UNIVERSITY OF LEEDS

Statistical Decision Making in Early Phase Clinical Trials

Andrew Martin Hall

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

Institute of Clinical Trials Research

School of Medicine

February 2024

Intellectual Property

Acknowledgements

I would like to express my sincere gratitude to my three supervisors, Prof. Sarah Brown, Dr. Duncan Wilson, and Dr. Stuart Barber. Their guidance, support, and expertise have been invaluable throughout my PhD. I am thankful for their time, encouragement and insights, which have played a crucial role in making me a better academic and the completion of this thesis.

I would also like to thank Prof. John Paul Gosling and Prof. David Cairns who acted as supervisors during the first two years and second year of my PhD respectively, and helped in part to shape the direction of the thesis.

Funding Acknowledgement

PhD funding provided by Myeloma UK, the only charitable organisation in the UK dealing exclusively with myeloma, dedicated to patients – making sure they get access to the right treatment at the right time.

Abstract

The main objective of dose finding trials is to find an optimal dose amongst a candidate set for further research. The trial design in oncology proceeds in stages with a decision as to how to treat the next group of patients made after every stage until a final sample size is reached or the trial stopped early.

The thesis applies a Bayesian decision theoretic approach to the problem focusing on the specification of a novel utility function and the role of correlation in the probability model. Utility independence axioms are used to give a simplified bivariate form to the utility function based on more easily assessed univariate utility functions. Utility functions for both efficacy and toxicity depend upon a reference point with different attitudes to risk depending upon whether above or below the point. A risk averse attitude (concave) is specified for perceived gains and risk prone (convex) for losses. A set of questions are posed for the utility function to be accurately elicited. A novel stopping rule derived from the utility function is also tested.

An inspection of copula theory and a simulation study demonstrate the difficulty in estimating correlation and recommend using a more parsimonious independent model. A simulation study demonstrates that the utility function has merit in further evaluation in this setting. The simulation results show that the decision criteria are more sensitive in detecting the optimal dose when candidate doses are around minimum efficacy and maximum toxicity thresholds.

The specification of the utility function is flexible to accommodate clinical beliefs allowing us to think about acceptable levels of patient risk. The work applies a broad framework to give insight to existing methods and potential to adapt to different endpoints and trial features.

Contents

List of Figures

List of Tables

Acronyms

APARA	The Arrow-Pratt measure of Absolute Risk Aversion
CDF	Cumulative Density Function
CRM	Continual Reassessment Method
CSS	Copula Simulation Study
CTCAE	Common Terminology Criteria for Adverse Events
DLT	Dose-Limiting Toxicity
DM	Decision Maker
EffTox	Efficacy Toxicity design
EffToxU	Efficacy Toxicity patient Utility design
FDA	Food and Drug Administration
FGM	Fairlie-Gumbel-Morgenstern copula
ICH	International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use
MTD	Maximium Tolerated Dose
NDS	percentage of trials with No Dose Selected
OD	Optimal Dose
PMF	Probability Mass Function
R2DT	Reference Dependent Decision Theoretic dose finding design
RIO	Rational Impartial Observer
VNM	Von Neumann–Morgenstern utility theory

Chapter 1

Introduction

Cancer develops when abnormal cells divide and multiply in an uncontrollable way, with many cancers eventually spreading into other tissues. Cancer continues to be the leading cause of death in the developed world [1]. Modern cancer treatment can be placed into four broad groups: surgery, radiotherapy, chemotherapy and targeted treatments [2]. Targeted treatments are able to target specific proteins or a faulty process within a cancer cell while chemotherapy or cytotoxic agents destroy both cancerous and healthy cells without distinction. Individual interventions or a combination of interventions from the four groups are used to treat the many different types of cancer.

The development pathway for an anti-cancer treatment starts at the discovery stage, when a biological idea is constructed and tested in cell and then animal models. Promising treatments are progressed to the toxicology stage where the effect of the agent on an organism and how the organism handles the agent are studied more closely in multiple animal models. Clinical trials for cancer treatments are almost exclusively conducted in patients rather than healthy volunteers given many of the serious side effects associated with treatment. The clinical development of a treatment for cancer has traditionally followed four sequential stages, Phases I-IV:

Phase I trials primarily assess the safety, toxicity and pharmacology of a treatment usually at multiple doses. Phase I trials are conducted in a low number of patients (typically 10-30) with the sample size highly dependent on the number of doses to be tested. The development of a treatment will stop at this stage if a treatment cannot be considered safe.

Phase II trials test to see if the regimen has sufficient activity to warrant further large scale study. The endpoints for efficacy are typically accessed over a shorter period of time in comparison to phase III. Phase II trials traditionally were conducted without comparison to a concurrent control (single arm trials) instead comparing to historic control data for the current standard of care in the specific setting. Randomisation is becoming more prominent in this setting to provide a more representative concurrent control [3]. Phase II studies are usually the first formal look at the efficacy of an agent, but may also look at further toxicity and safety. Sample sizes for phase II studies range from tens to low hundreds.

Phase III trials are large randomised trials to assess if the treatment can offer a benefit against the current standard of care. Typically this is assessed using a longer term efficacy measure. An improvement in efficacy is usually sought but non-inferiority can also be an objective if the new treatment is cheaper to deliver or has a more favourable toxicity profile. The successful completion of a phase III trial would typically see the agent being adopted into routine clinical practice. The sample size for a phase III trial are often hundreds or thousands of patients.

Phase IV trials are post marketing trials and evaluate the long term effects of the treatment and can change how a treatment is delivered in practice. It is also possible that the treatment may have its license removed if new unfavourable evidence emerges.

The first two phases can be described as early phase trials, the third a confirmatory trial and the fourth a post marketing trial. Although this appears to be a linear pathway, in reality there is overlap and trials can be conducted in parallel in different settings. The International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH E8) guideline on general considerations for clinical studies recognises the difficulties in classifying trials according to a traditional pathway whilst supporting the need for a step wise series of trials where information from early studies is used to support and plan later stage studies [4]. A classification system that is structured around the trial's objective is encouraged. For example, a study that seeks to find a suitable dose may be referred to as a phase I study in the traditional pathway whereas a dose finding study is a less ambiguous label.

Drug discovery and development is a long and costly process, estimates of cost range from \$314 million to \$2.7 billion to get a drug licensed [5, 6]. An estimate for the time it takes to

develop a compound before licensing is 7.3 years (range, 5.8 - 15.2 years) [7]. The large range of time is dependent upon disease area and the nature of the drug in addition to whether the compound was part of Fast Track, Breakthrough Therapy, Accelerated Approval, or Priority Review statuses from the Food and Drug Administration (FDA) [8]. A drug company will traditionally file for a patent around the discovery stage of a compound. Patents are normally granted for 20 years which allow the patent owner to exclusively sell the drug for this period. After the patent expires a huge proportion of sales are lost to generic competitors and the price for the drug plummets [9]. It is imperative for the pharmaceutical company that a drug is licensed in a timely manner in order to recoup the costs of development.

The high cost of drug development is in part due to the high failure rate during drug development, it is estimated that 70% of research and development budget is spent on failed projects [10]. Over the period of 2000-2015 only 3.4% of cancer drugs that made it to clinical testing received a license [11]. Breaking this figure down using transition probabilities; approximately 40% fail to transition to phase II testing, with less than a third of compounds tested at phase II progressing to phase III. Of the compounds reaching phase III only 35% successfully obtain a license. The main reasons for failure in phase III testing are lack of efficacy (64%), safety (12%) and commercial reasons (24%) [12]. Similar proportions, efficacy (52%), safety (24%) and commercial reasons (24%) are the estimated failure rates at phase II [13]. A major factor in the optimisation, and subsequently the effects, of the drug is the dose of the drug selected in the earlier stages of development. The dose of a treatment is intrinsically linked to both the efficacy and toxicity profile of a treatment with dose optimisation seen as a major factor in improving success rates for phase II and phase III studies [14].

The dose selection paradigm in oncology has been shaped historically by the prognosis of diagnosis, lack of effective treatments and the properties of cytotoxic treatments coming through development [15]. The effectiveness and the toxicity associated with a cytotoxic agent steeply increase with increasing dose. The highest dose of a cytotoxic agent that can be tolerated by patients, is considered optimal to progress to phase II testing. The trial design that developed to find the maximum tolerated dose (MTD) was to treat small groups of patients at increasing doses over a short period of time. The staged design is referred to as a best intention study design where a decision for the dose to treat the

next group of patients is made in the best interest of patients entering the study. Project Optimus is an FDA initiative to reform the dose optimization and dose selection paradigm in oncology drug development in response to the increased proportion of targeted treatments coming into development [16]. The project is wide ranging to help improve overall success rates of oncology treatments. This includes trial designs for later stage randomised studies comparing different doses. This thesis focuses on non-randomised best intention designs with the unique feature of optimising dosing at each stage for the benefit of patients entering the study. This type of design is referred to as a dose finding design from hereon.

There are multiple objectives for a dose finding study in oncology. The main scientific objective of a dose finding trial is to determine a dose for the treatment of patients in the future [17]. This dose is referred to as the optimal dose (OD). Given the potential for serious side effects when treating with untried cancer treatments, it is ethical to recruit patients rather than healthy volunteers to the studies. Patients who enter the trial are typically also seeking a therapeutic advantage. The main ethical objective therefore is to ensure patients studied within the trial are not exposed to excessive toxicity or doses with minimal efficacy. Lower level objectives associated with the trial design and delivery concern efficiency and reliability: the trial should utilise the minimum number of patients and be capable of finding an OD with a degree of statistical accuracy [18]. A dose finding trial design is a balance between meeting each of the stated objectives.

The ability of a trial to successfully achieve its objectives lies in appropriately robust clinical trial design. Different designs are relevant to different objectives or phases of the trial. Adaptive designs are one specific type of clinical trial design where the key feature is to modify an ongoing trial in a pre-planned manner after reviewing accrued data at interim analyses. In a dose finding trial as the trial proceeds, information accumulates that reduces uncertainty regarding optimal treatment dose; adaptive clinical trials are designed to take advantage of this accumulating information, by allowing modification to the dose in response to accumulating information and according to predefined rules. The use of an adaptive design in dose finding allows a staged approach to trial design where a decision about what dose to give the next group of patients is made after each stage. By sequentially adapting the ethical objective of treating patients optimally at each stage can be met, while still achieving the main scientific objective in an efficient and reliable manner. A robust adaptive design is

intended to both enhance flexibility and efficiency without undermining the study's integrity and validity [19]. Efficient adaptive trial designs in early phase oncology are imperative to maximise the scientific benefit whilst reducing the risk:benefit ratio for patients entering such studies [20].

There are two main statistical approaches applicable to the design and analysis of a clinical trial, frequentist (or classical) and the Bayesian approach. The majority of earlier clinical trial methodology followed the frequentist approach with the use of the Bayesian approach rapidly becoming more prominent [21]. Bayesian methods have and continue to make a significant contribution to the field of health research [22]. The Bayesian approach allows a formal mathematical approach to incorporating prior information into trial design, analysis and decision making. The Bayesian approach is argued to be more flexible and more efficient in the use of data in contrast to the traditional approach [23]. This flexibility approach lends itself to adaptive design procedures [24], although many frequentist procedures also exist [19].

Within a dose finding trial there are multiple points where a decision needs to be made concerning the treatment of patients. Bayesian decision analysis is closely linked with Bayesian inference but the two disciplines are distinct [25]. The Bayesian decision theoretic approach is a statistical method to determine an optimal action from a set of possible actions when there is uncertainty [26]. There are two main components: A Bayesian model representing the structure of a system and its associated uncertainty and a utility function that is capable of measuring preferences relating to the consequence of taking a particular course of action [25]. A fully decision theoretic approach to statistical decision making is scientifically sound, providing coherent decisions when each of the two main components can sufficiently be determined [27].

The main premise of this thesis is a novel Bayesian decision theoretic approach to dose finding in oncology. It is hoped that by using a scientifically more robust method in contrast to a more ad-hoc procedure that the objectives of a dose finding trial can be more closely met. A key component of this is the patient objective of ensuring optimal dose allocation. In achieving improved trial design in this setting this adds to the much bigger objective of improving patient care through drug development. The work is split into four chapters and a discussion chapter that are briefly introduced below.

Chapter 2 reviews the current state of the art for statistical designs in the setting. The chapter starts with a preliminary section to give a more detailed overview of how the objectives associated with dose finding can be met with a staged design. A Bayesian approach to analysis and decision making is highlighted as most existing designs in the literature exhibit some features of the approach. Statistical concepts relating to defining an optimal dose with conflicting objectives are introduced. A review of the literature highlights that designs do not follow the Bayesian decision theoretic approach and that designs are a simplification of clinical preferences.

Chapter 3 gives a closer inspection of modelling in the setting of dose finding as a key component of a decision theoretic approach. Specifically how binary toxicity and efficacy endpoints, typically used in this setting, are jointly modelled to account for any interaction, with the possibility that an efficacy event becomes more or less likely in the presence of a patient toxicity event. The approach is tested to understand its influence on dose finding studies.

Chapter 4 reviews the relevant statistical literature for decision making. The purpose of this is to understand how clinical preferences with multiple competing objectives can be encoded into a utility function. This is necessary so that a decision theoretical approach can be undertaken. This literature hasn't previously been applied in the setting of dose finding and a strong theoretical foundation yields a consistent method that can be revised and adapted as needed for a particular application with confidence. This work has a practical element so that a statistician can ask appropriate questions to a clinician in order to elicit their preferences for an OD.

Chapter 5 proposes a novel Bayesian Decision theoretic approach to dose finding trials. The method builds upon the work in the previous chapter. The new method has an accompanying set of questions to ask clinical experts so that the method can be practically applied. The method is contrasted against a prominent alternative design in a worked example with some initial evidence that significant improvement is seen on some key metrics.

Chapter 6 is a discussion chapter summarising the work in the thesis and placing it into a wider context. The merits of the work are evaluated and recommendations for future work are made.

Chapter 2

Statistical designs for dose finding studies in oncology

The literature on statistical designs for dose finding studies in studies for cancer therapeutics is reviewed in this chapter with a particular focus on how decisions are made to determine which dose to treat patients entering the trial and the recommendations for patients in the future. There are two categories defined in this chapter to split designs based upon whether a design utilises a single toxicity endpoint or both a toxicity and an efficacy outcome. Designs that utilise a single toxicity endpoint are typically referred to as phase I trials while trials that utilise both a toxicity and an efficacy endpoint are referred to as phase I-II or phase I/II designs. There is an initial preliminary section to the chapter where a general approach to dose finding trials is introduced with a closer inspection of what constitutes an OD. The general notation for a Bayesian design to dose finding is given as it is the method utilised by most of the designs reviewed as part of the literature review. The chapter concludes with a motivating dose finding example for the thesis.

2.1 Preliminaries

2.1.1 An algorithm for dose finding

To satisfy the ethical objective of treating patients at doses that are not overly toxic and (or) not efficacious the trial proceeds in stages. The design is to treat cohorts of patients in a sequential manner with a decision as to what dose to treat the next cohort of patients

made after each stage. An algorithm that encapsulates a framework for study design for dose finding in oncology is defined in this section.

A drug is intended to have some effect on the body, with the effect changing depending upon the dose administered. The effects of a drug at a particular dose can be considered to affect the cancer in a positive manner or to affect the body negatively. The positive effect is described as efficacy and the negative effect toxicity. An endpoint is defined as part of the trial design that measures each of these effects for an individual patient. The design type described earlier as phase I that uses a single toxicity measure only makes an assumption about efficacy that is stated in the next section when exploring what defines an OD. The basic algorithm for trial design remains the same however. Let $D \in \{d_1 < d_2 < \cdots < d_k\}$ be a set of k pre-defined doses to be studied within a dose finding trial. Let $Y = (Y_E, Y_T)$ where

$$Y_E = \begin{cases} 1 & \text{if efficacy} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad Y_T = \begin{cases} 1 & \text{if toxicity} \\ 0 & \text{otherwise} \end{cases} \tag{2.1}$$

be Bernoulli random variables representing an efficacy and toxicity event for a patient respectively following treatment at a particular dose. Binary variables are typically used to measure the effects of treatment in this setting because they can be specified to obtain short term measures of effect to allow the trial to adapt to accumulating evidence and complete in an acceptable time window. Additionally it can be possible to encapsulate several measures into a single binary effect [28]. There are alternative measures of effect that are considered later in the chapter.

Each event definition will depend on the particular clinical setting, for example efficacy may be measured by response or progression-free survival at a particular time point. The toxicity event is typically described as a dose-limiting toxicity (DLT), a severe toxicity event. Such toxicities are assessed according to the National Cancer Institute's Common Terminology Criteria for Adverse Events (CTCAE) classification [29], and usually encompass all grade three or higher toxicities with some pre-specifed exceptions that can be managed [18].

The trial design for a dose finding trial will stipulate how each of the items below are specified in the trial protocol.

1. An initial cohort of patients are treated at a starting dose
2. Effects of the treatment for each participant in the cohort are measured by an outcome representing toxicity, or efficacy and toxicity.
3. The study then continues iteratively as follows:
    - A decision process to specify which dose to treat the next cohort of patients at is made with respect to the effects from the previous cohort(s).
    - The next cohort of patients is treated at the dose recommendation from the decision process.
4. The trial continues until a maximal sample size is achieved or an additional rule called a stopping rule is initiated. This stopping rule may indicate an OD has been found or that it is no longer beneficial to continue development of the trial due to excess toxicity and/or lack of efficacy.
5. The OD will be declared at the end of trial to treat patients in future trials. If the design stops early because excess toxicity and/or lack of efficacy then no dose is recommended for further study. In many designs the recommendation of an OD will follow the previously defined decision process and the dose for the next cohort of patients will be the OD. In other designs there may be separate criteria to the decision process used to evaluate which dose is most appropriate to take forwards.

2.1.2 Defining an optimal dose

Traditional designs for dose finding in oncology utilised only a toxicity endpoint. This class of design relies on an implicit assumption that more dose of the drug relates to increased anti tumour activity. The OD in a trial that is suitable for a toxicity only design will be the highest dose with an acceptable toxicity profile defined with respect to the trial endpoint [30, 31]. The highest dose with an acceptable toxicity profile is referred to as the Maximum Tolerated Dose (MTD). Designs that seek to find the MTD as the OD for further study are referred to as phase I trial designs. The exact definition of the MTD is defined as part of the trial design with a number of definitions existing [28]. A cytotoxic agent is toxic or deadly to cells and is associated with early compounds in the development of cancer therapeutics [32]. Cytotoxic agents likely have a narrow range of doses that are both sufficiently efficacious

and acceptably toxic with the OD coinciding with the MTD [33]. As such a toxicity outcome alone can capture the OD.

A statistical design for a phase I design will define the MTD with respect to a population level summary for the probability that a patient will experience the DLT event [18]. Let $\pi_T = \Pr(Y_T = 1|d)$ denote the probability of experiencing a toxicity event at a dose, d. A dose toxicity relationship inspects the probabilities of efficacy and toxicity across a range of doses. For the dose toxicity relationship the relationship is typically monotonic, i.e. the chance of a toxic event increases with dose [34]. Assuming that the risk of toxicity increases with dose one definition of the MTD is the dose level that results in the maximum risk of toxicity that is no bigger than an acceptable level, $\bar{p}_{MTD}$.

$$MTD = \max\{d \in D : \ \pi_T \leq \bar{p}_{MTD}\} \tag{2.2}$$

Alternative definitions of the MTD may be the dose with a probability of toxicity closest to some predefined toxicity called the target toxicity level. Assuming that efficacy also increases with dose the minimally effective dose is the dose that gives a minimum amount of efficacy. Doses between the minimally effective dose and the MTD are described as the therapeutic window. When the MTD isn't minimally effective, the therapeutic window does not exist and the agent isn't viable.

When extending the dose finding design to incorporate efficacy endpoints, statistical designs will define the OD with respect to a population level summary relating to the chance of an efficacy event occurring offset against a toxicity outcome. Designs that utilise both an efficacy and toxicity outcome are referred to as phase I-II or phase I/II designs. Concepts relating to defining an OD with respect to event probabilities are explored in this section. Assuming that both efficacy and toxicity endpoints are measured over a similar time period. Let $\pi_E = \Pr(Y_E = 1|d)$ denote the probability of experiencing an efficacy event at a dose, d, and $\pi_T = \Pr(Y_T = 1|d)$ denote the probability of experiencing a toxicity event at a dose, d. The two probabilities are a measure of effect in the wider population and are used as a measure to decide the merits of a particular dose. The dose toxicity relationship is typically monotonic [34]. Increasing the dose will increase the chance of a toxic event. The efficacy relationship is less likely to be monotonic for many classes of compound and could plateau or decrease with increasing dose [35].

Given the risk of mortality due to cancer [36], some toxicity is deemed acceptable if the treatment is efficacious. The OD can be considered as a compromise between the effects of efficacy and toxicity. The definition of what constitutes "optimal" forms part of the trial design with multiple examples given later in the chapter.

A statistical approach to defining an OD could be achieved by a mathematical function of efficacy and toxicity. The joint probability space for efficacy and toxicity is defined over the unit square $[0, 1]^2$. An objective function $O(\pi_E, \pi_T)$ defines a ranking for all possible combinations of efficacy and toxicity. A larger magnitude of the objective function is preferred to one that is smaller. The OD is the dose that ranks the highest. There are two properties of any objective function that should be self-evident. More efficacy is preferred to less, if considering two doses with the same risk of toxicity the dose with more efficacy should be preferred. Similarly less toxicity is preferred to more. If two doses have the same efficacy then the dose with less toxicity is preferred. Any definitions of optimal in a trial design should possess these properties.

One example definition of the objective function could be to define the patient outcome of efficacy without any toxicity event as the only truly successful outcome for a patient [37]. As such, assuming that the two events are independent the dose that has the highest probability of a successful outcome is the OD:

$$O(\pi_E, \pi_T) = \pi_E(1 - \pi_T) \tag{2.3}$$

with the OD satisfying

$$\arg\max_d O(\pi_E, \pi_T). \tag{2.4}$$

It is possible to plot objective functions, with possible values of efficacy on the x axis and possible values of toxicity on the y axis. Values that the objective function takes creates a surface with larger values representing a combination of efficacy and toxicity that is preferred. The best possible combination of efficacy and toxicity is when there is perfect efficacy and zero chance of toxicity, i.e when $\pi_E = 1$ and $\pi_T = 0$. The worst possible case is if a dose has zero chance of efficacy and guaranteed chance of toxicity i.e when $\pi_E = 0$ and $\pi_T = 1$. Any function that doesn't give these two points as the best and worst will violate one or both of the self evident principles. All possible combinations of efficacy and toxicity

for a dose will give a value for the objective function between these two extreme points. As the objective function is in two dimensions there will combinations where the objective function gives combinations of efficacy and toxicity that are equally desirable, these points of indifference can be displayed by a line called a contour. When plotting objective functions an arbitrary number of contours are plotted to facilitate visualising the surface that the objective function creates, Figure 2.1.

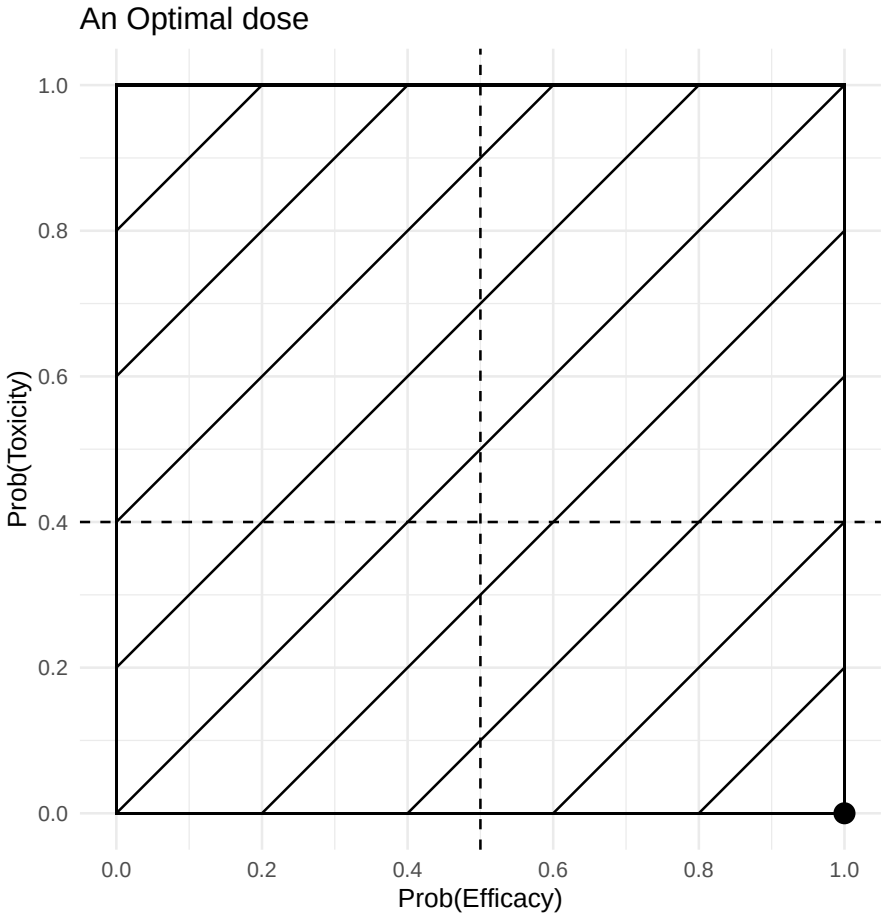

Figure 2.1: Example objective function. The probability space for efficacy and toxicity is defined over the unit square $[0, 1]^2$.Contours (solid lines) describe combinations of efficacy and toxicity that are equally desirable. Eight contours are arbitrarily drawn to represent the surface. In this example the surface is a ramp with the point at $(1, 0)$ representing perfect efficacy without toxicity as the most desirable, i.e every patient has an efficacy outcome and no patients have a toxicity outcome. Parallel contours nearer to lower right point are preferred to ones further away Admissibility criteria defined by $\bar{\pi}_E = 0.5$ and $\bar{\pi}_T = 0.4$ are represented by dashed lines. The admissibility criteria split the decision space into four quadrants. Doses that are constrained within the lower right quadrant are admissible. Doses outside of the quadrant cannot be considered to be optimal.

There may be additional constraints placed upon the objective function to define the OD. The OD in this situation becomes the dose that maximises the objective function subject

to satisfying a set of conditions called admissibility criteria. Doses that do not satisfy the criteria are described as inadmissible doses. One condition relating to the probability of toxicity would be the maximum amount of toxicity that would be deemed acceptable, $\overline{\pi}_{addT}$. Any dose with toxicity above this value would be deemed inadmissible. Similarly there may be a minimum amount of efficacy required for a dose to be considered acceptable. Any dose with efficacy below this threshold, $\overline{\pi}_{addE}$ is deemed inadmissible. An example of the two criteria is given in Figure 2.1. Doses where $\pi_E \leq \overline{\pi}_{addE}$ or $\pi_T \geq \overline{\pi}_{addT}$ are inadmissible and excluded when considering an OD.

This section provides a general framework for considering an OD statistically for a dose finding trial. Many of the statistical designs that are reviewed later in the chapter will follow these key ideas.

2.1.3 Bayesian approach to dose finding

The most common statistical approach to a trial with dose finding objectives is the Bayesian approach. The first design to do this was the phase I design, the continual reassessment method (CRM) [38] reviewed later in the chapter. This section gives the notation and short overview for the Bayesian method. The section ends by defining evidence levels from the Bayesian model to declare doses inadmissible as defined in the last section.

Features that are unknown about the external world are modelled by unknown states of nature $\theta \in \Theta$. In dose finding this may be parameters π_E and π_T for the probability of efficacy and toxicity at each dose or parameters associated with a dose response curve. In the case of a toxicity only design this may be π_T or parameters used to model the dose toxicity curve. The observation Y is drawn from a distribution $p(y|\theta)$ called the likelihood. Prior knowledge of $\theta \in \Theta$ is incorporated via a prior $p(\theta)$. This is updated through Bayes theorem in light of the observation(s), to give the posterior

$$p(\theta|y) \propto p(y|\theta) \times p(\theta). \tag{2.5}$$

This component is referred to as the statistical model with the details given as part of the design. The posterior distribution is used to make decisions at each stage and at the end of the trial. A utility function $u(d, \theta)$ specifies the utility of treating at dose $d \in D$ if the state of nature is $\theta \in \Theta$. This is similar to $O(d)$ defined previously in that it creates an

order of all combinations of efficacy and toxicity but the magnitude of the utility denotes a measure of preference. A more precise description of utility and how it differs to an objective function is the subject of Chapter 4. For this chapter, when referring to a design that maximises the expectation of a function, "utility" is used, otherwise the term objective function is used for consistency. There are a number of other terms that refer to functions used for making decisions in this setting including "Gain", "Loss" and "desirability", all have a similar interpretation.

The Bayes action (or decision) $d^* \in D$ is the action that maximises the posterior expected utility:

$$d^*(y) = \arg\max_d(E_\theta[u(d,\theta)|y]). \tag{2.6}$$

There are a number of designs that will specify a Bayesian statistical model but differ to the decision theoretic approach in the decision making. An alternative two stage approach used in a number of designs is to first find the posterior mean from the statistical model and then to use the estimate as an argument for the objective function. The dose that maximises the objective function based upon a posterior mean estimate defines the OD at each stage. i.e.

$$d^*(y) = \arg\max_d(O(d, E(\theta|y)). \tag{2.7}$$

The method is described as a hybrid Bayesian approach to decision making and is commented upon in the discussion of Chapter 4.

Admissibility criteria were introduced as part of defining an OD, these criteria are typically applied as admissibility or stopping rules to limit the inclusion of doses in the decision making. The admissibility criteria were defined earlier for π_E and π_T as minimum levels of efficacy and maximum levels of toxicity respectively for a dose to be considered acceptable. Admissibility rules formally use the Bayesian model and the criteria $\overline{\pi}_{addE}$ and $\overline{\pi}_{addT}$ to exclude doses from the decision process. This is done to avoid exposing patients to excessively toxic doses or doses that are in-efficacious. The posterior distribution given in Equation 2.5 is defined for each dose $d \in D$. It is possible to obtain marginal distributions for $\pi_E(d,\theta)$ and $\pi_T(d,\theta)$ denoting the probability of efficacy at a particular dose and the probability of toxicity. For brevity $\pi_E(d,\theta) = \pi_E$ and $\pi_T(d,\theta) = \pi_T$. Two further constants are specified to denote evidence levels p_E and p_T for a dose to be considered efficacious and safe.

Admissibility rules for efficacy and toxicity respectively are defined by

$$\Pr\{\pi_E < \overline{\pi}_{addE} \mid y\} > 1 - p_E \tag{2.8}$$

and

$$\Pr\{\pi_T > \overline{\pi}_{addT} \mid y\} > 1 - p_T. \tag{2.9}$$

If either rule is satisfied then the dose is considered inadmissible and is excluded when determining the OD.

2.1.4 Assessing a designs performance

A specified statistical design needs to demonstrate that it is capable of meeting its objectives. The primary method of doing this is by specifying some known "truths" around a trial process and assessing what the specified design does [39]. For dose finding the primary objective is to find an OD. How likely the design is to select the OD given some known dose effect relationships cannot be calculated exactly and simulation is used. A single simulation iteration mimics the proposed study but the data is generated using pseudo-random sampling with a known probability distribution. The use of pseudo-random sampling is referred to as Monte Carlo simulation. A simulation study conducts multiple iterations of the trial from simulated data and assesses the performance of the design using summary metrics; these are referred to as a designs operating characteristics. The use of simulation allows us to compare metrics from different designs, or different specifications of the same design, to demonstrate a study is capable of meetings its objectives. When different specifications of the same design are compared to select the design with the best operating characteristics this is described as the calibration of a design.

A data simulation model will specify the model associated with generating the binary patient responses in Equation 2.1. The prominent method in the dose finding literature is to specify a Bernoulli distribution with fixed probabilities for each dose. Specifying fixed effects and inspecting long run frequencies as described above is an inherently frequentist concept, but these can be used to assess properties of methods, even if the methods are Bayesian [40]. The fixed probabilities, known as a scenario, are chosen to represent a clinically plausible relationship between the doses. The process is repeated for a range of scenarios. In practice

little is known about the dose effect relationship; the number and range of scenarios should reflect this ignorance [41].

The main metrics used in the dose finding literature to assess design operating characteristics are described in this paragraph. The primary purpose of the simulation study in dose finding is to assess if it is capable of determining an OD. Secondary objectives for dose finding are linked to treating patients optimally on trial. Treating a large number of patients in and around the OD is desirable. To assess the ability of the design to meet the primary objective the metric of probability (or equivalently percentage) of correct selection is used. This will be the proportion of times from the high number of replicates that a design correctly chooses the dose that is optimal, predetermined from the scenario. Understanding the proportions of selection across different doses will also be used to understand what the design does when it doesn't select the OD. For example in the toxicity only setting a design that frequently selects a dose that is above the MTD is considered differently to one that has a tendency to select the dose below the MTD [42]. The average number of patients treated at each dose is used to assess how the design assigns patients to each dose level over the course of the trial. The patient objective is to treat patients optimally at each decision point. When trying to determine if this objective is achieved it is important to assess the proportions of selected doses at different decision points during the trial in addition to overall. For example a design may have a tendency to expose patients early in the trial to excessively toxic doses but overall perform similarly to one that is slower to escalate initially.

2.2 Toxicity-based designs

This sections reviews existing designs in the literature with a toxicity only endpoint. Phase I designs could fit into three broad categories: algorithm or rule based designs, model based designs and model assisted designs [43]. The key features of each of these designs with examples from the different categories of designs are explored in this section. A summary between the designs is given at the end of the section.

2.2.1 Rule-based designs

Rule based designs assume no statistical model for the outcome $Y = Y_T$ and do not have an explicit function to define the MTD. The classical design, or 3+3, involves starting at the

lowest dose $d = 1$ and following the rules of the schematic in Figure 2.2 until an MTD is identified [44]. The trial stops when there are no more dose levels to escalate or deescalate from. If the recommendation is to deescalate from the lowest available dose, $d = 1$ the trial will stop without recommending a dose for further study. The design has evolved over time through pragmatic considerations from treating clinicians rather than an underlying statistical basis. A systematic review of the use of phase I oncology trial designs with dose finding objectives found that 95% of trials in the previous two decades used the 3+3 design [28].

It is a misconception that the 3+3 design finds a dose with a 33% or 1/3 expected toxicity level. Simulation studies have demonstrated that the expected target toxicity of the MTD following the 3+3 design to be monotonically decreasing from about 30% to 0% as the number of dose levels, $d \in D$, increase from three to infinity [45]. There are many other variants on up-down designs with fixed rules for dose escalation or de-escalation (also referred to as A+B designs) [46]. Rule based designs are considered to be inferior to alternative statistical designs (model based and model assisted) given in the next section [47]. Rule based designs for phase I trials have inferior operating characteristics to alternatives with little or no ability to be flexible or the scope for extension [28].

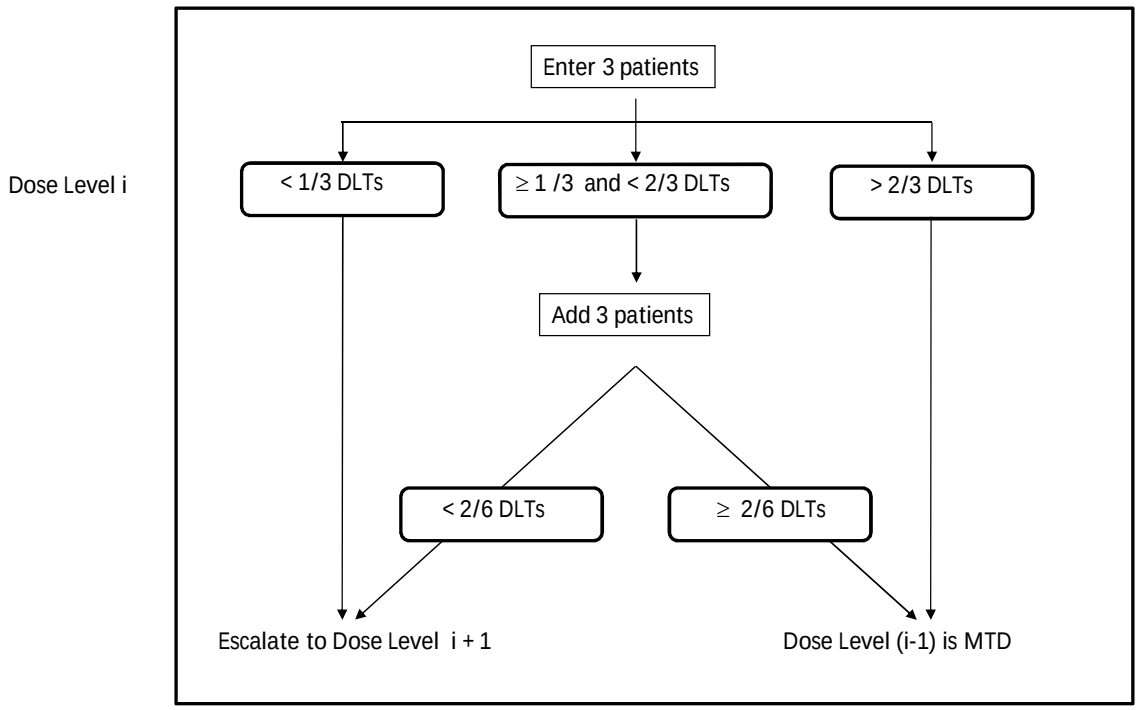

Figure 2.2: Design schematic of the classic 3+3 design

2.2.2 Model-Based designs

Model based designs have an underlying statistical model for the dose toxicity relationship. Decisions for each cohort and at the end of the trial will be derived from the statistical model. One of the first statistical designs in dose finding was the continual reassessment method (CRM) [38]. As a Bayesian design, CRM models the probability of toxicity for a given dose covariate. The model takes the form of a number of monotonically increasing functions such as the empirical function with the probability of toxicity at each dose equalling

$$\pi_T(x, \beta) = x^{exp(\beta)} \quad \text{for } 0 < x < 1 \tag{2.10}$$

where x represents a numeric covariate for each dose level called the "skeleton". The values of x at each dose are determined by initial estimates or values of the probability of DLT at each of the dose levels prior to commencing. These are back substituted into Equation 2.10 to determine the covariate value, x. The prior for β is normal with mean α and a standard deviation σ. The mean α is arbitrary because of the back substitution when assigning the skeleton and is set to zero.

The design originally proposed recruiting patients in cohorts of one and giving the next patient the dose believed to be closest to the target toxicity level, denoted by π_{MTD}. To achieve this a loss function is specified as follows,

$$L(\pi_T) = (\pi_T - \pi_{\text{MTD}})^2. \tag{2.11}$$

The loss function reverses the direction of the previously defined objective function in that a smaller value is preferred. As such it is the dose that minimises the loss function that is selected as optimal. There are two approaches to determine the dose that minimises the loss function, A "plug-in" approach or the Bayesian method. The Bayesian method minimises the expected posterior loss, while the plug in estimate will find the mean estimate of posterior distribution for β and plug this back into Equation 2.10 to give an estimate of the rate of toxicity at each dose. The dose that minimises the loss function is the dose to treat the next patient. The trial terminates after a fixed number of patients (N) have been recruited with the dose that would be used to treat the $(N + 1)$ patient as per the previous decision process declared the MTD.

There have been multiple extensions to the original CRM design. Most notably, and often referred to as the modified CRM (mCRM), pre-specified ad hoc rules of not skipping doses in escalation and recruiting in different size cohorts[48, 49]. These two features are common place in the dose finding literature, these are often imposed to reflect what would happen in practice and reflect compromises in complexity for models and decision criterion at the start of the trial with minimal data. Stopping rules have also been proposed to stop a trial early when all doses are too toxic or when there is sufficient evidence for a dose being the MTD [50, 51, 52]. The two stage CRM [53] allows a pre-specified initial escalation sequence until a DLT is observed. The specification of the prior skeleton has also been investigated extensively to give robust specification of the prior and skeleton [54].

The short comings of the rule based methods are strengths for the CRM method in that they are flexible with scope for extension [28]. Comprehensive reviews of the CRM in contrast to the 3+3 demonstrate improved ability to determine the OD and treating more patients in and around the MTD in nearly all scenarios [55].

Decision theoretic approaches have been applied to the problem [56, 57]. A two parameter logistic model was proposed by the authors as believed to give a wide class of sigmoidal curves appropriate in many clinical settings.

$$\pi_T(d) = \text{logit}^{-1}\eta(\theta, d) \tag{2.12}$$

with

$$\eta(\theta, d) = \theta_1 + \theta_2 \log(d) \tag{2.13}$$

A number of different utility functions were explored, in particular a patient and a variance utility function. The patient utility function corresponded to that used in the CRM, although maximises expected utility. The variance utility sought to minimise the variance at each decision stage. The notion of making decisions to gain information through minimising variance could be considered to violate the ethical principles of treating patients optimally set out earlier in the chapter.

