*((1/cl)+(t1/v))-(x/cl^2)

                             }

#derivative w.r.t. cl

#####################################################
#########  One-compartment PK model with  #########
#########  first-order absorption         #########
#####################################################

################## Model #####################

f2<-function(x,v,ke,ka,t) {

((x*ka)*(exp(-ke*t)- exp(-ka*t)))/(v*(ka-ke))

                          }

#'x' represents dose

############ AUC and its variance ###############

auc2<-function(v,ke,ka,t1,x) {

((x*ka)/(v*(ka-ke)))*((1-exp(-ke*t1))/ke-(1-exp(-ka*t1))/ka)

                             }

#t1 is the upper sampling time
```

```
c.v.f2<-function(v,ke,ka,t1,x) {
z1<-1-exp(-ke*t1)
z2<-1-exp(-ka*t1)
(-(x*ka)/((v^2)*(ka-ke)))*(z1/ke-z2/ka)
                                }
c.ke.f2<-function(v,ke,ka,t1,x) {
z1<-1-exp(-ke*t1)
z2<-1-exp(-ka*t1)
((x*ka)/(v*((ka-ke)^2)))*(z1/ke-z2/ka)+
((x*ka)/(v*(ka-ke)))*(-z1/ke^2+(t1*exp(-ke*t1))/ke)
                                 }
c.ka.f2<-function(v,ke,ka,t1,x) {
z1<-1-exp(-ke*t1)
z2<-1-exp(-ka*t1)
(-(x*ke)/(v*((ka-ke)^2)))*(z1/ke-z2/ka)+
((x*ka)/(v*(ka-ke)))*(z1/ka^2-(t1*exp(-ka*t1))/ka)
                                 }
############### Cmax and its variance ###############
cmax2<-function(v,ke,ka,tmax,x) {
((x*ka)/(v*(ka-ke)))*(exp(-ke*tmax)-exp(-ka*tmax))
                                }
cmax.v.f2<-function(v,ke,ka,tmax,x) {
-((x*ka)/(v^2*(ka-ke)))*(exp(-ke*tmax)-exp(-ka*tmax))
                                    }
cmax.ke.f2<-function(v,ke,ka,tmax,x) {
((x*ka)/(v*(ka-ke)^2))*(exp(-ke*tmax)*(2-ka*tmax)-
  exp(-ka*tmax)*(1+(ka*(1-ke*tmax))/ke))
                                     }
cmax.ka.f2<-function(v,ke,ka,tmax,x) {
(x/(v*(ka-ke)^2))*(ke*exp(-ke*tmax)*(-1+ka*(ka*tmax-1))+
 exp(-ka*tmax)*(ke-ka*(ke*tmax-1)))
                                     }
#########################################################
########  Functions representing the probabilities #######
########  of trinomial dose-response outcomes      #######
#########################################################
psi0<-function(theta1,theta2,theta3,theta4,x) {
```

```
z1<-exp(theta1+theta2*x)
z2<-exp(theta3+theta4*x)
1/((1+z1)*(1+z2))
                                                    }
psi1<-function(theta1,theta2,theta3,theta4,x) {
z1<-exp(theta1+theta2*x)
z2<-exp(theta3+theta4*x)
z1/((1+z1)*(1+z2))
                                                    }
psi2<-function(theta3,theta4,x)  {
z2<-exp(theta3+theta4*x)
z2/(1+z2)
                                    }
##########################################################
#########  Likelihood function for the CR model #########
##########################################################
lf1<-function(w) {
   v<-1
   w1<-w[1]
   w2<-w[2]
   w3<-w[3]
   w4<-w[4]
   for (i in 1:length(alloc.dose)) {
       dose.i<-alloc.dose[i]
       r0.i<-r0[i]
       r1.i<-r1[i]
       r2.i<-r2[i]
       z1<-exp(w1+w2*dose.i)
       z2<-exp(w3+w4*dose.i)
       psi0<-1/((1+z1)*(1+z2))
       psi1<-z1*psi0
       v<-v*(psi0^r0.i)*(psi1^r1.i)*((1-psi0-psi1)^r2.i)
                                      }
   return(v)
                  }
lf1.c<-cmpfun(lf1)       #compiling the likelihood function
```