A pragmatic paper built upon the Bayesian approach using two parameter logistic regression to support clinical decision making by presenting full posterior distributions [58]. The two parameter logistic regression model was utilised although the numerical dose component

was expressed as the log of the ratio between the dose and an arbitrary reference dose. This was done to give the intercept the more intuitive meaning of toxicity at the reference dose. Decisions were made according to the chance of being within four intervals to represent under dosing, targeted dosing, excessive toxicity and unacceptable toxicity. Each of these probabilities are inspected at each dose and depending on a decision rule the next dose selected. One example of the decision rule from the paper is the dose that maximises the probability of being within the targeted toxicity interval while satisfying a probability threshold of less than 25% chance of both excessive or unacceptable toxicity. The proposal was that full posterior distributions should be presented with an appropriate discussion with the clinical team enabling sensible dosing decisions to be made. Full decision theoretic utility functions were also applied to the four regions but considered to be challenging to specify.

Escalation with overdose control (EWOC) has been proposed where a piece wise loss function around the MTD was applied [59]:

$$L(\pi_T) = \begin{cases} (1-\lambda)(\pi_{\text{MTD}} - \pi_T) & \pi_T \leq \pi_{\text{MTD}} \\ \lambda(\pi_{\text{MTD}} - \pi_T) & \pi_T > \pi_{\text{MTD}} \end{cases} \tag{2.14}$$

where $\lambda > 1$. The two cases for the function are referred to as underdose, and overdose. The parameter λ is set to control how conservative an escalation sequence should be. The motivation was to have a parameter that was able to control the rate of escalation. The method specified a two parameter logistic regression function for the probability model.

This approach is similar to additional utility functions specified to be more conservative in the additional materials (but not evaluated) in the decision theoretic approach [56]:

$$L(\pi_T) = \begin{cases} (\pi_T - \pi_{\text{MTD}})^2 & \pi_T \leq \pi_{\text{MTD}} \\ C(\pi_T - \pi_{\text{MTD}})^2 & \pi_T > \pi_{\text{MTD}} \end{cases} \tag{2.15}$$

where $C > 1$, or

$$L(\pi_T) = \begin{cases} (\pi_T - \pi_{\text{MTD}})^4 & \pi_T \leq \pi_{\text{MTD}} \\ (\pi_T - \pi_{\text{MTD}})^2 & \pi_T > \pi_{\text{MTD}}. \end{cases} \tag{2.16}$$

The pragmatic Bayesian design, EWOC and the decision theoretic designs specify 2 parameter logistic regression models. There is some variation between how the models are parameterised and subsequent prior distributions but all specify a more complex model than is needed for the CRM. The one parameter model in the case of CRM is a reasonable approximation locally around the MTD in order to make robust decisions [60]. The pragmatic Bayesian approach to dose finding argued that many of the ad hoc rules necessary for CRM to work were as a result of overly simplistic decision models and dose–toxicity models. A two parameter logistic model was specified in this case to give a more accurate reflection of the full posterior distribution for the probability of a DLT at each dose to allow for the more complex decision making.

A relationship between conservatism in escalation at the cost of treating fewer patients at or around the MTD has been demonstrated for designs such as escalation with over dose control and the pragmatic Bayesian logistic regression design in comparison to CRM [59, 58]. Treating fewer patients in or around the MTD may seem something that is undesirable but the critical element is *when* these patients are treated. The designs that control escalation are slower to escalate and therefore when considering summary statistics for all trial patients they treat fewer patients at higher doses closer to the MTD. The two designs consider that when there is less information available to make decisions it is important to account for uncertainty in the decision for the next cohort of patients. This links to the ethical objective of not exposing patients to doses that are unacceptably toxic.

2.2.2.1 Model-assisted designs

This class of design could be considered a sub-set of model based designs with the major difference in that no intrinsic relationship is assumed between the doses [61]. They can be referred to as "short memory" designs with decisions based upon the observed data at any given dose. Doses are ordered in terms of potential for toxicity and the decision at each stage is made with respect to a statistical model that incorporates data from the current dose only. There are three potential actions, escalate to the dose above in the ordered list, remain at the current dose or de-escalate to the dose below in the ordered list. Further rules are introduced to exclude doses due to safety and to define the OD at the end of the trial.

Two examples of a model assisted design are the modified toxicity probability interval design

(mTPI-2) [62], and the Keyboard design [63]. Both designs were developed separately but are equivalent. They assume Bernoulli data with independent parameters for each dose and uniform prior distributions. Decisions are based upon specifying three regions to represent under, correct and over dosing. The decision at each stage is deescalate, remain or escalate. Decisions are made by segmenting the posterior into intervals of equal width to the region of correct dosing, the highest and lowest intervals may have smaller widths. The Keyboard design refers to the intervals as "keys". The unit probability mass is defined by the area under the curve contained within the interval divided by its width. The decision is determined by the location of the interval with highest unit probability mass. If the highest unit probability mass interval lies in the region of under, correct or overdosing the decision is to escalate, remain or de-escalate respectively (Figure 2.3).

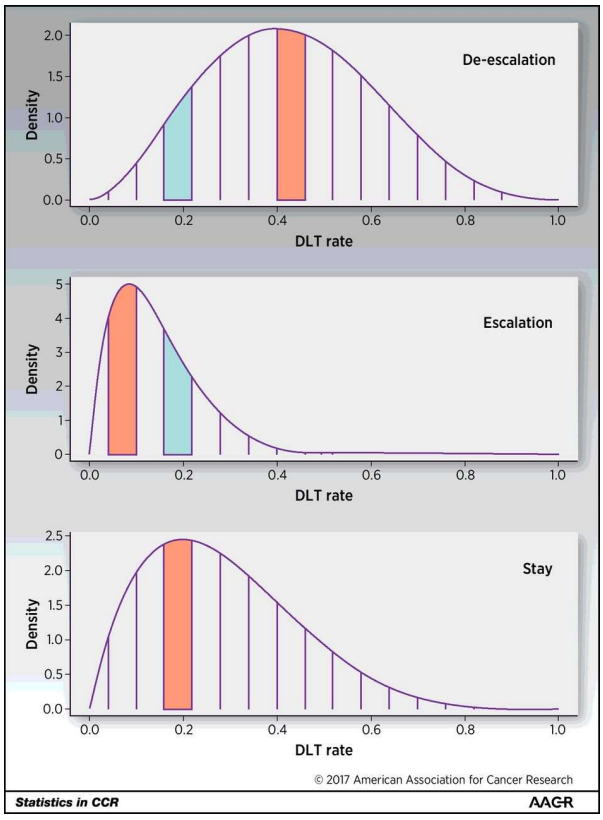

Figure 2.3: Posterior beta density plots of the DLT rate at the current dose level and the dose escalation/deescalation rules of the keyboard design [63]. The decision is made in relation to the positioning of the strongest key (red) in relation to the target key (blue).

To ensure patients are not treated at unsafe doses an admissibility rule is applied to the

DLT rate at each dose as per equation 2.9. If all doses are unsafe the trial will stop. The end of the trial, assuming at least one safe dose, will be declared after a fixed sample size has been reached. The decision to decide upon an OD is determined by isotonic regression [64]. Bayesian Optimal interval or BOIN design is similar to the Keyboard design with the decision at each stage based upon the observed DLT rate [65]. Parameters are set to construct an interval around the target toxicity level. Decisions then correspond with whether the mean observed DLT rate at a dose falls within the target interval (remain) below (escalate) or above (deescalate).

An advantage of model assisted designs is that they are simple to implement as decisions can be pre-specified in the protocol as a tabulation of number of patients treated and the number of DLTs observed or a simple diagram in the case of BOIN [66]. The FDA have granted the BOIN design the fit-for-purpose designation for dose finding, which has increased its significance and utilization in drug development programs [67]. Larger simulation studies demonstrate that mTPI, and BOIN designs have comparable overall performance finding the MTD reliably and treating a high number of patients in and around the MTD similar to the CRM [68, 69]. Model assisted designs use only the data at the current dose and ignore data from patients treated at alternative doses at each decision point. Safety rules mean that the design will limit exposure for patients to excessively toxic doses but by only considering data from the current dose level, it is challenging to argue that patients are treated optimally at each stage.

2.2.2.2 Phase I - is more better?

This section has reviewed designs for phase I trials with only a toxicity endpoint. Rule based designs were found to be unsuitable to achieve study objectives. The model based and model assisted designs were found to be suitable designs in the setting of phase I trials. A major considerations is whether more dose equates with preferable outcomes for patients. The Oncology Center of Excellence Project Optimus is an FDA initiative to reform the dose optimization and dose selection paradigm in oncology drug development [16]. This is in response to the changing drug-dosing conundrum in oncology that with modern agents with new mechanisms of action should not define the MTD as the OD [70]. A recent systematic review of dose effect relationships in oncology found that the assumption of more dose equating with improved efficacy outcomes to be violated [34]. Careful consideration should

be given at the design stage as to whether a trial using a toxicity outcome alone is suitable.

2.3 Efficacy and toxicity dose-finding designs

Clinical trials with dose finding objectives incorporating both an efficacy and toxicity endpoint are often referred to as phase I-II or phase I/II designs. Many of the designs are an adaptation of earlier dose finding studies incorporating a toxicity endpoint. The designs can similarly be split into rule based, model based or model assisted designs. A particular focus in this section is upon how the OD is defined.

The review of phase I designs found rule based designs to be unsuitable. Extensions to incorporate efficacy have similar shortcomings when compared to model based approaches and are briefly mentioned here for completeness. One of the first rule based dose finding designs was an application to bone marrow transplantation where the objective was to enable non-rejection (efficacy) without graft versus host disease (toxicity) [71]. Three different algorithms or A+B designs based upon number of efficacy or toxicity events in a given cohort were proposed and assessed through simulation. There was large variation between operating characteristics of the three contrasted designs, of particular note was that the first design based upon well-meaning heuristics performed poorly with modifications needed in the other two designs. In a situation when it is suitable to assume that no significant toxicity will occur other A+B designs have been proposed [72].

2.3.1 Model-based designs

2.3.1.1 Extensions to CRM

Ivanova proposed a design with an additional constraint to A+B rules incorporating the CRM [73]. Given the most recent patient treated at d_j, dose selection for the next patient is as follows: d_{j-1} is selected if there was a toxic event, d_j if there was a response without toxicity and d_{j+1} in the absence of a response and toxicity. A CRM model was fitted to the toxicity data as previously described and the next dose is selected as the lowest dose from either the CRM or the algorithm. A similar design was proposed by Hardwick, described as the directed walk design with the rule component also incorporating d_{j-2} [74]. The dose toxicity and dose efficacy curves were proposed by a number of parametric and non parametric functions with the dose maximising an estimate of the product of efficacy and

no toxicity selected for future patients.

The CRM was extended in the setting of a phase I/II trial in a new antiretroviral treatment for children infected with HIV [75]. The trial sought an OD based upon a maximum likelihood estimate of the probability of efficacy given no toxicity, $\hat{\pi}_{E|T'}$ and the probability of no toxicity $(1 - \hat{\pi_T})$

$$OD = \arg\max_d \{\hat{\pi}_{E|T'} * (1 - \hat{\pi_T})\} \tag{2.17}$$

The likelihood for the probability for toxicity followed the CRM with the dose covariates based upon the "skeleton". Efficacy given no toxicity followed a similar one parameter model. A two stage decision process was used to define doses with acceptable toxicity and a series of hypothesis tests to determine dosing at each stage.

An extension to CRM has also been proposed in the context of a trial of allogenic cell transplantation for high-risk Leukemia patients [76]. The joint distribution for efficacy and toxicity was defined as follows:

$$\pi(y_T, y_E|d) = k(\pi_E, \pi_T, \psi)\pi_E^{y_E}(1 - \pi_E)^{1-y_E}\pi_T^{y_T} \times (1 - \pi_T)^{1-y_T}\psi^{y_T y_E}(1 - \psi)^{1-y_E y_T} \tag{2.18}$$

where ψ represents the correlation between efficacy and toxicity and is assumed to be constant across doses. The constant $k(\pi_E, \pi_T, \psi)$ is a normalising constant. The prior for ψ is Uniform$(0, 1)$. The posterior mean of each of the parameters was calculated in order to estimate probabilities of efficacy and toxicity at each dose, $\hat{\pi}_E$ and $\hat{\pi}_T$. The trial sought to minimise the following criterion at each stage with respect to desirable efficacy and toxicity constants π_E^* and π_T^*:

$$OD = \arg\min_d \{w(\hat{\pi}_T - \pi_T*)^2 + (1 - w)(\hat{\pi}_E - \pi_E^*)^2\} \tag{2.19}$$

where $w \in [0, 1]$ describes the payoff between efficacy and toxicity (Figure 2.4). Admissibility rules for toxicity were part of the decision process. Here, the definition of the OD is inconsistent with a fundamental concept that more efficacy is preferred to less. The design implies that a dose with perfect efficacy and zero toxicity is less desirable than the target at the centre of the circular contours.

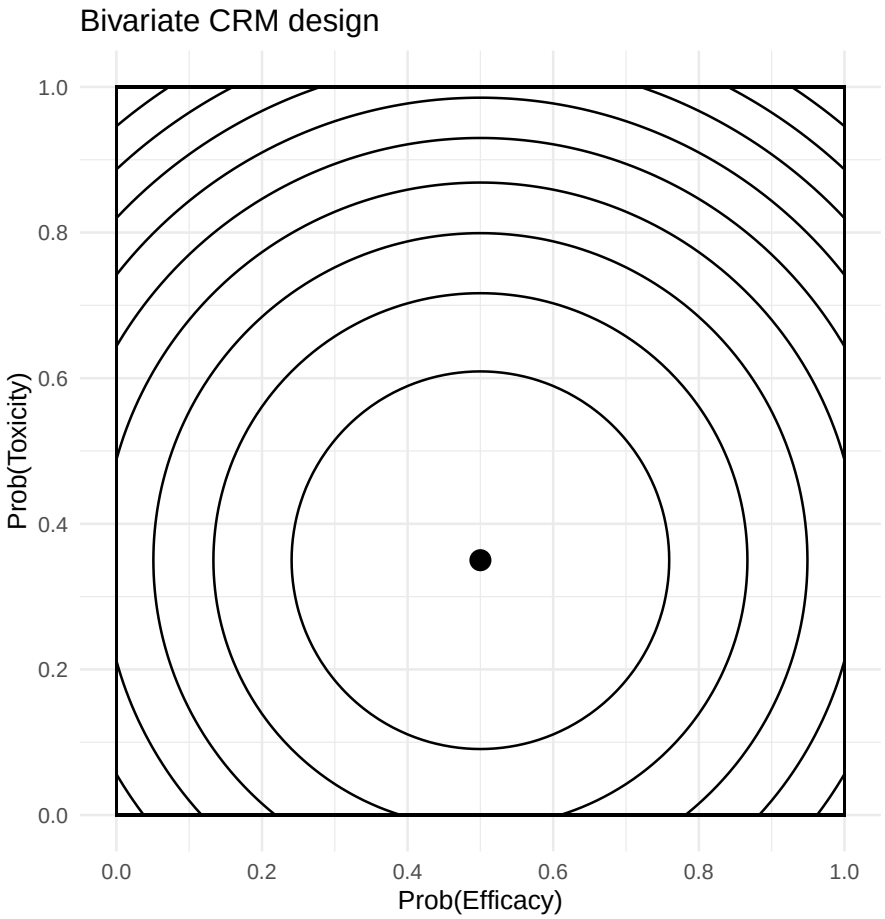

Figure 2.4: Braun bivariate continual reassement method with $(\pi_E^*, \pi_T^*) = (0.5, 0.35)$ and $w = 0.5$. Lines describe equal desirability, contours closer to the point (0.5,0.35) more desirable. This violates the fundamental idea that more efficacy is preferred to less and less toxicity is preferred to more, with a dose having 50% efficacy and 35% toxicity preferred to a dose with perfect efficacy and zero toxicity

2.3.1.2 EffTox design

Thall proposed a Bayesian adaptive design utilising correlated binary efficacy and toxicity endpoints called the EffTox model [77]. The design forms the basis for multiple extensions described in the book: Bayesian Designs for Phase I-II Clinical Trials [78]. The design was one of the first to incorporate efficacy and toxicity endpoints modelled as binary endpoints given in Equation 2.1. Many subsequent designs all possess the same basic structure of three components. These are the probability model, the constraints to determine an admissible set of doses and a objective function described as efficacy-toxicity trade off contours. The objective function specifies the risk:benefit ratio of a given dose and is used to select the dose after each stage and the recommended dose at the end of the trial.

The probability model is defined as follows: a set of numeric doses, $d \in \mathbb{R}_{>0}$, are transformed

by centering around the geometric mean.

$$x_j = \log(d_j) - \frac{1}{k}\sum_{r=1}^{k}\log(d_r) \qquad j = 1 \dots k \tag{2.20}$$

Marginal probabilities for efficacy and toxicity at each dose are defined with an inverse-logit link function

$$\pi_{E,j} = \text{logit}^{-1}\{\mu_E + \beta_{E1}x_j + \beta_{E2}x_j^2\} \tag{2.21}$$

$$\pi_{T,j} = \text{logit}^{-1}\{\mu_T + \beta_T x_j\}. \tag{2.22}$$

A bivariate distribution for the probability of any event $Y = (Y_E = a, Y_T = b)$ is defined using the Fairlie-Gumbel-Morgenstern (FGM) copula [79, 80].

$$\pi_{a,b} = (\pi_E)^a(1-\pi_E)^{1-a}(\pi_T)^b(1-\pi_T)^{1-b} + (-1)^{a+b}\pi_E(1-\pi_E)\pi_T(1-\pi_T)\psi \tag{2.23}$$

where

$$\psi = \frac{e^\phi - 1}{e^\phi + 1}. \tag{2.24}$$

Model parameters for the design are defined by $\boldsymbol{\theta} = (\mu_E, \beta_{E1}, \beta_{E2}, \mu_T, \beta_T, \phi)$ and data by $\mathcal{D}_n = (Y_i, x_i)$, the posterior for patients $i = 1, \dots, n$ is given via Bayes theorem, with the prior for $\boldsymbol{\theta}$ following independent normal distributions with corresponding hyper parameters for the mean and variance.

To create an ordering of preference over all possible combinations of π_E and π_T, a family of trade-off contours is defined. The corner $(1, 0)$ of the probability space is most desirable, i.e every patient has an efficacy outcome and no patients have a toxicity outcome. In the cited paper this was achieved using a quadratic function to define a contour and a Euclidean distance to create an order. Subsequent findings from the author found that this method led to undesirable operating characteristics and suggested that the contour be created using L^p vector norms [81]. An objective function using L^p norms is described as follows:

$$O(\pi_E, \pi_T) = 1 - \left(\left(\frac{\pi_E - 1}{\pi^*_{1,E} - 1}\right)^r + \left(\frac{\pi_T}{\pi^*_{2,T}}\right)^r\right)^{\frac{1}{r}} \tag{2.25}$$

where $r > 0$ (Note the Euclidean distance is a special case with $r = 2$). An initial target contour, $\mathcal{C}$, is used to define the constants and is calculated from three elicited probability pairs

$(\pi^*_{1,E}, 0)$, $(1, \pi^*_{2,T})$ and $(\pi^*_{3,E}, \pi^*_{3,T})$ and solving $O(\pi^*_{3,E}, \pi^*_{3,T}) = 0$ to find r. The objective function defines an ordering for any combination of efficacy and toxicity in the unit square. An example with $(\pi^*_{1,E}, 0) = (0.5, 0)$, $(1, \pi^*_{2,T}) = (1, 0.65)$ and $(\pi^*_{3,E}, \pi^*_{3,T}) = (0.75, 0.25)$ to give $r = 0.848$ is plotted in Figure 2.5 with an arbitrary number of parallel contours.

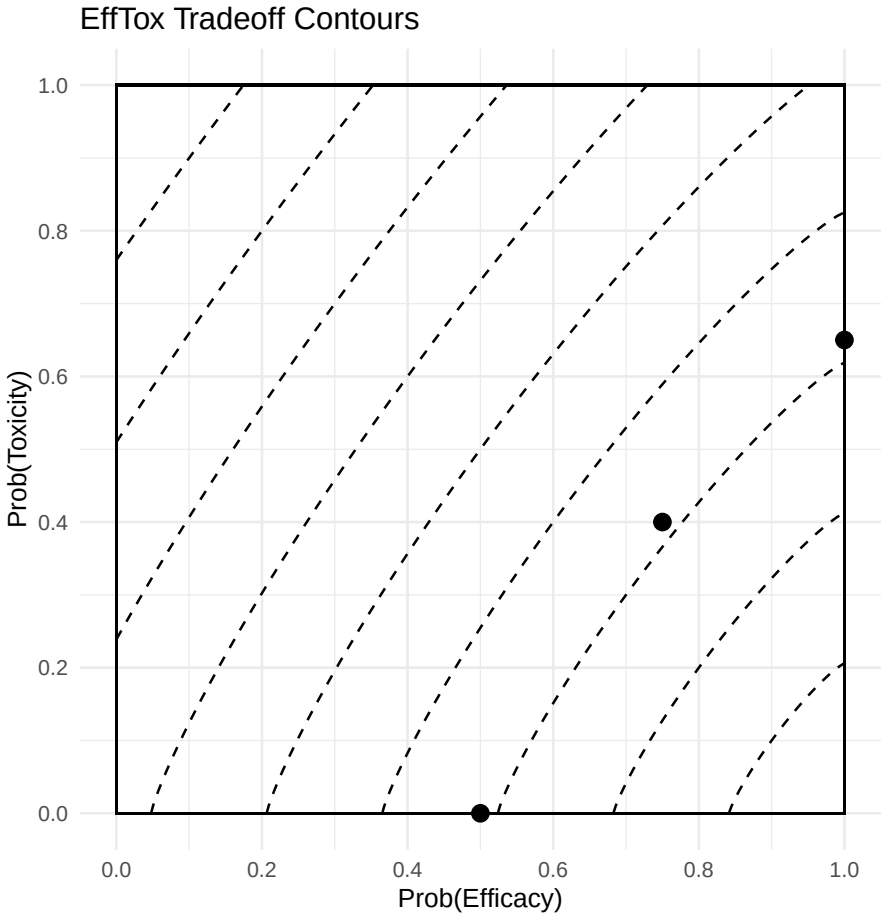

Figure 2.5: EffTox trade-off contours with $(\pi^*_{1,E}, 0) = (0.5, 0)$, $(1, \pi^*_{2,T}) = (1, 0.65)$ and $(\pi^*_{3,E}, \pi^*_{3,T}) = (0.75, 0.25)$ (plotted points) to give $r = 0.848$. Lines describe equal desirability, contours closer to bottom right corner (1,0) more desirable

At each decision point, admissibility rules are applied to all doses as per Equation 2.8 and Equation 2.9. Those doses that are inadmissible are excluded from the decision making process. The most desirable dose is determined from the posterior mean probabilities and selecting the maximum amongst the acceptable doses according to the objective function given in Equation 2.25, i.e.

$$\max_j O(E\{\pi_{E,j}(\mu_E, \beta_{E1}, \beta_{E2})|\mathcal{D}_n\}, E\{\pi_{T,j}(\mu_T, \beta_T)|\mathcal{D}_n\}) \tag{2.26}$$

for $j = 1, \ldots, k$. Note that this is not the Bayesian approach of maximising expected utility.

The expectation is evaluated for the parameters of the probability model first, then the objective function is minimised using estimates of each parameter.

2.3.1.3 EffTox utility design

Utilities designs based upon the four elementary patient outcomes have previously been proposed as an alternative to trade off contours [82, 78]. The probability model is the same as EffTox in the case of two binary responses but uses a decision theoretic decision process. The utility function is defined as follows:

$$u(Y_E = a, Y_T = b) = \begin{cases} K(1,1), & \text{for } a = 1 \text{ and } b = 1 \\ K(0,0), & \text{for } a = 0 \text{ and } b = 0 \\ K(1,0), & \text{for } a = 1 \text{ and } b = 0 \\ K(0,1), & \text{for } a = 0 \text{ and } b = 1 \end{cases} \tag{2.27}$$

Constants $K(a, b)$ are described as the numerical utilities of a patient achieving one of the elementary outcomes, $Y = (Y_E = a, Y_T = b)$. Given maximising expected utility is invariable to linear transformations, the worst and best of the four possible outcomes can be given a utility $K(0, 1) = 0$ and $K(1, 0) = 1$ with $K(0, 0)$ and $K(1, 1)$ to be defined. The calculation of expected utility at each dose is

$$E(u(Y)) = \int_\theta \sum_{a=1}^{1} \sum_{b=1}^{1} K(a, b)\pi(a, b) \tag{2.28}$$

The utilities are elicited from the clinician as a score relative to $K(0, 1) = 0$ and $K(1, 0) = 1$ to quantify the risk-benefit trade-off under each outcome. Based on experience, clinicians can easily comprehend the meaning of utility scores and provide specifications that align with clinical judgments [83]. It is also suggested that $K(1, 1) > K(0, 0)$ to reflect that achieving a response is typically more clinically beneficial.

The correlation component of the probability model is incorporated into the utility function as the joint probability of each patient outcome depends on this parameter (Equation 2.23). Note that this approach is more akin to a Bayesian decision theoretic approach as the decision rule is to maximise the expected utility at each stage.

2.3.1.4 Robust Bayesian design

A robust Bayesian design modelling toxicity and efficacy using a flexible non-parametric dynamic model has been proposed [84]. Here, the model borrows some information across doses without imposing a more stringent parametric form. The utility function incorporates aspects of the admissibility rules by creating a penalty when a toxicity threshold is surpassed as follows:

$$U(\pi_E, \pi_T) = \pi_E - w_1\pi_T - w_2\pi_T I(\pi_T > \overline{\pi}_{addT}) \tag{2.29}$$

The weights w_1 and w_2 are constants to be determined as part of the design work up. The function is not continuous at the threshold point with a jump in utility as shown in an example with $w_1 = 0.33$, $w_2 = 1.09$ and $\overline{\pi}_{addT} = 0.35$, Figure 2.6.

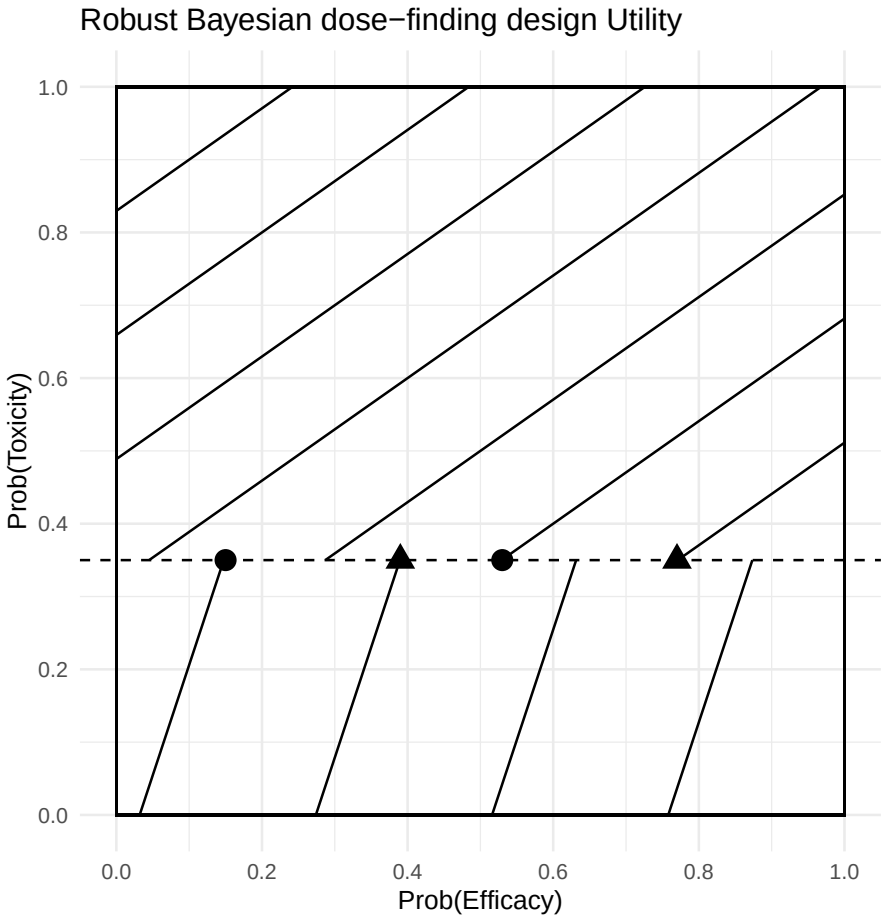

Figure 2.6: Robust Bayesian Approach Utility function with $w_1 = 0.33$, $w_2 = 1.09$ and $\overline{\pi}_{addT} = 0.35$. Lines describe utility equal to (0.1,0.2,..0.9), utility at (1,0) is 1. The points indicate the jump in utility

2.3.1.5 Decision theoretic designs

The decision theoretic approach from toxicity only designs were extended by the same corresponding author to include efficacy endpoints by utilising the ordinal outcome model [85, 86]. The probability model for toxicity was the same as the toxicity only design given in Equation 2.12. The efficacy component was conditional upon not having a toxicity event and modelled using a logistic regression model. The decision rule was based upon gaining information (variance) at each stage and a final trial stopping rule to stop the trial once a precise enough estimate of the MTD is obtained. Stopping rules limited exposure to doses with excessive toxicity. A patient gain function was also proposed in the latter publication to maximise the probability of efficacy without toxicity. A similar probability model was utilised with an objective function to balance interests of the patients both within and outside the trial, with different weights applied to patients inside and outside of the trial [87]. The interests were defined with respect to a reward function based around the probability of efficacy without toxicity, and with toxicities at each dose. The utility design described by four patient outcomes earlier could be considered a decision theoretic design as it maximises the expected utility [82, 78]. A similar design to patient level outcomes has also been utilised but in this case limiting the decision space to one of three possible actions: escalate, remain or de-escalate from the current dose level [88]. The probability model was Dirichlet at at each dose with four parameters for each of the potential patient outcomes, with the resulting design requiring a 3×4 utility table to be specified. A consequence function associated with continuous thresholds that cause each of the patient outcomes has also been proposed as a decision theoretic design in this setting [89].

2.3.2 Model assisted designs

The mTPI/Keyboard design has been adapted to include efficacy endpoints in the toxicity and efficacy interval design (TEPI) [90]. The toxicity rate is split into four intervals to denote low, moderate, high and unacceptable toxicity. The efficacy is similarly split into four intervals to denote low, moderate, high and superb efficacy. A (4×4) table splits the joint outcome space into sixteen distinct regions with a corresponding decision elicited from the clinician for each square. The possible decisions at each stage are escalate, remain or deescalate. The joint outcome is modelled using the product of independent beta distributions for efficacy and toxicity. The decision rule corresponds with a region having the highest joint

unit probability mass defined as the ratio of the probability of being within a region divided by its area. Admissibility rules for efficacy and toxicity are imposed to restrict treatment at a dose too toxic or not efficacious as per the admissibility rules in Equation 2.8 and 2.9. The decision at the end of the trial is based upon the admissible dose set that maximises a segmented joint utility function (Figure 2.7):

$$u(\pi_E, \pi_T) = u(\pi_E)u(\pi_T) \tag{2.30}$$

where

$$u(\pi_E) = \begin{cases} 0, & \pi_E \in [0, e_1] \\ \frac{\pi_E - e_1}{e_2 - e_1}, & \pi_E \in (e_1, e_2) \\ 1, & \pi_E \in [e_2, 1] \end{cases} \tag{2.31}$$

and

$$u(\pi_T) = \begin{cases} 1, & \pi_T \in [0, t_1] \\ \frac{\pi_T - t_1}{t_2 - t_1}, & \pi_T \in (t_1, t_2) \\ 0, & \pi_T \in [t_2, 1] \end{cases} \tag{2.32}$$

with e_1, e_2, t_1 and t_2 constants to define intervals which are specified as part of the design process. An example is plotted in Figure 2.7, with constants set to $e_1 = 0.5$, $e_2 = 0.8$, $t_1 = 0.15$ and $t_2 = 0.45$, representing a range of plausible efficacy and toxicity. Part of each contour has a segment of vertical or horizontal lines. The interpretation in the case of a vertical line is indifferent to more toxicity provided it is below 0.15 and in the case of a horizontal line there is no merit in additional efficacy above 0.8. The utility is maximised, $u(\pi_E, \pi_T) = 1$, for all combinations of efficacy and toxicity in the region where $\pi_T \leq 0.15$ and $\pi_E \geq 0.8$. The utility function does not strictly satisfy the self evident rules of preference for any combination of values for e_1, e_2, t_1 and t_2, that is, there will be regions of the utility function where more efficacy is not preferred to less, and regions where less toxicity is not preferred to more. The TEPI design has been adapted to change how the intervals are specified to improve upon operating characteristics without a change to the utility function [91].

An extension of the TEPI design is the utility-based toxicity probability interval design [92]. In this design the 4 patient utility function, given in Equation 2.28, is maximised to

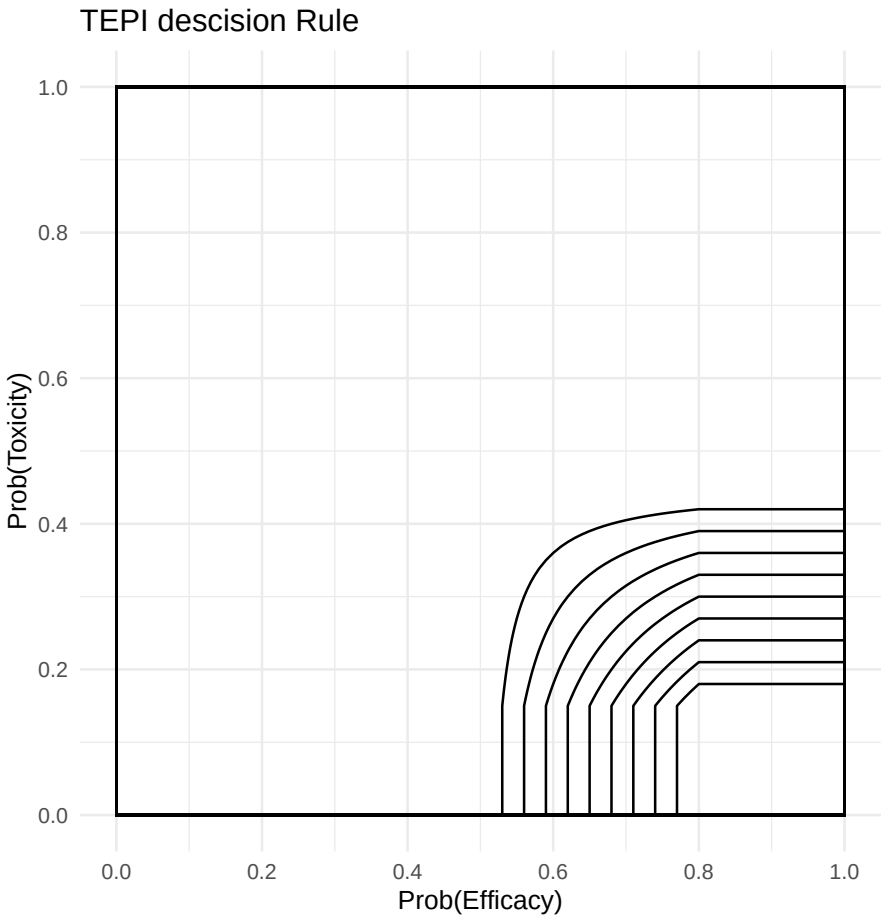

Figure 2.7: Toxicity and efficacy interval design final dosing decision utility function (Equation 2.30) with $e_1 = 0.5, e_2 = 0.8, t_1 = 0.15, t_2 = 0.45$. Lines are equal utility at (0.1,0.2,...,0.9), contours closer to the point (1,0) have larger utility.

make decisions subject to acceptable toxicity. The design leads to superior and more robust operating characteristics over the TEPI design. A notable feature of the design is the ability to tabulate all possible decisions in a clinical protocol.

The BOIN design has also been extended to include an efficacy endpoint. with the observed efficacy rate at a dose providing further rules as to whether to escalate, stay or deescalate from the current dose level. The design is referred to as BOIN-ET [93]. The STEIN design is very similar to BOIN-ET but differs slightly in how the interval to determine escalate, de-escalate and stay at the same is constructed [94]. The OD is selected at the end of the trial by non-parametric modelling for both outcomes; doses below the estimated MTD define a subset from which the dose that maximises efficacy is selected as the OD.

Further extension to the BOIN method for efficacy and toxicity outcomes include the Utility-BOIN or U-BOIN method [95]. The method is a two stage design, whereby a toxicity only

dose finding design, using BOIN, is specified before preceding to a second stage, utilising efficacy and toxicity endpoints. The use of two stage approaches is commented upon further in Section 2.4. There are three decision functions that are specified as part of the design that are of interest for the chapter. Decisions are made according to a hybrid Bayesian approach whereby the expectation of parameters from the probability model is found and inputted into the utility function. The use of "utility" in describing the method is not consistent with the rest of this chapter, where the term has only been used when maximum expected utility is used to make decisions. There are three decision making functions specified. The first two were specified in earlier papers associated with parametric models and are described in Equations 2.28 and 2.29. The last uses a single parameter to describe a payoff between efficacy and toxicity.

$$u(\pi_E, \pi_T) = u(\pi_E) - \omega u(\pi_T) \tag{2.33}$$

This is shown to be a special simplified case of Equation 2.28 in the paper. The utility function has a utility contour plot very similar to Figure 2.1 with the parameter ω responsible for the gradient of the linear contours. Simulation studies indicate that U-BOIN is more accurate in identifying the optimal dose and exhibits greater robustness compared to a more complex model-based phase I/II design described in Section 2.3.1.2.

It was highlighted, when looking at model-assisted designs for toxicity only, that the dosing decision may not be optimal at each stage. This is amplified in the setting of bivariate dose finding where there isn't a strict ordering of preference for the doses. The designs define what is optimal for the dosing decision at each stage, and what is optimal for the decision at the end of the trial, differently. Considering the objective of treating patients optimally at each stage of a dose finding trial the proposed model assisted designs do not meet this objective. This is because at any given stage the decision only considers data from the current dose, it may be the case that by considering the totality of the data a different recommendation could be made. This effect has been investigated numerically in a further extension to the BOIN design [96], the BOIN12 design. The BOIN12 design incorporates the utility function described Equation 2.28) to make decisions between doses at each stage. Using the toxicity only BOIN design, if the observed toxicity rate at the current cohort is contained or below the interval then the next cohort of patients is selected according to the utility function.

2.4 Extensions

The basic design follows the initial staged algorithm for binary efficacy and binary toxicity, there are a number of extensions or changes to the algorithm that are detailed here. There are designs that relax the component of treating each cohort of patients at the OD each time. The declaration of Helsinki states "While the primary purpose of medical research is to generate new knowledge, this goal can never take precedence over the rights and interests of individual research subjects" [97]. Prioritising the scientific objective of OD optimisation at the expense of patients on study may be challenging to justify to an ethics committee in practice. Careful consideration is needed around the trial objectives and the evidence base. The decision theoretic design introduced earlier gave the main approach of minimising variance at each stage [85, 86]. This is an example of a design that could be considered unethical due to not treating patients optimally at each stage. The original design paper gave an example of a trial for healthy volunteers in the setting of inflammatory diseases where this may be more acceptable due to the chronic nature of the disease. Safety rules were additionally stipulated as part of the design. Penalised optimal designs similarly seek to maximise information at each stage with penalties for unsafe and non-efficacious doses [98].

The review of statistical designs in this chapter has assumed binary efficacy and toxicity events observed over a similar relatively short time period to measure the effects at each dose. Continuous efficacy events with a binary toxicity event is one example of an alternative outcomes model [99]. Measuring toxicity as a ordinal variable is another possibility [100]. When events are observed over a longer period of time it is often not possible to conduct sequential recruitment over a feasible time period. A version of the CRM that weights events according to followup is one approach to allow decision making with incomplete data [101]. Similar approaches have been proposed for efficacy [102].

The designs considered so far have assumed that a patient's response is dependent upon the dose given. A patient's response could also additionally depend upon a covariate, such as tumor stage. The EffTox approach was extended to allow the OD to also depend upon the covariate [103]. Historical data in the form of informative priors helped to guide decisions with the design allowing the elicited contour and admissibility rules to vary based upon the covariate. An alternative design has been applied for trials where there are multiple similar

strata, this could be indications, regions or subgroups [104]. Assuming the number of strata is small, borrowing information across different strata could be of interest, particularly within rare populations. The phase I approach jointly modelled patient data assuming that patient could be both homogeneous or heterogeneous units simultaneously with a mixing component. An example of designs accounting for additional sub groups or covariates could be seen as a step towards a personalised dosing approach [105].

2.5 Motivating example and thesis case study

The motivation for the work in this thesis came from designing a dose finding study through Leeds Institute of Clinical Trials Research. The study was in relapsed-refractory multiple myeloma, a cancer of the plasma cells with an aim of investigating four doses of a treatment in combination with fixed dose standard of care therapies. The study is used as a case study throughout the rest of this thesis. To determine an OD it was felt that the higher doses, if tolerated, wouldn't necessarily give significantly more efficacy, with the interpretation that the efficacy dose response curve isn't monotonic. A phase I-II design was deemed appropriate with the toxicity endpoint a binary indicator of whether a DLT is experienced in the first two, four-week cycles. The efficacy endpoint was binary as to whether the patient achieved a "very good partial response" or not within the same time period. The EffTox design detailed earlier in the chapter was considered [77]. As part of the consultation with the clinician the objective function sought answers to the following questions specified in the paper in order to elicit three points to define the constants in Equation 2.25. The questions were proposed to the clinical team as follows:

1. The smallest efficacy probability considered desirable if toxicity were impossible (i.e. with no toxicity, what is the lowest chance of efficacy you would accept?).