```
lf1.cond<-function(w) {  #likelihood function with
lf1.c(w)*(w[1]>=w[3])    #the condition that theta1>theta3
                       }
theta1.f1<-function(w) {
(w[1])*lf1.c(w)*(w[1]>=w[3]) #function in the numerator when we find
                         }     #the posterior estimate for theta1
theta2.f1<-function(w) {     #function in the numerator when we find
(w[2])*lf1.c(w)*(w[1]>=w[3]) #the posterior estimate for theta2
                         }
theta3.f1<-function(w) {     #function in the numerator when we
(w[3])*lf1.c(w)*(w[1]>=w[3]) #find the posterior estimate for theta3
                         }
theta4.f1<-function(w) {    #function in the numerator when we
(w[4])*lf1.c(w)*(w[1]>=w[3]) #find the posterior estimate for theta4
                          }
#in the likelihood function and thereafter,
#w[1], w[2], w[3] and w[4] represent theta1, theta2,
#theta3 and theta4, respectively
############################################################
###########   Penalty function for the CR model   #########
############################################################
penalty<-function(theta1,theta2,theta3,theta4,x) {
((psi1(theta1,theta2,theta3,theta4,x))^(-CS))*
((1-psi2(theta3,theta4,x))^(-CT))
                                                  }
############################################################
#################  FIM for the CR model  ###################
############################################################
I11<-function(theta1,theta2,theta3,theta4,x) {
(psi1(theta1,theta2,theta3,theta4,x)*(1-psi2(theta3,theta4,x)))/
(psi0(theta1,theta2,theta3,theta4,x)*((1+exp(theta1+theta2*x))^2))
                                               }
I12<-function(theta1,theta2,theta3,theta4,x) {
(x*I11(theta1,theta2,theta3,theta4,x))
                                               }
I13<-0
I14<-0
#I21<-I12
```

```
I22<-function(theta1,theta2,theta3,theta4,x) {
((x^2)*I11(theta1,theta2,theta3,theta4,x))
                                          }
I23<-0
I24<-0
I31<-0
I32<-0
I33<-function(theta3,theta4,x) {
(psi2(theta3,theta4,x))*(1-psi2(theta3,theta4,x))
                               }
I34<-function(theta3,theta4,x) {
(x*I33(theta3,theta4,x))
                               }
I41<-0
I42<-0
#I43<-I34
I44<-function(theta3,theta4,x) {
((x^2)*I33(theta3,theta4,x))
                               }
#####################################################################
######  Defining the associated functions for the Cox Model  #######
#####################################################################
psi00<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
z1<-exp(theta1+theta2*x)
z2<-exp(theta3+theta4*x)
z3<-exp(theta5+theta6*x)
1/(1+z1+z2+z3)
                                                             }
psi01<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
z1<-exp(theta1+theta2*x)
z2<-exp(theta3+theta4*x)
z3<-exp(theta5+theta6*x)
z1/(1+z1+z2+z3)
                                                             }
psi10<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
z1<-exp(theta1+theta2*x)
```

```
z2<-exp(theta3+theta4*x)
z3<-exp(theta5+theta6*x)
z2/(1+z1+z2+z3)
                                                                    }
psi11<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
z1<-exp(theta1+theta2*x)
z2<-exp(theta3+theta4*x)
z3<-exp(theta5+theta6*x)
z3/(1+z1+z2+z3)
                                                                    }
psi1.<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
#marginal prob. of efficacy
psi10(theta1,theta2,theta3,theta4,theta5,theta6,x)+
psi11(theta1,theta2,theta3,theta4,theta5,theta6,x)
                                                                    }
psi.1<-function(theta1,theta2,theta3,theta4,theta5,theta6,x) {
#marginal prob. of toxicity
psi01(theta1,theta2,theta3,theta4,theta5,theta6,x)+
psi11(theta1,theta2,theta3,theta4,theta5,theta6,x)
                                                                    }

#########################################################################
#############  Likelihood function for the Cox model  ###################
#########################################################################

lf2<-function(w) {
   v<-1
   w1<-w[1]
   w2<-w[2]
   w3<-w[3]
   w4<-w[4]
   w5<-w[5]
   w6<-w[6]
   for (i in 1:length(alloc.dose)) {
       dose.i<-alloc.dose[i]
       r0.i<-r0[i]
       r1.i<-r1[i]
       r2.i<-r2[i]
       r3.i<-r3[i]
```

```
        z1<-exp(w1+w2*dose.i)
        z2<-exp(w3+w4*dose.i)
        z3<-exp(w5+w6*dose.i)

        psi00<-1/(1+z1+z2+z3)
        psi01<-z1*psi00
        psi10<-z2*psi00
        v<-v*(psi00^r0.i)*(psi01^r1.i)*(psi10^r2.i)*
            ((1-psi00-psi01-psi10)^r3.i)
                                    }
   return(v)
                  }
lf2.c<-cmpfun(lf2)
# compiling the likelihood function
theta1.f2<-function(w) {
(w[1])*lf2.c(w)
                        }
theta2.f2<-function(w) {
(w[2])*lf2.c(w)
                        }
theta3.f2<-function(w) {
(w[3])*lf2.c(w)
                        }
theta4.f2<-function(w) {
(w[4])*lf2.c(w)
                        }
theta5.f2<-function(w) {
(w[5])*lf2.c(w)
                        }
theta6.f2<-function(w) {
(w[6])*lf2.c(w)
                        }
#in the likelihood function and thereafter, w[1], w[2],
#w[3] and w[4] etc. represent theta1, theta2, etc.