2. The maximum toxicity probability considered acceptable if we have perfect efficacy (i.e. with perfect efficacy, what is the maximum chance of toxicity you would accept?)

3. A point that is "equally desirable to the first two" but between them (i.e. has some chance of toxicity, and less than perfect efficacy, but is just as desirable as the above scenarios; what combination of toxicity and efficacy would be desirable?)

The elicitation exercise generated the three points for (efficacy, toxicity): $(0.25, 0), (1, 0.4)$

and $(0.5, 0.25)$, as answers to each of the questions. Additionally admissibility constants were $\overline{\pi}_{addE} = 0.25$ and $\overline{\pi}_{addT} = 0.3$ defining the minimum efficacy and maximum toxicity acceptable to treat patients (Equations 2.8 and 2.9). The initial objective function is plotted in the left pane of Figure 2.8. The performance of the design was assessed by simulating a number of scenarios with different dose efficacy and dose toxicity relationships. The operating characteristics gave a high number of recommendations for low doses when the higher doses were considered optimal and few patients treated at higher doses. The specified design performed poorly with a tendency to get stuck at lower doses in scenarios where the highest and second highest doses from the four doses were determined to be the OD.

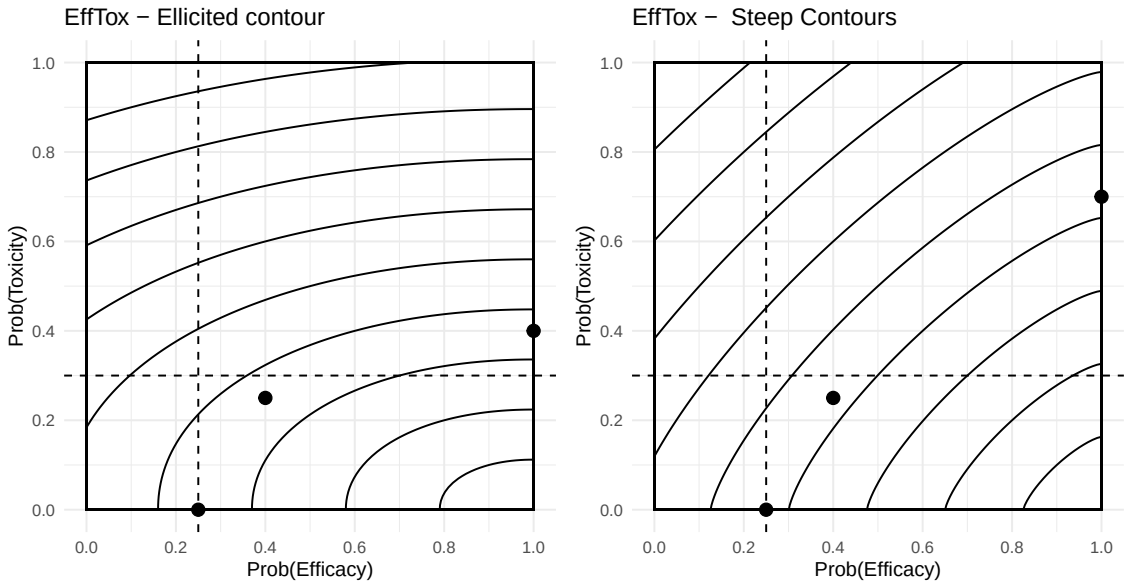

Figure 2.8: Left Pane: EffTox trade-off contours (solid lines) with elicited points $(\pi^*_{1,E}, 0) = (0.25, 0)$, $(1, \pi^*_{2,T}) = (1, 0.4)$ and $(\pi^*_{3,E}, \pi^*_{3,T}) = (0.5, 0.25)$. Lines describe equal desirability, contours closer to bottom right corner (1,0) more desirable. Dashed lines represent admissibility criteria $\overline{\pi}_{addE} = 0.25$ and $\overline{\pi}_{addT} = 0.3$. Right Pane: EffTox trade-off contours (solid lines) with points $(\pi^*_{1,E}, 0) = (0.25, 0)$, $(1, \pi^*_{2,T}) = (1, 0.7)$ and $(\pi^*_{3,E}, \pi^*_{3,T}) = (0.5, 0.25)$. Lines describe equal desirability, contours closer to bottom right corner (1,0) more desirable. Dashed lines represent admissibility criteria $\overline{\pi}_{addE} = 0.25$ and $\overline{\pi}_{addT} = 0.3$.

In the EffTox method and the utility extension the authors stress the importance of contour specification. Contours that are "Insufficiently steep" will lead to "pathological behaviour". Pathological behaviour describes the tendency of a design to repeatedly recommend a low dose without exploring higher doses. From a visual inspection of the contours in the left pane of Figure 2.8, the gradient of the right edge of the contours isn't very steep in contrast to the

left edge. This would constitute an "Insufficiently steep" contour. The individual elicited contour described the situation where the clinician was indifferent, but further contours are then extrapolated as part of the design to represent the objective function. The contours in the lower right quadrant would constitute doses that are admissible. In this region the flat right hand edge of the contour represents that a small increase in toxicity is worse than the same increase in efficacy. If asked to consider two points on any given contour in this region the clinician had a strong preference for doses towards the right hand side of the curve. This suggests that the contours do not represent lines of equal preference. Indifference contours in the lower right quadrant were better represented as steeper as plotted in the right pane of Figure 2.8. When considering the upper right quadrant of this design however, the contours do not represent clinical preferences, here considering doses on a given contour in the upper right quadrant there is a strong preference for doses to the left of the contour where toxicity is lower and more acceptable.

This created a dichotomy where the choice of which specification to choose was not clear. It would be desirable to design a trial that captured clinical preferences in this setting.

2.6 Discussion

In this chapter statistical trial designs to meet dose finding objectives in oncology were summarised. Broad concepts of what an OD is and the trial objectives associated with finding the OD were given. Given the low evidence base at this early stage of development of a compound it is imperative that patients are treated optimally. An overarching adaptive and staged approach was stated in order to meet trial objectives. A more formal statistical, specifically Bayesian, approach was detailed which most designs follow. The chapter has split the literature into two major classes of dose finding designs, phase I designs with a univariate toxicity endpoint and phase I-II designs incorporating both an efficacy and a toxicity endpoint. The choice of phase I design depends on whether toxicity alone can be assumed to be a surrogate endpoint for efficacy and subsequently an OD.

While this chapter focuses on the statistical designs there are pragmatic considerations of implementing a design. A recent review of the slow uptake of novel dose finding designs found that the main barriers to further uptake are lack of expertise in the trial and clinical team and limited resources for study design [47]. This increased complexity is necessary

to ensure the scientific validity of the trial and importantly ensure that patients within the trial are treated ethically. The traditional 3+3 is arguably the most straightforward design to implement and the reason it remains the most utilised designs [106]. This pragmatism is offset by its vastly inferior capability of meeting the scientific and patient objectives of the study. Model assisted designs are proposed to bridge the gap between simplicity and more complicated statistical designs as encouragement for more practitioners to use designs that better meet study objectives [66]. The importance of designs that are easy to implement can be seen with the rapidly expanding literature for model assisted designs as is seen in Section 2.3.2.

There is a rapidly expanding field of designs to meet dose finding objectives, in particular designs with efficacy and toxicity endpoints to meet the modern drug paradigm. From this chapter there isn't a single "best" design to recommend for use but a collection of designs that favour some objectives and settings more than others. The assessment involves a careful consideration of the trial objectives and a statistical design to meet those objectives. To meet the objectives, designing a dose finding trial is a complex and lengthy process requiring collaboration between statistician and the clinical team. There are two major components of the statistical designs for dose finding that feature in the designs reviewed: the probability model and the decision process.

There are a number of approaches taken in the literature that model the relationship between dose and toxicity, and dose and efficacy. A parametric model such as logistic regression gives a closed form for the relationship with a number of parameters. Non parametric approaches make fewer assumptions about the dose effect relationship. Any Bayesian model will include priors that need to be specified before the study is initiated. In the setting of dose finding where minimal data is available the role of the prior is important [78]. With any given design, in particular the parametric designs there will be a number of assumptions as to how the dose effect relationship modelled, it is important that a design is robust in performance when misspecified.

One element of the probability model in the dose finding literature that is less well understood is the role of correlation. Individual outcomes for binary toxicity and efficacy from a given patient could be correlated. In the case of positive correlation observing both efficacy and toxicity or neither becomes more likely. Some designs ignore correlation and simply

model the two outcomes independently. The use of a copula as in the EffTox design (equation 2.23) seems desirable as it allows the specification of a dose efficacy relationship, a dose toxicity relationship and a copula to model the correlation between them. The rationale and the effect of specifying the copula isn't justified as part of the approach. A model that includes correlation may be more reflective of how the data is generated and lead to better inferences and decisions. Chapter 3 investigates the role of copulas in dose finding.

The performance of statistical dose finding designs with multiple endpoints is more challenging to compare. Each new design will compare to some previous designs with some improvement in a limited set of scenarios. Different designs propose different strategies for how the OD is defined statistically. Given a proposed scenario one design may define one particular dose as optimal while another design a different dose. There is a question as to which design performs better as success is defined differently. Both designs should recommend their respective OD. This point has been made in a recent review with the recommendation to carefully consider and select a decision procedure that best aligns with the specific dose-toxicity and dose-efficacy scenarios during the design stage [107]. It is also worth noting that the optimum dose is defined with respect to both admissibility rules and the objective function. There isn't an agreed upon approach to defining the OD, there are however a number of heuristics described in the next paragraph that feature in most designs. These are described as heuristics as they are common features without a consistent statistical approach.

The OD is typically described with respect to population level parameters. This means that inferences at any stage are made with respect to the whole patient population entering the study, although more personalised dosing designs exist [108]. There is a limit to the amount of toxicity that is acceptable; toxicity only designs will explicitly target some level of toxicity while efficacy designs will typically have a cut off using admissibility rules. This limit is typically somewhere between 20-40% and is defined with respect to the disease area of intervention and previous studies [109]. Many designs will also have a similar threshold level for efficacy. How the OD is defined and how this related to the decision making a major component of the thesis and was seen through the motivating example.

The issue highlighted as part of the motivating example isn't a problem with the manner in which the contour is elicited. It would be possible to reformulate the questions so that the

initially elicited contour would be contained in the lower right quadrant. In this instance this would produce steeper contours associated with improved operating characteristics. The contour that was initially elicited however still represents the clinical situation. The objective function specified as part of the EffTox method is a simplification of the situation described in the motivating example. The methods highlighted as part of the review all specify objective functions that cannot be specified to reflect the clinical situation in the motivating example, this includes the 4 outcome utility function specified by a number of recent designs (Equation 2.28). Many authors consider the objective function as part of the statistical design with a set of components that need tuning through simulation and less formal clinical consultation [103]. This allows specification of an objective function that doesn't fully reflect the clinical situation so long as the design has good operating characteristics.

The Bayesian decision theoretical approach would suggest that preferences should be clinical and be encoded within the utility function. Additional ad hoc rules are typically imposed in the literature to prevent unethical choices for patients; these are included in the previous proposed Bayesian decision theoretic approaches to dose finding also. The rules include admissibility criteria to define an evidence level for an estimated minimum amount of efficacy and maximum amount of toxicity in order for a dose to be considered in the decision. When there is minimal evidence available in particular at the start of the trial the rules won't be efficient in excluding doses that appear to be quite toxic or not efficacious. The admissibility rules are necessary to compensate for a simplified objective function. The use of admissibility rules constitute a two stage approach to decision making, by restricting the decision space according to the admissibility rules before optimising an objective function that doesn't fully capture the situation. This approach falls short of a fully decision theoretic approach [89].

There is a large body of applied statistical literature surrounding the formulation of decision analysis with multiple competing objectives [110]. The intention is to gain a good understanding of this work and propose a general structural form for the utility function that is capable of better capturing the clinical beliefs in the setting. It is hoped that a utility function that more closely captures the situation will lead to improved performance in terms of of correctly selecting the OD in a wider range of scenarios. Given all designs,

whether model based or model assisted, in this setting use a function to describe clinical preference in order to make decisions, the work could be applied to other designs. Elicitation methods to reliably obtain any proposed utility function is also included in this work. The work is restricted to efficacy and toxicity dose finding designs to reflect the current nature of dose selection paradigm in oncology drug development [16]. Chapter 4 reviews the statistical decision making literature with multiple objectives. Chapter 5 proposes a novel utility function and assesses the operating characteristics in comparison to a more established design to understand potential benefits and limitations.

The main objectives of this thesis are:

- Investigate correlation for binary endpoints through copulas to to gain insights and make recommendations about modelling in PI-II dose finding studies
- Gain a comprehensive understanding of the applied statistical literature on decision analysis with multiple competing objectives, that would be capable of capturing clinical beliefs in this context.
- Propose a novel utility function and assess its operating characteristics in comparison to a more established design, in order to understand its potential benefits and limitations.
- Develop an elicitation protocol to capture the preferences of a single decision-maker or a team of key opinion leaders for inclusion in a utility function through a structured set of questions.

Chapter 3

Copula models for dose finding

3.1 Introduction

Model based approaches for phase I-II dose-finding map the relationship between response (achieving efficacy and/or experiencing toxicity) and dose for a patient using a parametric model. One approach is to assume that efficacy and toxicity are independent and to model the outcomes separately with respect to dose. It is possible that there may be some dependence between the two outcomes with the chance of an individual patient having an efficacious response changing if they also have a toxic event, for example. One approach, utilised by the EffTox method, as in the thesis motivating example (Section 2.5), is to model efficacy and toxicity separately and combine the two using a copula. A copula model is a flexible framework to allow the specification of a dependence structure between separate marginal distributions [80]. In the case of joint efficacy and toxicity modelling, in dose finding, this allows the specification of a dose efficacy relationship, a dose toxicity relationship and a copula to model the correlation between the two random variables. The main advantage of the copula approach is that each of the separate marginal distributions are retained which is helpful because they have an easily interpretive meaning to make decisions concerning dose allocation. The definition of a copula function with a number of examples will be introduced as part of this chapter. The over arching purpose of this chapter is to inspect the statistical properties of copulas for discrete data to gain insight into their use in dose finding.

A simulation study looking specifically at the use of copula models in phase I-II clinical

trials has previously been reported [111]. The paper is referred to as the copula simulation study (CSS) from herein for ease of brevity. The primary goal of the paper was to assess operating characteristics of the copula model in the dose finding setting with data simulated with correlation. The CSS assessed the Braun [76] and the EffTox [77] models, introduced in Chapter 2. Both designs utilise a function for a joint distribution based upon probabilities of efficacy and toxicity at each dose, Equations 2.25 and 2.18 for the Braun and Efftox models respectively. The CSS contrasted the two proposed designs with an independent model ($\pi_{ET} = \pi_E \pi_T$ i.e. no correlation). The decision rule to decide dosing at each stage followed that specified in EffTox, Equation 2.25, with admissibility rules given in Equations 2.8 and 2.9. A number of scenarios with fixed vectors for the probability of efficacy and toxicity with different degrees of correlation between efficacy and toxicity were evaluated for each of the three probability models. The correlation was specified by a range of parameter values assuming the copula probability model specified for the two designs. Specifically when assuming the EffTox joint probability model, $\psi = 0, 0.4, 0.8$ with positive values of ψ representing positive correlation and $\psi = 0$ the independent model.

The CSS results suggested that the copula models gave very similar operating characteristics to a simple model assuming independence. The results were consistent even when simulating data with larger values of the correlation parameter, representing strong positive correlation between endpoints. In some instances marginally worse operating characteristics were obtained for the copula models. Possible explanations for this counter intuitive result were i) that "the likelihood may contain little information about the correlation parameter", and ii) any benefit to accounting for correlation was at the cost of less precision in other parameters. This was linked to the small sample sizes typically used in dose finding studies and an inability to identify a "correct" copula. The CSS discussed the overall purpose of dose finding and whether it was sufficient to use a more parsimonious model, such as the independent model in order to achieve the aim of selecting an OD. The CSS concludes that given the empirical results, copula models are not that useful in dose finding.

The CSS did not give an interpretation as to how parameter values relating to correlation from different models should be interpreted. Additionally, the paper did not acknowledge that the copula used as part of EffTox belongs to a much wider class of potential copula functions. Conclusions from the CSS are difficult to extrapolate to other phase I-II designs

with a high degree of confidence as there isn't theory relating to copulas underpinning the simulation study. The CSS provides empirical results for the performance of the specific copula model and decision function used in EffTox without a justification as to why it may extend to other settings. This chapter adds to the literature by providing a theoretical understanding of copulas and their use for binary data. In doing so the reasons for the findings in the CSS paper are justified with a more general analytic (algebraic) understanding. The chapter provides insight to the use of the wider class of copula models. Additionally this provides an understanding of the potential impact of the copula model on different decision functions in dose finding. This is achieved by introducing a consistent measure of correlation, Kendal's Tau, to compare the performance of different models.

The chapter starts by introducing how correlation for two binary variables can be measured. The notation for probability and cumulative density functions are then introduced to allow the mathematical definition of a copula. Some of the statistical properties of copulas in the case of discrete data are given to gain insight into their use in dose finding. A simulation study is conducted looking at the impact of different copula models on the ability to estimate correlation and corresponding marginal distributions. The CSS paper findings are placed in the context of the chapter with further inferences made.

3.2 Kendall's tau

In the dependency theory, correlation is defined as a measure of dependence or statistical relationship between two random variables [112]. Many measures of correlation depend upon two concepts of concordance and discordance. In the general sense concordance is when the magnitude of the random variables coincide, when one goes up the other also is likely to go up or vice versa. For discordance the opposite is true. Kendall's tau [113] is a measure of correlation between two random variables. Let (X_1, Y_1) and (X_2, Y_2) be two independent realisations from the joint distribution of (X, Y). The population version of Kendall's tau is defined by

$$\tau(X, Y) = \mathrm{P}[(X_1 - X_2)(Y_1 - Y_2) > 0] - \mathrm{P}[(X_1 - X_2)(Y_1 - Y_2) < 0] \tag{3.1}$$

A concordant pair is when both X and Y for one pair is bigger (or smaller) than another pair given by the first part of the equation. Discordant pairs are when one variable is bigger while the other smaller (second part of equation). The difference between the probability of concordant and discordant pairs is defined as Kendall's tau which is defined on the interval $(-1, 1)$. For positive correlation, the chance of concordance (similarity) increases with the opposite true for negative correlation. When $\tau = 0$ the chance of concordance and discordance are equal.

Kendall's coefficient initially was proposed for continuous variables where concordance and discordance are the only possibility for pairs of random variables. In the discrete case there is the additional chance of pairs of observations taking the same value. A pair $\{(X_1, Y_1), (X_2, Y_2)\}$, is said to be tied if $X_1 = X_2$ or $Y_1 = Y_2$. An adjusted version τ_b is defined by

$$\tau_b(X, Y) = \frac{\tau(X, Y)}{\sqrt{\mathrm{P}(X_1 \neq X_2)\mathrm{P}(Y_1 \neq Y_2)}} \tag{3.2}$$

This allows in the discrete case, for τ_b to be in the range $[-1, 1]$, corresponding to perfect negative and positive correlation. With continuous random variables the denominator would equal 1 and $\tau(X, Y) = \tau_b(X, Y)$. Kendall's τ_b is used as a consistent measure of correlation throughout the rest of this chapter.

3.3 Copulas

3.3.1 Marginal distributions

This section introduces key properties for continuous marginal distributions and the differences for discrete distributions. This enables a more formal definition of a copula to be given in the next section, which relies on properties of marginal distributions.

Take a random variable Y where $Y \in \mathbb{R}$. The cumulative density function (CDF), F is defined by

$$\Pr(Y \leq y) = F(y) \tag{3.3}$$

The function maps from the domain of Y to probabilities in the interval $[0, 1]$. The inverse CDF maps the opposite way so that:

$$F^{-1}(u) = y. \tag{3.4}$$

The distribution of $U = F(Y)$ is Uniform$(0, 1)$ i.e.

$$P[F(Y) \leq u] = u, \quad u \in [0, 1]. \tag{3.5}$$

The cumulative distribution for a discrete random variable, is defined by

$$F(Y) = P(Y \leq y) = \sum_{y_i \leq y} P(Y = y_i) = \sum_{y_i \leq y} f(y_i), \tag{3.6}$$

where f is the probability mass function (PMF). The distribution is not continuous and has plateaus for distinct values of y_i.

The inverse also isn't unique mapping 1:1 with plateaus at each distinct value of y. We therefore define a left-continuous generalised inverse function for natural numbers as follows

$$F^{\leftarrow}(u) = \inf\{y \in \mathbb{N} : F(y) \geq u\} \tag{3.7}$$

or the lowest value of y that satisfies $F(y) \geq u$.

The PMF is linked to the CDF in the discrete case by

$$Pr[Y = y] = F_Y(y) - F_Y(y\prime) \tag{3.8}$$

where $y\prime$ is the largest value of y such that $y\prime < y$. For $Y \in \mathbb{N}$ then $y\prime = y - 1$.

3.3.2 Definitions

Copulas are functions that allow us to model correlation between multiple random variables. This is achieved by decomposing the multivariate CDF into univariate marginal distributions and then a copula that captures the dependence structure.

Consider a multivariate random variable $\mathbf{Y} = (Y_1, \ldots, Y_m.)$, where $Y_i \in \mathbb{R}$. A m-variate copula C, mapping $[0, 1]^m \to [0, 1]$, is the CDF of a random vector $(U_1, ..., U_m)$ with uniform margins

$$C(\mathbf{u}) = \mathrm{P}[U_1 \leq u_1, ..., U_d \leq u_m], \quad U_j \sim \text{Uniform}(0, 1) \tag{3.9}$$

Sklar [114] showed that if C is a m-variate copula and $F_1, ..., F_m$ are univariate CDFs then

a function H exists such that

$$H(\mathbf{y}) = C(F_1(y_1), ..., F_m(y_m)) \tag{3.10}$$

where H is a m-variate CDF with margins $F_1, ..., F_m$. In addition Sklar's second theorem stated that if H is an m-variate CDF with univariate CDF's $F_1, ..., F_m$, then there exists a copula C such that the equation above holds and that C is unique and equal to

$$C(\mathbf{u}) = G(F_1^{\leftarrow}(u_1), ..., F_m^{\leftarrow}(u_m)) \tag{3.11}$$

and that $H(.) = G(.)$.

The Frèchet-Hoeffding Theorem states that lower and upper limits exist for the copula function [80].

$$\max\{1 - m + \sum u_i, 0\} \leq C(u_1, ..., u_m) \leq \min\{u_i\} \tag{3.12}$$

This provides a set of conditions that a copula function needs to abide by in order for Sklar's theorem to hold. So for example, if $u_1 = 0$, then any joint distribution incorporating u_1 must also equal zero which is given by the limits above.

When using a copula function when at least one of the marginal distributions are discrete the left-continuous generalised inverse function given in Equation 3.11 isn't a unique 1:1; instead there are a range of possible values of the CDF for a given uniform distribution. It is still possible however to specify a copula using a parametric form with Sklar's theorem still holding [115].

There are several classes of parametric copula functions in the literature that satisfy the axioms of Sklar's theorem [80]. The bivariate case is considered from this point as it corresponds with the dose finding setting. The presented material is however extendable to the multivariate setting. Three examples of parametric forms with a single parameter θ, are given below and with some properties explored further in the next section. The bivariate Farlie-Gumbel-Morgenstern Copula (FGM) [79, 80] is defined by

$$C_\theta(u_1, u_2) = u_1 u_2(1 + \theta(1 - u_1)(1 - u_2)), \quad \theta \in [-1, 1]. \tag{3.13}$$

The Plackett copula [116] defined by

$$C_\theta(u_1, u_2) = \frac{1 + (\theta - 1)(u_1 + u_2) - \sqrt{[1 + (\theta - 1)(u_1 + u_2)]^2 - 4\theta(\theta - 1)u_1 u_2}}{2(\theta - 1)} \quad \theta \in (0, \infty), \tag{3.14}$$

and the bivariate Gaussian copula given by

$$C_\theta(u_1, u_2) = \Phi_\theta\big(\Phi^{-1}(u_1), \Phi^{-1}(u_2)\big), \tag{3.15}$$

where Φ^{-1} is the inverse of the cumulative standard normal and Φ_θ is a bivariate cumulative normal distribution with mean vector of zeros and covariance matrix

$$\Sigma = \begin{bmatrix} 1 & \theta \\ \theta & 1 \end{bmatrix} \quad \theta \in [-1, 1]. \tag{3.16}$$

The range of the parameters in each function has been determined so that the Frèchet-Hoeffding bounds hold for all possible values of u_1 and u_2.

3.3.2.1 Copula model identifiability

The copula model itself is unidentifiable in the discrete setting, as it is not possible to ascertain empirically which model data is generated from [117]. Table 3.1 demonstrates this through an example. All of the copula models in the example from the table are equally valid, all able to model the correlation. The parameter values for each copula are unique and do not have an interpretative meaning outside of the context of the copula model. Later in the chapter a formula to obtain Kendall's tau is given. All three models in this instance have the same Kendall's tau, which highlights the necessity of a consistent measure. In this small example it doesn't matter which copula model was used to model the correlation.

Table 3.1: Consider a pair of Bernoulli random variables (Y_1, Y_1) with $\mathrm{P}(Y_1 = 0) = p$ and $\mathrm{P}(Y_2 = 0) = q$ with a joint probability of $r = \mathrm{P}(Y_1 = 0, Y_2 = 0) \in [\max(0, p + q - 1), \min(p, q)]$, the Frèchet-Hoeffding bounds

Bernoulli $\max(0, p + q - 1 \leq r \leq \min(p, q)$	Copula Model	θ
$p = 0.3, q = 0.4, r = 0.1704$	FGM, $\theta \in [-1, 1]$	$\theta = 1$
$p = 0.3, q = 0.4, r = 0.1704$	Plackett, $\theta \in (0, \infty)$	$\theta = 2.693764$
$p = 0.3, q = 0.4, r = 0.1704$	Gaussian, $\theta \in [-1, 1]$	$\theta = 0.3602642$

3.3.3 Bivariate logistic regression modelling using copulas

Consider Y_1, Y_2 dependent variables and $\mathbf{x} \in \mathbb{R}^p$ explanatory variables, which may be non mutually exclusive. In copula modeling, marginals $F_1(\cdot|\mathbf{x}), F_2(\cdot|\mathbf{x})$ are fitted and dependence induced through a copula - a bivariate distribution function with uniform margins on $[0, 1]$. i.e

$$Pr(Y_1 \leq y_1, Y_2 \leq y_2|\mathbf{x}) = C\{F_1(y_1|\mathbf{x}), F_2(y_2|\mathbf{x})\} \tag{3.17}$$

holds for a specific function C and for all values $(y_1, y_2) \in \{0, 1\}$ and $\mathbf{x} \in \mathbb{R}^p$ where p is the number of explanatory covariates.

The probability mass function of a discrete copula is defined by a range of values (Equation 3.8) i.e.:

$$\begin{aligned} p(y_1, y_2|\mathbf{x}) &= P[Y_1 = y_1|\mathbf{x}, Y_2 = y_2|\mathbf{x}] && (3.18)\\ &= C(F_1(y_1|\mathbf{x}), F_2(y_2|\mathbf{x})) - C(F_1(y_1 - 1|\mathbf{x}), F_2(y_2|\mathbf{x})) && (3.19)\\ &\quad - C(F_1(y_1|\mathbf{x}), F_2(y_2 - 1|\mathbf{x})) + C(F_1(y_1 - 1|\mathbf{x}), F_2(y_2 - 1|\mathbf{x})) \end{aligned}$$

In a setting where response variables are dichotomous and the marginal distribution, $F_j(\cdot|\mathbf{x})$ follows a logistic regression model, i.e. $\pi_j(\mathbf{x}) = \Pr(Y_j = 1|\mathbf{x})$ for each $j \in \{1, 2\}$ where

$$\pi_j(\mathbf{x}) = \frac{\exp(\mathbf{x}^T\beta_j)}{1 + \exp(\mathbf{x}^T\beta_j)}, \tag{3.20}$$

and β_j is a $p_j \times 1$ vector of parameters corresponding to the width of the design matrix $\mathbf{x}$. Both marginals follow a Bernoulli distribution with CDF:

$$F_j(y|\mathbf{x}) = \begin{cases} 0 & \text{if } y < 0 \\ 1 - \pi_j(\mathbf{x}) & \text{if } 0 \leq y < 1 \\ 1 & \text{if } y \geq 1 \end{cases} \tag{3.21}$$

defining $\bar{\pi}_j = 1 - \pi_j$ and noting that

$$C(\mathbf{u}) = 0 \quad \text{if any } u_j = 0, \tag{3.22}$$

$$C(1, u_j) = u_j \tag{3.23}$$

and using Equation 3.18 the PMF for the four possible outcomes are given as [118]:

	$Y_2 = 0$	$Y_2 = 1$
$Y_1 = 0$	$C(\bar{\pi}_1, \bar{\pi}_2)$	$\bar{\pi}_1 - C(\bar{\pi}_1, \bar{\pi}_2)$
$Y_1 = 1$	$\bar{\pi}_2 - C(\bar{\pi}_1, \bar{\pi}_2)$	$1 - \bar{\pi}_1 - \bar{\pi}_2 + C(\bar{\pi}_1, \bar{\pi}_2)$

The table above or equivalently Equation 3.18 gives the probability mass function for a joint distribution that is used to specify the likelihood when incorporating copulas into dose finding. The quoted equation for EffTox using the FGM copula, to find the joint probability for an observation $Y_E = a, Y_T = b$, is given by

$$\pi_{a,b} = (\pi_E)^a(1-\pi_E)^{1-a}(\pi_T)^b(1-\pi_T)^{1-b} + (-1)^{a+b}\pi_E(1-\pi_E)\pi_T(1-\pi_T)\theta. \tag{3.24}$$

Which is a tidy way of expressing the FGM copula in a single line (Equation 3.13) for two binary variables. The more general equation given in Equation 3.18 and the simplified contingency table given in Table 3.1 are equivalent. This is demonstrated below for $Y_E = 0, Y_T = 0$:

From EffTox equation above,

$$\pi_{0,0} = (1-\pi_E)(1-\pi_T) + \pi_E(1-\pi_E)\pi_T(1-\pi_T)\theta. \tag{3.25}$$

from Equation 3.18,

$$\pi_{0,0} = C(1-\pi_E, 1-\pi_T) - C(0, 1-\pi_T) - C((1-\pi_E), 0) + C(0,0) \tag{3.26}$$

$$= C(1-\pi_E, 1-\pi_T), \tag{3.27}$$

as Table 3.1. Using Equation 3.13, the FGM copula,

$$\pi_{0,0} = (1-\pi_E)(1-\pi_T)(1+\pi_E\pi_T)\theta, \tag{3.28}$$

which is equal to the EffTox equation.

When marginal distributions are specified by logistic regression different copula models will give a different correlation structure. Extending the example given in Table 3.1, where three different copula models were specified to the data $p = 0.3, q = 0.4$ and $r = 0.1704$. If there

was a second Bernoulli variable (i.e a further dose) with $p' = 0.5$ and $q' = 0.5$ then assuming each copula model was the correct model, this would give $r' = (0.3125, 0.3107, 0.3086)$ for the FGM, Placket and Gaussian copulas respectively. These are not hugely different in the particular example but it does demonstrate the copula is a model that implies a different correlation structure across different doses. As such the copula is a model choice.

3.3.3.1 The range of Kendall's tau for different copulas

Different copulas are capable of measuring different levels of correlation. In order to compare different copula models a consistent measure of correlation is needed. The population Kendall's tau for a discrete bivariate copula was given by Nikoloulopoulos [119] as follows: Let Y_i, $i = 1, 2$ be integer-valued discrete random variables whose joint distribution is H, with marginal CDFs F_i, PMFs $f_i, i = 1, 2$ and copula C. Then the population version of Kendall's tau for Y_1 and Y_2 is given by

$$\tau(Y_1, Y_2) = \sum_{y_1=0}^{\infty}\sum_{y_2=0}^{\infty} h(y_1, y_2)\{4C(F_1(y_1-1), F_2(y_2-1)) - h(y_1, y_2)\} + \sum_{y_1=0}^{\infty} f_1(y_1)^2 + \sum_{y_2=0}^{\infty} f_2(y_1)^2 - 1 \tag{3.29}$$

where $h(y_1, y_2)$ is the joint PMF as given previously. It is important to notice the measure of correlation is dependent on the magnitude of each of the marginal distributions. This is a general result rather than something that can be remedied by changing the measure of correlation [115]. In the case when each of the marginals is Bernoulli, τ also can be derived from the population definition

$$\begin{aligned}
\tau(Y_1, Y_2) &= \mathrm{P}[(Y_{11} - Y_{12})(Y_{21} - Y_{22}) > 0] - \mathrm{P}[(Y_{11} - Y_{12})(Y_{21} - Y_{22}) < 0] \\
&= 2\{\mathrm{P}[Y_1 = 0, Y_2 = 0]\mathrm{P}[Y_1 = 1, Y_2 = 1] - \mathrm{P}[Y_1 = 1, Y_2 = 0]\mathrm{P}[Y_1 = 0, Y_2 = 1]\} \\
&= 2\{h(0,0)h(1,1) - h(0,1)h(1,0)\} \\
&= 2\{r(1 - \bar{\pi}_1 - \bar{\pi}_2 + r) - (\bar{\pi}_1 - r)(\bar{\pi}_2 - r)\} \\
&= 2\{r - \bar{\pi}_1 r - \bar{\pi}_2 r + r^2 - \bar{\pi}_1\bar{\pi}_2 + \bar{\pi}_1 r + \bar{\pi}_2 r - r^2\} \\
&= 2\{r - \bar{\pi}_1\bar{\pi}_2\}
\end{aligned} \tag{3.30}$$

where $r = \mathrm{P}[Y_1 = 0, Y_2 = 0] = h(0, 0)$. The two formulas can be shown to be equivalent for bivariate Bernoulli data. Kendall's τ_b accounts for ties (Equation 3.2):

$$\tau_b = \frac{r - \bar{\pi}_1 \bar{\pi}_2}{\sqrt{\bar{\pi}_1(1 - \bar{\pi}_1)\bar{\pi}_2(1 - \bar{\pi}_2)}} \tag{3.31}$$

Different copula models will have different ranges for possible values of τ_b. The parameter range is restricted as part of the copula definition so that the Frechet-Hoeffding bounds hold for all values of $(u_1, u_2) \in [0, 1]^2$. Using the extreme parameter values for a given copula gives maximum and minimum values of r as upper and lower limits that the copula is capable of measuring at, given two probabilities of p and q. These are then converted into the corresponding measure of τ_b using Equation 3.31.

The bounds of Kendall's tau (τ_b) for the FGM and Gaussian copulas are plotted in Figures 3.1 and 3.2 by substituting the parameter extremes for a given copula into the formula for Kendall's τ_b. It is apparent that the FGM has a very limited range [120] unable to measure stronger correlation.

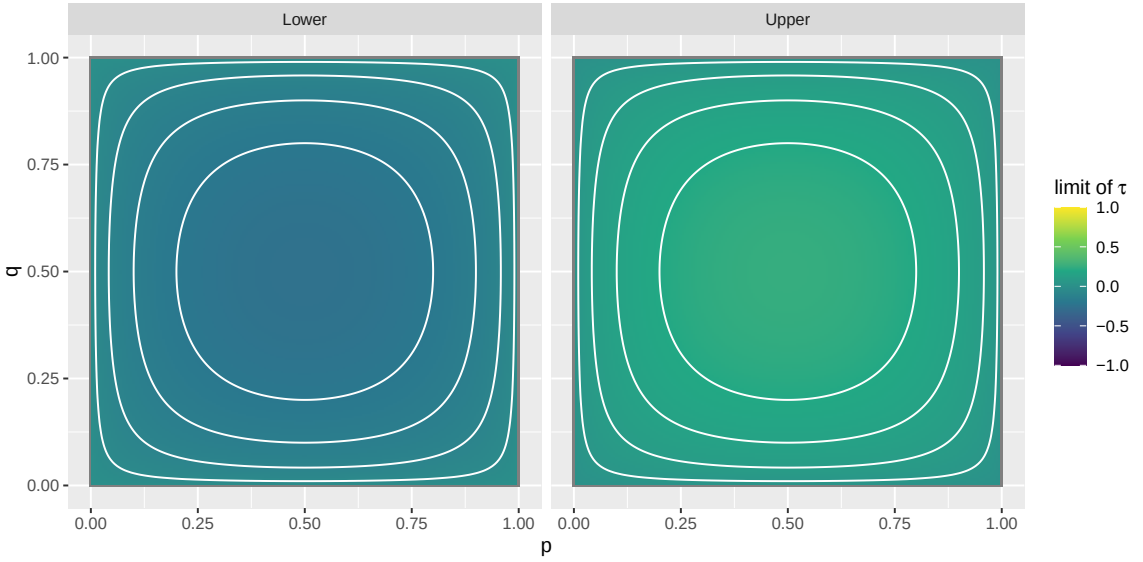

Figure 3.1: The range of Kendall's tau b for the FGM bivariate copula. Contours are lines of equal correlation, outer square has $\tau_b = 0$ with increments of 0.05 moving to centre with maximum and minimum values $\tau_b = 0.25$ and $\tau_b = -0.25$ at $p = q = 0.5$

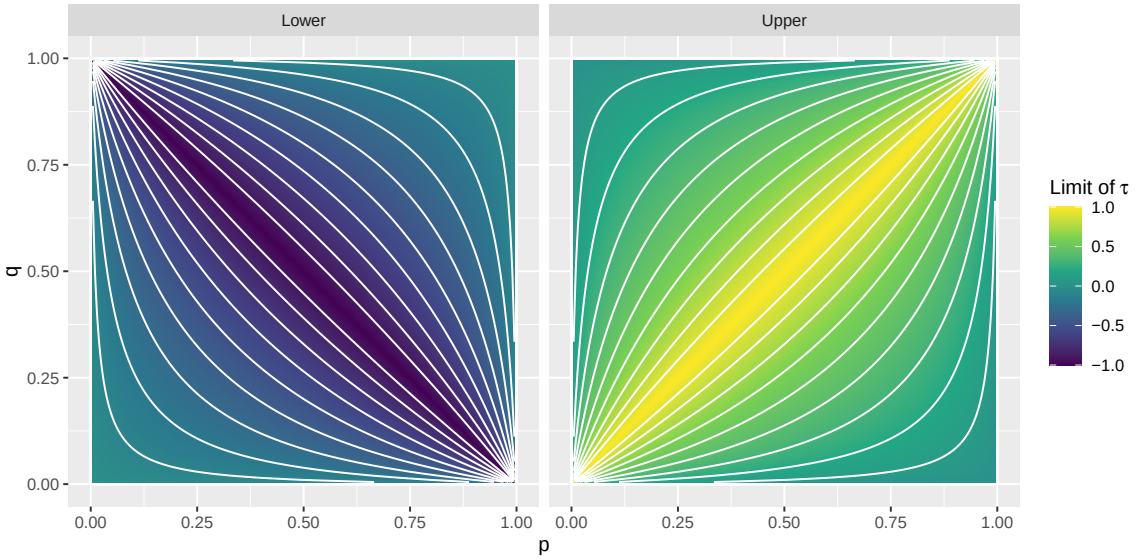

Figure 3.2: The range of Kendall's tau b for the Gaussian bivariate copula. Contours are lines of equal correlation, outer square has $\tau_b = 0$ with increments of 0.1 moving to centre with maximum and minimum values of $\tau_b = 1$ and $\tau_b = -1$ at $p = q$ and $p = (1 - q)$ respectively

3.4 Copulas for dose finding

3.4.1 Decision functions

The choice of dose to allocate a patient at each stage of a dose finding trial using a Bayesian approach depends upon maximising the expected utility (Equation 2.5). Many designs will additionally sub-set the decision space using admissibility rules (Equations 2.8 and 2.9). In the case of EffTox , as in the motivating example (Section 2.5), both of the components of the decision process (admissibility rules and desirability contours) depend only upon the probabilities π_E and π_T. Linking the more general copula notation from the previous section $\pi_1 = \pi_E$ and $\pi_2 = \pi_T$ and using Table 3.1,

$$\begin{aligned} \pi_E &= P(Y_E = 1, Y_T = 0) + P(Y_E = 1, Y_T = 1) \\ &= \bar{\pi}_T - C(\bar{\pi}_E, \bar{\pi}_T) + 1 - \bar{\pi}_E - \bar{\pi}_T + C(\bar{\pi}_E, \bar{\pi}_T) \\ &= 1 - \bar{\pi}_E \\ &= \pi_E \end{aligned} \tag{3.32}$$

The calculation of the marginal parameter π_E does not depend on the copula. A similar result can be shown for toxicity. When calculating the expectation of the decision function

the correlation parameter is redundant as it does not feature in the decision function. As such, an independent model and a copula model would have identical implied utility distributions at each dose, if the marginal distributions for efficacy and toxicity were the same. This statement is true for any copula function. Any impact upon the operating characteristics utilising a decision function composed of the marginal probabilities will be down to any differences in deriving the marginal posteriors from the same data. A simulation study is conducted in the next section to quantify the effect of a copula model upon marginal probabilities to explain any differences in operating characteristics.