##########################  End  ##############################
```

Bibliography

Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Designs, with SAS*. Oxford University Press, Oxford.

Atkinson, A. C., V. V. Fedorov, A. M. Herzberg, and R. Zhang (2014). Elemental information matrices and optimal experimental design for generalized regression models. *Journal of Statistical Planning and Inference 144*(1), 81–91.

Babb, J., A. Rogatko, and S. Zacks (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine 17*(10), 1103–1120.

Babb, J. S. and A. Rogatko (2001). Patient specific dosing in a cancer phase I clinical trial. *Statistics in Medicine 20*(14), 2079–2090.

Babb, J. S. and A. Rogatko (2004). Bayesian methods for cancer phase I clinical trials. In N. Geller (Ed.), *Contemporary Biostatistical Methods in Clinical Trials*, pp. 1–40. Marcel Dekker, New York.

Bazzoli, C., T. T. Nguyen, A. Dubois, E. Retout, S. Comets, and F. Mentré (2010). *PFIM 3.2 User Guide*. Université Paris Diderot and INSERM.

Beal, S. L. and L. B. Sheiner (1982). Estimating population kinetics. *Critical Reviews in Biomedical Engineering 8*(3), 195–222.

Beal, S. L. and L. B. Sheiner (1992). *NONMEM User's Guide*. NONMEM Project Group, University of California at San Francisco, San Francisco, CA.

Berger, M. P. and W. K. Wong (2009). *An Introduction to Optimal Designs for Social and Biomedical Research*. Wiley, New York.

Berry, S. M., B. P. Carlin, J. J. Lee, and P. Müller (2010). *Bayesian Adaptive Methods for Clinical Trials.* CRC Press, New York.

Bogacka, B., D. Uciński, and A. C. Atkinson (2014). Comparison of optimum adaptive experimental designs for dose-finding in early phase clinical trials. Preprint.

Bos, A., F. De Vos, E. De Vries, J. Beijnen, H. Rosing, M. Mourits, A. Van der Zee, J. Gietema, and P. Willemse (2005). A phase I study of intraperitoneal topotecan in combination with intravenous carboplatin and paclitaxel in advanced ovarian cancer. *European Journal of Cancer 41*(4), 539–548.

Braun, T. M. (2002). The bivariate continual reassessment method: Extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials 23*(3), 240–256.

Bryant, J. and R. Day (1995). Incorporating toxicity considerations into the design of two-stage phase II clinical trials. *Biometrics 51*(4), 1372–1383.

Cheng, J. D., J. S. Babb, C. Langer, S. Aamdal, F. Robert, L. R. Engelhardt, O. Fernberg, J. Schiller, G. Forsberg, R. K. Alpaugh, et al. (2004). Individualized patient dosing in phase I clinical trials: The role of escalation with overdose control in PNU-214936. *Journal of Clinical Oncology 22*(4), 602–609.

Chernoff, H. (1953). Locally optimal designs for estimating parameters. *The Annals of Mathematical Statistics 24*(4), 586–602.

Cheung, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika 92*(4), 863–873.

Cheung, Y. K. and R. Chappell (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics 56*(4), 1177–1182.

Chow, S. C. and J. P. Liu (2004). *Design and Analysis of Clinical Trials: Concepts and Methodologies,* 2nd Edition. Wiley, New York.

Chu, P.-L., Y. Lin, and W. J. Shih (2009). Unifying CRM and EWOC designs for phase I cancer clinical trials. *Journal of Statistical Planning and Inference 139*(3), 1146–1163.

Collins, J. M., C. K. Grieshaber, and B. A. Chabner (1990). Pharmacologically guided phase I clinical trials based upon preclinical drug development. *Journal of the National Cancer Institute 82*(16), 1321–1326.

Collins, J. M., D. S. Zaharko, R. L. Dedrick, and B. A. Chabner (1986). Potential roles for preclinical pharmacology in phase I clinical trials. *Cancer Treatment Reports 70*(1), 73–80.

Cox, D. R. (1970). *The Analysis of Binary Data*. Chapman and Hall, London.

Cunningham, D., Y. Humblet, S. Siena, D. Khayat, H. Bleiberg, A. Santoro, D. Bets, M. Mueser, A. Harstrick, C. Verslype, I. Chau, and E. V. Cutsem (2004). Cetuximab monotherapy and cetuximab plus irinotecan in irinotecan-refractory metastatic colorectal cancer. *New England Journal of Medicine 351*(4), 337–345.

Davidian, M. (2010). *An Introduction to Nonlinear Mixed Effects Models and PK/PD Analysis*. American Statistical Association Biopharmaceutical Section webinar. `http://www4.stat.ncsu.edu/~davidian/webinar.pdf`.