The EffTox utility design [82, 78] introduced in the previous chapter maximises the expectation of individual patient level outcomes given in Equation 2.28 (and again below) to make decisions.

$$E(u(Y)) = \int_{\theta} \sum_{a=1}^{1} \sum_{b=1}^{1} K(a,b)\pi(a,b) \tag{3.33}$$

where $K(a,b)$ were constants to be specified and $\pi(a,b) = P(Y_E = a, Y_T = b)$. Assuming the use of the FGM copula as specified in EffTox, and standardising with $K(0,1) = 0$ and $K(1,0) = 1$, as the best and worst of the four possible outcomes, the equation can be rewritten as a function of π_E, π_T and c:

$$E(u(Y)) = \int_{\theta} K(1,1)\pi_E + K(0,0)(1-\pi_T) + (1 - K(0,0) - K(1,1))(\pi_E(1-\pi_T) - c) \tag{3.34}$$

where $c = \pi_E(1-\pi_E)\pi_T(1-\pi_T)\theta$. This includes the parameter for correlation, θ. From the equation above specifying $K(0,0) + K(1,1) = 1$ would make the expectation of the utility function independent of the correlation component, c. If an independent model was specified with the decision function it is possible to plot the contours in two dimensions, Figure 3.3. The correlation component would add a further dimension to the plot. A design with $K(1,1) = 0.5$ and $K(0,0) = 0.3$ has been stated as suitable in many settings [78]. The differences in the expected utility function are plotted for extreme values of θ with these values of the utility decision function in Figure 3.3. The correlation component for the FGM copula is only capable of having a small impact on the utility function. This partly because the extreme values of θ correspond with small values of Kendall's tau for the FGM copula (Figure 3.1). As such when specifying an FGM copula with the EffTox utility design, the marginal probabilities of π_E and π_T will tend to dominate the correlation component in

decision making. Small changes in the specifying utility constants will have a greater effect on the decision making process in contrast to specifying the FGM copula or an independent model. The correlation component in Figure 3.3 gives the extreme of possible values. The conclusion here is that the correlation component as specified is going to make minimal impact upon decision making.

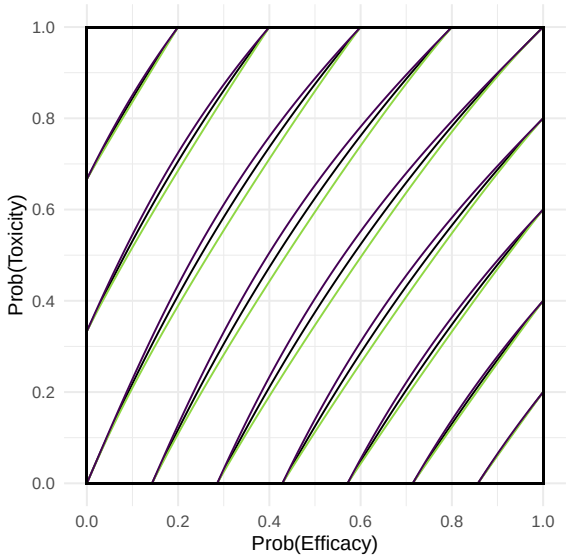

Figure 3.3: EffTox patient utility design [82, 78] utility contours with decision parameters $K(1,1) = 0.5$ and $K(0,0) = 0.3$ fit with a FGM copula. Black curve is the utility for an independent model ($\theta = 0$), green curve when $\theta = 1$ the maximum positive correlation possible in the model and purple curve when $\theta = -1$, the minimum negative correlation possible in the model

3.4.2 Correlation

The magnitude of correlation for two binary variables is dependent upon the marginal distributions (Equation 3.30). In dose finding little is known about the marginals prior to the trial beyond what is informed by the data collected in the trial. With the smaller samples sizes used in dose finding, imprecise posterior distributions for the marginals will give further imprecision for the correlation parameter. This suggests that the precision of correlation estimates will be limited in dose finding where the sample size is small. To get an idea of this, the simulation also seeks to quantify the variation in calculating the posterior of the correlation component. Given that the parameter in the copula has no meaning outside of the copula this is expressed in terms of Kendall's τ_b.

EffTox specifies a reparameterised version of the FGM copula parameter as in Equation

3.24, where $\theta = \frac{e^{\phi}-1}{e^{\phi}+1}$. This is done to allow the specification of a standard normal prior on ϕ. The reparameterisation and prior are mildly informative for independence when using $\phi \sim N(0,1)$ as is suggested for the method (Figure 3.4). The CSS paper and the simulation study in this chapter use an uninformative uniform prior over the range of θ.

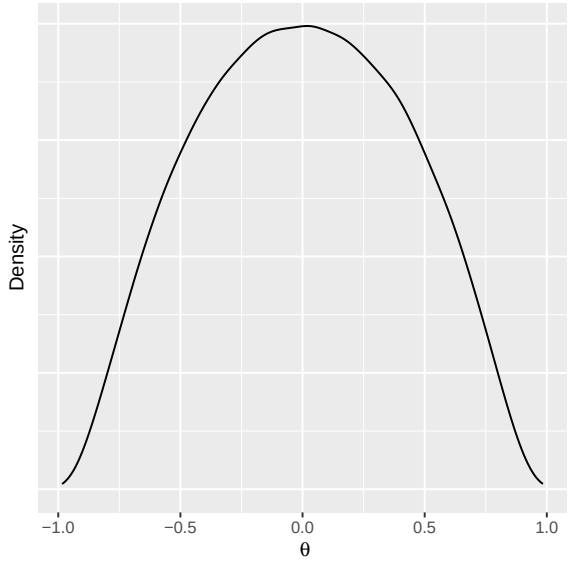

Figure 3.4: Implied prior for θ specified in EffTox method where $\phi \sim N(0,1)$ and $\theta = (e^{\phi}-1)/(e^{\phi}+1)$. Where θ is the parameter specified in the FGM copula (Equation 3.13)/ Priors used in the CSS and this chapter follow a uniform distribution

3.4.2.1 Exchangeability

An independent model and a copula model use different levels of information from the data. For a copula model, efficacy and toxicity pairs from patients treated at a given dose level are exchangeable. For example a change of ordering from the three patients outcomes treated at the same dose

$$\big((Y_{1E}, Y_{1T}), (Y_{2E}, Y_{2T}), (Y_{3E}, Y_{3T})\big) \quad \text{and} \quad \big((Y_{3E}, Y_{3T}), (Y_{1E}, Y_{1T}), (Y_{2E}, Y_{2T})\big) \tag{3.35}$$

will give identical posterior distributions. The independent model allows the the efficacy and toxicity responses to also be exchangeable. To extend the previous example any ordering of the patient outcomes

$$\big((Y_{2E}, Y_{1T}), (Y_{1E}, Y_{3T}), (Y_{3E}, Y_{2T})\big) \tag{3.36}$$

will give the same posterior distribution in the case of independence.

Multiple realisations from a Bernoulli distribution can equivalently be fit as a Binomial distribution with sufficient statistics for patients treated at a particular dose are $\sum Y_E$ and $\sum Y_T$. A sufficient statistic for the copula model needs to be considered jointly at each dose, i.e $\left(\sum\sum Y_E = a, Y_T = b\right)$ for $a = 0, 1$ and $b = 0, 1$ with one of the four possibilities derivable from the other three. In essence, for the independent model two binomial likelihoods are specified for a particular dose while for the copula model a four parameter multinomial likelihood is needed with the individual probabilities coming from the copula. For example, consider data from 6 patients treated at a dose, where

$$\boldsymbol{Y} = \big((0,0),(1,1),(0,1),(1,0),(0,0),(0,0)\big). \tag{3.37}$$

The likelihood is proportional to

$$\pi_E^{\sum Y_E}(1-\pi_E)^{n-\sum Y_E}\pi_T^{\sum Y_T}(1-\pi_T)^{n-\sum Y_T} = \pi_E^3(1-\pi_E)^3\pi_T^2(1-\pi_T)^4, \tag{3.38}$$

with an independent model, and

$$\pi_{00}^{\sum\sum Y_E=0,Y_T=0}\pi_{11}^{\sum\sum Y_E=1,Y_T=1}\pi_{10}^{\sum\sum Y_E=1,Y_T=0}\pi_{01}^{\sum\sum Y_E=0,Y_T=0} = \pi_{00}^3\pi_{11}\pi_{10}\pi_{01}, \tag{3.39}$$

with a copula model. The difference is a key point in reducing the computational burden of large simulation studies in Chapter 5. The purpose of highlighting the differences here is that the copula model uses additional information from the data in comparison to the independent model. This constitutes counts relating to the number of ties, discordant and concordant pairs in order to estimate correlation.

3.4.2.2 The range for copula models

In the section earlier it was shown that different copulas will give slightly different correlation structures across doses. The major difference in the FGM and the Gaussian copula selected however is in the range of possible values for Kendall's tau. The FGM copula has a limited range in contrast to the Gaussian copula. This would make the FGM model suitable to model weak correlation only. Uniform priors over the parameter ranges have been selected in the simulation in the next section. The transformed prior distribution for τ_b is over a greater range for the Gaussian copula due to different ranges of Kendall's tau. For example,

if we were to fix $\pi_E = 0.5$ and $\pi_T = 0.5$ then the prior for τ_b is uniform over $(-0.25, 0.25)$ and $(-1, 1)$ for the FGM and Gaussian copulas respectively.

3.5 Simulation study

To inspect the properties of copula models with binary data a simulation study was conducted. The main aim of the simulation study was to investigate the effect on posterior marginal distributions for efficacy and toxicity when fitting a copula model. The control for this question is a simpler independent model. It is possible that by correctly modelling correlation the variance of the marginal distributions is reduced. It is also possible that the copula would constitute a less parsimonious model. If there is little or no difference between marginal distributions when fit to the same data this would suggest that the copula will not impact upon decisions. This is because in dose finding trials many decision models depend only upon the marginal distributions for the probability of efficacy and toxicity. Given that a copula function separates marginal distributions and correlation, the effect on marginal distributions is expected to be negligible.

A secondary aim is to see the effect of fitting different copula models and to check whether the model can identify correlation as intended, and to quantify this. Little is likely to be known about the correlation prior to the start of the trial. The prior therefore is uninformative constituting the full range of possible values. Each value will correspond with a different measure of Kendall's τ_b with the measure at a particular dose depending upon the magnitude of marginal distributions. With small samples it is possible that little may be learnt about the correlation given the dependence on the marginal distributions; the simulation study will investigate this.

3.5.1 Single dose simulation

Multiple doses are evaluated in dose finding trials where marginal distributions can be estimated using logistic regression and dependence induced through a copula. The aim of the simulation study is to understand the effect between the marginals and the added correlation structure. To achieve this aim, the logistic regression component can be simplified to give greater focus on the copula component. This is achieved by looking only at a single dose. A simulation study was conducted on 20 patients. The sample size was chosen as an upper

estimate of the number of patients that would be treated at a single dose in a dose finding trial, similar to the motivating example (Section 2.5). Marginal parameters π_E and π_T represent the probability of any patient having an efficacy and a toxicity event respectively. Both parameters are assumed to have uninformative uniform priors. The simulation study will compare copula models to an independent model. The posterior distributions for the independent model are conjugate beta distributions i.e $\pi_E \sim \text{Beta}(1 + x_e, 1 + 20 - x_e)$ and $\pi_T \sim \text{Beta}(1 + x_t, 1 + 20 - x_t)$ where x_e and x_t are the total number of efficacy and toxicity events respectively.

Two copula models are fitted, the FGM copula given in Equation 3.13 and a Gaussian copula given in Equation 3.15. The copula parameter, $\theta \in [-1, 1]$, has a uniform prior distribution, $\theta \sim \text{Uniform}(-1, 1)$ for both copulas. The simulation study evaluates τ_b as a measure of correlation; the implied prior distribution for both copulas is plotted in Figure 3.5. The increased variance for the Gaussian copula is down to its structure as previously described. The posterior distributions for the copula models have no known closed form and are estimated using Markov Chain Monte-Carlo integration using Stan software [121].

Given the reduced complexity of a single dose it is possible to look at this via an exact method. This involves assessing every possible combination of the data. Given the ordering of patient responses will give the same posterior distributions, this reduces the number of possible combinations. This is because different orderings all give the same sufficient statistics. There are 1771 ways of choosing 20 patients from the possible four patient responses. The number of possible different data for the independent model is 441.

3.5.1.1 Single example

A reduced dataset is initially chosen to understand the aims of the simulation study. Let $x_e = 14$ and $x_t = 6$, this corresponds with a sample estimate of $\hat{\pi}_E = 0.7$ and $\hat{\pi}_T = 0.3$ taken from scenario 1 dose level 3 of the CSS, which is reproduced in the Appendix A.1.1. There are seven possible combinations of data that retain $x_e = 14$ and $x_t = 6$, these correspond with the seven possible number of patients with both an efficacy and a toxicity event. The number of concordant event pairs or ties ranges between 0 and 6. In the case of the independent model given the exchangeability of individual patients, the posterior distribution is the same for any number of ties when $x_e = 14$ and $x_t = 6$. Summary

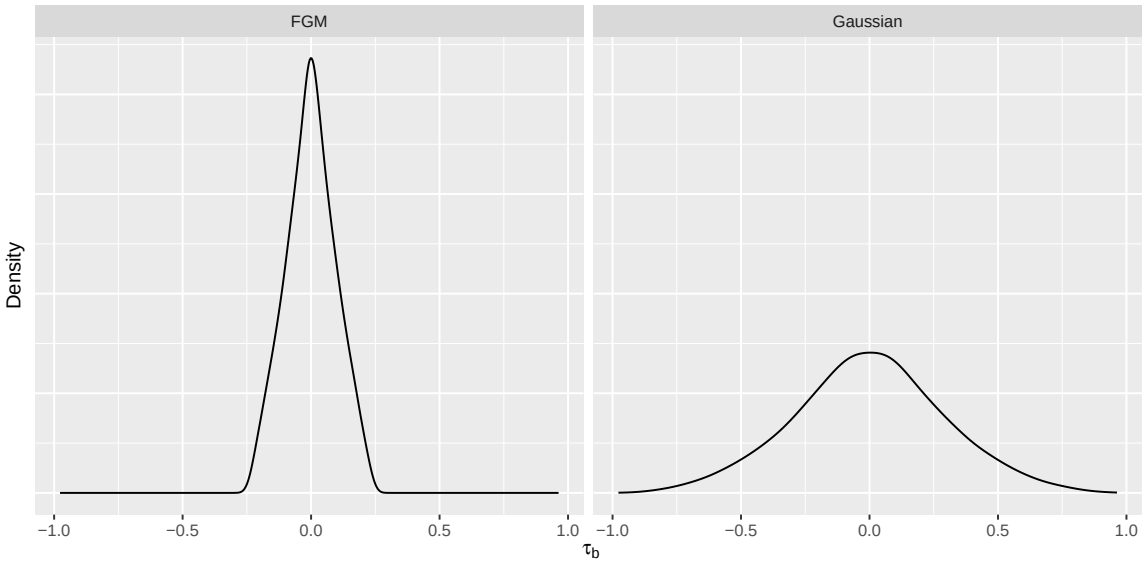

Figure 3.5: Implied prior for τ_b specified in the simulation study for each of the copula models. The parameter depends upon the uniform distributions specified for π_E, π_T and a uniform distribution over the parameter range in the model θ. Differences are down to limits that the copula model imposes, see Section 3.3.3.1

statistics for the posterior distributions for π_E, π_T and τ_b from the three models are given in Table 3.2.

Fitting a copula model has very little impact upon the marginal distributions with summary statistics being within 1 percentage point of the independent model for both copulas at all possible combinations of data. Efficacy and toxicity parameters medians lie in between the prior median of 0.5 and data estimates. The upper and lower range for the FGM copula limits its ability to estimate the correlation (Table 3.2). When the data suggest a stronger correlation the FGM estimates the correlation close to its limit (Figure 3.1). The FGM copula does appear to estimate the direction of correlation but makes little distinction when the correlation is stronger in the data. For the FGM copula when the number of ties is 1 or fewer $\tau_b = 0$ is excluded from the 90% credible interval, in all other instances $\tau_b = 0$ is included. The Gaussian copula is able to model stronger correlation with many of the credible intervals excluding $\tau_b = 0$. Even if the Gaussian copula is able to suggest the direction of correlation the credible interval is relatively wide. When there is little correlation as measured by the data τ_b the FGM copula has a narrower CI; this is due to the effect of the prior which is concentrated around independence (Figure 3.5).

The chance of the ties occurring given $x_e = 14$ and $x_t = 6$ is worth considering to place

Table 3.2: Posterior distribution summaries for all possible combinations of data from a single dose with 14 efficacy events and 6 toxicity events. Ties refers to the number of patients with both an efficacy and a toxicity event. The independent model has same fit independent of the number of ties. Data τ_b is the sample estimate of τ_b. Credible intervals (CI) are equal-tailed intervals

Model fit - data	Median π_E (90% CI)	Median π_T (90% CI)	Median τ_b (90% CI)	Data τ_b
Independent	0.69 (0.51, 0.83)	0.31 (0.17, 0.49)		
FGM - Ties=0	0.68 (0.50, 0.83)	0.32 (0.17, 0.50)	-0.16 (-0.23, -0.04)	-1.00
FGM - Ties=1	0.68 (0.51, 0.83)	0.32 (0.17, 0.50)	-0.15 (-0.22, -0.00)	-0.76
FGM - Ties=2	0.68 (0.51, 0.83)	0.32 (0.17, 0.49)	-0.13 (-0.22, 0.05)	-0.52
FGM - Ties=3	0.69 (0.51, 0.83)	0.31 (0.17, 0.49)	-0.09 (-0.21, 0.12)	-0.29
FGM - Ties=4	0.69 (0.51, 0.83)	0.31 (0.17, 0.49)	-0.02 (-0.19, 0.17)	-0.05
FGM - Ties=5	0.69 (0.51, 0.83)	0.31 (0.17, 0.49)	0.06 (-0.15, 0.20)	0.19
FGM - Ties=6	0.69 (0.52, 0.84)	0.31 (0.17, 0.48)	0.12 (-0.09, 0.21)	0.43
Gaus - Ties=0	0.69 (0.51, 0.83)	0.31 (0.17, 0.49)	-0.76 (-0.92, -0.49)	-1.00
Gaus - Ties=1	0.68 (0.51, 0.83)	0.32 (0.17, 0.49)	-0.58 (-0.80, -0.27)	-0.76
Gaus - Ties=2	0.68 (0.51, 0.83)	0.32 (0.17, 0.49)	-0.40 (-0.67, -0.08)	-0.52
Gaus - Ties=3	0.68 (0.51, 0.83)	0.32 (0.17, 0.49)	-0.22 (-0.52, 0.09)	-0.29
Gaus - Ties=4	0.68 (0.51, 0.83)	0.32 (0.17, 0.49)	-0.04 (-0.36, 0.25)	-0.05
Gaus - Ties=5	0.69 (0.51, 0.83)	0.31 (0.17, 0.48)	0.14 (-0.17, 0.39)	0.19
Gaus - Ties=6	0.70 (0.52, 0.84)	0.30 (0.16, 0.47)	0.33 (0.08, 0.53)	0.43

these results in context. If the true underlying data mechanism had $\pi_E = 0.7$, $\pi_T = 0.3$ and $\theta = 0.8$, for the FGM copula or $\tau_b = 0.168$ the chance of seeing the ties $= \{6, 5, 4, 3, 2, 1, 0\}$ would equal $\{0.27, 0.45, 0.23, 0.05, 0.00, 0.00, 0.00\}$ to 2 d.p., while if the data generating mechanism was independent this would be $\{0.08, 0.31, 0.39, 0.19, 0.04, 0.00, 0.00\}$.

3.5.1.2 Effect on marginal distributions

The single dose simulation study is now considered for all possible values of x_e, x_t and number of ties. The difference in means between the independent and the copula models for the probability of efficacy is is given in Figure 3.6. A negative value would constitute a larger mean for the independent model. The second column of the plot gives the ratio of standard deviations of the posterior distributions between the copula models and an independent model fit to the same data. A value greater than 1 occurs when the independent model has a smaller standard deviation. The equivalent plots from each of the copula models for the probability of toxicity are given in Figure A.1.1. There is only very minor effect on the marginal distributions by fitting a copula model as measured by the mean in all possible outcomes of data.

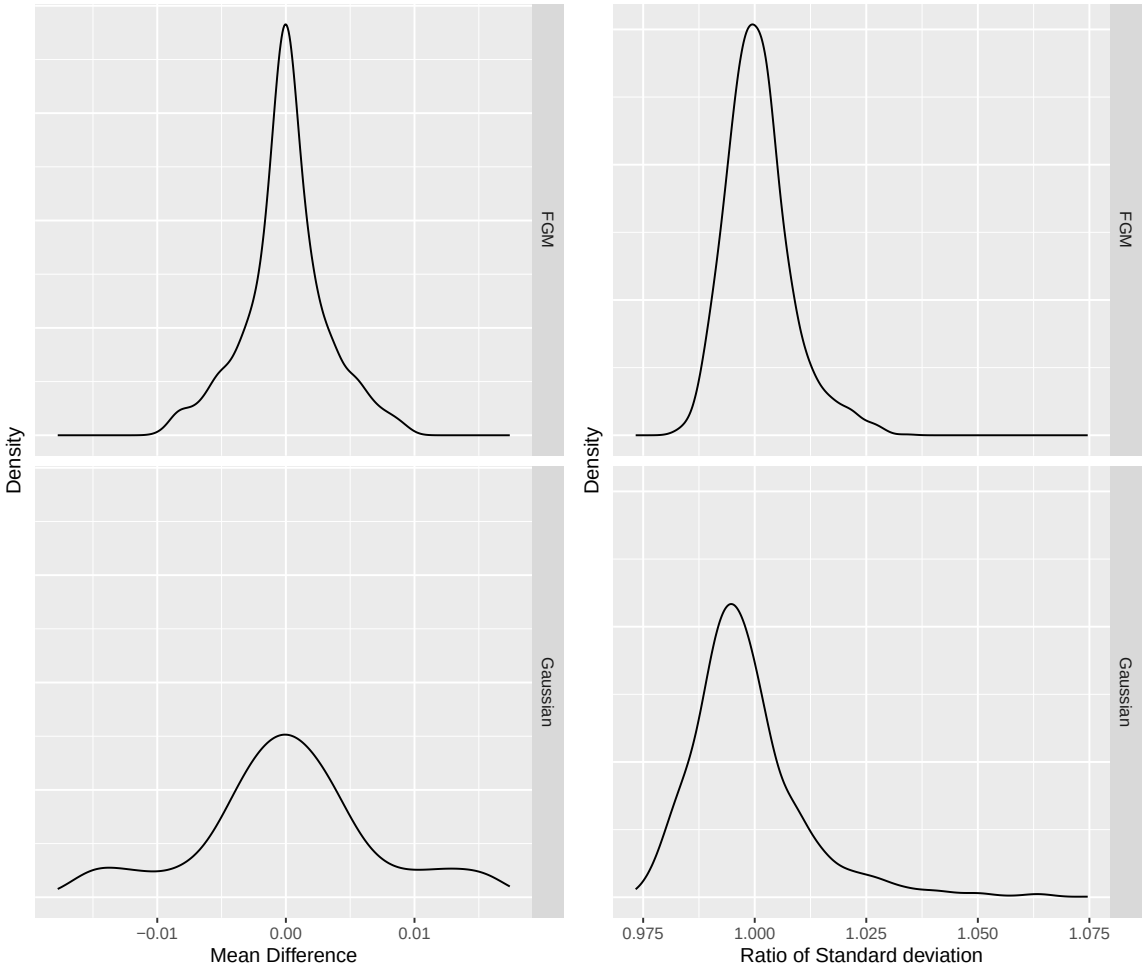

Figure 3.6: Difference in efficacy marginal distributions for all possible combination of data for 20 patients recruited at a single dose between copula models and independent models. First row of plots is from the FGM copula model and the second row from the Gaussian Copula. First column is the difference in means between the Copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

There is very strong similarity in the ratio of standard deviations. A small number of data sets result in a ratio of the standard deviation in excess of 1.05 for the Gaussian copula. This would suggest that the standard deviation for the copula model is larger than an independent model. An inspection of these data-points, seem to suggest this is when the correlation in the data is large ($|\tau_b| > 0.68$) and when both of the mean marginal probabilities are estimated to be large or small i.e in the range $[0, 0.15]$ or $[0.85, 1]$. Overall, the mean and the ratio of standard deviations are very similar.

This section has so far evaluated every possible combination of data and summarised. This will include a lot of data that is highly unlikely in a particular setting. Suppose if the true data generating mechanism is now defined by $\pi_E = 0.7$, $\pi_T = 0.3$ and $\theta = 0.8$ for the

FGM copula or $\tau_b = 0.168$ to correspond with scenario 1 of the CSS. The data generating process could also be defined with respect to a Gaussian Copula where $\tau_b = 0.168$ (See Table 3.1). The plots can be re-plotted with the density now proportional to the chance of any particular data set occurring under the true data generating mechanism. The efficacy probabilities and toxicity probabilities are plotted in Figures 3.7 and A.1.2 for efficacy and toxicity respectively. It is interesting that the skew on the means is different between the two copula models and the skew is in the opposite direction between toxicity and efficacy. A plot when the data generating process is independent is shown in Figures 3.8 and A.1.3 for efficacy and toxicity respectively. The Gaussian copula still has a small difference in the model fit while the FGM copula is very similar with a mean difference close to 0 and the ratio of standard deviations at 1.

Overall for this example the simulation study demonstrates that with small numbers of patients the copula model does not make any meaningful difference to the marginal distributions. The small effect on the marginal distributions for the FGM copula is smaller than the normal distribution. A possible concern of the conclusion is that extending the marginal distributions to account for the covariate of dose may have some effect. A small example from the CSS is recreated and simulated in Appendix A to check for any evidence for this. The results are consistent with the single dose simulation with the magnitude of difference between copula and independent models very similar.

3.5.2 Correlation

There are two aims of the simulation study relating to correlation, these are to see whether the copula model is capable of capturing the correlation structure and to what extent. The small example with $x_e = 14$ and $x_t = 6$ would suggest that copula models can at least capture the direction of the correlation. Figure 3.9 is a scatter plot comparing an estimate of the correlation in the data against the mean of the posterior distribution from that data from the two copula models. The FGM copula suggests a non-linear relationship unlike the Gaussian Copula. The relationship is caused by the limited range of the copula, where correlation as estimated by the data is beyond the limit imposed by the copula (Figure 3.1). As the correlation in the data increases beyond the copula limit, the copula estimates the correlation at its limit, to give what appears to be multiple small vertical asymptotes. The Gaussian copula is able to estimate the stronger correlation. The Gaussian copula has a

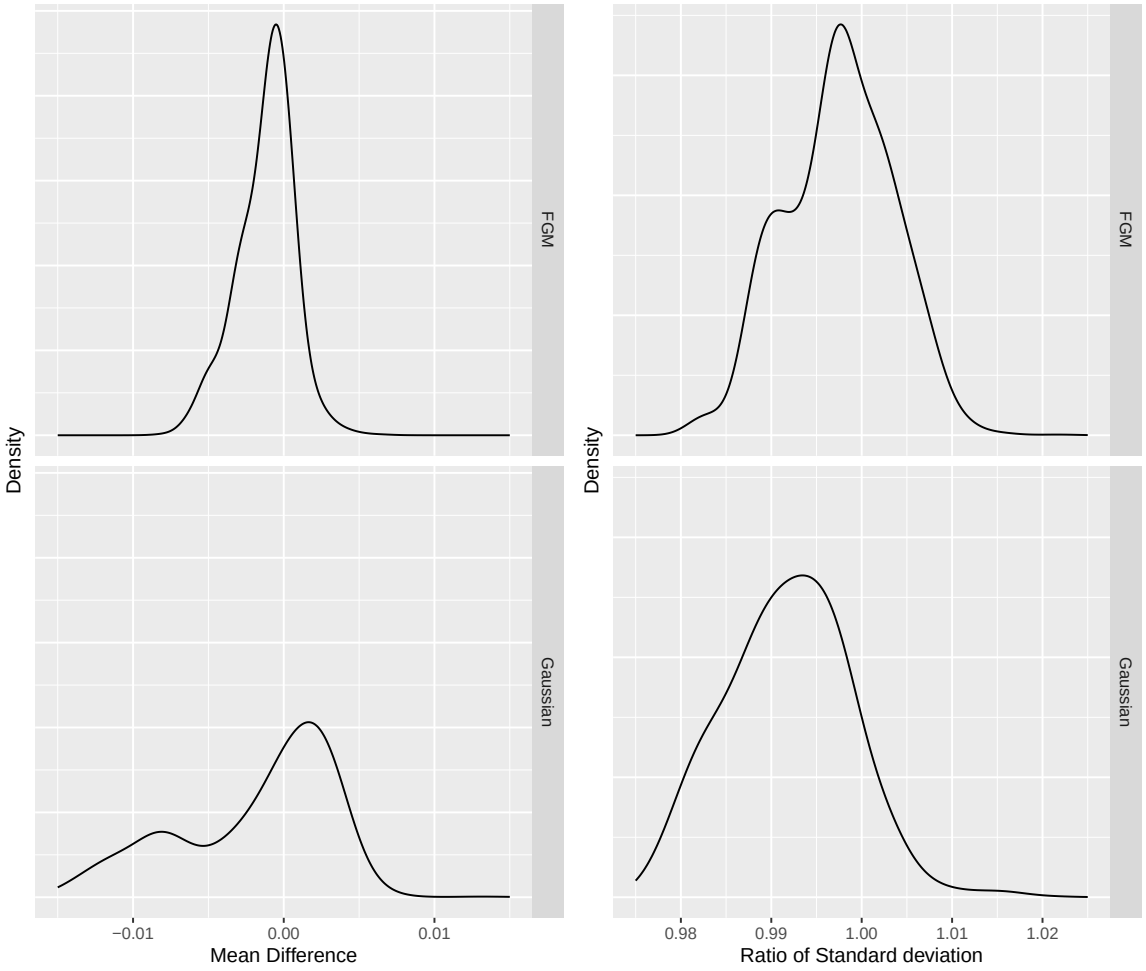

Figure 3.7: Difference in efficacy marginal distributions of data for 20 patients recruited at a single dose between copula models and independent models. Data is generated from a dose with $\pi_E = 0.7$, $\pi_T = 0.3$ and $\tau_b = 0.168$. First row of plots is from the FGM copula model and the second row from the Gaussian Copula. First column is the difference in means between the Copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

mean estimate nearer to zero in all instances in contrast to the data estimate. This is down to the effect of the prior which is uniform and centred around zero.

The range of credible intervals can be used to understand how accurate the estimation of correlation can be for a given model. Ninety percent equal tailed credible intervals are added to Figure 3.9 in Figure 3.10. For the FGM copula, in general it is only when the data $|\tau_b| > 0.5$ is $\tau = 0$ excluded from the credible interval. The range of the 90% credible interval for the Gaussian copula is approximately 0.5 with $|\tau_b| > 0.3$ generally excluding $\tau_b = 0$. This demonstrates that Gaussian copula model can identify correlation within the data although credible intervals for its estimation are quite wide.

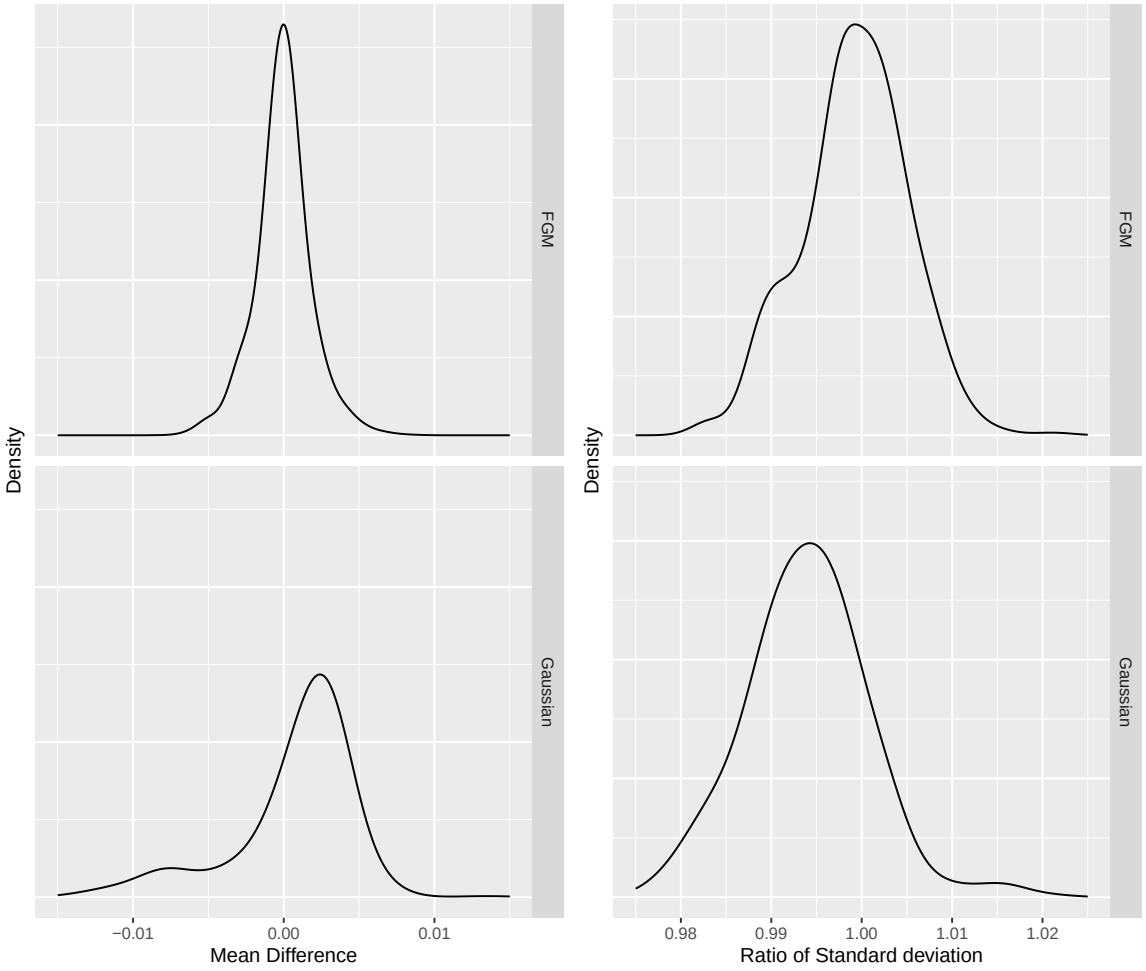

Figure 3.8: Difference in efficacy marginal distributions of data for 20 patients recruited at a single dose between copula models and independent models. Data is generated from a dose with $\pi_E = 0.7$, $\pi_T = 0.3$ and $\tau_b = 0$ or an independent data generating process. First row of plots is from the FGM copula model and the second row from the Gaussian Copula. First column is the difference in means between the Copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

To place this into context an example with a specified data generating process is needed. Using the example from previous and defining the data generating mechanism to $\pi_E = 0.7$, $\pi_T = 0.3$ and $\tau_b = 0.168$ (or $\theta = 0.8$ using an FGM copula). The probability of the credible interval excluding 0 is 0.13% and 16.54% for the FGM and Gaussian copulas respectively. When the data generating mechanism is independent with $\theta = 0$ this becomes 0.07% and 8.09% chance. For a data generating mechanism following the Gaussian copula with $\pi_E = 0.7$ and $\pi_T = 0.3$, $\theta > 0.69$ or $\tau_b > 0.36$ is needed to give at least a 50% chance of the credible interval excluding independence.

In summary there seems to be little merit in fitting the FGM copula to a limited data set

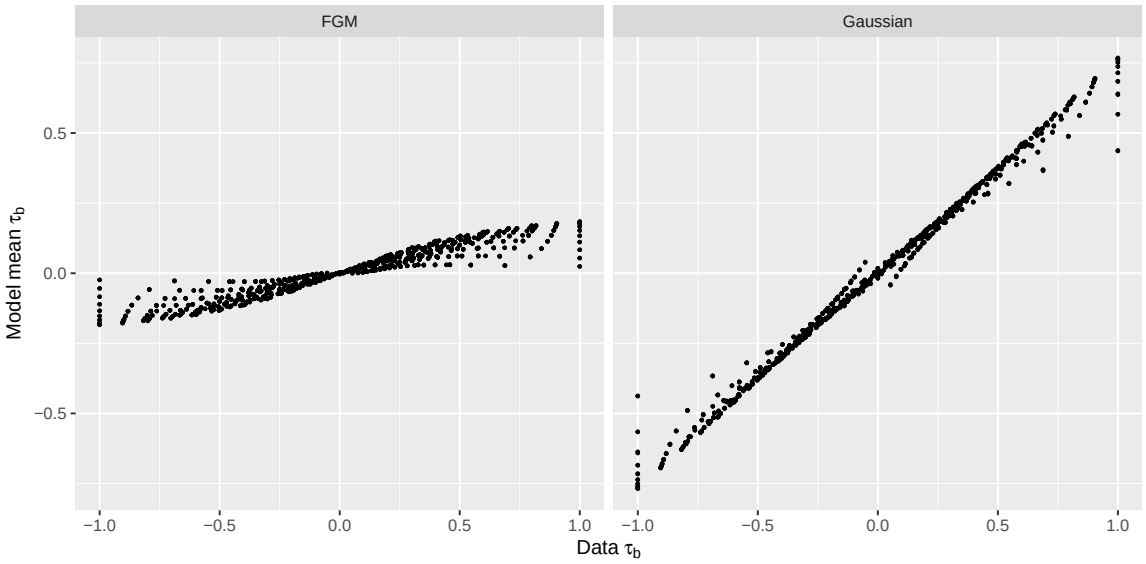

Figure 3.9: Correlation plot between the mean of the posterior distribution for τ_b and the estimate of τ_b from the data. Plotted for all values from the FGM copula and Gaussian copula

to gain insight into the correlation. When there is strong correlation in the data the copula doesn't account for it very well, when there is little correlation credible intervals remain quite wide relative to the prior. The small sample size isn't sufficient to provide insight over the FGM copulas limited range. The Gaussian copula provides better insight into the correlation and is able to suggest stronger correlation when it is a feature. The credible intervals are still quite wide and even for moderate correlation it is not able to provide good insight. If the correlation is present there is still a good chance it will not appear in the data.

3.6 Discussion

This chapter has drawn together copula theory for binary outcomes and applied it to the setting of dose finding. Previous work in the CSS provided empirical evidence in a limited setting that the fitting of a copula model doesn't seem to make a difference to operating characteristics in dose finding. There are a number of limitations to wider applicability and explanations as to why the results occur in the CSS. This chapter explains the findings of the CSS and adds a number of novel conclusions to the dose finding literature as follows:

1. Any study looking at correlation needs a consistent correlation measure in order to compare different models. Kendal's Tau has been suggested as a suitable measure.

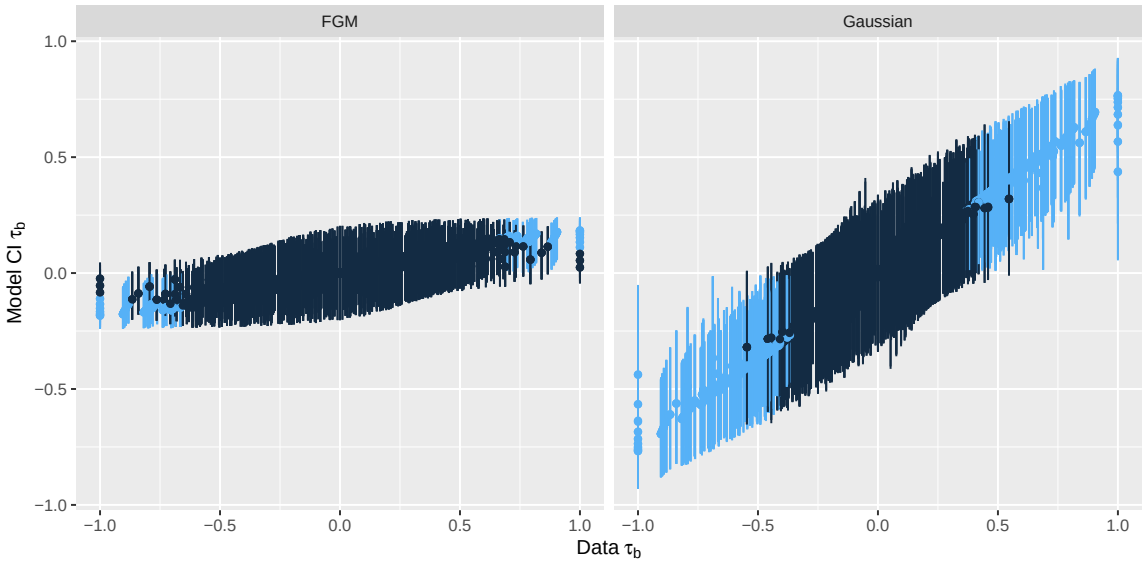

Figure 3.10: Correlation plot between the mean of the posterior distribution for τ_b and the estimate of τ_b from the data. Plotted for all values from the FGM copula and Gaussian copula. 90% equal tailed credible intervals are displayed. Light blue lines represent when the credible interval excludes $\tau_b = 0$, independence. Dark blue lines are when the credible intervals include $\tau_b = 0$, independence.

The interpretation of the correlation parameter in a copula function is specific to the model itself.

2. In dose finding if the decision function doesn't consider correlation, then a copula model will have very limited effect on operating characteristics in comparison to an independent model. As such it is suggested that the independent model is a more parsimonious model.

3. Different copulas are able to model different levels of correlation as measured by Kendal's Tau. The FGM copula that has previously been used in dose finding has a very limited ability to measure strong correlation. The Gaussian copula is more suitable to measure stronger correlation.

4. The precision of estimating correlation within a dose finding trial is limited, predominately due to the dependence on the marginal distributions in estimating any correlation.

The main advantage and appeal of copula models is the ability to have marginal distributions which have a intuitive meaning and then model correlation separately. The CSS describes the Braun model as a "copula". While the model does incorporate a parameter

to account for correlation between two patient outcomes it isn't a bivariate function that can be written in terms of the CDF of random vectors as given in Equation 3.9. The theoretical copula literature presented in this chapter doesn't apply to the Braun model. The efficacy component of the Braun model is conditional upon toxicity having not occurred. A correlation measure such as Kendall's tau is calculable for the model and will be needed to explore if considering it against other models using correlation.

This chapter has demonstrated that fitting an independent model in favour of a copula model has very little effect on deciding between doses. This is because marginal distributions that are used to make these decisions are extremely similar when fitting either models. A caveat would be if the decision model includes a correlation component then a more detailed inspection of how this affects the operating characteristics is needed. The EffTox individual patient utility is one such design. In this chapter the utility function was plotted for the extreme parameter values from the FGM copula. This was shown to have a very small effect on the utility function. The specification of the other two decision parameters has been shown to have a large effect on operating characteristics [78]. The utility plots from such changes are noticeably different while in the case of maximum and minimum correlation from the FGM copula they look similar. The correlation component is difficult to measure with accuracy as such conclusions for the EffTox utility design are similar in that the correlation component will have only a small effect with the marginal probabilities dominating the decision making process.

The work in this chapter provides a theoretical understanding as to why correlation is challenging in limited sample sizes. The parameter in a copula model can only be interpreted through the specified copula. The use of a standard, such as Kendall's tau allows comparison and evaluation of correlation across different models. The metric was chosen as it corresponded with much of the literature around binary copulas. Pearson's correlation coefficient is a similar measure that could be used. The correlation component appears to be estimated primarily through the number of ties in relation to what is expected from an independent model. Estimation of correlation is dependent upon the marginal distributions. These factors combine to yield a wide credible interval for estimating the correlation component. The model does however correctly identify correlation, as intended. The Gaussian and FGM copula explored in this chapter have the same parameter range but measure different

amounts of correlation. The FGM copula has a limited range for correlation in contrast to a different copula such as the Gaussian copula. The Gaussian copula is able to measure stronger correlation.

There are very small differences between posterior distributions from independent and the copula models. The term "dose ambivalence" has been introduced in the setting of dose finding when the expected utility between doses is very similar [41]. In a situation of dose ambivalence the error in model fitting from the same data could determine the dose recommendation. Similarly, in the setting of using a copula or not it is possible that small differences in posterior distributions could lead to a number of different decisions. In such instances there isn't a correct or incorrect decision; very minor differences in the specification of the utility function or priors may also similarly lead to different decisions. In the context of the the main objective of the study different factors such as the specification utility function or the admissibility rules will induce a much larger effect on the performance of the design.

Alternative approaches that allow modelling on the marginal scale such as a bivariate Probit model have been applied in the setting of dose finding [98]. The correlation in this model is applied to the link function on the continuous odds scale. The paper did not explore or contrast the correlation component in the model. A copula approach is a general approach not limited to two binary outcomes. Archimedean copulas have been applied to the dose finding setting with categorical toxicity and continuous efficacy [122]. There have been a number of uses of the structural form of a copula model for use in combination studies [123]. Here the structural form of a copula function allows MTD probabilities of individual drugs to be combined into a single probability for a single outcome. The role of Copulas in the setting of outcomes observed over a different time period has also been applied [124]. In theory correlation plays a more important role in this setting. When there is strong correlation, the observation of one component of the patient response can be more informative of the other to allow better decision making. In practice the limited ability of estimating the correlation as demonstrated in this chapter is likely to still dominate.

The independent model is computationally more efficient. The simulation study looking at a single dose for 20 patients was effectively instantaneous for the independent model. Whereas to fit the 1771 combinations took approximately half an hour (1 second per model fit) to

fit the FGM copula and 2 and a half hours (5 seconds per model fit) to fit the Gaussian copula due to additional evaluation of a normal integral. When evaluating a large dose finding simulation study the ability to fit an independent model allows significant savings in computational time (Appendix A). If the decision model for binary outcomes relies only upon the marginal posteriors a pragmatic solution would be to apply the independent model for the purpose of dose finding without significant impact on conclusions. A secondary or exploratory objective as part of the study could be used to evaluate the correlation. The remainder of this thesis utilises a more parsimonious independent model to capitalise on the gains in computational efficiency.

Chapter 4

Bayesian statistical decision theory

4.1 Introduction

This chapter reviews the methods of Bayesian statistical decision making. This covers both the theory and the rationale for maximising expected utility in addition to practical guidance in constructing and eliciting a utility function. The work is in the general setting of clinical trials although many examples are given in dose finding. The purpose of keeping the methods separate from the consistent theme of dose finding in this thesis is that the methods introduced are general and apply to many different settings. Later chapters describing the application to dose finding use a principled approach that can be used in many different disciplines. This can give us confidence that when adapting a dose finding trial design into a more specific individual trial setting it has strong foundations and a consistent method that can be revised and adapted.

A Bayesian decision theoretic approach to decision making can help to decide upon an action from a set of possible alternatives when the outcome is uncertain [27]. There are two main components: a Bayesian model representing the structure of a system and its associated uncertainty, and a consequence function to measure the merit of taking an action when the future outcome is known [25]. Specifying the consequence function is the main topic of this chapter. The problem of efficacy and toxicity dose finding involves making decisions with two competing factors, namely maximising efficacy while having an acceptable toxicity profile. Multi attribute decision making builds upon the key concepts that apply to the simpler single attribute setting. The chapter introduces these key concepts before moving

onto dealing with the bivariariate setting. Chapter 5 applies the content laid out in this chapter.

The layout of the chapter is as follows: A preliminary stage of scoping the decision problem is considered. This is a practical set of questions to ask in a decision analysis in order to refine it into a mathematical framework. The formal definition of maximising expected consequence is then restated. Formal mathematical notation of preference is given alongside an example of the perils of maximising the expected consequence function with an inappropriate scale. A consistent scale, utility, is then introduced to overcome the difficulties, with accompanying notation and a method to elicit it from a decision maker. Setting a parametric utility function is explored for a single attribute in detail by defining attitudes of a decision maker when faced with an uncertain choice. Two approaches to defining utility when there is more then one factor are then given.

4.2 Decision scope

A decision analysis can be used whenever a choice is to be made between at least two courses of possible action. There are a number of qualitative factors that need to be determined before undertaking a decision analysis. The aim of this is to understand what about the problem is important and how a decision analysis can help. The broad specification of the decision problem involves finding provisional answers to a number of general questions that are later encoded using the language of mathematics [25]. Much of the initial scoping of the decision problem is similar to the concepts in more general frameworks of creating and defining clinical trial questions.

4.2.1 Decision terminology

A number of terms are defined in this section that are subsequently used in the rest of the chapter associated with Bayesian decision analysis. Definitions all overlap with general terminology used within a clinical trials setting; as such each definition is accompanied with an example from a phase I dose finding trial. The example is a dose finding trial with binary efficacy and toxicity endpoints using the EffTox method, as for the motivating example in Section 2.5. The method models π_E and π_T, the probability of efficacy and toxicity events at each dose, modelled through a copula with logistic regression for each marginal distribution

(Equation 2.23). The decision to choose a dose for the next cohort at each stage is made with respect to an objective function,

$$O(\pi_E, \pi_T) = 1 - \left(\left(\frac{\pi_E - 1}{\pi^*_{1,E} - 1} \right)^r + \left(\frac{\pi_T}{\pi^*_{2,T}} \right)^r \right)^{\frac{1}{r}} \tag{4.1}$$

with $\pi^*_{1,E}$, $\pi^*_{2,T}$ and r constants pre-specified as described in 2.25.

Definition (**action**). *An action describes one of the possible choices in a decision analysis.*

In the example this would be the choice of dose at any stage of the trial. Note that the language in the decision literature is not entirely consistent and the term, "alternative" is often used to describe actions.

Definition (**state of nature**). *States of nature are the different possible conditions or scenarios in the decision analysis that may exist in the future. It represents the various states or situations that are beyond the control of the decision-maker and are uncertain or unpredictable.*

A probability density function gives the probability of each state of nature. When this is combined with data as part of a Bayesian model this is referred to as the posterior density function. In the example the parameters of the probability model all represent states of nature. The probability model for EffTox is described in 2.23. The prominent framework of setting a research question in clinical trials is the PICO (population, intervention, control, and outcomes) framework [125]. Part of the PICO framework is establishing objectives and outcome measures. The term outcome is typically used to denote a measured variable, for example a DLT event (yes/no). Endpoints refer to the analysed parameter that will be part of the clinical trial decision.

Definition (**attribute**). *In a decision analysis an attribute is used in a function to decide what the optimal action is.*

The definition coincides with the use of endpoints in clinical trials to establish whether the objectives are met. The example is multi-attributed with the chance of experiencing a toxicity event π_T, and an efficacy event π_E being the attributes. A level of an attribute describes the situation when an attribute takes a particular value, for example $\pi_E = 0.5$. An attribute will be a subset or derived from the states of nature. The use of the term attribute

allows us to refer to different components of the states of nature that will be associated with different objectives of the decision analysis. In the example the states of nature included in the probability model include parameters for correlation between outcomes as well as the two attributes.

Definition (**consequence**). *Consequence in decision analysis is a measure of the impact of enacting one of the actions when the true state of nature takes a particular level.*

In this chapter consequence is used to denote something positive with more consequence preferred to less. A consequence function describes a function that maps all levels of an attribute to a corresponding measure of consequence. This function is often referred to as an "objective", "loss", "gain", "value" or "utility" function within the clinical trials literature. Formal definitions of value and utility, which are both consequence functions, are given later in this chapter. In this example the consequence function is denoted by the objective function which gives a numeric value for every possible combination of the attributes π_E and π_T.

4.2.2 A unitary decision maker

Definition (**Decision Maker (DM)**). *A single decision maker (DM) is responsible for the decision analysis. In practice this means there is a single consequence function.*

In dose finding all the designs reviewed in Chapter 2 had a unitary DM. This thesis follows a normative decision theory defined by how a DM's beliefs should be structured if they were to follow certain elementary consistent rules which would be expected of a rational individual [26]. A normative theory is suitable for clinical trials where decisions need to be prescribed and justified in a trial protocol. The alternative is a descriptive decision theory; that models how decisions are made in practice. The goal of a clinical study is objectivity and as such a normative theory is the only sensible choice. The normative approach isn't however an objectively "correct" approach dictating to a DM how to act. The interpretation is that a greater understanding of the problem is developed through a decision analysis allowing the DM to make an informed choice.

There may however be multiple people that may input and share some of the responsibility for making a particular decision. In addition to this there may be further experts that may need consulting and stakeholders who will be impacted by the decision analysis. Aggregating

multiple opinions into a single set of preferences analytically is desirable in principle but in practice adds a further level of complexity to the problem, which cannot easily be resolved. Given there may be a conflict of opinion, who's opinion is more valid and who makes this judgement? Arrow's impossibility theorem in essence demonstrates that it is impossible to have both a completely democratic and rational decision with a group having any conflict of opinion [126].

The unitary DM can be thought as a concept whereby preferences are a compromise or consensus representing the situation written in the trial protocol which encodes the choices of multiple parties. In a slightly different setting of expert elicitation this has been encoded and given the term of a rational impartial observer (RIO) [127]. Multiple experts are gathered to elicit expert opinion, accepting that differences may occur, RIO sits outside of the group and gives a single impartial judgement accounting for all of the discussion that has been had. Scoping who the key people and stakeholders are in a decision analysis can help in making sure that the consequence function meets its objectives. The role of the statistician is to facilitate the discussion and obtain preferences from experts to encode this into mathematics in the form of a decision analysis.

Patients are one of the key stake holders in dose finding. A patient may find it difficult to comprehend the consequences of experiencing a DLT given that such an event has never been experienced [78]. This would make it difficult for a patient to provide a suitable utility function. In general however there is a growing acceptance that there may be disagreement between patient reported outcomes and clinician reported outcomes [128, 129]. A number of recent designs with a toxicity endpoint only have been proposed whereby a patient reported outcome is considered in tandem with the clinician reported toxicity [130, 131]. The design considers two separate attributes, the clinician reported outcomes and a further one based upon a validated patient questionnaire. The MTD is defined with respect to each of the endpoints, one patient MTD and one clinician MTD. The decision rule considers both perspectives as equally valid with the goal of finding a dose that satisfies both viewpoints.

4.2.3 Maximising expected consequence

The decision theoretic approach was stated in Chapter 2 and is repeated here for completeness. The fundamental principle in decision making in essence is to choose the action that

is most likely to give the DM the greatest reward. This is encoded into mathematics following a Bayesian statistical paradigm as follows. A decision problem can be defined when a DM needs to make a single action $a \in A$ from a set of possible actions. Features that are unknown about the situation before making the decision are modelled by an unknown state of nature, $\theta \in \Theta$. A consequence function $c(a, \theta) \in C$ specifies the consequence of making the decision $a \in A$ if the future outcome is $\theta \in \Theta$. Before making the decision an outcome $Y = y$ may be observed which depends on the unknown state θ. Prior knowledge of $\theta \in \Theta$ is incorporated via a prior $p_\theta(\cdot)$ and this is updated through Bayes theorem in light of the observation(s), to give the posterior distribution

$$p_\theta(\theta|y) \propto p_Y(y|\theta) \times p_\theta(\theta). \tag{4.2}$$

The Bayes decision $a \in A$ is the decision that maximises the posterior expected consequence:

$$a(y) = \arg\max_a(E_\theta[c(a, \theta)]). \tag{4.3}$$

This method is as for the decision theoretic approaches in dose finding as given in Chapter 2, Section 2.3.1.5. Note that if consequence is initially defined so that more consequence is worse than less then the Bayes decision is one that minimises the posterior expected consequence. For example, the CRM objective function, Equation 2.11, minimises an objective function. In principle the method is straight forward, but this misses the subtlety of defining the consequence function and difficulties in establishing the probability model. The rest of the chapter assumes that a DM has set up the parts of the decision analysis detailed earlier in the chapter. This includes a set of actions, defined attributes and a probability model. The specification of the consequence function from attributes is the focus of the rest of the chapter.

4.3 Value

A core tenant of the theory of decision making is preference. The notation of describing a DM's preferences is given in this section. This is used to give a numerical quantity to represent this preference; this is called a value function. Consider two levels x_1 and x_2 of an

attribute from the domain X describing all possible attribute levels. Consider the notation

$$x_1 \succcurlyeq x_2, \tag{4.4}$$

which is used when the DM considers x_1 to be at least as good as x_2. A stronger condition,

$$x_1 \succ x_2 \tag{4.5}$$

is when a DM strictly prefers x_1 to x_2. When

$$\{x_1 \succcurlyeq x_2 \text{ and } x_2 \succcurlyeq x_1\} \Rightarrow x_1 \sim x_2, \tag{4.6}$$

representing that a DM is indifferent between x_1 and x_2.

Transitivity is the property that with a third level x_3, in the domain X,

$$\{x_1 \succcurlyeq x_2 \text{ and } x_2 \succcurlyeq x_3\} \Rightarrow x_1 \succcurlyeq x_3. \tag{4.7}$$

The interpretation is that a DM has a fixed set of preferences that that do not change.

The property of *completeness* is when all levels within the domain of alternatives can be expressed using these relationships (also called *comparable*). The corollary of this is that a DM is decisive, in that given any pair of alternatives a preference (or indifference) can be stated.

If the conditions of completeness and transitivity hold for a DM then a real valued *value* function $v(\cdot)$ can be defined that represents preferences [27]:

$$\forall x_1, x_2 \in X, v(x_1) \geq v(x_2) \Leftrightarrow x_1 \succcurlyeq x_2. \tag{4.8}$$

A value function is not unique; any monotonic transformation will convey the same preference structure. In many settings the identity function of an attribute will constitute a value function. For example in the domain of clinical trials in oncology, when assessing only one attribute, more survival is invariably preferred to less, and therefore the proportion of patients alive at 1 year would constitute a value function. An attribute may have a natural preference structure but be non-monotonic. For example in a trial measuring blood

insulin levels, there is a peak of the value function within an acceptable range. Defining a value function between attribute levels above and below this range would need careful consideration.

The value function is a numerical representation of the ranking of preference and does not infer a measure relating to the idea of strength of preference. Consider three levels of an attribute of survival at one year, $t_1 = 100\%$, $t_2 = 80\%$ and $t_3 = 50\%$. In the previous paragraph it was stated that the raw attribute or identity function for the attribute constituted a value function as $v(t_1) \geq v(t_2) \geq v(t_3)$ to describe the preference $t_1 \succcurlyeq t_2 \succcurlyeq t_3$. It is not possible to infer anything from the magnitude of the differences however. It may be that $v(t_1) - v(t_2) < v(t_2) - v(t_1)$ but it isn't necessarily true that t_1 is preferred to t_2 less than t_2 is preferred to t_3.

The value function is on an ordinal scale; as calculating averages of ordinal variables is meaningless, calculating the expectation of a value function has no intrinsic meaning. Measuring the strength of preference so that a value function can be defined on a ratio scale is not an idea that is easily articulated or elicited [27]. Any value function that is able to measure strength of preference would be defined in the context of certainty. In the next subsection the classic example of the St Petersburg paradox is given to demonstrate that a DMs preferences also change when faced with an uncertain situation.

4.3.1 The role of uncertainty

Maximising expected value in order to make decisions in the presence of uncertainty can lead to poor decisions. This is made famous by the so called the St Peterburg paradox which was first described in the 18th century [132]. The example uses the attribute of money, where we note that a rational DM would prefer more money to less. Maximising the expected value of different actions in order to make decisions would intuitively seem rational and sensible in this case.

The paradox describes a game where a fair coin is flipped an indefinite number of times until a head appears. The prize is then $\$2^n$ where n is the number of times the coin is flipped. The decision to be made is the maximum price that a DM should pay to enter the game. The expected value of the game is infinite, calculated by summing each possible reward multiplied by the chance of it happening. The paradox is that according to a strategy of

maximising expected value a DM should pay any amount of money to enter the game even though it is almost certain that they will win a small amount.

There are a number of detailed explanations and technical details surrounding the history and development of St Peterburg paradox and accompanying development of utility theory that are beyond the scope of this thesis [133]. Maximising expected value only makes sense if using an appropriate scale, in the paradox the identity or raw attribute is inappropriate. Early solutions to the St Peterburg paradox formulated the initial stake in terms of a DM's overall net worth [132]. This accounted for the idea that many DM's would enter the game for a small stake but would not if the stake was high. Any scale that is chosen to measure strength of preference should account for the fact that a DM's preferences may change when faced with an uncertain outcome. Utility is defined in the next section assuming preferences for uncertain situations from the outset. This also allows us to convey a strength of preference or a ratio scale for consequence.

4.4 Utility

Von Neumann–Morgenstern utility theory [134] (VNM) gives four axioms of rational behaviour in the presence of uncertain outcomes. If these are accepted, the existence of a utility function follows. The theory accounts for uncertain outcomes.

The theory is set up using probability density functions for a continuous attribute x. The term lottery is used in this thesis for a finite subset of an attribute with all definitions still holding for the subset. Let $\mathbf{x} = x_1, x_2, \ldots, x_m$ be a set of m possible levels of a single attribute with $x_i \in X$ for $i = 1, \ldots, m$, where X denotes the domain of all possible values of x. A vector of probabilities, $\mathbf{d} = d_1, d_2, \ldots, d_m$, where $\sum d_i = 1$, represents the chance of $\mathbf{x}$ occurring. A *degenerate lottery* is used when $m = 1$ i.e, a single attribute level with certainty.

Let p, q and r represent three possible probability distributions. The four axioms of Von Neumann–Morgenstern utility state that a DM should be able to represent preferences over the distributions as follows:

Axiom 1 (Completeness) - A DM can express any distribution p, q in the form $p \succcurlyeq q$ or $q \succcurlyeq p$.

The DM is decisive and able to give a preference for any distribution.

Example: Suppose a DM is choosing between two treatments, p and q, where p has a 90% survival at one year with probability 0.5 and a 70% survival at one year with probability 0.5, and q has an 80% survival at one year with certainty. The DM should be able to state either $p \succcurlyeq q$ or $q \succcurlyeq p$ based on their preferences

Axiom 2 (Transitivity) - If $p \succcurlyeq q$ and $q \succcurlyeq r$ then $p \succcurlyeq r$.

The DM is consistent in stating preferences between distributions.

Example: Following the previous example, If a DM prefers p (90% survival at one year with probability 0.5 and 70% survival at one year with probability 0.5) over q (80% survival at one year with certainty), and prefers q over r (80% survival at one year with probability 0.5 and 60% survival at one year with probability 0.5), then the DM should also prefer p over r.

It is possible to mix probability distributions with a mixing probability $\alpha \in [0, 1]$ to form a new distribution. The notation $\alpha p + (1 - \alpha)r$ represents a distribution mixing p with probability α and r with probability $1 - \alpha$. Note that the mixing component must sum to one to be a further probability distribution.

Axiom 3 (Continuity) - If $p \succcurlyeq q \succcurlyeq r$, then there exists a probability $\alpha \in [0, 1]$ with

$$\alpha p + (1 - \alpha)r \sim q. \tag{4.9}$$

This axioms states that there is a tipping point between distributions.

Example: Consider a related but different example to above, If a DM prefers a treatment p (90% chance of survival at one year) over treatment q (80% chance of survival at one year), and prefers treatment q over treatment r (70% chance of survival), then there should be a probability α where the DM is indifferent between q and a mix of p and r. For instance, the DM might be indifferent between treatment q and a mix of p with probability 0.6 and r with probability 0.4.

Axiom 4 (Independence) - Given a probability $\beta \in [0, 1]$ then

$$p \succcurlyeq q \Rightarrow \beta p + (1 - \beta)r \succcurlyeq \beta q + (1 - \beta)r. \tag{4.10}$$

That is, preferences between p and q are independent of the presence (or absence) of other probability distributions.

Example: Consider a further example, If a DM prefers treatment p (70% chance of survival) over treatment q (60% chance of survival), then if we introduce another treatment r (30% chance of survival), the DM should still prefer a mix of p and r with probability β over a mix of q and r with the same probability β. For instance, if $\beta = 0.5$, the DM should prefer a mix of getting p with probability 0.5 and r with probability 0.5 over a mix of getting q with probability 0.5 and r with probability 0.5.

The main theory of VNM is that if a DM satisfies the four axioms then there is a function u that assigns a real value $u(x)$ to each possible value of x such that for any two distributions p and q,

$$p \succ q \iff E_p(u(x)) > E_q(u(x)) \tag{4.11}$$

Note that in Equation 4.11 the expectation relation still holds with some linear transformations of u. As such a utility function is unique up to a positive linear transformation. Utilities can be applied over any real number interval but more typically over the interval $[0, 1]$ (or $[0, 100]$) with $u(x_0) = 0$ and $u(x_*) = 1$ where x_0 and x_* are the minimum (least preferable) and maximum (most preferable) levels an attribute, x can take. In the case of unbounded continuous attributes a maxima and minima need choosing.

The axioms of VNM allow a DM to make coherent decisions by maximising expected utility. The rest of this section seeks to find a utility function that satisfies the VNM axioms.

4.4.1 Elicitation

The purpose of this section is to introduce a number of methods in order to elicit a utility function from a DM. This is achieved by obtaining preferences over a small number of lotteries. These are used to establish a more general preference structure in order to define a utility function at all points. As preliminary notation consider a number of levels of the attribute denoted by x_i with $x_i \in X$, the domain of all possible attribute levels. The elicitation assessment methods will consider lotteries with two components in addition to a mixing component α, with $0 \leq \alpha \leq 1$. The notation of angle brackets $\langle x_1, \alpha, x_2 \rangle$, is used to denote a lottery between x_1 and x_2 with $p(x_1) = \alpha$ and $p(x_2) = 1 - \alpha$. The relation

is abbreviated to $\langle x_1, x_2 \rangle$ when denoting an equal lottery with $\alpha = 0.5$. The assessment in utility elicitation used in this section will be of a comparison between the lottery and a certain outcome:

$$\langle x_1, \alpha, x_3 \rangle \; R \; x_2 \tag{4.12}$$

where R is one of the relations, $\{\prec, \sim, \succ\}$ to denote the direction of preference or indifference. An underline is used to denote the object that is being elicited. There are four main methods defined by which object of the lottery is being elicited.

1. preference comparison: $\langle x_1, \alpha, x_3 \rangle \; \underline{R} \; x_2$
2. probability equivalence: $\langle x_1, \underline{\alpha}, x_3 \rangle \sim x_2$
3. value equivalence: $\langle \underline{x_1}, \alpha, x_3 \rangle \sim x_2$
4. certainty equivalence: $\langle x_1, \alpha, x_3 \rangle \sim \underline{x_2}$

The latter three methods use axiom 3 of VNM as there exists a value of α that satisfies the relation. The method for preference comparison is to ask which option is preferred by the DM. The method for the latter three is to find a suitable value for the elicited component until the DM is indifferent between the two options. Preference comparisons will not give a magnitude for the utility; however by proposing a large number of preference comparisons these can be used in order to limit the utility function to be within quite tight bounds; the method when used in this manner is described as a discrete choice experiment [135]. Elicitation involving lotteries are routinely used in economic analysis for health states [136, 137].

The main determinants of deciding which method to use in a given situation are the ease of obtaining a utility whilst reducing any error or bias. Any elicitation is subject to bias or error using appropriate methods and an awareness of the main sources of bias ensures that a utility function is as accurate as possible. There are a number of factors that are known to induce measurement error from the behavioural sciences including how the question is asked [138]. Earlier in the chapter it was stated that the statistician is responsible for conducting the elicitation to obtain preferences from a DM. A greater awareness of the major sources of biases ensures that preferences of the DM are as accurate as possible and unintentional biases from the statistician are not introduced.

When choosing attribute levels they should be reasonably close together in a space that is well understood. This is not a strict rule but deviations from this principle increase the elicitation error [139]. For example, for an attribute representing the probability of efficacy for a new treatment consider a lottery between perfect efficacy and zero efficacy. Such levels are rarely encountered in practice and are hypothetical choices which are at the extremes of the attribute space, likely to yield a biased result. Another bias is the difficulty most people have in appreciating probabilities that are very small or very large. Assessing lotteries where $\alpha > 0.9$ or $\alpha < 0.1$ is likely to increase the measurement error. Each method could produce a slightly different bias, relying on one technique alone to elicit a number of relations could result in a more systematic bias.

The midpoint method is a non-parametric method to ascertain a complete utility function when the attribute is increasing. Define the maxima and minima as x_0 and x_* of the attribute range respectively. The midpoint method initially uses the certainty equivalent method to elicit the point $x_{.50}$ defined by the lottery

$$\langle x_0, x_* \rangle \sim \underline{x_{.50}}. \tag{4.13}$$

further points are then elicited with certainty equivalents with the range defined by the previous elicitation

$$\langle x_0, x_{.50} \rangle \sim \underline{x_{.25}} \tag{4.14}$$

and

$$\langle x_{.50}, x_* \rangle \sim \underline{x_{.75}}. \tag{4.15}$$

Two steps may be a reasonable approximation to the function or a further four evaluations (and so on) may be needed until a reasonable approximation to the utility function is generated [25]. Points on the continuous function can be connected by interpolation to define the function at all points. For example, consider treatments with percentage of survival at one year. Minima and maxima are $x_0 = 0\%$ and $x_* = 100\%$. A certainty equivalent was posed to a DM to find where they would be indifferent to receiving a treatment with certainty or an equal lottery to a treatment with perfect one year survival and zero survival

at one year. The DM gave 80% for the certainty equivalent,

$$\langle 0\%, 100\% \rangle \sim \underline{80\%}. \tag{4.16}$$

No further steps in the elicitation were taken. Points are interpolated to define a utility function over the entire range, which is plotted in Figure 4.1.

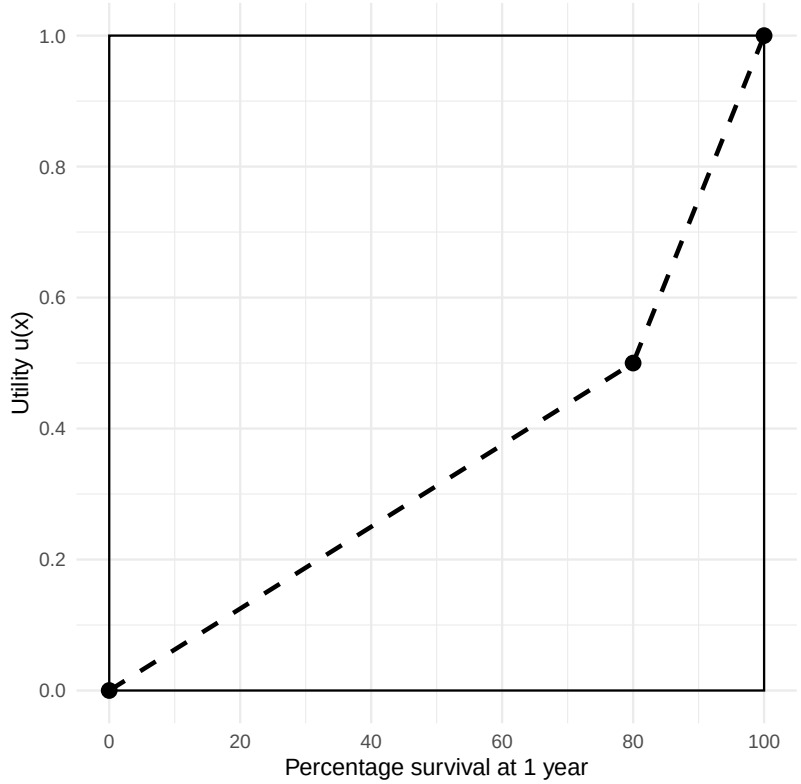

Figure 4.1: Midpoint method of elicitation for an attribute percentage of survival at 1 year. Minima and maxima of the attribute are $x_0 = 0\%$ and $x_* = 100\%$. A certainty equivalent elicitation, $\langle 0\%, 100\% \rangle \sim \underline{80\%}$ is plotted for one step. Points are interpolated to define a utility function over the entire range.

The midpoint method is a simple method of defining a utility function. The main issue is however in the likely quite high elicitation error. The first elicitation is a lottery over the entire domain space, which is known to not be best practice. Subsequent lotteries all rely on the initial lottery. As in the earlier example elicitation in Equation 3.24, asking a DM to weigh up treatments with perfect or zero efficacy is challenging, as they are not encountered in practice.

An alternative to eliciting a non parametric utility function is to assign a parametric functional form to define a utility function. This is particularly useful for a continuous attribute

where interpolation is known to be an approximation. Further benefits are that it is easier to describe a DMs attitude to risk and the use of a function makes it easier to calculate expected utility. Parameters that define the function can be elicited by using the techniques of this section as points elicited following the parametric function. How to decide whether a function form captures the situation is the topic of the next section.

4.5 Parametric utility functions

The purpose of this section is to inspect a number of basic preference structures that in turn will imply a functional form or suitable shape for the utility function. These are considered in turn with a number of functional forms considered at the end of the section.

In many applications the raw attribute constitutes a natural ordering or a value function. Any utility function has the property of also being a value function with a numerical representation of preference. The basic structuring can be useful in the first steps in defining a parametric form for a utility function. Take for example the probability of surviving one year following diagnosis with multiple myeloma. More survival is invariably preferred to less. Take two levels of a value function x_1 and x_2, where $x_2 > x_1$. The utility function for the attribute must also be *monotonically increasing* so that if

$$\langle x_2 > x_1 \rangle \Leftrightarrow \langle u(x_2) > u(x_1) \rangle. \tag{4.17}$$

For a monotonically decreasing attribute

$$\langle x_2 > x_1 \rangle \Leftrightarrow \langle u(x_2) < u(x_1) \rangle, \tag{4.18}$$

This initial step greatly reduces the number of possible functions. The next section goes through structuring a monotonically increasing utility function from a small number of preferences. A monotonically decreasing function uses the same ideas and concepts but many of the definitions flip; this is highlighted in the section proceeding.

4.5.0.1 Monotonically increasing utility

The general process is to inspect how a DM's preferences change when faced with uncertainty. This is described as a DM's attitude to risk. How a DM's attitude to risk changes

at different points in the domain of the attribute will imply different parametric forms for the utility function. Three attitudes to risk can be defined with respect to certainty equivalents defined in the continuity axiom of VNM utility theory for an attribute x that is monotonically increasing.

Risk aversion is when a DM prefers the expected consequence of any non degenerate lottery to that lottery.

Risk neutrality is when a DM is indifferent between the expected consequence of the lottery and the lottery.

Risk prone is when a DM prefers the lottery to the expectation of the lottery.

Consider the example elicitation for the midpoint method given earlier, where the attribute was survival at one year,

$$\langle 0\%, 100\% \rangle \sim \underline{80\%}. \tag{4.19}$$

The expectation of the lottery is 50%. The elicitation describes Risk aversion. If the certainty equivalent was 50% this would describe risk neutrality. If the certainty equivalent was less than 50% this would describe risk prone.

It is important to be aware each of the definitions are made with respect to a raw attribute that has been chosen as part of the decision problem work up. It is possible that a DM may have a different attitude to risk if a non linear transformation of the attribute is applied. Risk is a mathematical definition with respect to an attribute and is useful to give a consistent approach to defining utility. It should not be confused with the everyday usage of the words that imply an ethical judgement of a DMs preferences.

Different attitudes to risk will imply a different shape for the utility function. This can be shown graphically in Figure 4.2, where a probability equivalent elicitation,

$$\langle x_0, x_* \rangle \sim \underline{x} \tag{4.20}$$

is shown for a number of scenarios. The maximum, x_*, and minimum. x_0 levels of the attribute are plotted by green dots. The attribute has been standardised for simplicity, with the lowest level $x_0 = 0$ and the highest level of the attribute $x_* = 1$. The expected value of the equal lottery is 0.5. The definitions above define that the DM is risk neutral

if the probability equivalent is 0.5 (black dot), risk averse if the equivalent is less than 0.5 (red dot) and risk prone if the equivalent is more than 0.5 (blue dot).

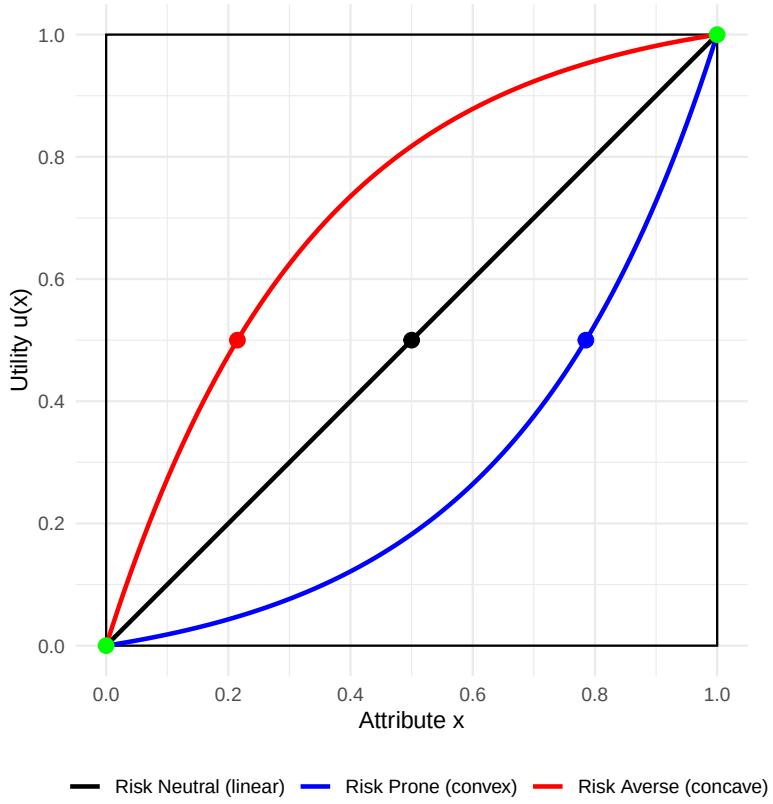

Figure 4.2: Plot of the mapping of an attribute x (x axis) onto $u(x)$ (y axis) for a monotonically increasing attribute. Green points represent the simple lottery $\langle x = 0, x = 1\rangle \sim \underline{x}$ with $\underline{x}$ the certainty equivalent. When the DM is risk neutral, the certain equivalent is equal to the expectation of the lottery. Risk aversion is when the certainty equivalent is less than the expected consequence of the lottery. Risk prone is when the certainty equivalent is more than 0.5. Each of these relations imply a corresponding shape for the utility. If the DM is risk averse at all points the shape of the utility must be concave. A risk prone DM has a convex utility function.

The risk premium of a lottery can be defined as the difference between the expected value that the lottery takes and the certainty equivalent. By eliciting a small number of certainty equivalences over the domain of the attribute, it is likely in many settings that the DM will have a consistent risk premium across the domain of the attribute [110]. If the risk premium is consistently zero the DM preferences will be described by a linear function. If consistently positive the DM is consistently risk averse and a concave function will describe preferences. While if the risk premium is consistently negative (risk prone) a convex function will describe preferences.

In order to inspect the properties of a parametric form of a utility and link back to a DMs

preferences it is necessary first to give a further measure of risk aversion. The Arrow-Pratt measure of absolute risk aversion (APARA) [140, 141], at a point $x = b$ is given by

$$r(b) = \frac{-u''(b)}{u'(b)}. \tag{4.21}$$

where u' and u'' are the first and second order differentials of the function with respect to the attribute, x. For a monotonic parametric utility function the sign of $r(b)$ denotes the attitude to risk, and the magnitude denotes a measure of departure from risk neutrality.

When a DM is risk neutral a DM's utility function is linear, $u(x) = x$. The linear function has a second order differential of $u''(b) = 0$ at all points, so APARA will be zero. Note that the utility is indifferent to linear transformations so $u(x) = mx + c$, where m and c are constants will yield the same result.

The following can be shown for risk aversion and risk seeking [110]:

Risk aversion has positive APARA and gives a concave shape.

Risk prone has negative APARA and gives a convex shape.

Two common utility functions are introduced with a single parameter, a. The first is the exponential function, which has the property that a DM has a constant APARA at all points. The second is the power function which has APARA proportional to the raw attribute b or that $br(b) = a$. The exponential utility function is defined as:

$$u(x) = \begin{cases} (1 - e^{-ax})/a & a \neq 0 \\ a & a = 0 \end{cases} \tag{4.22}$$

Note that $r(b) = a$ at all points for the exponential utility function. As such a DM will be risk neutral when $a = 0$, risk averse when $a > 0$ and risk prone when $a < 0$. The constant proportional risk aversion utility function or power function is defined as

$$u(x) = \begin{cases} x & a = 0 \\ (x^{1-a}) & a < 1, a \neq 0 \\ ln(x) & a = 1 \\ -x^{-(a-1)} & a > 1 \end{cases} \tag{4.23}$$

Care should be taken with the sign of the raw attribute, x as APARA is now inversely proportional. The attitude to risk at any point now depends on the sign of the raw attribute and the sign of constant a. For a positive attribute the attitude to risk at any point has the same interpretation as the exponential function, but the magnitude now increases with increasing x.

A number of preliminary elicitations were described at the start of the section that could be used to get a sense of the attitude to risk and also a non parametric measure, the risk premium, for a number of lotteries across the domain of the attribute. Inspecting the pattern of the risk premium with respect to the attribute and comparing to the APARA measure for a function allows us to obtain an appropriate corresponding function [110]. In the reference this is quoted as an iterative and heuristic approach rather than strictly quantitative, where a suitable function is proposed and fitted before consistency checks to see if it fits with the preliminary elicitations. For example, a positive and constant risk premium across the attribute domain would suggest that an exponential function is appropriate with the parameter $a > 0$.

Parameters associated with a function can be obtained by using a single probability equivalent, a certainty equivalent or value equivalent method and then back substituted to find the parameter. For example if $u(x) = exp(-ax)$, where a is the constant to be found and the relation

$$\langle x_1, \alpha, x_3 \rangle \sim \underline{x_2}, \tag{4.24}$$

would be found using a certainty equivalent, then the parameter a is the value that satisfies the equation

$$\alpha e^{-ax_1} + (1 - \alpha)e^{-ax_3} = e^{-ax_2}. \tag{4.25}$$

Bias in the elicitation can be reduced by choosing fixed values that minimise the magnitude between the upper and lower levels in the lottery [25]. If using a probability equivalence elicitation the fixed values can be chosen so that α isn't too small (or too large) in order to minimise any elicitation error. The preferences implied by the utility at other points should be checked via consistency checks in relation to the preliminary elicitations.