Derendorf, H., L. J. Lesko, P. Chaikin, W. A. Colburn, P. Lee, R. Miller, R. Powell, G. Rhodes, D. Stanski, and J. Venitz (2000). Pharmacokinetic/pharmacodynamic modelling in drug research and development. *Journal of Clinical Pharmacology 40*(12), 1399–1418.

Dixon, W. J. and A. M. Mood (1948). A method for obtaining and analyzing sensitivity data. *Journal of the American Statistical Association 43*(241), 109–126.

Dragalin, V. and V. Fedorov (2006). Adaptive designs for dose-finding based on efficacy–toxicity response. *Journal of Statistical Planning and Inference 136*(6), 1800–1823.

Dubois, A., J. Bertrand, and F. Mentré (2011). *Mathematical Expressions of the Pharmacokinetic and Pharmacodynamic Models Implemented in the PFIM Software*. Université Paris Diderot and INSERM.

Faries, D. (1994). Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics 4*(2), 147–164.

Fedorov, V. V. (1972). *Theory of Optimal Experiments*. Academic Press, New York.

Fedorov, V. V., N. Flournoy, Y. Wu, and R. Zhang (2011). Best intention designs in dose-finding studies. Preprint.

Fedorov, V. V. and P. Hackl (1997). *Model-Oriented Design of Experiments*. Lecture Notes in Statistics. Springer, New York.

Gabrielsson, J. and D. Weiner (2000). *Pharmacokinetic/Pharmacodynamic Data Analysis: Concepts and Applications*. Routledge, London.

Gehan, E. A. (1961). The determination of the number of patients required in a preliminary and a follow-up trial of a new chemotherapeutic agent. *Journal of Chronic Diseases 13*(4), 346–353.

Gelmon, K., D. Stewart, K. Chi, S. Chia, C. Cripps, S. Huan, S. Janke, D. Ayers, D. Fry, J. Shabbits, et al. (2004). A phase I study of AMD473 and docetaxel given once every 3 weeks in patients with advanced refractory cancer: A national cancer institute of Canada-clinical trials group trial, IND 131. *Annals of Oncology 15*(7), 1115–1122.

Giles, F. J., M. S. Tallman, G. Garcia-Manero, J. E. Cortes, D. A. Thomas, W. G. Wierda, S. Verstovsek, M. Hamilton, E. Barrett, M. Albitar, et al. (2004). Phase I and pharmacokinetic study of a low-clearance, unilamellar liposomal formulation of lurtotecan, a topoisomerase 1 inhibitor, in patients with advanced leukemia. *Cancer 100*(7), 1449–1458.

Goodman, S. N., M. L. Zahurak, and S. Piantadosi (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine 14*(11), 1149–1161.

Gooley, T. A., P. J. Martin, L. D. Fisher, and M. Pettinger (1994). Simulation as a design tool for phase I/II clinical trials: An example from bone marrow transplantation. *Controlled Clinical Trials 15*(6), 450–462.

Govindarajulu, Z. (1988). *Statistical Techniques in Bioassay*. Karger, Basel.

Graham, M. A. and P. Workman (1992). The impact of pharmacokinetically guided dose escalation strategies in phase I clinical trials: Critical evaluation and recommendations for future studies. *Annals of Oncology 3*(5), 339–347.

Hardwick, J., M. C. Meyer, and Q. F. Stout (2003). Directed walk designs for dose-response problems with competing failure modes. *Biometrics 59*(2), 229–236.

Heyd, J. M. and B. P. Carlin (1999). Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine 18*(11), 1307–1321.

Hooker, A. and P. Vicini (2005). Simultaneous population optimal design for pharmacokinetic-pharmacodynamic experiments. *AAPS Journal 7*(4), E759–E785.

Ishizuka, N. and Y. Ohashi (2001). The continual reassessment method and its applications: A Bayesian methodology for phase I cancer clinical trials. *Statistics in Medicine 20*(17-18), 2661–2681.

Ivanova, A. (2006). Dose-finding in oncology - nonparametric methods. In N. Ting (Ed.), *Dose Finding in Drug Development*, pp. 49–58. Springer, New York.

Johnson, S. G. and B. Narasimhan (2013). *Cubature: Adaptive Multivariate Integration Over Hypercubes*. R package version 1.1-2.

Jung, S.-H., M. Carey, and K. M. Kim (2001). Graphical search for two-stage designs for phase II clinical trials. *Controlled Clinical Trials 22*(4), 367–372.

Jung, S.-H., T. Lee, K. M. Kim, and S. L. George (2004). Admissible two-stage designs for phase II cancer clinical trials. *Statistics in Medicine 23*(4), 561–569.

Kiefer, J. and J. Wolfowitz (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics 12*(3), 363–366.

Korn, E. L., D. Midthune, T. T. Chen, L. V. Rubinstein, M. C. Christian, and R. M. Simon (1994). A comparison of two phase I trial designs. *Statistics in Medicine 13*(18), 1799–1806.

Kuhn, E. and M. Lavielle (2005). Maximum likelihood estimation in nonlinear mixed effects models. *Computational Statistics and Data Analysis 49*(4), 1020–1038.