Consider the example elicitation for the midpoint method given earlier, where the attribute was survival at one year. The certainty equivalent could be made closer to what is en-

countered in clinical practice. In doing so the elicitation bias is reduced in contrast to the midpoint method. The following elicitation was obtained:

$$\langle 40\%, 60\% \rangle \sim \underline{55\%}. \tag{4.26}$$

Assuming that the exponential utility function was appropriate, Equation 4.25 could be used, where $x_1 = 40\%$, $x_3 = 60\%$, $x_2 = 55\%$ and $\alpha = 0.5$, to find the parameter a of the utility function. The utility function is defined over the entire domain, further checks via elicitation beyond $x_1 = 40\%$ and $x_3 = 60\%$ should be made to ensure this is a reasonable approximation.

4.5.0.2 Monotonically decreasing utility

This short section highlights the differences when the attribute is monotonically decreasing. The definitions of risk previously given still hold in the case of a monotonically decreasing function; a DM is risk averse (neutral, seeking) when they prefer (indifferent to, do not prefer) the expected consequence of any non degenerate lottery to that lottery. The utility function representing these preferences must of course be reflective of this. The shape of the utility function (concave or convex) now changes so that the definitions of risk hold, Figure 4.3. The measure of risk aversion, APARA also changes. For simplicity define mAPARA as

$$mr(b) = (-1)r(b), \tag{4.27}$$

the negation of APARA. When working with a monotonically decreasing attribute substitution of APARA with mAPARA will ensure that all the definitions for the parametric utility functions in the previous section still hold.

4.6 Bivariate utility

In this section it is demonstrated that definitions for univariate attributes still hold, with a number of additional considerations needed for two attributes. The intent is to define a utility function when there are two attributes. As a preliminary introduction to the notation and the problem, consider two attributes X and Y, we define a point (x, y), within

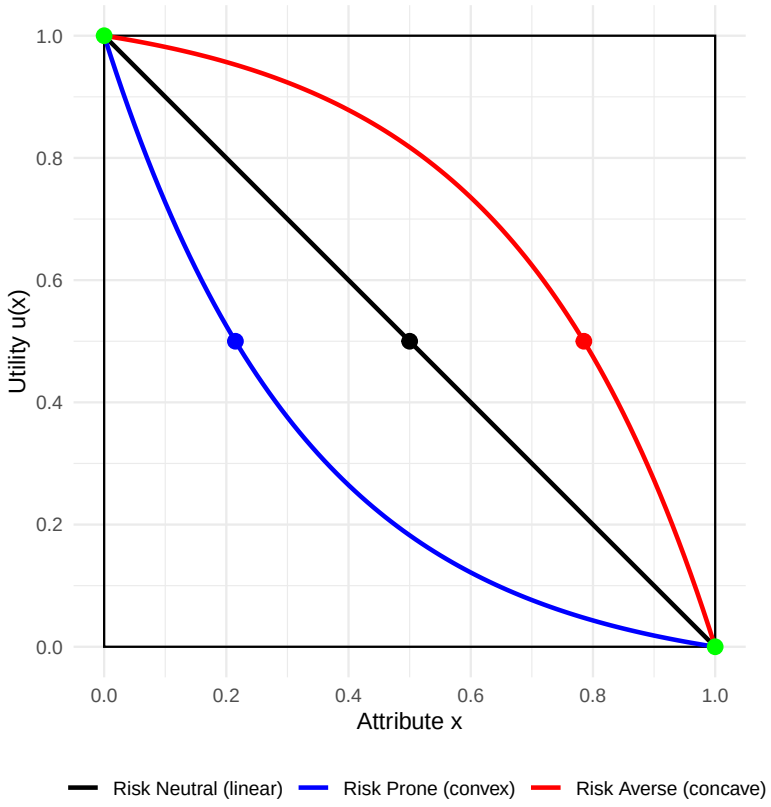

Figure 4.3: Plot of the mapping of an attribute x (x axis) onto $u(x)$ (y axis) for a monotonically decreasing attribute. Green points represent the simple lottery $\langle x = 0, x = 1\rangle \sim \underline{x}$ with $\underline{x}$ the certainty equivalent. When the DM is risk neutral, the certain equivalent is equal to the expectation of the lottery. Risk aversion is when the certainty equivalent is more than the expected consequence of the lottery. Risk prone is when the certainty equivalent is less than 0.5. Each of these relations imply a corresponding shape for the utility. If the DM is risk averse at all points the shape of the utility must be concave. A risk prone DM has a convex utility function.

the domain of all possible outcomes $X \times Y$, such that

$$x_0 \leq x \leq x_* \quad \text{and} \quad y_0 \leq y \leq y_* \tag{4.28}$$

where x_0 and y_0 are the minima and x_* and y_* the maxima of each domain.

Following the notation in the univariate setting to represent uncertainty, define lotteries, p and q by a probability distribution between possible outcomes, now defined by points (x, y), within the consequence space C, now defined over the surface $X \times Y$. The main theory of VNM utility theory still holds,

$$p \succ q \iff E_p(u(x, y)) < E_q(u(x, y)), \tag{4.29}$$

and a DM will act rationally by maximising expected utility [27]. The bivariate utility function is defined as $u(x, y)$ from hereon.

It is possible to elicit the utility, $u(x, y)$, at any point using the elicitation methods for a simple lottery between the most and least preferable points in the consequence space. One option could be to split the consequence space into a grid, assuming that the utility at all points within the grid are equal. Specifying a relatively modest number of intervals for each attribute, such as four would still result in 16 separate evaluations, with further necessity for consistency checks. Two issues arise: the function isn't particularly easy to assess in a manner that reduces elicitation error, and what you end up with is in most cases an over simplification of the problem [110]. The method doesn't exploit any features of the attributes, such as monotonicity, that may be known.

There are two broad approaches introduced in this section to yield a bivariate utility function. The methods have a similar approach to those proposed in the univariate setting where basic preferences and attitudes to risk are elucidated to specify a broad form that is then easily elicited from the DM. The first method tests to see if the joint utility can be given functional form, f formed from separate utility functions, u_X and u_Y of the two attributes , i.e.

$$u(x, y) = f(u_X(x), u_Y(y)). \tag{4.30}$$

Certain conditions are necessary for this to be the case and imply a specific function. The method exploits the relative ease of specifying univariate utility functions. The second approach, which is not as prominent and is included for completeness, maps the two attributes to a single scalar value function which in turn can be transformed into utility using the methods described previously for a single attribute.

4.6.1 Utility independence

In this section a formal definition of utility independence is given. If utility independence can be demonstrated then a simple linear relationship is implied between marginal utility functions [142, 110]. Utility independence can be tested with a small number of elicited lotteries as given in the previous section. Constants of the utility can also be determined by elicitation.

Following Keeney, utility independence requires the inspection of uni-dimensional *conditional utility functions*. These are defined by fixing one of the attributes and inspecting the utility function across the other. The conditional utility function of x given $y = y_1$, a point in the domain of y, is denoted by $u(x|y_1)$ and similarly the conditional utility function of y given $x = x_1$ is denoted by $u(y|x_1)$. This can be represented in a diagram, Figure 4.4. Define four points in the attribute space $X \times Y$ by $(x_1, y_1), (x_2, y_1), (x_1, y_2)$ and (x_2, y_2). Two conditional utility functions of x on y can be defined as $u(x|y_1)$ and $u(x|y_2)$ (horizontal dashed lines in the Figure). Similarly there are two conditional utility functions of y on x, $u(y|x_1)$ and $u(y|x_2)$ (vertical dashed lines in the Figure).

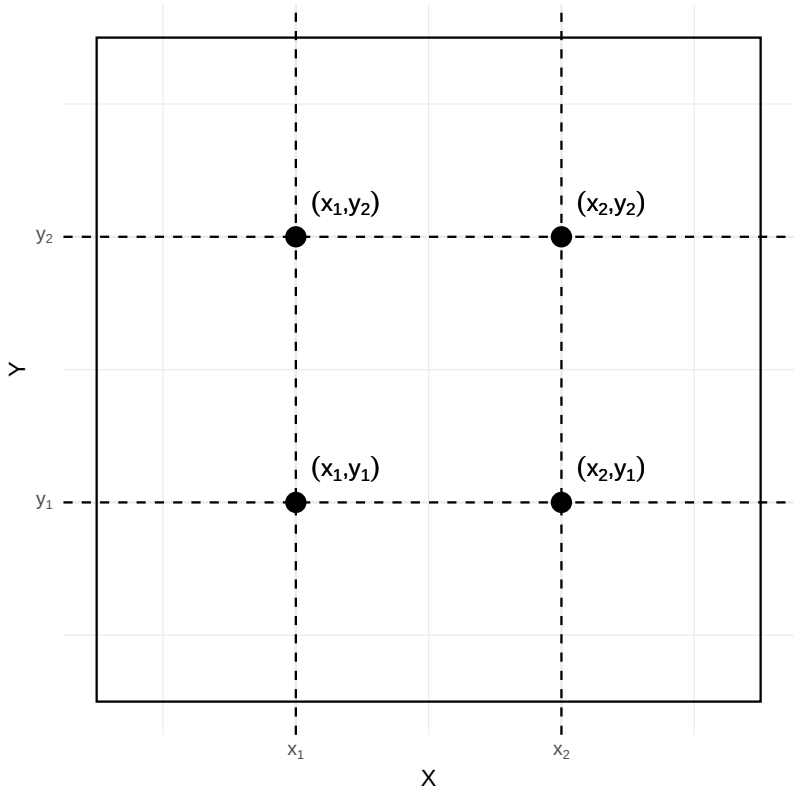

Figure 4.4: Plot of the joint attribute space $X \times Y$. Four points in the joint outcome space are plotted, $(x_1, y_1), (x_2, y_1), (x_1, y_2)$ and (x_2, y_2). Horizontal dashed lines are the domains of two conditional utility functions $u(x|y_1)$ and $u(x|y_2)$. Vertical dashed lines are the domains of the conditional utility functions of y on x.

Considering the utility functions $u(x|y_1)$ and $u(x|y_2)$, these can considered to be equivalent if the utility function does not change with y. If this is true for all values of y then X can be considered utility independent of Y. More formally, any attribute X is *utility independent* of another attribute Y, when conditional preferences for lotteries on X given Y do not depend on the particular level of y. In practice it is likely that marginal conditional utility

functions won't have been established at the preliminary stage of a decision analysis. To establish whether X is utility independent of Y consider the four points and the two utility functions $u(x|y_1)$ and $u(x|y_2)$ in Figure 4.4. Consider a certainty equivalent elicitation and a further preference relation comprising of components of the four points as follows:

$$\begin{aligned} \langle (x_1, y_1), (x_2, y_1) \rangle &\sim \underline{(x', y_1)} \\ \langle (x_1, y_2), (x_2, y_2) \rangle &\ \underline{R}\ (x', y_2). \end{aligned} \tag{4.31}$$

This represents the conditional utility functions along the horizontal dashed lines and assessing certainty equivalents between the two points on each line. If the second relationship can be considered equivalent then this would imply that the conditional utility functions are equivalent. Establishing this relation over further points and regions of the joint domain would be convincing that X is utility independent of Y. The certainty equivalents can be adapted to assess whether Y is utility independent of Y. For example, consider the previous elicitation in Equation 4.26, for an attribute of survival at one year, described by X, and a second attribute, the chance of experiencing a severe infection within the same time period, described as Y. Establishing whether X is utility independent of Y can be achieved by taking the certainty equivalent relation and establishing whether it is appropriate when $y = 20\%$ and $y = 60\%$.

When X and Y are mutually utility independent we can express the utility function $u(x, y)$ in a multi-linear (bi-linear) form as follows [142]:

$$u(x, y) = k_X u_X(x) + k_Y u_Y(y) + k_{XY} u_X(x) u_Y(y) \tag{4.32}$$

where

1. $u_X(x)$ is a conditional utility function on Y normalised by $u_X(x_0) = 0$ and $u_x(x_*) = 1$.
2. $u_Y(y)$ is a conditional utility function on X normalised by $u_Y(y_0) = 0$ and $u_y(y_*) = 1$.
3. $k_X = u(x_*, y_0)$.
4. $k_Y = u(x_0, y_*)$.
5. $k_{XY} = 1 - k_X - k_Y$.

If mutual utility independence can be demonstrated this greatly simplifies the problem

of deriving joint utility functions. The marginals can be found using the methods in the previous section, leaving only two constants k_X and k_Y to be found. These can be established with a minimum of two elicitations (value, certainty or probability equivalent) of simple lotteries and back substituted to find the two constants. Given the clear importance of the two constants on overall utility it is a number of consistency checks are advised.

The ratio $k_X : k_Y$ can be considered a linear payoff of the two attributes when $k_{XY} = 0$, i.e how much is an incremental increase in one utility worth on the scale of the other utility. The parameter k_{XY} makes the interpretation more challenging in general however. Given the normalising of the conditional utility functions the parameters k_X and k_Y will give the utility at the point where one attribute is maximised and the other minimised (conditions 3 and 4).

Given that both conditional utility functions are normalised, $-1 \leq k_{XY} \leq 1$. The constant k_{XY} comprising of the other two constants is referred to as an interaction component between the two utility functions. The interpretation of $k_{XY} > 0$ is a positive interaction; when u_X gets larger, the effect of u_Y on the overall utility gets larger and vice versa. Alternatively positive interaction would imply that both of the attributes need to have high utility in order to consider the overall utility high. When $k_{XY} < 0$ this is a negative interaction; interpreting a negative interaction in the context of marginal effects is essential to distinguish between a ceiling effect and a qualitative difference. A ceiling effect indicates that the combined impact of u_X and u_Y is less than their additive effects due to a natural limit, while a qualitative difference means that one variable reduces the effect of the other, fundamentally altering the interaction. When $k_{XY} = 0$ the utility independence equation reduces to a simpler additive form.

$$u(x, y) = k_X u_X(x) + k_Y u_Y(y). \tag{4.33}$$

This is described as *additive utility*. In order for this to be true a further condition of additive independence also needs to hold. Using the four points in Figure 4.4, the following preference needs to hold

$$\langle (x_1, y_1), (x_2, y_2) \rangle \sim \langle (x_1, y_2), (x_2, y_1) \rangle, \tag{4.34}$$

for all values of $x = (x_1, x_2)$ and $y = (y_1, y_2)$. This means the overall utility is the sum of its consistent parts. If mutual utility independence cannot be established it is still possible that there is utility independence for one attribute. Alternative formulas can be considered in this case see Keeny and Raiffa [110].

4.6.2 Multivariate value

A multivariate value function maps two attributes to a single scalar (value) which in turn can be transformed into utility. The first part of the section defines the value function to demonstrate how this can be used to give joint utility. The second part establishes a method of indifference curves to obtain a value function.

The univariate definition of a value function given earlier in the chapter simply extends to the bivariate setting by redefining the two points $a = (x', y')$ and $b = (x'', y'')$ in the domain of all possible outcomes $A = X \times Y$

$$\forall a, b \in A, v(a) \geq v(b) \Leftrightarrow a \succcurlyeq b. \tag{4.35}$$

The bivariate value function takes the two attributes as arguments and yields a single scalar that is monotonic (as per definition). The unidimensional value function can be considered as a single "attribute" and used to define the joint utility,

$$u(x, y) = u(v(x, y)). \tag{4.36}$$

The actual scalar quantity that the value function takes will not have a interpretive meaning, only an order. This is because the attributes X and Y are measuring different things. When taking the next step of eliciting the utility, equivalence methods will need to be evaluated using points on the original scale.

Two attributes can be displayed graphically with each attribute corresponding with the horizontal and vertical axes. Indifference curves are a graphical technique requiring a DM to define a curve of indifference within the joint attribute space, $X \times Y$. A family of indifference curves implies an ordering of preference. This technique is not used in the application of dose finding the next chapter. The method is touched upon here as in some simpler cases it is natural to think of rates of payoff between attributes. A more detailed evaluation of

the technique including multiple parametric forms is given by Keeny and Raiffa [110]. The technique is predominately used for decision making under certainty. An example of an application of this method in dose finding is the EffTox method [77] which was reviewed in Chapter 2, (Figure 2.5). Note the step of defining utility was not applied in this method.

4.7 Bivariate utility function summary

A method for assessing a bivariate utility has been proposed in this chapter. The method is briefly summarised as follows. A preliminary decision scope allows the formulation of a research question; this includes specification of actions to decide between, finding suitable attributes capable of measuring the trial objectives and a suitable probability model. Some basic preference structures can initially be inspected, this includes whether the attributes have a monotonic value function. This allows the preliminary assessment whether mutual utility independence holds using elicitation techniques comprising of simple lotteries. If the condition holds the individual utility functions can be created for each attribute. Each attribute can have a number of preliminary relations inspected for simple lotteries. These in turn will imply attitudes to risk over the attribute space. A suitable parametric form based on the preliminary assessments can then be proposed and elicited. The process is repeated for the second attribute. The parameters associated with the multiplicative utility function can be ascertained with a minimum of two further elicitations. Consistency checks are necessary to provide reassurance that the parametric utility is appropriate. It is at this stage that the utility function is suitable for inclusion in a clinical trial protocol. The entire process can be documented. At the analysis stage the decision corresponding with the alternative that maximises the expected utility is the Bayes decision.

4.8 Discussion

This chapter has reviewed the statistical literature for decision making. One of the main findings of this work is that maximising the expected consequence function with an inappropriate scale can lead to poor decision making. Preferences concerning a decision are different when faced with an uncertain situation. A consistent scale, utility, overcomes the difficulties by setting up the problem in terms of preferences for simple lotteries. No design in the dose finding literature has been defined with respect to VNM utility theory in terms

of preferences for lotteries. The chapter has reviewed the relevant statistical literature so that problems incorporating multiple attributes as is the case for dose finding with efficacy and toxicity attributes may be deconstructed and elicited from a DM.

This chapter has focused upon the Bayesian method for decision analysis. Decision analysis not using the Bayesian method is a possibility. The change in statistical paradigm requires a different optimisation procedure to decision making [27]. An example of this procedure has been applied in dose finding described as the likelihood CRM [143]. The likelihood approach is set up similarly to this chapter with a model and a consequence function but the decision is two stage. The parameter(s) describing the states of nature is deterministic within the frequentist paradigm and is estimated according to the maximum likelihood estimate before the parameter estimate is inputted into a consequence function. The dose that ranks the highest is selected for the next cohort. The method needs to have a DLT observed in order to yield a likelihood estimate. The approach is similar to the hybrid Bayesian approach defined in Equation 2.7. In contrast to the Decision theoretic approach uncertainty is first aggregated by finding the mean estimates of the parameters in the model before maximising the consequence function. The approach cannot incorporate different preferences in the presence of uncertainty.

The reason why a frequentist analysis might be argued for is that the introduction of a prior and the use of utilities introduces subjectivity into the analysis. Subjectivity isn't however absent in other procedures. Any decision with two attributes will need some form of subjective assessment regarding preference. Take for example two treatment options, one that is both high toxicity and high efficacy and one that is slightly lower efficacy and lower toxicity. It is not possible to ascertain which is the preferred, even if the attribute levels of the two treatments are known, without some subjectivity. Equation 2.3 gave one possible definition without the need for specifying a parameter. Even in this case there is a subjective choice that the simple payoff adequately describes the situation. Adding further conditions such as a maximum amount of toxicity are also further subjective choices. The method described in this chapter fully acknowledges the inherent subjectivity in any given situation and is based on methods to reduce any bias with transparency in recording choices. This chapter has demonstrated that making decisions in an uncertain situation is different to when the states of nature are known. This uncertainty cannot easily be resolved outside

of a subjective decision theoretic framework.

The next chapter applies the methodology and practical approach of this chapter in the setting of dose finding in oncology when there is both an efficacy and a toxicity attribute.

Chapter 5

Decision theoretic dose finding

5.1 Introduction

The previous chapter reviewed the literature for a Bayesian decision theoretic approach to statistical decision making. The purpose of the approach is to choose an optimal action from a set of possible actions when the outcome is uncertain [27]. There are two main components: a Bayesian model representing the structure of a system and its associated uncertainty; and a consequence or utility [25]. Utility is a numerical measure of consequence that follows an axiomatic basis for rational decision making. A decision theoretic approach to statistical decision making is scientifically sound, providing coherent decisions when each of the two main components can sufficiently be determined.

This chapter proposes a dose finding method for phase I trials that follows the Bayesian decision theoretic approach accounting for uncertainty and is referred to as Reference Dependent Decision Theoretic dose finding (R2DT) from hereon. Note that reference dependence is defined for the first time later in this chapter. The motivation for this work is to propose a utility function that better reflects clinical preferences to avoid reliance upon a two staged approach to decision making. The two stage approach comprised of first restricting the decision space with admissibility rules then maximising an over simplified objective function. A full justification is given in the discussion of Chapter 2.

The rest of this chapter is structured as follows: The decision theoretic approach specific to R2DT is restated. Utility functions based upon attitudes to uncertainty with individual

attributes are then defined. Multivariate utility theory gives a broad form for the utility function with constants to be set. A detailed elicitation protocol is then given in order to elicit the constants as part of the proposed utility function. The R2DT method is then applied to an example in multiple myeloma and evaluated using simulation.

5.2 Decision theoretic dose finding

The general Bayesian method to dose finding was given in Chapter 2. The approach is restated here in this short section with specific detail for the R2DT method.

Let $D = \{d_1 < d_2 < \cdots < d_k\}$ be a set of k pre-defined doses to be studied and $Y = (Y_E, Y_T)$ where

$$Y_E = \begin{cases} 1 & \text{if efficacious} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad Y_T = \begin{cases} 1 & \text{if toxic} \\ 0 & \text{otherwise} \end{cases} \tag{5.1}$$

are Bernoulli random variables representing an efficacy and toxicity event respectively. Features that are unknown about the external world, namely the probability of efficacy, π_E and toxicity at each dose, π_T are modelled by unknown states of nature $\theta \in \Theta$, where θ represents the parameters associated with data generation. The observation Y is drawn from a distribution $p_Y(y|\theta)$. Prior knowledge of $\theta \in \Theta$ is incorporated via a prior $p(\theta)$.

This is updated through Bayes theorem in light of the observation(s), to give the posterior

$$p(\theta|y) \propto p(y|\theta) \times p(\theta). \tag{5.2}$$

The probability model $p(y|\theta)$ for the R2DT method follows independent logistic regression models for efficacy and toxicity and are expanded upon in the next section.

A utility function $u(d, \theta)$ specifies the utility of making the decision to treat at dose $d \in D$ if the state of nature is $\theta \in \Theta$. The potential actions at each stage are to assign a dose d to treat the next cohort. The Bayes action (or decision) $d^* \in D$ is the action that maximises the posterior expected utility:

$$d^*(y) = \arg\max_d (E[u(d, \theta)|y]). \tag{5.3}$$

The trial recruits in cohorts of size c with the posterior formed from data from each cohort. The Bayes decision determines the dose for the next cohort. No skipping of untried doses in escalation is stipulated as an additional safety rule outside of the probability model to account for model misspecification in earlier cohorts [77], an additional rule typical of many phase I designs. Specifically, if the Bayes decision is more than one dose above the highest dose already treated at, the dose for the next cohort will be the dose above the highest dose currently treated at. The trial continues until a maximum sample size is reached with the Bayes decision following the final cohort determining the OD. The two attributes used to make decisions through $u(d, \theta)$ in addition to any ad hoc procedures are π_E and π_T, the probability of an efficacy event and the probability of a toxicity event at each dose.

A novel stopping rule for R2DT as well as the conventional rules to prevent treatment at doses with unacceptable levels of toxicity or efficacy are detailed for the design in Section 5.3.4.

5.2.1 Probability model

The probability model used in R2DT follows previous work in this setting, namely logistic regression for efficacy and toxicity [77]. Following the conclusions of Chapter 3, a bivariate distribution for the probability of any event $Y = (Y_E = a, Y_T = b)$ is defined independently for efficacy, π_E, and toxicity π_T.

$$\pi_{a,b} = (\pi_E)^a(1 - \pi_E)^{1-a}(\pi_T)^b(1 - \pi_T)^{1-b} \tag{5.4}$$

The covariate for a set of numeric doses, $d \in \mathbb{R}_{>0}$, are transformed by centering around the geometric mean.

$$x_j = \log(d_j) - \frac{1}{k}\sum_{r=1}^{k}\log(d_r) \qquad j = 1 \dots k \tag{5.5}$$

Marginal probabilities for efficacy and toxicity at each dose are the attributes used in the utility function, for ease in deriving these quantities an inverse-logit link function is defined for efficacy and toxicity,

$$\pi_{E,j} = \text{logit}^{-1}\{\mu_E + \beta_{E1}x_j + \beta_{E2}x_j^2\} \tag{5.6}$$

$$\pi_{T,j} = \text{logit}^{-1}\{\mu_T + \beta_T x_j\}. \tag{5.7}$$

The additional squared term, β_{E2} in the efficacy model allows the possibility that efficacy may not be monotonic in dose. Model parameters for the design are defined by $\boldsymbol{\theta} = (\mu_E, \beta_{E1}, \beta_{E2}, \mu_T, \beta_T)$ and data by $\mathcal{D}_n = (Y_i, x_i)$. With the prior for $\boldsymbol{\theta}$ following independent normal distributions with corresponding hyper parameters for the mean and variance.

5.2.1.1 EffTox model priors

In Bayesian analysis an uninformative prior can be constructed for a parameter by assigning a large variance representing little information. In logistic regression however uninformative priors on the odds scale can lead to very informative priors on the probability scale [144]. This is particularly a problem in dose finding because decisions are based upon posteriors containing little data. Pathological priors in this setting are described when decision making is dominated by the prior and this leads to unusual decisions such as getting stuck at one dose level or stopping the trial independent of accumulating data [78].

An algorithm for establishing priors using effective sample sizes was proposed to simplify the specification of parametric priors [145]. The motivation for the method is that it is easier to elicit a prior on a familiar scale and then transform into parameter values. Prior means are specified on the probability scale for efficacy and toxicity at each dose. The Effective Sample size (ESS) relates the information contained within the prior to the equivalent information gained from a number of patients and is used to assign the variance component of priors.

The measure of ESS is calculated by simulating data for a range of sample sizes and calculating a posterior from a vague prior. These posteriors are compared to the proposed prior by a measure of similarity between distributions. The effective sample size of the prior is the sample size of the posterior most similar. In a simple setting a $beta(a, b)$ distribution has an $ESS = a + b$. This is argued from a beta binomial conjugate model for a binomial outcome, with number of successes, X out of n trials and success probability π. Assuming a beta prior gives $p(\pi|x, n) \sim beta(a + x, b + n - x)$. The prior may be identified with a similar $beta(c+x, d+m-x)$ posterior arising from a previous $beta(c, d)$ prior having a very small amount of information, with m the the effective sample size.

The algorithm was elaborated upon as follows for the EffTox model [146]. Set the second order parameter term for efficacy to be $\beta_{E,2} \sim N(0, 0.2)$. Note the method also details setting the correlation parameter $\psi \sim N(0, 1)$; but R2DT doesn't use the parameter as independent is assumed. Elicit values for the mean toxicity and efficacy at each dose on the probability scale and then determine mean and variance parameter hyper parameter estimates for a given ESS. The final choice of prior is determined from simulation with good choices of ESS between 0.5 and 1.5. The idea is to give a large enough ESS so that early decisions are guided by sensible prior choices without pathological behaviour, but not too strong that the posteriors are dominated by the prior irrespective of accumulating data. The algorithm was implemented in this thesis using `trialr` software [147].

5.3 R2DT Utility specification

The review of dose finding methods in Chapter 2 demonstrated that objective or utility functions are often used but not elicited carefully to reflect clinical preferences. By setting up utilities which better reflect clinical preferences by applying methods detailed in Chapter 4 it is hoped that the proposed R2DT design will improve upon operating characteristics.

R2DT assumes some conditions to define the utility function as $u(\pi_E, \pi_T) = f(u_E(\pi_E), u_T(\pi_T))$ with $f(\cdot)$ a linear function, u_E a marginal utility function of π_E , and u_T a marginal utility function of π_T. In allowing the parametric joint form more easily assessed marginal utility functions can be considered. The marginal utilities are first defined with attitudes to uncertainty and reference dependence. The two functions are combined in the next section accounting for how the two utilities interact. The role of utility in stopping the trial in light of all doses being overly toxic and/or efficacious is then expanded upon to give an additional novel stopping rule.

5.3.1 Marginal utility functions

5.3.1.1 Attributes

In decision making, objectives are characterised by attributes; in this setting, measures of efficacy and toxicity corresponding with the population that the treatment is intended for are used. These are π_E and π_T, the probabilities of efficacy and toxicity at each dose. The attributes are able to measure an OD at the end of the trial as well as able to guide

decision making for an individual patient. Any utility function for efficacy must be strictly monotonically increasing, since more efficacy is invariably preferred to less. If π_E^* and π_E^{**} are two levels of the efficacy attribute then

$$[\pi_E^* > \pi_E^{**}] \Leftrightarrow [u(\pi_E^*) > u(\pi_E^{**})]. \tag{5.8}$$

Similarly, any toxicity utility function must be a strictly monotonically decreasing function as more toxicity is worse than less. If π_T^* and π_T^{**} are two levels of the toxicity attribute then

$$[\pi_T^* > \pi_T^{**}] \Leftrightarrow [u(\pi_T^*) < u(\pi_T^{**})]. \tag{5.9}$$

When considering whether the ethical objective of not exposing patients to non efficacious and toxic doses, it is not possible to ascertain this from the attributes alone, without further constants. Admissibility criteria are typically used in the setting to define a threshold to give context and split each attribute into regions of acceptability or unacceptability with the intention of meeting the ethical objective. As such any utility function that aims to incorporate the ethical objective must have some context or an additional parameter in addition to the raw attribute. R2DT considers attributes for toxicity and efficacy against a single reference point; $\overline{\pi}_E$ and $\overline{\pi}_T$ for efficacy and toxicity respectively.

In the R2DT method the efficacy reference, $\overline{\pi}_E$, is suggested to correspond with the current efficacy estimates for standard of care rather than an aspirational level of efficacy associated with the continued development of the drug, i.e. a minimum efficacy threshold. The toxicity reference, $\overline{\pi}_T$ is suggested to be thought as a target toxicity level, typically associated with toxicity-only dose finding designs [109].

The attributes for efficacy and toxicity at each dose are first transformed as $\pi_E - \overline{\pi}_E$ and $\overline{\pi}_T - \pi_T$ respectively to be arguments to the marginal R2DT utility functions. Note the transformation of the toxicity attribute to satisfy Equation 5.9 ($[1-\pi_T]-[1-\overline{\pi}_T] = \overline{\pi}_T - \pi_T$). A negative transformed attribute is labeled a "loss" and a positive value a "gain" upon the reference. With the labels "loss" and "gain" reflecting the basic interpretation of a bad and a good level of the attribute respectively. For example an improvement from the reference for efficacy is beneficial for the patient i.e a gain, with $\pi_E - \overline{\pi}_E > 0$.

The merit of an incremental increase in either attribute is considered differently depending

on whether it is considered a "gain" upon the reference or a "loss". Increasing the distance from the reference for a gain is more beneficial for the patient than a smaller distance with the opposite is true for a loss. For example, considering gains, increasing efficacy beyond the reference is increasingly beneficial. While for toxicity reducing toxicity below the reference is similarly increasingly beneficial. The main premise is that the ethical objective of patient benefit for each attribute is considered with respect to departure from either side of a reference point, and this is called reference dependence. Creating a utility function for each attribute with reference dependence allows us to incorporate the ethical objectives directly into determination of the OD rather than with separate admissibility criteria for each attribute.

Attitude to risk describes how a DMs preferences change in the presence of uncertainty. Considering either attribute as reference dependent, it is likely that attitudes to risk may be dependent upon whether considering the level of the attribute as a gain or a loss. Both marginal utility functions for R2DT are defined using a piece-wise function that splits the attribute into gains and losses. Prospect theory, reviewed in the next section, takes a similar approach although uses a descriptive decision model.

5.3.1.2 Prospect theory

A consequence function with reference dependence was first described in prospect theory [148]. The authors, Daniel Kahneman and Amos Tversky proposed the method as an alternative to expected utility theory as a descriptive model of decision making in economics. The method isn't directly applicable to the setting of dose finding, as in Chapter 4, it was stated that only a method that prescribed rationale decisions i.e. a normative theory would be suitable for the dose finding setting. The idea to define a piece-wise function for the marginal utility functions that splits the attribute into gains and losses came from prospect theory. The functional form of each of the piece-wise components, utility independence to join the two marginal distributions and maximising expected utility, all part of R2DT method does not feature in prospect theory. The method is expanded upon further in this section to highlight differences.

The first part of prospect theory is that value is assigned to gains and losses with respect to some reference rather than the attribute. Prospect theory has a second component to

account for how individuals perceive probabilities, the probability weighting function. The theory explains that small probabilities are perceived to be higher than they actually are, with medium to large probabilities perceived to be less. The reason for highlighting the theory is that the value function in prospect theory takes the same shape to what has been proposed for the R2DT method. The primary reason for the similarity is that both methods exhibit reference dependence for a monotonically increasing attribute with broad concepts of "gain" and "loss" coinciding. The reference dependence in prospect theory is seen as a psychological heuristic. Individuals feel the pain of loss more intensively than the equivalent pleasure of a gain and this is accounted for in a descriptive theory.

The development of R2DT utilises some of the language, ("losses", "gains") and broad ideas from prospect theory. There is however a clear distinction between a method that is a description of how people make decisions and a normative theory of R2DT that prescribes rationale dosing choices as introduced in the previous chapter. Prospect theory has been previously applied in the domain of health research [149, 150, 151, 152]. In the domain of evaluating additional life years participants were found to give preferences to risk suggesting reference dependence with attitudes to risk similar to that of prospect theory [150]. The use of parametric models to describe attitude to risk over the entire domain, i.e. without reference dependence was seen to be implausible[149]. Examples were applied as a descriptive theory, but is highlighted here as evidence of reference dependence in the health domain.

5.3.2 R2DT marginal utility functions

This section gives a marginal utility function for efficacy and toxicity which should be more capable of representing clinical preferences than the alternative designs from Chapter 2. A parametric function allows a broad form for the utility function to be specified in different settings without the need to restudy the basis for the function such as reference dependence each time. Specific parameter values in the utility function can be elicited as is expanded upon in Section 5.4.

The attitude to risk for both efficacy and toxicity utility functions is proposed to be different depending on whether considering the attribute a gain or a loss. The power function is specified for each segment of the utility functions as it is a parametric utility function where the absolute risk aversion index is dependent upon the distance from the reference point

(Equation 4.23). The power utility is a commonly used utility function when an attribute is measured relative to a reference [110].

The efficacy utility is defined as

$$u_E(\pi_E) = \begin{cases} g\big((\pi_E - \overline{\pi}_E)^{\alpha_{GE}}\big) & \pi_E \geq \overline{\pi}_E \\ g\big(-\lambda_E|\pi_E - \overline{\pi}_E|^{\alpha_{LE}}\big) & \pi_E < \overline{\pi}_E, \end{cases} \tag{5.10}$$

with $\lambda_E \geq 0$, $\alpha_{GE} \geq 0$, $\alpha_{LE} \geq 0$ and $g(u) = [u - u(0)]/[u(1) - u(0)]$. The normalising function, g, places the utility function in the range $[0, 1]$. Note that utility functions are indifferent to linear transformations. In R2DT, the scaling is necessary to ensure the utility is on the same scale as the toxicity utility function.

The loss aversion index, λ_E, considers the merit of "gains" with respect to "losses" (Figure 5.1(B)). It has been argued that gains and losses are considered differently. Loss neutral, $\lambda_E = 1$, considers gains and losses as equally important. Increasing the loss aversion index so that $\lambda_E > 1$ represent an increasing preference of avoiding losses more so than pursuing gains. In dose finding loss aversion corresponds with the ethical objective of avoiding exposing patients to in-efficacious doses. This component of utility function is similar in form to escalation with overdose control in the setting of toxicity only dose finding (Equation 2.14) [59].

The parameters α_{LE} and α_{GE} specify the attitude to risk for losses and gains respectively; $\alpha_\cdot = 1$ would indicate risk neutrality. It is proposed $\alpha_{GE} < 1$ represents an incremental increase in the attribute, becoming less desirable the further away from the reference; this gives a concave (risk-averse) utility function for the gain segment. It is proposed $\alpha_{LE} < 1$ represents a convex (risk-prone) utility function for losses. This represents a preference that an incremental increase nearer the reference has more impact than one further away. The magnitude of the parameters coincides with admissibility rules, with the extreme when $\alpha_{GE} = 0$ and $\alpha_{LE} = 0$, the utility function becomes a step function that is similar to the admissibility rules.

An example of possible shapes for the utility function is plotted in Figure 5.1(A).

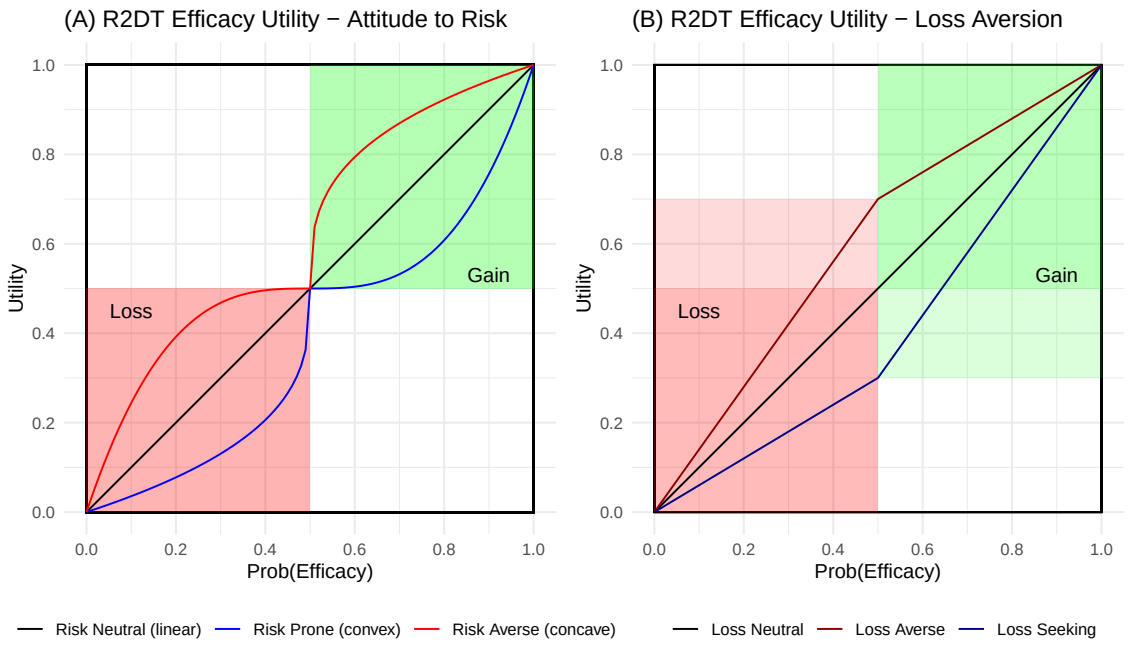

Figure 5.1: Attitudes to risk and loss aversion for efficacy utility function (A): Constantly increasing utility for efficacy attribute (probability of efficacy event) with reference probability of 0.5 defining whether a "Loss" or "Gain". Convex shape is risk prone, concave shape is risk averse. Figure depicts different attitudes to risk depending on the reference. R2DT proposes a sigmoidal utility, convex for "Losses" and concave for "Gains". (B) Loss aversion is depicted with a risk neutral utility for efficacy attribute with reference probability of 0.5 to define whether a "Loss" or "Gain". R2DT proposes loss aversion (stretches loss region) to reflect ethical objective of avoiding exposure to non-efficacious doses

The following utility function is proposed for toxicity.

$$u_T(\pi_T) = \begin{cases} h\big((\overline{\pi}_T - \pi_T)^{\alpha_{GT}}\big) & \pi_T \leq \overline{\pi}_T \\ h\big(-\lambda_T|\overline{\pi}_T - \pi_T|^{\alpha_{LT}}\big) & \pi_T > \overline{\pi}_T \end{cases} \tag{5.11}$$

with $\lambda_T \geq 0$, $\alpha_{GT} \geq 0$, $\alpha_{LT} \geq 0$ and $h(u) = [u - u(1)]/[u(0) - u(1)]$. The normalising function, h, places the utility function in the range $[0, 1]$. It is proposed that $\alpha_{GT} < 1$, $\alpha_{LT} < 1$ and $\lambda_T > 1$ with similar interpretation and attitudes to risk to the efficacy utility function. Due to the initial transformation of the attribute the toxicity utility is proposed to mirror the efficacy utility i.e. an inverted sigmoidal shape, Figure 5.2.

Individual utility functions have been proposed for efficacy and toxicity. These are combined into a joint form in the next section.

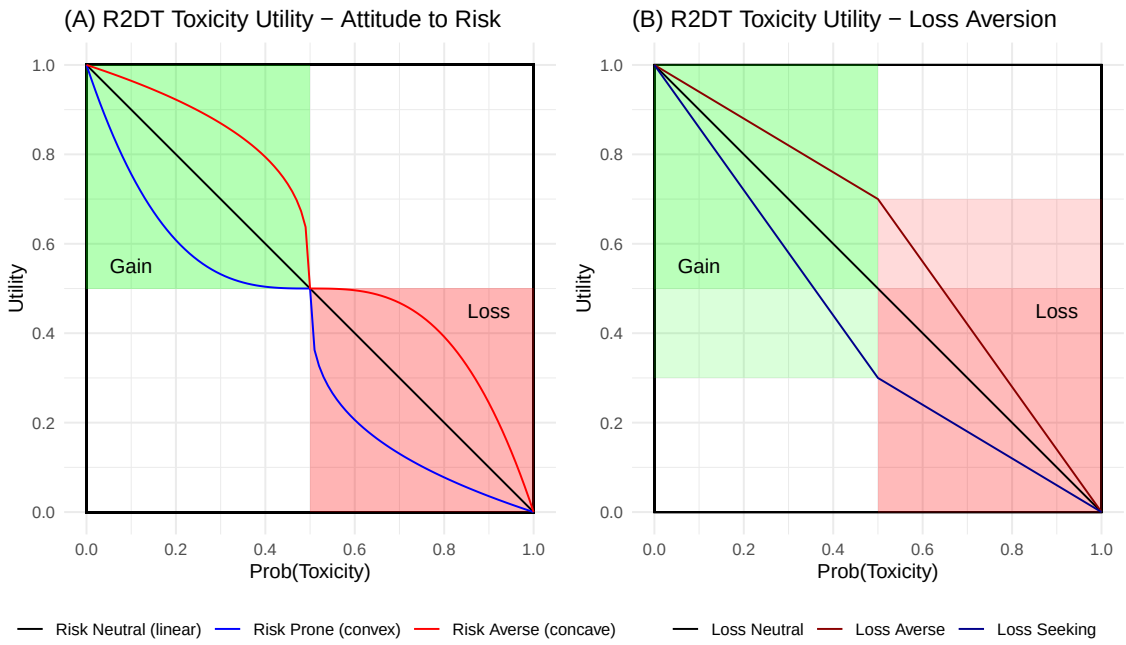

Figure 5.2: Attitudes to risk and loss aversion for toxicity utility function: (A): Constantly decreasing utility for toxicity attribute (probability of toxicity event) with reference probability of 0.5 defining whether a "Loss" or "Gain". Convex is risk prone, Concave is risk averse. Figure depicts different attitudes to risk depending on the reference. R2DT proposes an inverted sigmoidal utility, convex (red) for "Losses" and concave (blue) for "Gains". Loss aversion is considered to be neutral for the purpose of the figure (B) Attitudes to "Losses" are depicted with a risk neutral utility for toxicity attribute with reference probability of 0.5 to define whether a "Loss" or "Gain". R2DT proposes Loss aversion (stretches loss region) to reflect ethical objective of avoiding exposure to overly toxic doses. Attitudes to risk are considered to be neutral in the figure

5.3.3 Joint utility

R2DT assumes a number of conditions to define the utility function in the form $u(\pi_E, \pi_T) = f(u_E(\pi_E), u_T(\pi_T))$ with $f(\cdot)$ a linear function, u_E a marginal utility function of π_E , and u_T a marginal utility function of π_T. These conditions, the functional form and interpretation of additional parameters defined in $f(\cdot)$ are given in this section.

With attributes π_E, efficacy, and π_T toxicity, consider a point (e, t), within the domain of all possible levels $\pi_E \times \pi_T$, such that

$$0 \leq e \leq 1 \quad \text{and} \quad 0 \leq t \leq 1 \tag{5.12}$$

Consider two conditional utility functions $u(e', \cdot)$ and $u(e'', \cdot)$ from two points e' and e''. Defining a lottery from the conditional utility function $u(e', \cdot)$ concerning two points t_1

and t_2 and associated certainty equivalent $\hat{t}$. We then contrast this with the certainty equivalent from the same lottery from the conditional utility function $u(e'', \cdot)$. If the certainty equivalent, $\hat{t}$ does not shift we can say that the two are strategically equivalent. This is expanded upon in section 4.6.1 of previous chapter.

Efficacy is utility independent of toxicity when conditional preferences for lotteries on π_E given π_T do not depend on the particular level of t. When efficacy and toxicity are mutually utility independent we can express the utility function $u(e, t)$ in a multi-linear (bilinear) form as follows, Equation 4.32, [142]:

$$u(\pi_E, \pi_T) = k_E \; u_E(\pi_E) + k_T \; u_T(\pi_T) + k_{ET} \; u_E(\pi_E) \; u_T(\pi_T) \tag{5.13}$$

where

1. $k_E > 0$ and $k_T > 0$, and
2. u_E is a conditional utility function on π_T.
3. u_T is a conditional utility function on π_E.
4. $k_{ET} = 1 - k_E - k_T$.

The conditional utility functions u_E and u_T do not depend on the level of the other attribute, as per the condition of mutual utility independence, as such these are referred to as efficacy and toxicity marginal utility functions for simplicity. The constant k_{ET} represents an interaction between the two attributes. A smaller sum of k_E and k_T would constitute a greater interaction and $k_{ET} = 0$ no interaction.

When combining two measures of consequence through a function to give a single measure of consequence as is the case here it is necessary to have an understanding of what the function is achieving. The key to this is the interaction term, k_{ET}. An interpretation was given in Chapter 4 surrounding the explanation of the utility independence equation (Equation 4.32). Here the interpretation is specific to the attributes of efficacy and toxicity.

The simplest case is the independent case, this is also called additive utility independence. With additive utility independence there is no interaction term ($k_E T = 0$) and the relationship between the two attributes is a simple linear payoff. There is only a single parameter that needs specifying since $k_T = 1 - k_E$. A small incremental increase in efficacy utility is

directly proportional to a increase in toxicity utility (lower toxicity) with the magnitude of the constant dictating how much a small incremental increase in efficacy utility is worth in terms of the same increase toxicity utility. This simple payoff remains constant at all levels of efficacy and toxicity.

A positive interaction is when $k_E + k_T < 1$ and would imply that the higher the efficacy utility, the greater (more positive) the effect of toxicity utility (reduction in toxicity) on overall utility. Similarly, the higher the toxicity utility, the greater (more positive) the effect of efficacy utility on overall utility. The opposite being true of a negative value for the interaction parameter. This description of each possible interpretation for the interaction term is plotted in Figure 5.3. It can be seen for the plot with no interaction that the slope for toxicity with respect to efficacy is a constant at points within the joint domain. For a positive interaction the slope for toxicity with respect to efficacy is initially steep at the left hand end of the contour and reduces moving left to right. This suggests that as toxicity increases the effect of efficacy is reduced. The interpretation is synonymous with the clinical situation described for the motivating example in Chapter 2, with the effect of additional efficacy when there is high toxicity being minimal. A negative interaction describes the opposite to the situation in that the slope gets progressively steeper or the effect of additional efficacy becomes greater with more toxicity.

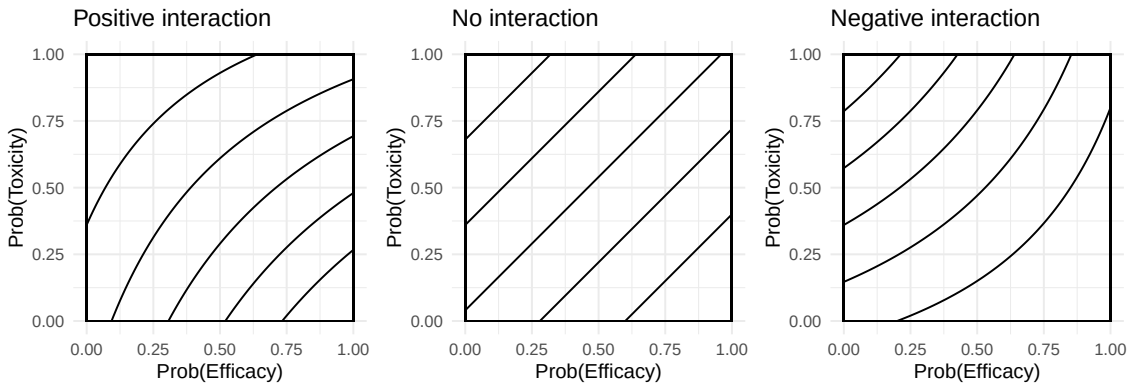

Figure 5.3: Example to visualise the effect of the interaction component of the joint utility function. All plots follow the utility independent relation in 5.13 with simple risk neutral marginal utility functions i.e. $u_E = \pi$ and $u_T = 1 - \pi$. For the Positive interaction plot, $k_E = 0.25$, $k_T = 0.25$ and $k_{ET} = 0.5$. For the no interaction plot, $k_E = 0.5$, $k_T = 0.5$ and $k_{ET} = 0$. For the negative interaction plot, $k_E = 0.75$, $k_T = 0.75$ and $k_{ET} = -0.5$. It can be seen how the slope of contour changes with utility for each of the different interactions.

The Figure represents a simplification of the marginal utility functions, before stating that the positive interaction is the only choice that fits with the clinical situation for this setting a further example is considered. Define two 50-50 lotteries $\langle A = (e_1, t_2), C = (e_2, t_1), \rangle$ and $\langle B = (e_1, t_1), D = (e_2, t_2) \rangle$ with points straddling the reference point, i.e $e_1 < \overline{\pi}_E < e_2$ and $t_1 < \overline{\pi}_T < t_2$. In general for any two utilities following the utility independent equation if $\langle A, C \rangle \sim \langle B, D \rangle$ then this would mean that there is no interaction. $\langle A, C \rangle \succ \langle B, D \rangle$ a positive interaction and $\langle A, C \rangle \prec \langle B, D \rangle$ a negative interaction (Chapter 4).

Interpretation in this example in the context of marginal R2DT utilities for the points is as follows: A is a loss for both efficacy and toxicity, B is loss for efficacy but a gain for toxicity, point C is a gain for both efficacy and toxicity and D is a gain for efficacy but a loss for toxicity. In dose finding it is proposed that the lottery involving point C, two gains, is preferred.

$$\langle A, C \rangle \succ \langle B, D \rangle \Leftrightarrow k_{ET} > 0 \tag{5.14}$$

The interpretation is that both attributes need to be 'good' for the overall utility to be considered likewise. In terms of losses, if one attribute is a loss this is almost as bad as if both attributes are losses - in both cases neither would likely be suitable to treat the wider population. With respect to the reference points the increase in efficacy cannot compensate for high toxicity and given low efficacy this cannot be compensated with low toxicity, this means that there is an interaction and from Equation 5.13 the direction of this interaction implies:

$$k_E + k_T < 1 \tag{5.15}$$

This is thought to be the case in oncology dose finding settings where the payoff becomes more beneficial when both attributes improve. The parameters k_E and k_T and subsequently the interaction term are ascertained by elicitation as detailed later in the chapter. One design that is shown to be analogous to the multiplicative utility function is the EffTox utility design.

5.3.3.1 EffTox utility design

The EffTox utility design [82, 78] introduced in Chapter 2 is used as a comparator to assess the R2DT in the simulation study. This section shows that the EffTox utility design can be formulated as a special case of R2DT that assumes simple risk neutral marginal utility

functions.

The EffTox Utility design specifies a discrete utility function on the four possible individual patient level outcomes, $Y = (Y_E = a, Y_T = b)$, as follows

$$u(Y_E = a, Y_T = b) = \begin{cases} K(1,1), & \text{for } a = 1 \text{ and } b = 1 \\ K(0,0), & \text{for } a = 0 \text{ and } b = 0 \\ K(1,0), & \text{for } a = 1 \text{ and } b = 0 \\ K(0,1), & \text{for } a = 0 \text{ and } b = 1 \end{cases} \tag{5.16}$$

Where $K(a, b)$ are constants to be specified. Given that utility is indifferent to linear transformations, $K(1, 0) = 1$ and $K(0, 1) = 0$ can be specified as the best and worst outcomes respectively. Expected utility is calculated by averaging the utility function over the chance of a state of nature (each patient outcome) happening. For the EffTox utility design the expectation is given by

$$E(u(Y(Y_E = a, Y_T = b)) = \int \sum_{a=0}^{1} \sum_{b=0}^{1} K(a, b)\pi_{a,b}, \tag{5.17}$$

where π_{ab} represent the probability of an event happening. Assuming independence with $\pi_{11} = \pi_E\pi_T$, $\pi_{00} = (1-\pi_E)(1-\pi_T)$, $\pi_{10} = \pi_E(1-\pi_T)$, $\pi_{01} = (1-\pi_E)\pi_T$, and standardising with $K(0, 1) = 0$ and $K(1, 0) = 1$, the expectation equation can be rewritten as a function of π_E and π_T:

$$E(u(Y)) = E(u(\pi_E, \pi_T)) = \int_\theta K(1,1)\pi_E + K(0,0)(1-\pi_T) + (1 - K(0,0) - K(1,1))\pi_E(1-\pi_T)d\theta$$

The expected utility equation can be written as a function of the population level parameters for the probability of an event at each dose. The specific equation has been written in this form as it is analogous to the utility independence equation, Equation 5.13 with $K(1, 1) = k_E$, $K(0, 0) = k_T$, $u_E = \pi_E$ and $u_T = 1-\pi_T$. The marginal utility functions are the identity function or the degenerate case of R2DT, $\lambda_E = \lambda_T = \alpha_{GE} = \alpha_{LE} = \alpha_{GT} = \alpha_{LT} = 1$, with $\overline{\pi}_T$ and $\overline{\pi}_E$ becoming redundant in this special case due to the normalisation function. This demonstrates that the EffTox utility design can be formulated as a special case of R2DT

that assumes simple risk neutral marginal utility functions with interpretation from the perspective of population level parameters.

When the utility independence equation was stated in the last chapter (Equation 4.32) the constants were specified as the utility at the minimum of one attribute and the maximum of the other. The EffTox utility method elicits the numerical consequence of each of the four patient level outcomes. An alternative way of thinking about this in the context of event probabilities as attributes is that after the outcome for a patient has been observed, there is no uncertainty. The chance of the event happening therefore equals zero or one depending on whether it happened or not. This corresponds with $u(\pi_E = 1, \pi_T = 1) = K(1,1) = k_E$ and $u(\pi_E = 0, \pi_T = 0) = K(0,0) = k_T$. In Section 5.4.3 the interpretation of utility is from the perspective of lotteries rather than numerical consequence and this is used to define an alternative elicitation method.

5.3.4 Stopping rule

The intention of R2DT is to move away from reliance on ad hoc admissibility rules. The stopping rules are responsible for stopping the trial when all doses are unsuitable in addition to limiting the decision space at each stage. The initial specification of the design still utilises the admissibility rules, with a novel stopping rule explored as part of the design that is able to capture dependency on both of the interplay between attributes when stopping the trial. For example, more toxicity may be acceptable for a treatment if the efficacy is very good in comparison to treatment with poor efficacy.

The admissibility criteria as is typical in this setting are defined separately in relation to reference cut points $\overline{\pi}_{addE}$ and $\overline{\pi}_{addT}$, below. If either criteria is met the dose will be excluded from the set D.

$$\Pr\{\pi_E(x, \boldsymbol{\theta}) < \overline{\pi}_{addE} | \mathcal{D}\} > 1 - p_E \tag{5.18}$$

$$\Pr\{\pi_T(x, \boldsymbol{\theta}) > \overline{\pi}_{addT} | \mathcal{D}\} > 1 - p_T \tag{5.19}$$

If all doses are excluded from the set D then the trial stops.

The decision theoretic method chooses the decision that maximises the expected utility of a set of actions, with stopping rules an addition to the theory. Within the Bayesian decision theoretic framework, d^* maximises the utility of all potential dosing decisions at the decision

point. In a dose finding trial the potential actions are choosing a dose $d \in D$ or the action s to stop the trial and treat no further patients. The action s could be given a utility scale that corresponds with the actions of choosing a dose d. The specification of a separate utility function would however be needed as the implications of stopping the trial are different to simply choosing a different dose. There may be a need to incorporate additional attributes beyond those specified to limit the choice of doses, such as the cost of setting up the trial and abandoning future development of the treatment [153]. The specification is technically possible but may be overly complex in the context of a single trial to be practically feasible. R2DT builds upon the work in setting up the R2DT utility function proposing a novel stopping rule based upon the utility function. The stopping rule is applied similarly to the admissibility rules in equations 5.18 and 5.19, in that it is an addition and separate to maximising the expected utility to choose a dose.

The admissibility rules for efficacy and toxicity attributes used in the decision making process are applied separately with the dose excluded if either are initiated. Constants for admissibility criteria can be specified by considering both efficacy and toxicity at the same time however [41]. In the cited application of Efftox, the threshold for efficacy (Equation 5.18), $\overline{\pi}_{addE}$, was 5% above the upper bound of where the alternative standard of care treatment was believed to be. The toxicity threshold, (Equation 5.19), $\overline{\pi}_{addT}$, was chosen to be the highest toxicity that would be acceptable based upon the aspirational efficacy level. The utility function of R2DT is specified to more closely meet the joint preferences of toxicity and efficacy in any given setting. The trial stopping decision(s) should be guided by the same preferences between the two attributes that guide dosing decisions. A utility function is also a value function able to rank all combinations of the attribute. By using this property, this allows us to consider an unacceptable contour that accounts for different levels of efficacy and toxicity. In principle the approach is to use a single contour and evidence threshold to define unacceptability in contrast to two separate rules, Figure 5.4. The major difference in this example is to incorporate that the toxicity threshold varies depending upon the amount of efficacy a particular dose has.

The following stopping rule is referred to as the utility admissibility rule: If the probability of the utility being below the equal utility contour surpasses a predefined threshold, $1 - p_u$, the dose is excluded from the set D. If D is an empty set, i.e. no dose satisfies the equation,

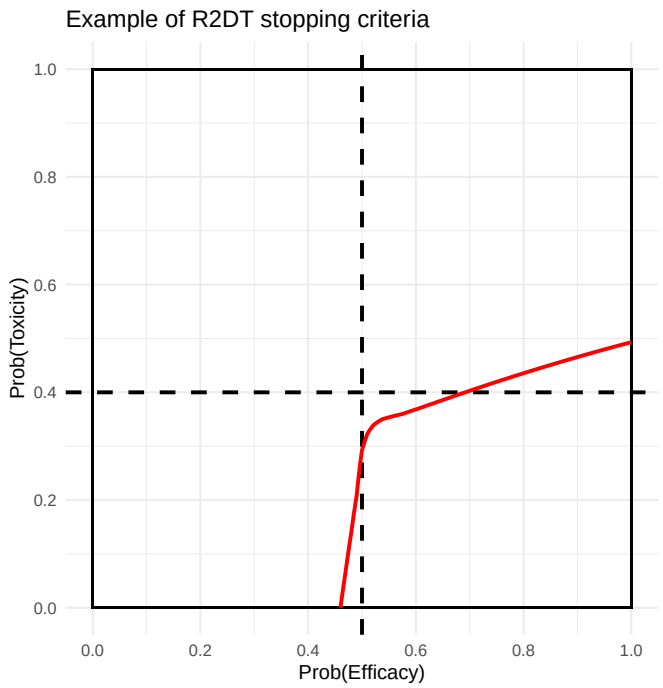

Figure 5.4: Example to visualise the use of utility admissibility criteria. The utility admissibility criteria allows us to define a single contour from the utility function (solid red line) with any dose above deemed unacceptable. This is contrasted with the two black dashed lines representing the conventional admissibility criteria with any dose not in the lower right quadrant unacceptable. The quadrant is formed for efficacy as any dose to the left of the vertical line at $\overline{\pi}_{addE} = 0.5$ and for toxicity any dose above the line at $\overline{\pi}_{addT} = 0.4$ is deemed unacceptable. The contour allows us to accept different levels of an attribute depending on the level of the other attribute. The varying rate of unacceptable toxicity depending on efficacy is most prominent in the example. Note the example coincides with admissibility criteria given in *R2DT (1)* and the utility admissibility rules in*R2DT (3)(ii)* given later in chapter for the simulation study

the trial is stopped:

$$\Pr\{u(\pi_E, \pi_T) < u(\overline{\pi}_{UaddE}, \overline{\pi}_{UaddT})\} > 1 - p_u, \tag{5.20}$$

where $\overline{\pi}_{UaddE}$ and $\overline{\pi}_{UaddT}$ are constants that specify any point on the unacceptable contour. Linking to the overarching trial objectives the contour should be the limit between what would be unacceptable and what would be acceptable for a patient, in terms of efficacy and toxicity, to be treated on trial. Specifying an aspirational level of efficacy may limit exposure of patients to unacceptable doses but the higher bar may have a tendency to stop the trial early based upon little data. One choice of efficacy and toxicity values for the unacceptable contour could be the reference points $\overline{\pi}_{UaddE} = \overline{\pi}_E$ and $\overline{\pi}_{UaddT} = \overline{\pi}_T$ with the intuitive interpretation that anything beyond the contour would constitute a loss. This isn't entirely

correct as the point $\overline{\pi}_T$ is seen as a target toxicity level here, rather than toxicity levels associated with other standard of care treatments. It might be the case, with an alternative treatment that is really well tolerated, that a higher utility threshold is needed. This is considered as part of the elicitation process later in the chapter.

To remain closer to the decision theoretic method an alternative stopping rule is also proposed: the utility trial stopping rule. Here Equation 5.20 is used to stop the trial only, rather than to limit the doses for consideration at each stage. If the criterion is surpassed for all k doses the trial will stop, otherwise the set D contains all k doses at each stage. Differences will only occur between the two proposed decision rules when the dose that is the Bayes optimal action is excluded using the stopping rule while other doses are not. In most instances it is likely when the Bayes optimal action is excluded, all other doses will be excluded and the trial will be stopped.

5.4 R2DT elicitation

This section talks through a number of steps using the methods proposed in Chapter 4 in order to elicit the R2DT utility function. A single DM is responsible for the decision analysis. In practice, this means that there is a single utility function. The DM would likely be the chief investigator of the study but may be a wider team of key opinion leaders. The methods in Section 4.2.2 of the previous chapter described a method to come to a consensus for a group. A DM is referred to as a single entity for simplicity in the rest of this section.

The purpose of the elicitation is to obtain suitable values for parameters in the utility function and stopping rules specified in the last section. This is achieved through a conversation with the DM with a number of precise questions detailed in this section designed to elicit a single parameter for each question. The two attributes are the probability of efficacy, π_E and π_T. Prior to the elicitation exercise it would be helpful to obtain a greater understanding of the attributes and the interpretation of various levels in the setting. This understanding would be gained through an evidence dossier from publications in the setting corresponding with attributes for the group of patients.

The parameters are obtained using simple lotteries. As a brief reminder of the notation and

method given fully in the previous chapter,

$$\langle x_1, \alpha, x_3 \rangle \sim x_2 \tag{5.21}$$

denotes indifference between a simple lottery between two levels x_1 and x_3 of an attribute with mixing component α and the level x_2 with certainty. The purpose of an elicitation is to fix all but one of the components of the lottery and to find the value where indifference is satisfied. Three different methods were described depending upon which component was being elicited. A probability equivalence method elicits α, a value equivalence method elicits one of the attribute levels in the lottery (x_1 or x_3) and a certainty equivalence method elicits x_2.

For consistency, in this section we assume $x_1 < x_2 < x_3$. The question for each elicitation method is to find a suitable value for the elicited component until the DM is indifferent between the two options. For example, to elicit x_2 for the certainty equivalence method, the question may be "what level of efficacy would you be indifferent to receiving with certainty, compared to a 50-50 lottery between a treatment with 40% efficacy and 60% efficacy?".

When choosing attribute levels they should be reasonably close together in a space that is well understood, this is to reduce bias. Additionally assessing lotteries where $\alpha > 0.9$ or $\alpha < 0.1$ is likely to increase the measurement error. Using any one of the elicitation methods above from Equation 5.21, the following relation is a consequence of the Von Neumann–Morgenstern utility axioms in the last chapter and is used to establish parameters in the utility function

$$\alpha u(x_1) + (1 - \alpha)u(x_3) = u(x_2). \tag{5.22}$$

The first step of the elicitation process is to ascertain whether the utility independence axiom is sufficient. The process is described in detail in the previous chapter by considering a number of lotteries to ascertain whether conditional utility functions can be considered equivalent. The preliminary work should consider lotteries at values of toxicity and efficacy that are likely to be seen as feasible in this setting. If this condition holds then we can consider the rest of this section in order to elicit the parameters of the utility function.

5.4.1 Efficacy utility

This section goes through the elicitation of the parameters in the efficacy utility function

$$u_E(\pi_E) = \begin{cases} g\big((\pi_E - \overline{\pi}_E)^{\alpha_{GE}}\big) & \pi_E \geq \overline{\pi}_E \\ g\big(-\lambda_E|\pi_E - \overline{\pi}_E|^{\alpha_{LE}}\big) & \pi_E < \overline{\pi}_E, \end{cases} \tag{5.23}$$

The method proposes a number of individual relations in order to elicit parameters from single lotteries given in Equation 5.22. There are four parameters that need to be ascertained; the reference point $\overline{\pi}_E$, Attitudes to risk α_{GE} and α_{LE} in the gain and loss domains respectively and the loss aversion parameter λ_E. The function g is then ascertained by standardising values already elicited so that $g(u) = [u - u(0)]/[u(1) - u(0)]$. Part of the workup in demonstrating utility independence is that it shouldn't matter what the toxicity level is, when eliciting the conditional efficacy utility. It might still however be useful to keep the conversation solely focused on the efficacy component by explicitly stating at the start that toxicity is acceptable and similar to the current standard of care.

5.4.1.1 Reference point

The first task is to obtain the reference point for efficacy, $\overline{\pi}_E$, in the utility function. This parameter doesn't need a lottery to be established as it is the tipping point of preference where a dose of the new drug is preferred to what is available outside of the trial. This conversation will be informed by evidence from the literature regarding efficacy rates associated with current standard of care treatments.

A suitable question would be "At what efficacy level is the current standard of care?". This could then be followed up with a similar question: "If toxicity profiles were similar between a dose of the intended trial drug and what is available for this group of patients elsewhere, so just considering efficacy, at what level would you be indifferent between the two treatment options".

5.4.1.2 Risk for gains

The next step is to elicit the attitudes to risk, denoted by the parameters α_{GE} and α_{LE}. This is first considered for α_{GE} by considering lotteries restricted to values above the previously elicited reference point, i.e. $x_1 \geq \overline{\pi}_E$. As an example consider the certainty equivalent

method. An equal lottery ($\alpha = 0.5$) could have the lower point, x_1 as the reference point and a reasonable improvement such as $x_3 = \overline{\pi}_E + 20\%$. for example,

$$\langle \overline{\pi}_E, \overline{\pi}_E + 20\% \rangle \sim \underline{x_2}. \tag{5.24}$$

It is expected that x_2 would be less than the midpoint of this lottery to reflect the risk averse attitude. The segment of the utility function when efficacy is higher than $\overline{\pi}_E$ (a gain) is given by

$$u(x) = (x - \overline{\pi}_E)^{\alpha_{GE}}. \tag{5.25}$$

Any constants that are part of the function g will cancel at the next step. Using the utility function, substituting into Equation 5.22 and rearranging gives

$$\alpha_{GE} = \frac{\log(1-\alpha)}{\log(x_2 - \overline{\pi}_E) - \log(x_3 - \overline{\pi}_E)} \tag{5.26}$$

Note that this is only true when $x_1 = \overline{\pi}_E$. For the case $x_1 > \overline{\pi}_E$,

$$\alpha_{GE} = \frac{\log(\alpha) + \log(1-\alpha)}{\log(x_2 - \overline{\pi}_E) - \log(x_1 - \overline{\pi}_E) - \log(x_3 - \overline{\pi}_E)}. \tag{5.27}$$

This is because $u(\overline{\pi}_E) = 0$ from the equation above and the log transformation isn't possible.

5.4.1.3 Risk for losses

When considering the attitude to risk in the loss domain to find the α_{LE}, a further elicitation exclusively in the loss domain is needed. This is similar to the elicitation of α_{GE} and could also incorporate the reference point as the upper point in the lottery, i.e. $x_3 = \overline{\pi}_E$. The method would suggest that the point x_2 would be above the midpoint of the other two lotteries if the mixing component is $\alpha = 0.5$ to denote risk prone behaviour. The utility function for lotteries in the loss domain is given by

$$u(x) = -\lambda_E |x - \overline{\pi}_E|^{\alpha_{LE}}, \tag{5.28}$$

substituting into Equation 5.22 and rearranging gives

$$\alpha_{LE} = \frac{\log(\alpha)}{\log(\overline{\pi}_E - x_2) - \log(\overline{\pi}_E - x_1)}. \tag{5.29}$$

The modulus function has been replaced since $|x - \overline{\pi}_E| = (\overline{\pi}_E - x)$ when $x \leq \overline{\pi}_E$. Again, this is only true when $x_3 = \overline{\pi}_E$. For the case $x_3 < \overline{\pi}_E$ we have

$$\alpha_{LE} = \frac{\log(\alpha) + \log(1-\alpha)}{\log(\overline{\pi}_E - x_2) - \log(\overline{\pi}_E - x_1) - \log(\overline{\pi}_E - x_3)}. \tag{5.30}$$

5.4.1.4 Loss aversion

In order to elicit the loss aversion parameter a further lottery is needed such that $x_1 < \overline{\pi}_E$ and $x_3 > \overline{\pi}_E$. The utility function at x_1 and x_3 is given by Equations 5.28 and 5.25 respectively. There are two possibilities for the utility at x_2 depending on whether the level is considered to be a loss or gain. If $x_2 < \overline{\pi}_E$ (a loss) then

$$-\lambda_E = \frac{(1-\alpha)(x_3 - \overline{\pi}_E)^{\alpha_{GE}}}{(\overline{\pi}_E - x_2)^{\alpha_{LE}} - \alpha(\overline{\pi}_E - x_1)^{\alpha_{LE}}}, \tag{5.31}$$

whereas if $x_2 \geq \overline{\pi}_E$ (a gain) then

$$-\lambda_E = \frac{(x_2 - \overline{\pi}_E)^{\alpha_{GE}} - (1-\alpha)(x_3 - \overline{\pi}_E)^{\alpha_{GE}}}{\alpha(\overline{\pi}_E - x_1)^{\alpha_{LE}}}. \tag{5.32}$$

The parameters α_{GE} and α_{LE} have previously been found for the attitudes to risk in the loss domain and need to be used in order to ascertain the value of λ_E.

5.4.2 Toxicity utility

The marginal utility function for toxicity is given by

$$u_T(\pi_T) = \begin{cases} h\big((\overline{\pi}_T - \pi_T)^{\alpha_{GT}}\big) & \pi_T \leq \overline{\pi}_T \\ h\big(-\lambda_T |\overline{\pi}_T - \pi_T|^{\alpha_{LT}}\big) & \pi_T > \overline{\pi}_T \end{cases} \tag{5.33}$$

The intent is to elicit the reference point $\overline{\pi}_T$, attitudes to risk α_{GT} and α_{LT} in the gain and loss domains respectively and the loss aversion parameter λ_{LT}. The method to obtain these parameters is broadly the same as that specified for the efficacy utility. The notation is changed to refer to a level of the toxicity attribute as y. Simple lotteries are elicited as previously described using one of the three methods and are of the form

$$\langle y_1, \alpha, y_3 \rangle \sim y_2 \tag{5.34}$$

where $y_1 < y_2 < y_3$. The utility function is of the form $u(y_1) > u(y_2) > u(y_3)$. This needs to be considered when trying to find indifference in the elicitation as the attitude to risk is different (Figure 4.3). For example, if the certain outcome y_2 is preferred this would indicate that a larger value of y_2 is needed for equivalence. If the uncertain outcome is preferred then a smaller value of y_2 would be needed to find indifference (this is the opposite to efficacy). When indifference is achieved, equation 5.22 is valid, but now is expressed in terms of the attribute y,

$$\alpha u(y_1) + (1 - \alpha)u(y_3) = u(y_2), \tag{5.35}$$

The reference point $\overline{\pi}_T$ is specified as a target toxicity, corresponding with target toxicity levels that are specified in phase I toxicity-only designs. This is an acceptable level of toxicity and isn't necessarily the point at which the toxicity becomes non-viable or unethical. The level depends on the population under study, treatment options available outside of the trial and specific toxicities associated with treatment under study [109]. Using the concepts of loss and gain that are defined as part of the transformation of the attribute, the reference point is the point at which the attitude to risk changes and the perception of the attribute changes.

Here is a dialogue that could be used to find the reference point. The clinician will likely have a perception as to whether treatment available outside of the trial is considered to be quite toxic, reasonably well tolerated, or have a good toxicity profile. The purpose is to elicit the point at which the clinician starts to consider a treatment as increasingly toxic (but not necessarily unacceptable). This could be achieved by initially agreeing an interval from an initial discussion where the lower bound is a level that the DM is confident constitutes a well tolerated treatment and an upper bound constituting a level where there is uncertainty as to whether the treatment has an acceptable toxicity profile. In this interval there is a point (the reference point) at which the thinking shifts and there is indifference as to whether a particular rate is well tolerated or toxic. The tipping point corresponds with $\overline{\pi}_T$.

The three required parameters to be elicited follow the same rationale as the efficacy utility function. For the parameter α_{GT}, elicit a simple lottery in the gain domain with $y_3 = \overline{\pi}_T$.

Using the utility function $u(y) = (\overline{\pi}_T - y)^{\alpha_{GT}}$ and substituting into equation 5.35 gives,

$$\alpha_{GT} = \frac{\log(\alpha)}{\log(\overline{\pi}_T - y_2) - \log(\overline{\pi}_T - y_1)}. \tag{5.36}$$

For the parameter, α_{LT}, elicit a simple lottery in the loss domain with $y_1 = \overline{\pi}_T$. Using the utility function $u(y) = (y - \overline{\pi}_T)^{\alpha_{LT}}$ yields

$$\alpha_{LT} = \frac{\log(1-\alpha)}{\log(y_2 - \overline{\pi}_T) - \log(y_3 - \overline{\pi}_T)}. \tag{5.37}$$

For the loss aversion parameter λ_T, a lottery is needed such that $y_1 < \overline{\pi}_T$ and $y_3 > \overline{\pi}_T$.

If $y_2 \leq \overline{\pi}_T$, (a gain), then

$$-\lambda_T = \frac{(\overline{\pi}_T - y_2)^{\alpha_{GT}} - \alpha(\overline{\pi}_T - y_1)^{\alpha_{GT}}}{(1-\alpha)(y_3 - \overline{\pi}_T)^{\alpha_{LT}}}. \tag{5.38}$$

whereas if $y_2 \geq \overline{\pi}_T$ (a loss) then

$$-\lambda_T = \frac{\alpha(\overline{\pi}_T - y_1)^{\alpha_{GT}}}{(y_2 - \overline{\pi}_T)^{\alpha_{LT}} - (1-\alpha)(y_3 - \overline{\pi}_T)^{\alpha_{LT}}}. \tag{5.39}$$

5.4.3 Joint utility

The joint utility function for the two attributes of π_E and π_T has the following form:

$$u(\pi_E, \pi_T) = k_E \; u_E(\pi_E) + k_T \; u_T(\pi_T) + (1 - k_E - k_T) \; u_E(\pi_E) \; u_T(\pi_T) \tag{5.40}$$

The conditional utility functions u_E and u_T have been established in the preceding subsections with two parameters k_E and k_T to be determined.

It was shown earlier that the EffTox utility method had an equivalent form as the utility independence formula (Section 5.3.3.1). The method of obtaining the two parameters for utility Efftox is to ask the DM a question based upon the four patient outcomes. This involved the DM trying to ascertain a value as a measure of weighted preference between patient outcomes. An alternative formulation based upon uncertainty and utility could be

to elicit two probability equivalences

$$\begin{aligned}\langle(\pi_E = 0, \pi_T = 1), \underline{\alpha_1}, (\pi_E = 1, \pi_T = 0)\rangle &\sim (\pi_E = 1, \pi_T = 1)\\ \langle(\pi_E = 0, \pi_T = 1), \underline{\alpha_2}, (\pi_E = 1, \pi_T = 0)\rangle &\sim (\pi_E = 0, \pi_T = 0),\end{aligned} \tag{5.41}$$

to obtain $k_E = 1 - \alpha_1$ and $k_T = 1 - \alpha_2$. This elicitation however involves lotteries that are at extreme values that in nearly all situations will be hypothetical. An alternative method and one that is recommended here would be to to find points of indifference. If the decision maker is indifferent between any two points $(\pi_E = x_1, \pi_T = y_1)$ and $(\pi_E = x_2, \pi_T = y_2)$, the utility of both points must also be equal so that:

$$u(x_1, y_1) = u(x_2, y_2) \tag{5.42}$$

where

$$\begin{aligned}u(x_1, y_1) &= k_E\ u_E(x_1) + k_T\ u_T(y_1) + (1 - k_E - k_T)\ u_E(x_1)\ u_T(y_1)\\ u(x_2, y_2) &= k_E\ u_E(x_2) + k_T\ u_T(y_2) + (1 - k_E - k_T)\ u_E(x_2)\ u_T(y_2)\end{aligned} \tag{5.43}$$

A further equivalence relation would yield a second equation which could be solved simultaneously to obtain the two unknown parameters. The simplest way to do this is to rearrange the equation above so that it is of a simple linear form and then solve simultaneously,

$$Ak_E + Bk_T = C \tag{5.44}$$

where

$$\begin{aligned}A &= u_E(x_1) - u_E(x_1)u_T(y_1) - u_E(x_2) + u_E(x_2)u_T(y_2)\\ B &= u_T(y_1) - u_E(x_1)u_T(y_1) - u_T(y_2) + u_E(x_2)u_T(y_2)\\ C &= u_E(x_2)u_T(y_2) - u_E(x_1)u_T(y_1).\end{aligned} \tag{5.45}$$

The joint reference point $(\overline{\pi}_E, \overline{\pi}_T)$ could be a suitable point to establish equivalence. The DM can be asked what would constitute a substantial improvement in efficacy over the reference for the group of patients under study. The next stage would be to offset this improvement in efficacy against toxicity so that the DM is indifferent between the efficacious but toxic option and the reference. This would yield two points of equal utility for the first linear

equation. The second equation could be gained by considering the joint reference point again but this time to enquire about a reduction in toxicity and to offset this against a reduction in efficacy until there is indifference.

5.4.4 Stopping rule

The R2DT stopping rule involves specifying a threshold contour beyond which any treatment with a lower utility would be considered unacceptable. Given that the full elicitation of the parametric utility has already happened all that is needed for this elicitation is a single point. The single point (x, y) corresponds with constants $\overline{\pi}_{UaddE} = x$ and $\overline{\pi}_{UaddT} = y$ given in equation 5.20. A simple way of doing this is to consider an efficacy level that is seen as both feasible and constitutes a significant step in improving outcomes for patients. The question is then what is the maximum amount of toxicity that would be considered acceptable for this level of efficacy? I.e. beyond this, toxicity would be considered unacceptable. The value of utility at this point constitutes the utility threshold.

5.4.5 Consistency checks

The elicitation methods described in this section give the minimum number of simple lotteries or points of indifference that need to be elicited in order to obtain the parameters of the utility function. At this stage a complete utility function and trial stopping rule has been specified. The utility function by definition of a continuous attribute implies an infinite number of other possible simple lotteries within the joint attribute space; some of these should be checked to ensure consistency. It was suggested throughout the elicitation that the magnitude of the difference between x_1 and x_3 in any lottery be kept quite small in order to reduce elicitation bias. Using the specified parametric form for each marginal utility allowed extrapolation beyond the range of the elicitation lottery. The implied lotteries from the utility function should be considered, i.e when $x < x_1$ and when $x > x_3$ to see if the parametric form is acceptable in regions beyond what was initially elicited.

The utility function was used in order to define the contour for utility stopping rule. This is a particular feature that should be used for consistency checks. If the contour doesn't describe the bounds for acceptability this would imply that the utility function needs adjusting. This process is described by an example. The joint utility space can be considered by four regions defined for gains and losses for efficacy and toxicity. The issue identified in this example is

when there is a gain in efficacy, and toxicity is excessive or a loss. The contour lines in this region are too steep, and do not give a sensible contour for the stopping rules, as per the motivating example given in Chapter 2.

The problem with contours too steep in one region could imply a number of small adjustments to aspects of the utility (or combinations of) to make the contour a better approximation to the situation. One parameter that may need adjusting is the interaction component, $k_{ET} = (1 - k_E - k_T)$, of the joint utility. The example would imply that it could be insufficiently large; i.e. the low toxicity utility doesn't interact sufficiently to counteract the large efficacy utility. The adjustments of any parameter require revisiting the corresponding initial elicitation to see if a shift in the elicited component of the lottery is acceptable. A further possibility is the toxicity utility function needs adjusting. In this region the loss aversion parameter, λ_T and the attitudes to risk α_{LT} define the marginal toxicity function for a loss. The example would imply that both parameters are too small. The attitude to risk for efficacy, α_{GE}, similarly may be too small. Adjusting any part of the utility function will have an effect on other regions of the utility function and these should have further checks.

If consistency checks suggest some aspect of the specified parametric form of the utility function doesn't hold then this may need changing. This should be possible to adapt as per the general methods given in Chapter 4. While R2DT is proposed as a closer resemblance to the general situation of dose finding in oncology, it is still a simplification. Consideration should be given as to whether the utility is a close enough approximation to this general situation; the purpose of the utility is to facilitate effective decision making rather than to obtain a utility that perfectly captures the situation.

5.5 R2DT elicitation pilot study

The elicitation protocol prospectively specified in the previous section, to obtain the parameters of R2DT, has not previously been conducted in this setting. A novel and untried aspect of this elicitation is that questions are framed around lotteries for uncertain outcomes. A further consideration is that the proposed elicitation protocol has more questions than the motivating example given in Section 2.5. This section reports on a real elicitation exercise or pilot study to assess the feasibility of the proposed method. The key metrics to assess

in the pilot study are how easy the questions are to understand and the estimated time required for a complete elicitation. Following the study, several practical improvements for the elicitation design are proposed to enhance the design before its potential full real-world implementation.