Kwan, K. C., G. O. Breault, E. R. Umbenhauer, F. G. McMahon, and D. E. Duggan (1976). Kinetics of indomethacin absorption, elimination, and enterohepatic circulation in man. *Journal of Pharmacokinetics and Biopharmaceutics 4*(3), 255–280.

Le Tourneau, C., J. J. Lee, and L. L. Siu (2009). Dose escalation methods in phase I cancer clinical trials. *Journal of the National Cancer Institute 101*(10), 708–720.

Leung, D. H.-Y. and Y.-G. Wang (2001). Isotonic designs for phase I trials. *Controlled Clinical Trials 22*(2), 126–138.

Lévy, V., S. Zohar, R. Porcher, and S. Chevret (2001). Alternate designs for conduct and analysis of phase I cancer trials. *Blood 98*(4), 1275–1275.

Lindstrom, M. J. and D. M. Bates (1990). Nonlinear mixed effects models for repeated measures data. *Biometrics 46*(3), 673–687.

Mander, A. P. and S. G. Thompson (2010). Two-stage designs optimal under the alternative hypothesis for phase II cancer clinical trials. *Contemporary Clinical Trials 31*(6), 572–578.

Mander, A. P., J. Wason, M. J. Sweeting, and S. G. Thompson (2012). Admissible two-stage designs for phase II cancer clinical trials that incorporate the expected

sample size under the alternative hypothesis. *Pharmaceutical Statistics 11*(2), 91–96.

McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models,* 2nd Edition. Chapman and Hall, London.

Mielke, T. (2012). *Approximation of the Fisher Information and Design in Nonlinear Mixed Effects Models.* Ph. D. thesis, Otto-von-Guericke University Magdeburg.

Møller, S. (1995). An extension of the continual reassessment methods using a preliminary up-and-down design in a dose finding study in cancer patients, in order to investigate a greater range of doses. *Statistics in Medicine 14*(9), 911–922.

Murtaugh, P. (1989). *Simultaneous Analysis of Efficacy and Toxicity in Dose Ranging Trials with New Drugs.* Ph. D. thesis, University of Wahsington.

Nelder, J. A. and R. Mead (1965). A simplex method for function minimization. *The Computer Journal 7*(4), 308–313.

Oehlert, G. W. (1992). A note on the delta method. *The American Statistician 46*(1), 27–29.

Okamoto, I., A. Hamada, Y. Matsunaga, J.-i. Sasaki, S. Fujii, H. Uramoto, H. Yamagata, I. Mori, H. Kishi, H. Semba, and H. Saito (2006). Phase I and pharmacokinetic study of amrubicin, a synthetic 9-aminoanthracycline, in patients with refractory or relapsed lung cancer. *Cancer Chemotherapy and Pharmacology 57*(3), 282–288.

O'Quigley, J. (2002). Continual reassessment designs with early termination. *Biostatistics 3*(1), 87–99.

O'Quigley, J., M. D. Hughes, and T. Fenton (2001). Dose-finding designs for HIV studies. *Biometrics 57*(4), 1018–1029.

O'Quigley, J., M. Pepe, and L. Fisher (1990). Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics 46*(1), 33–48.

O'Quigley, J. and E. Reiner (1998). A stopping rule for the continual reassessment method. *Biometrika 85*(3), 741–748.

O'Quigley, J. and L. Z. Shen (1996). Continual reassessment method: A likelihood approach. *Biometrics 52*(2), 673–684.

O'Quigley, J. and S. Zohar (2006). Experimental designs for phase I and phase I/II dose-finding studies. *British Journal of Cancer 94*(5), 609–613.

Piantadosi, S. and G. Liu (1998). Improved designs for dose escalation studies using pharmacokinetic measurements. *Statistics in Medicine 15*(15), 1605–1618.

Pinheiro, J. C. and D. M. Bates (1995). Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics 4*(1), 12–35.

Pinheiro, J. C. and D. M. Bates (2000). *Mixed-Effects Models in S and S-PLUS*. Springer, New York.

Prieur, D. J., D. M. Young, R. D. Davis, D. A. Cooney, E. R. Homan, R. L. Dixon, and A. M. Guarino (1973). Procedures for preclinical toxicologic evaluation of cancer chemotherapeutic agents: Protocols of the laboratory of toxicology. *Cancer Chemotherapy Reports. Part 3 4*(1), 1–39.

Pronzato, L. (2000). Adaptive optimization and D-optimum experimental design. *Annals of Statistics 28*(6), 1743–1761.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria.

Ratain, M. J., R. Mick, R. L. Schilsky, and M. Siegler (1993). Statistical and ethical issues in the design and conduct of phase I and II clinical trials of new anticancer agents. *Journal of the National Cancer Institute 85*(20), 1637–1643.

Reigner, B. G. and K. S. Blesch (2002). Estimating the starting dose for entry into humans: Principles and practice. *European Journal of Clinical Pharmacology 57*(12), 835–845.

Reiner, E., X. Paoletti, and J. O'Quigley (1999). Operating characteristics of the standard phase I clinical trial design. *Computational Statistics and Data Analysis 30*(3), 303–315.

Retout, S., S. Duffull, and F. Mentré (2001). Development and implementation of the population Fisher information matrix for the evaluation of population pharmacokinetic designs. *Computer Methods and Programs in Biomedicine 65*(2), 141–151.

Riviere, J. E. (2011). *Comparative Pharmacokinetics: Principles, Techniques and Applications,* 2nd Edition. Wiley-Blackwell, New York.

Rosenbaum, S. E. (2011). *Basic Pharmacokinetics and Pharmacodynamics: An Integrated Textbook and Computer Simulations.* Wiley, New York.

Sheiner, L. B. and S. L. Beal (1980). Evaluation of methods for estimating population pharmacokinetic parameters. I. Michaelis-Menten model: Routine clinical pharmacokinetic data. *Journal of Pharmacokinetics and Biopharmaceutics 8*(6), 553–571.

Sheiner, L. B., B. Rosenberg, and K. L. Melmon (1972). Modelling of individual pharmacokinetics for computer-aided drug dosage. *Computers and Biomedical Research 5*(5), 441–459.

Sheiner, L. B. and J. L. Steimer (2000). Pharmacokinetic/pharmacodynamic modelling in drug development. *Annual Review of Pharmacology and Toxicology 40*(1), 67–95.

Simon, R. (1989). Optimal two-stage designs for phase II clinical trials. *Controlled Clinical Trials 10*(1), 1–10.

Simon, R., L. Rubinstein, S. G. Arbuck, M. C. Christian, B. Freidlin, and J. Collins (1997). Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute 89*(15), 1138–1147.

Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics 45*(3), 925–937.

Storer, B. E. (2001). An evaluation of phase I clinical trial designs in the continuous dose-response setting. *Statistics in Medicine 20*(16), 2399–2408.

Thall, P. F. and J. D. Cook (2004). Dose-finding based on efficacy–toxicity trade-offs. *Biometrics 60*(3), 684–693.

Thall, P. F. and H. Q. Nguyen (2012). Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics 22*(4), 785–801.

Thall, P. F., H. Q. Nguyen, and E. H. Estey (2008). Patient-specific dose finding based on bivariate outcomes and covariates. *Biometrics 64*(4), 1126–1136.

Thall, P. F. and K. E. Russell (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics 54*(1), 251–264.

Tighiouart, M., A. Rogatko, and J. S. Babb (2005). Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine 24*(14), 2183–2196.

Vonesh, E. F. (1996). A note on the use of Laplace's approximation for nonlinear mixed-effects models. *Biometrika 83*(2), 447–452.

Walker, S. (1996). An EM algorithm for nonlinear random effects models. *Biometrics 52*(3), 934–944.

Wang, J. (2007). EM algorithms for nonlinear mixed effects models. *Computational Statistics and Data Analysis 51*(6), 3244–3256.

Wason, J. M., A. P. Mander, and T. G. Eisen (2011). Reducing sample sizes in two-stage phase II cancer trials by using continuous tumour shrinkage endpoints. *European Journal of Cancer 47*(7), 983–989.

Whitehead, J., Y. Zhou, L. Hampson, E. Ledent, and A. Pereira (2007). A Bayesian approach for dose-escalation in a phase I clinical trial incorporating pharmacodynamic endpoints. *Journal of Biopharmaceutical Statistics 17*(6), 1117–1129.

Wolfinger, R. (1993). Laplace's approximation for nonlinear mixed models. *Biometrika 80*(4), 791–795.

Wolfinger, R. D. and X. Lin (1997). Two Taylor-series approximation methods for nonlinear mixed models. *Computational Statistics and Data Analysis 25*(4), 465–490.

Wynn, H. P. (1972). Results in the theory and construction of D-optimum experimental designs. *Journal of the Royal Statistical Society Series B 34*(2), 133–147.

Yuh, L., S. Beal, M. Davidian, F. Harrison, A. Hester, K. Kowalski, E. Vonesh, and R. Wolfinger (1994). Population pharmacokinetic/pharmacodynamic methodology and applications: A bibliography. *Biometrics 50*(2), 566–575.

Zhang, W., D. J. Sargent, and S. Mandrekar (2006). An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine 25*(14), 2365–2383.

Zhou, Y., J. Whitehead, P. Korhonen, and M. Mustonen (2008). Implementation of a Bayesian design in a dose-escalation study of an experimental agent in healthy volunteers. *Biometrics 64*(1), 299–308.