The minimum number of questions that need asking to obtain all the parameters of R2DT is eight; three for each marginal utility function and a further 2 questions for the parameters of Equation 5.13. This does not include the reference points or any subsequent consistency checks. This is five more questions than the EffTox method used in the motivating example and six more than the patient utility design detailed in Section 5.3.3.1. It was reported with the motivating example that neither of these methods were capable of capturing the clinical preferences. Whether the increased effort is justified, in an attempt to capture something closer reflecting clinical judgment, links to any benefit in operating characteristics, reported in Section 5.6.

To achieve the two aims of the pilot study only the elicitation of the efficacy marginal utility function was planned to be captured. The questions to obtain the toxicity utility function are of an identical structure, only relating to a different endpoint. The questions to capture the joint parameters are similar to previous work in this setting. As such, this component can already be concluded to be feasible; this is expanded upon in the discussion of this chapter. An estimate of the total length of time to conduct the R2DT elicitation in its entirety can also be obtained from the proposed study. Assuming that the elicitation of the toxicity marginal utility function takes a similar amount of time.

To assess the feasibility of the R2DT elicitation protocol, the motivating example was adapted. Dr. Chris Parish, a consultant haematologist and associate professor at the University of Leeds Clinical Trials Research Unit, volunteered to attend a meeting to conduct the elicitation. The feasibility study is illustrated using a hypothetical example from primary double-refractory multiple myeloma, based on the motivating example presented in Chapter 2. In this context, the toxicity endpoint was defined as a binary indicator of whether a dose-limiting toxicity (DLT) occurs within the first two four-week cycles. Efficacy is measured by a binary variable indicating whether the patient achieves a "very good partial response" within the same time frame. Further details of the study, such as an intervention and number of doses are not needed for the elicitation.

5.5.1 Pilot study report

Study set up

The endpoints and population for the application study was agreed with the clinician prior to the meeting. The clinician was also asked to consider what response rate, defined as the percentage of patients achieving the efficacy endpoint, could be achieved for this patient group. The meeting was planned to be in-person and conducted over an hour period. Dr. Duncan Wilson, PhD supervisor, acted as an observer for the meeting. The actual elicitation of any equal certainty equivalent, $x_2 \sim \langle x_1, x_3 \rangle$, was conducted on a white board with a visual representation as given in Figure 5.5 drawn out and specific values wiped out each time.

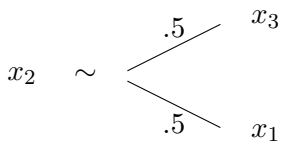

Figure 5.5: Pictorial representation of the lottery $\langle x_1, x_3 \rangle \sim x_2$

A short slide presentation was used to provide background on the purpose of the exercise. It included the objectives of a dose-finding study in oncology and the study design detailed in Section 2.1.1. The presentation covered endpoints and patient population, and explained how an optimal dose for these endpoints could be defined as a payoff between the two. It also showed how a contour plot was used under certainty to decide between doses. If the elicitation was to be used in an actual trial, the introduction will likely be known from preliminary discussions. A simple example of whether or not to play a lottery with a very small chance of a winning a large amount was used to explain the idea of maximising expected consequence, which was defined using money in the example. In a dose finding setting, it was pointed out that a measure of consequence combining effects of efficacy and toxicity was not easily defined. The primary purpose of the elicitation was to break this down into a simpler set of questions to measure the consequence and aid in decision making. The following lottery incorporating money was given to the clinician,

$$\langle -\$5000, \$5000 \rangle \underline{R} \langle -\$10, \$10 \rangle. \tag{5.46}$$

Given that both had the same expected return, the clinician was asked if there was a

preference. The clinician agreed that there was a preference for the lottery incorporating winning or losing $10. This was used to demonstrate that there wasn't an objectively correct way of assessing the answer to such a question and that choices made under uncertainty, similar to those made in dose finding, needed to account for uncertainty. The definition of a certainty equivalent was stated and that this was going to be the basis for all of the questions today. The following practice example was obtained from the clinician.

$$\langle -\$5000, \$5000 \rangle \sim \underline{-\$100}. \tag{5.47}$$

5.5.1.1 Elicitation exercise

Reference point

The current standard of care was estimated to be between 30%-40% for this group of patients by the clinician. The question was asked that if the new drug had an identical toxicity profile to the standard of care and had a known 40% response rate, would you choose the standard of care or the new intervention? The response of the "new intervention" suggested that the response rate for the standard of care was too high. Following some discussion a response rate of 30% was concluded to be reflective of the standard of care for this group of patients and designated as the reference point, $\overline{\pi}_E$.

A gain upon the reference

The clinician was asked to give a large but plausible improvement over the reference point to inform levels of the equal lottery $x_1 = 30\%$ and $x_3 = 60\%$. The clinician was asked whether they would be happy giving a certainty equivalent, or whether an initial preference comparison was preferred. The expected consequence of the lottery was chosen as an initial preference comparison,

$$\langle 30\%, 60\% \rangle \ \underline{R} \ 45\%. \tag{5.48}$$

A preference for the certain 45% was given, suggesting that the clinician was risk averse and the certainty equivalent is less than 45%. A further preference comparisons was used with $x_2 = 40\%$, again with the certain prospect chosen to give

$$\langle 30\%, 60\% \rangle \sim \underline{37.5\%}, \tag{5.49}$$

as a certainty equivalent.

A loss upon the reference

A question of a large and plausible reduction in efficacy was asked, giving $x_1 = 20\%$ and $x_3 = 30\%$ as components of the lottery. More familiar with the process, the clinician declared that they would be risk averse and would need something quite close to 30% to justify the lottery. The following relation,

$$\langle 20\%, 30\% \rangle \sim \underline{28\%}, \tag{5.50}$$

was settled upon. At this point it was pointed out that the clinician was actually conveying a risk seeking attitude rather than risk aversion. This sat uncomfortably, the clinician felt that perhaps they had initially made an error in judgement and the certainty equivalent was revised to

$$\langle 20\%, 30\% \rangle \sim \underline{24\%}. \tag{5.51}$$

A mixed lottery A lottery incorporating the extremes of the previous two lotteries was used as a mixed lottery that would be used to obtain the loss aversion parameter.

$$\langle 20\%, 60\% \rangle \sim \underline{45\%}. \tag{5.52}$$

Consistency checks Three further lotteries were given to the clinician as consistency checks. The gain lottery made $x_3 = 50\%$, the loss lottery $x_1 = 10\%$ and the mixed lottery incorporated the two new values $x_1 = 10\%$ and $x_3 = 50\%$. The clinician was consistent with attitudes to risk for these lotteries.

5.5.1.2 Feedback

The following feedback was provided by the clinician as answers to the set of questions following the elicitation:

Did you understand the questions asked and how confident were you in answers given? I was confident I understood what was being asked and reasonably confident of the answers given. Considering potential losses is more mentally taxing, as everyday practice typically focuses on improving patient outcomes.

How variable do you think the values would be if the elicited from other clinicians? There would be some variability among different clinicians. Some might be more willing to gamble on potentially better treatments, while others would avoid treatments potentially worse than the standard of care. This variability is less likely when considering treatments that might be worse than the standard of care, as avoiding harm is a common principle in clinical practice.

If repeating the exercise with a group of clinicians and patients do you think you could arrive at a consensus if there was a difference in opinion? Including a diverse group of carefully selected clinicians it would be feasible to reach a consensus. Involving patients in the elicitation process could be valuable, though the complexity might be challenging for some. Adjusting the framing of questions and more training could make the process more accessible for patients.

Thoughts about answering a similar exercise for the toxicity endpoint? Repeating the exercise for dose-limiting toxicities is definitely feasible and is potentially easier to think about.

5.5.2 Pilot study discussion

The internal consistency and feedback from the clinician were very positive. Suggesting that the pilot study was a success with the conclusion that the method of elicitation proposed in this thesis is feasible. The initial slides, practice, elicitation and consistency checks was achieved in an hour meeting. Extrapolating would suggest that the elicitation could be achieved in two hours. This is a more involved process than what has been proposed elsewhere in the setting. There are a number of small enhancements to the elicitation protocol that are detailed in subsequent paragraphs that were learnt through the exercise.

Training is important, the explanations were well understood, but the practice exercise was easier to articulate than the lotteries in the elicitation. Speed and understanding of what was happening seemed to improve as the meeting progressed. A further practice example in the health domain would be useful before commencing. The pilot study only included a single clinician, it would be useful to conduct elicitations with more than one clinician and to conduct these separately. This would allow a better understanding of variability and provide confidence that the utility function was reflective of the wider community. The staged approach of preference relations in order to obtain a certainty equivalent was helpful

in moving towards a point of indifference. The initial mixed lottery was too spread out, with $x_1 = 20\%$ and $x_3 = 60\%$, it would be preferable to consider different values to make these closer together. Getting an idea of the reference from the literature prior to the meeting would enable a comprehensive plan of lotteries. The values in the lottery were created in the meeting in response to answers from the clinician. This was challenging, if feasible it may be preferable to consult a clinician not involved in the elicitation so that these can be pre-planned.

Elicitation in the loss domain was more challenging than mixed lotteries in the gain domain. This was evidenced by the change in attitude to risk in the elicitation. The assessment involved making judgements involving something that the clinician wouldn't be comfortable making in practice, i.e. settling for a guaranteed "loss" upon the reference. The lower option in the initial lottery was 10% lower than the reference. Considering the initial certainty equivalent of 28% and the latter value of 24%, in the context of the initial variability in the reference rate, it would suggest that both values are close to clinical equivalence. A lottery more spread apart would have been preferable. The idea of acting in a "Risk seeking" manner, when pointed out, sat uncomfortably with the clinician and this seemed to have an influence that the clinician perhaps feeling they "should" give a value below the reference, to be acting in an appropriate risk averse manner. Upon reflection, the definitions are not helpful in the elicitation. From revisiting what was said, it is plausible that the clinician was framing the lottery with respect to "losses" rather than the percentage of patients achieving the efficacy endpoint [154]. The clinician appeared to be seeking to minimise losses, by giving a value closer to 30%, rather than accepting the consequences of a lottery. If the lottery in Equation 5.52 changed so that the attribute measured a relative "loss" i.e. $\overline{\pi}_E - \pi_E$. So that $x_3 = 30\%$ becomes 0%, $x_1 = 20\%$ becomes 10% and $x_2 = 24\%$ becomes 6% then the clinician is conveying a risk seeking attitude with respect to the loss attribute as follows,

$$\langle 10\%, 0 \rangle \sim \underline{6\%}. \tag{5.53}$$

Given the increased difficulty in the loss domain there are a number of recommendations for future work in eliciting the R2DT. Use a practice example involving losses. The statistician should also avoid introducing terminology associated with attitudes to risk. The statistician should also to be more prepared to explain the potential for different attitudes to risk,

particularly in the loss domain should these discussion arise, without casting an ethical judgement. Changing the order of the session by initially eliciting two mixed lotteries and then conducting the lottery in the loss domain would allow for greater consistency checking for the attitude to risk in the loss domain.

5.6 Simulation

The concept of using simulation to evaluate how well a proposed design works was introduced in Chapter 2. The merits of the R2DT design are explored utilising simulation with comparison against the efficacy toxicity utility design *EffToxU*, which was described earlier in the chapter. The main hypothesis is to see if the R2DT design offers an improvement over the more established designs and to confirm that having a consequence function that closer resembles preferences in practice is beneficial. Secondary questions are to assess the performance of the R2DT stopping rule.

The designs are applied to a fictitious example in primary double-refractory multiple myeloma reflecting the motivating example given in Chapter 2. The toxicity endpoint in this setting is a binary indicator of whether a DLT is experienced in the first two, four-week cycles. Efficacy will be a binary variable as to whether the patient achieved a "very good partial response" within the same time period. The trial will investigate four doses of an investigational medicinal product with units mg/kg, $D = (20, 30, 40, 50)$.

Fixed probability vectors $\tilde{\pi}_E(D)$ and $\tilde{\pi}_T(D)$ are specified for 10 clinically plausible scenarios (Figures 5.6 and 5.7). The scenarios have been specified to understand the merits of the R2DT as well as potential short comings, rather than being a favourable set of scenarios where the R2DT design is likely to exclusively excel compared to other designs. Simulated data comprising of efficacy and toxicity outcomes is generated for each scenario for all patients at each dose for 2000 repeated trials. Outcomes for dummy patients are drawn according to $Y_E \sim B(\tilde{\pi}_E(d_j))$ and $Y_T \sim B(\tilde{\pi}_T(d_j))$ where B is a Bernoulli distribution. Different trial designs are applied to the simulated data with the performance of designs assessed by operating characteristics, defined by the percentage of selection across the 2000 replicates and the average number of patients treated at each dose. An efficient simulation technique was utilised avoiding repetition in model fitting, and is described fully in Appendix A.

The different trial designs in terms of the decision functions are described in subsequent text and listed in Table 5.3. Utility contour plots for each design are plotted with stopping rules in Figure 5.9. The trial will start at the 20mg/kg dose (the first dose level). Successive cohorts of size $c = 3$ will be recruited to the trial until a pre-defined maximum sample size of 45 is achieved or the trial stopped early. The impact of sample size on the R2DT design is explored as part of the simulation study. The patient group is expected to have a 50% response rate if treated outside of the trial with the standard of care established agent. The target toxicity rate of 35% has been established.

The same probability model and priors have been specified for each of the different designs to better understand the decision element. Efficacy and toxicity are modelled independently as specified in Section 5.2.1.1. All of the different designs have prior hyper-parameters as defined in Table 5.2. Priors have been specified according to a mean vector at each dose and equivalent sample size (ESS), as per Section 5.2.1.1. The mean vector was chosen as the mean of the first six scenarios; a range of ESS values were explored for the *EffTox* design utilising *EffTox* software [155]. The chosen ESS was optimal from a range of 0.5 : 1.5 in increments of 0.1 giving acceptable operating characteristics across all 10 scenarios.

5.6.1 R2DT simulation

The R2DT method, labelled *R2DT (1)* is used to make decisions at each stage. This is specified using the marginal efficacy function, marginal toxicity function and joint utility function..

The marginal toxicity utility is determined by Equation 5.10 with parameters as specified in Table 5.1 and plotted in Figure 5.8A. The marginal toxicity utility is determined by Equation 5.11 with parameters as specified in Table 5.1 and plotted in Figure 5.8B. The joint utility combines the two marginal utility functions following Equation 5.13, with constants specified in Table 5.1. The joint utility function plotted in Figure 5.8C.

Admissibility rules for efficacy and toxicity are applied as per Equations 5.18 and 5.19, with $\overline{\pi}_{addE} = 0.5$, $p_E = 0.075$, $\overline{\pi}_{addT} = 0.4$ and $p_T = 0.075$.

Table 5.1: Short description of each of the different constants and interpretation specified in *R2DT (1)*

Constant	Description
	Efficacy utility function 5.10
$\overline{\pi}_E = 0.5$	Reference point where attitude to risk changes. $\pi_E \leq 0.5$ is described as a loss and $\pi_E > 0.5$ a gain
$\lambda_E = 2$	Loss aversion parameter, specified so that losses are twice as impactful as gains
$\alpha_{GE} = 0.7$	risk averse attitude to risk above the reference point (a gain)
$\alpha_{LE} = 0.7$	risk seeking attitude to risk below the reference point (a loss)
	Toxicity utility function 5.11
$\overline{\pi}_T = 0.35$	Reference point where attitude to risk changes. $\pi_T \geq 0.35$ is described as a loss and $\pi_T < 0.35$ a gain
$\lambda_T = 2$	Loss aversion parameter, specified so that losses are twice as impactful as gains
$\alpha_{GT} = 0.7$	risk averse attitude to risk below the reference point (a gain)
$\alpha_{LT} = 0.7$	risk seeking attitude to risk above the reference point (a loss)
	Joint utility function 5.13
$k_E = 0.25$	utility when $\pi_E = 1$ and $\pi_T = 1$
$k_T = 0.15$	utility when $\pi_E = 0$ and $\pi_T = 0$
$(1 - k_E - k_T) = 0.6$	positive interaction between marginal utility functions

The comparison is the *EffTox* patient outcome utility design described earlier, labelled *EffToxU (2)*, with $K(1,1) = 0.25$ and $K(0,0) = 0.15$ (Figure 5.8D). Note from earlier this a degenerate case of the R2DT design with $K(1,1) = k_E = 0.25$, $K(0,0) = k_T = 0.15$ and $u_E = \pi_E$ and $u_T = 1 - \pi_T$ i.e $\lambda_E = \lambda_T = \alpha_{GE} = \alpha_{LE} = \alpha_{GT} = \alpha_{LT} = 1$, with $\overline{\pi}_T$ and $\overline{\pi}_E$ becoming redundant in this special case due to the normalisation function. Stopping rules are applied as per *R2DT (1)*.

Table 5.3 summarises each of the different methods or decision functions including the stopping rules. Contour plots for *R2DT (1)* and *EffToxU (2)* are plotted in Figure 5.8, all other decision functions are plotted in Figure 5.9 and described in the proceeding paragraphs.

The effect of the R2DT stopping rules is explored by specifying the same utility function as *R2DT (1)* and *EffToxU (2)* and adapting the stopping rule. *R2DT (3)* and *EffToxU (5)* apply the utility admissibility rule. That is, at each decision point doses are excluded from choosing the maximum utility if there is insufficient evidence that a dose has acceptable levels of combined efficacy and toxicity. *R2DT (4)* applies the utility trial stopping rule. This maximises the expected utility for all doses at each decision point with the trial stopping

if all doses fail to satisfy the stopping rules given for *R2DT (3)* and *EffToxU (5)*. The threshold is $p_u = 0.1$ for all designs according to Equation 5.20. For *EffToxU (5)* using the specified utility function, $u(0.5, 0.35) = 0.42$. The value of 0.42 has been used to define the acceptable contour. A number of contours are explored for R2DT since the stopping rule is not directly comparable with the admissibility rules as part of the main comparison. Three acceptability contours that have been specified for the R2DT design:

(i) $u(u_E(\overline{\pi}_E), u(u_T(\overline{\pi}_T)) = u(0.5, 0.35) = 0.58$ this has the additional label of *(i)*.

(ii) $u(0.7, 0.4) = 0.62$ this has the additional label of *(ii)*.

(iii) $u(u_E(\overline{\pi}_E), u(u_T(\overline{\pi}_T)) = u(0.9, 0.4) = 0.69$ this has the additional label of *(iii)*.

The unacceptable contours for the admissibility rules are plotted in Figure 5.9. The admissibility rules accept lower utility for *R2DT (3i)* and *R2DT (3ii)* in comparison to the contour of *R2DT (3iii)*, which declares higher utility unacceptable. Linking these differences to scenarios, some doses will be designated as acceptable according to one stopping rule while another stopping rule may say they are not acceptable.

There are two sensitivity analyses applied to the comparative *EffToxU* method to demonstrate that conclusions are not just the result of a poorly specified comparator. The design with $K(1, 1) = k_E = 0.5$ and $K(0, 0) = k_T = 0.3$ has been stated as suitable in many settings [78] and is specified in *EffToxU (7)*. The ratio of of $(k_E : k_T)$ is the same as *EffToxU (2)* but the magnitude of $k_E + k_T = 0.8$ is increased suggesting a smaller interaction for k_{ET} from Equation 5.13. *EffTox (6)* applies the method of trade off contours [77], with specification of the design contour corresponding with the *EffToxU (2)* equal utility contour passing through the reference point defined in *R2DT (1)* and the points on the contour which have no toxicity and perfect efficacy.

Subsequent pages contain all of the tables and figures associated with the set up of the simulation study.

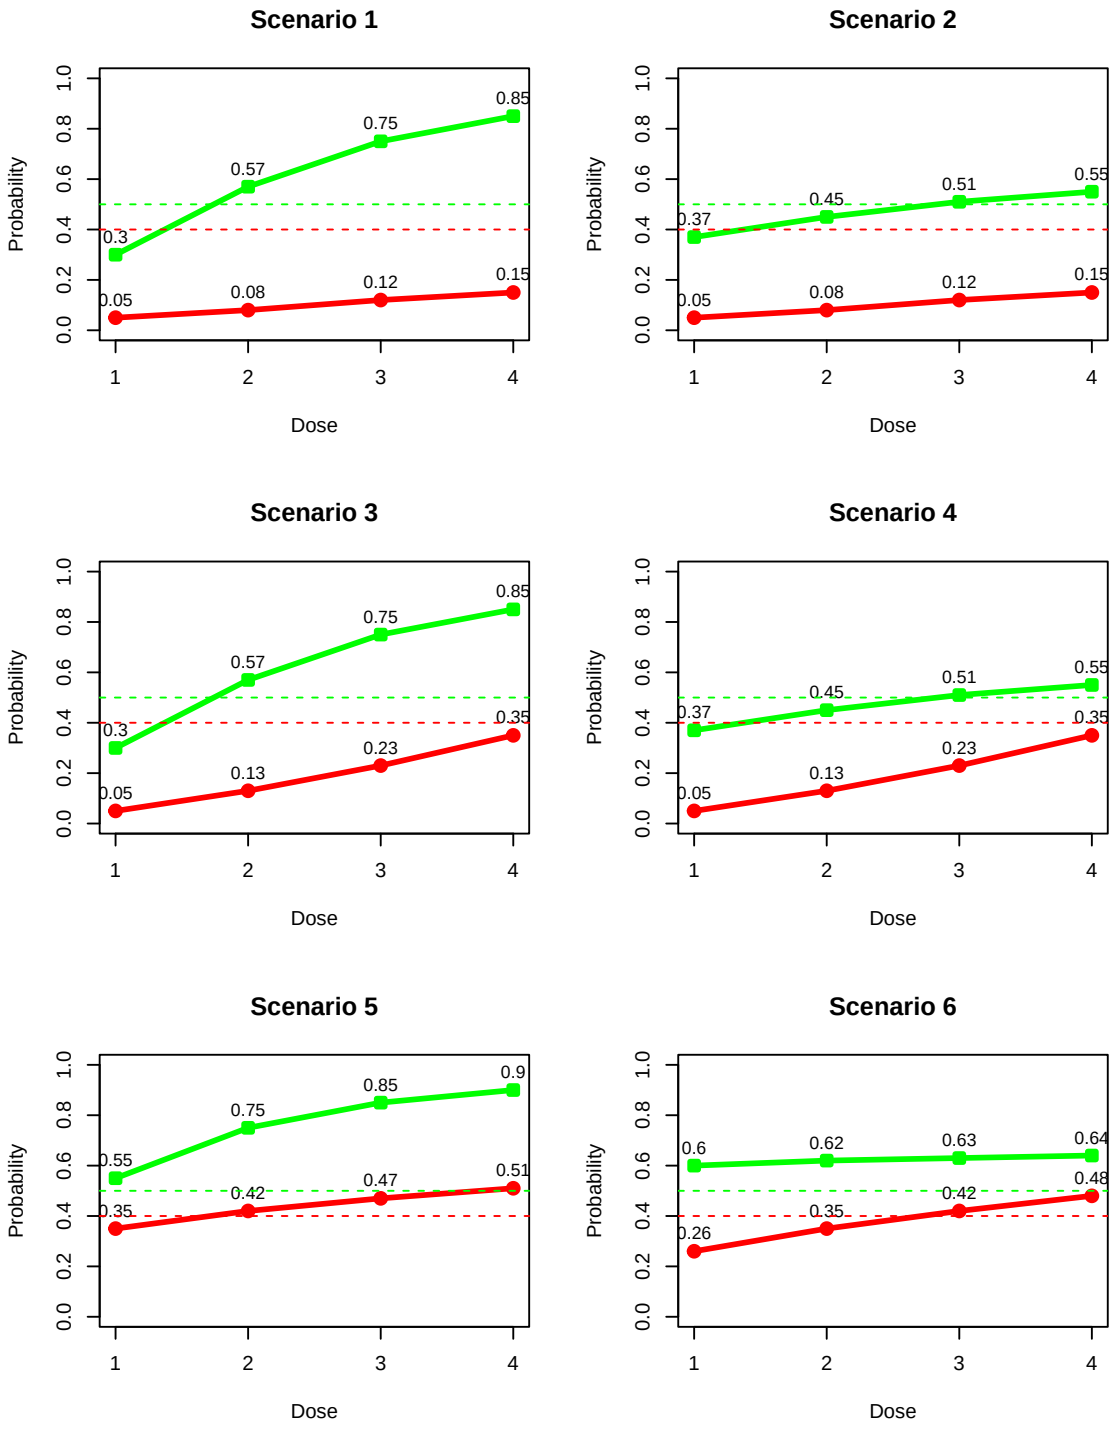

Figure 5.6: Scenarios 1:6. Green line is the fixed probabilities for efficacy ($\tilde{\pi}_E(D)$) and red line for toxicity ($\tilde{\pi}_T(D)$). Dashed lines represent the cut points for the admissibility rules given in *R2DT (1)*

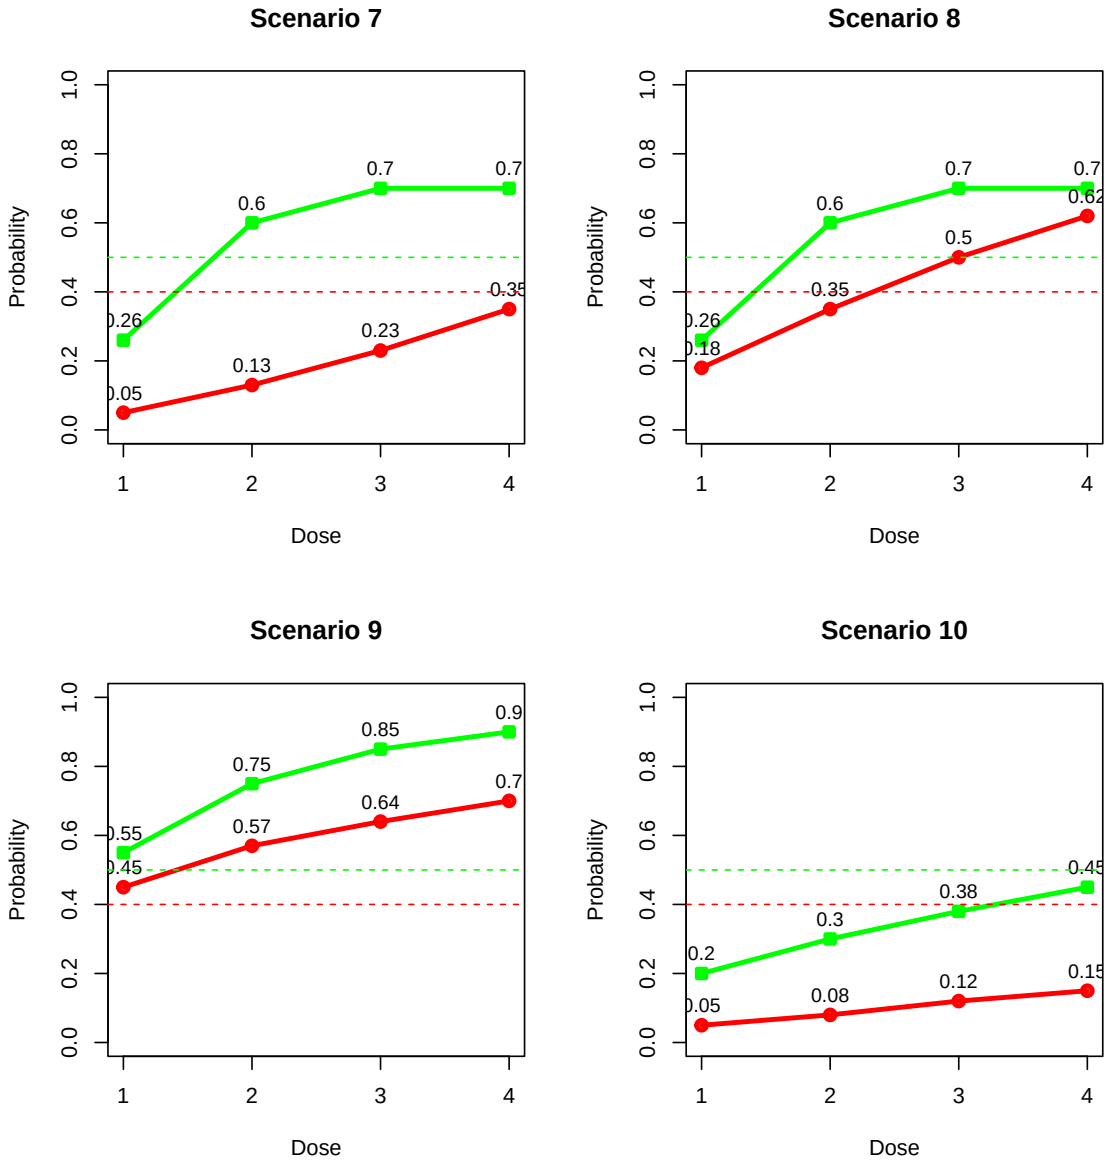

Figure 5.7: Scenarios 7:10. Green line is the fixed probabilities for efficacy ($\tilde{\pi}_E(D)$) and red line for toxicity ($\tilde{\pi}_T(D)$). Dashed lines represent the cut points for the admissibility rules given in *R2DT (1)*

Table 5.2: Listing of each of the probability and fixed trial parameters for simulation study

Notation	Value	Interpretation
D	$[20, 30, 40, 50]$	actual doses
x	$[-0.5, -0.1, 0.19, 0.41]$	transformed doses
x^2	$[0.25, 0.01, 0.04, 0.17]$	square of transformed doses
α_T	$N(-3.17, 2.88)$	toxicity intercept
β_{1T}	$N(-3.56, 2.79)$	toxicity slope
α_E	$N(0.73, 2.44)$	efficacy intercept
β_{1E}	$N(-0.11, 2.34)$	efficacy slope
β_{2E}	$N(0, 0.2)$	efficacy squared slope
$\tilde{\pi}_E(D)$	$[0.42, 0.57, 0.67, 0.72]$	Efficacy prior probabilities
$\tilde{\pi}_T(D)$	$[0.14, 0.2, 0.26, 0.33]$	Toxicity prior probabilities
	$[1, 1]$	ESS toxicity and efficacy
	20	Starting dose
N	45	Max Sample Size
	3	Cohort Size
	2000	Number of simulation repetitions

Table 5.3: Short description of each of the different methods in simulation study

Label	Description
R2DT (1)	Sigmoidal and inverted sigmoidal shaped efficacy and toxicity marginal utility functions respectively. Joint utility $k_E = 0.25$ and $k_T = 0.15$. Admissibility rules applied as separate step functions at each dose
EffToxU (2)	Marginal utilities are linear. Joint utility and admissibility rules applied as R2DT (1)
R2DT (3i)	R2DT (1) but single admissibility rule based upon contour,u(0.5,0.35)= 0.58
R2DT (3ii)	as (i) but contour includes u(0.7,0.4)= 0.62
R2DT (3iii)	as (i) but contour includes u(0.9,0.4)= 0.69
R2DT (4i)	R2DT (1) but single trial stopping rule based upon u(0.5,0.35)= 0.58. All doses considered at each stage but trial stops if all doses are admissible
R2DT (4ii)	as (i) but contour includes U(0.7,0.4)= 0.62
R2DT (4iii)	as (i) but contour includes U(0.9,0.4)= 0.69
EffToxU (5)	EffToxU (2) single admissibility rule based upon u(0.5,0.35)= 0.42
EffTox (6)	EffTox method applied defined from equal contour passing u(0.5,0.35)= 0.42. Admissibility rules applied as separate step functions at each dose
EffToxU (7)	EffToxU (2) but with $k_E = 0.5$ and $k_T = 0.3$

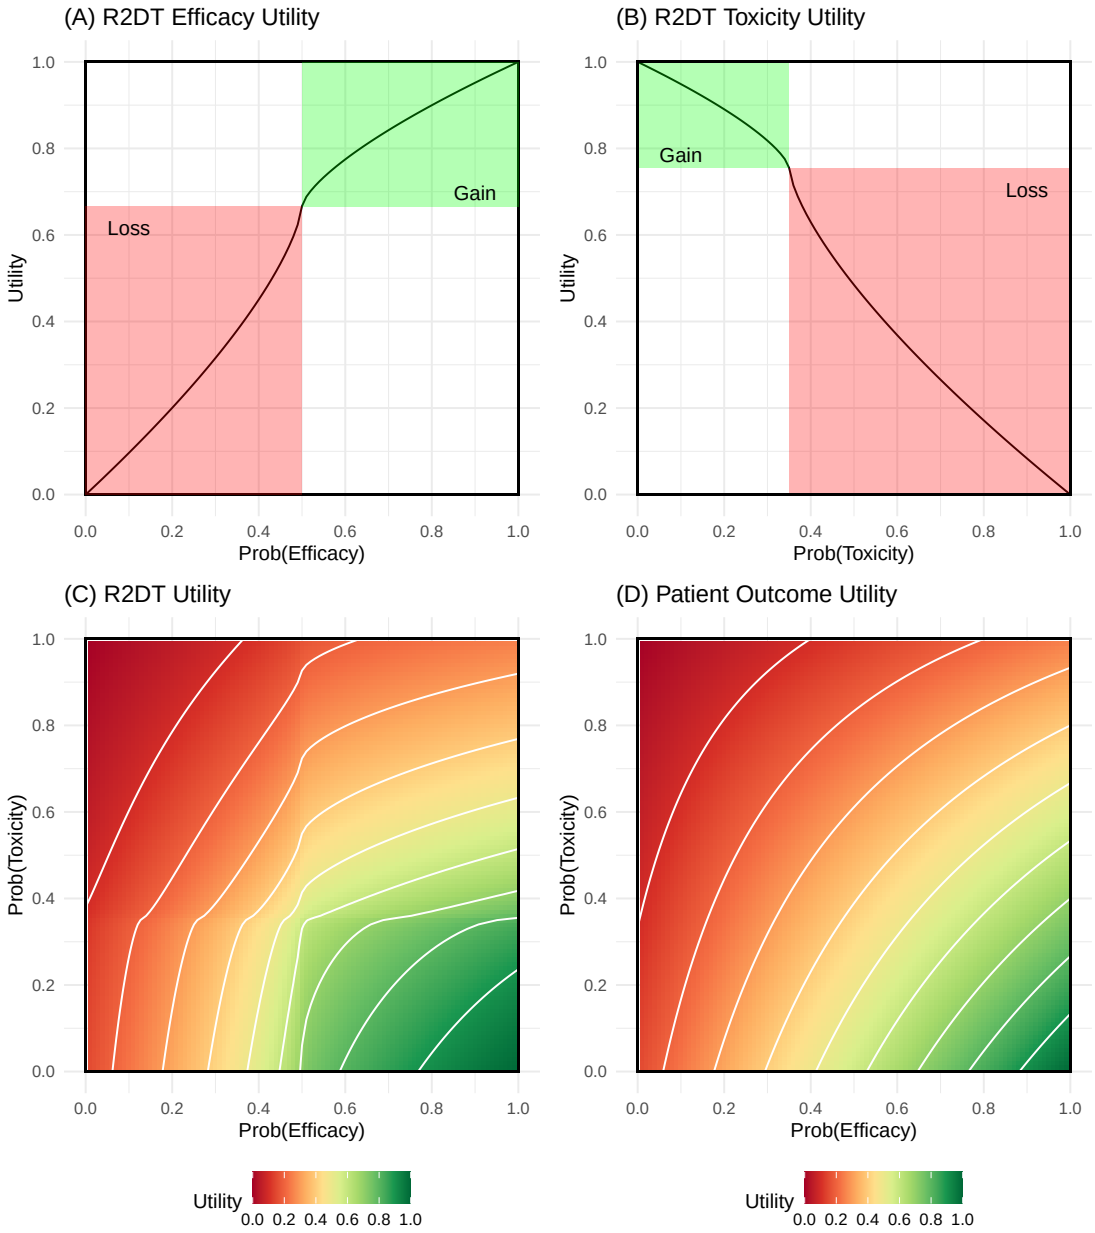

Figure 5.8: R2DT Utility function, Contours in the joint utility represent equal utility at 0.1,0.2,...,0.9 with the the point at guaranteed efficacy and no toxicity having utility of 1. A,B,C depict *R2DT (1)* method in simulation study. D depicts joint utility function of *EffToxU (2)*

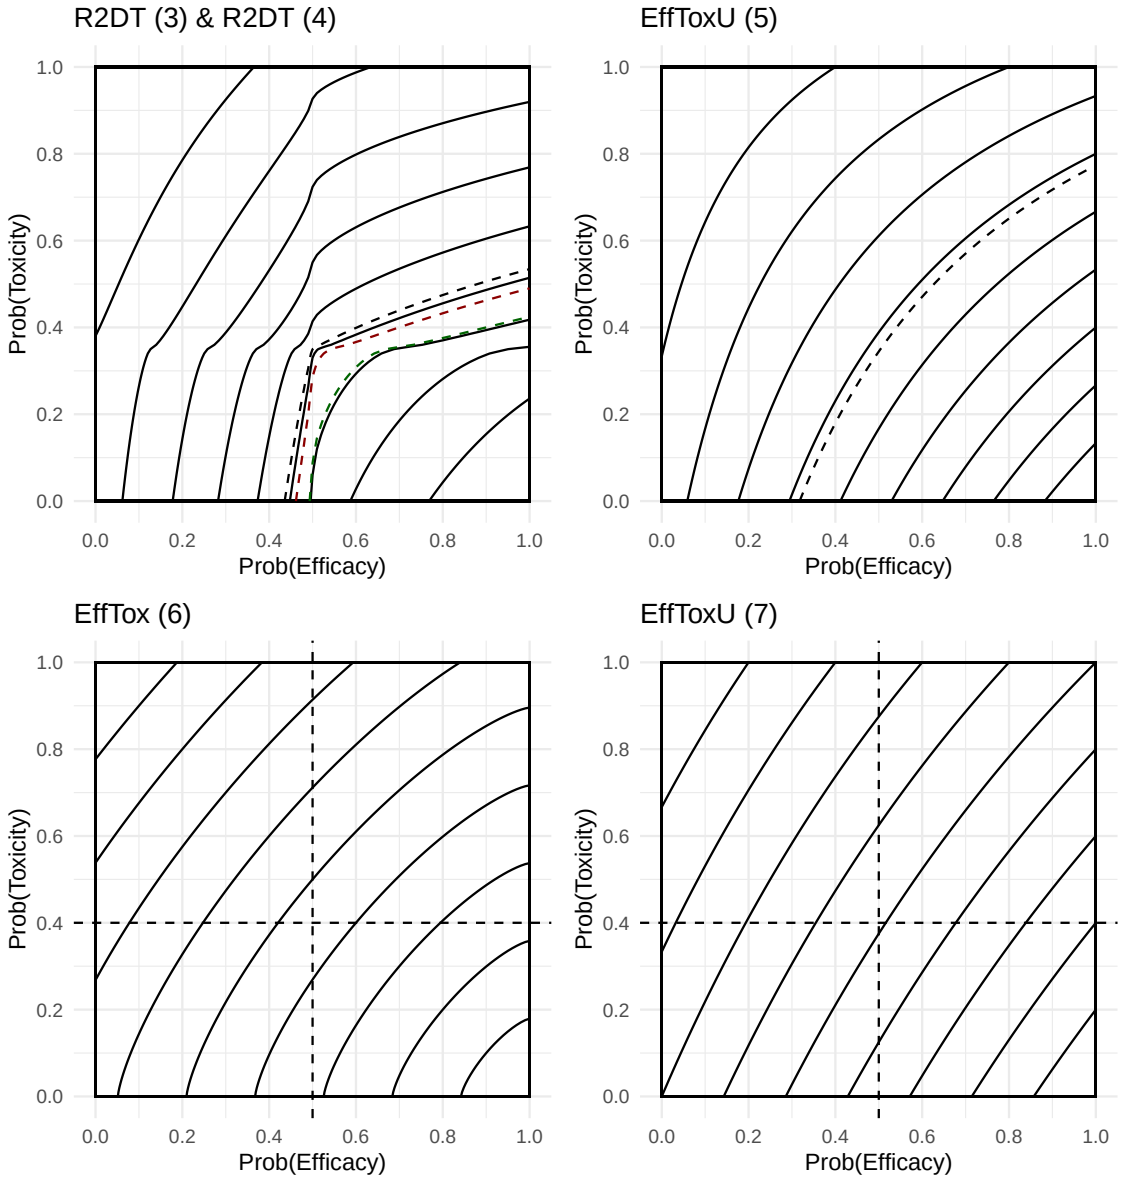

Figure 5.9: Simulation Utility Functions: Contours in the joint utility represent equal utility at 0.1,0.2,...,0.9 with the the point at guaranteed efficacy and no toxicity having utility of 1. Dashed lines are limits for admissibility rules. The contour plot for *R2DT (3)* and *R2DT (4)* gives the stopping rule *(i)* in black ($u(0.5, 0.35) = 0.58$), *(ii)* in red ($u(0.7, 0.4) = 0.62$) and *(iii* in Green ($u(0.9, 0.4) = 0.69$)

5.6.2 Results

The results of the simulation study proposed in the previous section is presented in tables and figures at the end of this section with interpretation and commentary preceding. Tables and Figure present the following:

- Table 5.4 compares the operating characteristics between *R2DT (1)* and *EffToxU (2)* for 10 scenarios.
- Figure 5.10 presents the probability of selection of each dose for Scenarios 2 - 7 for*R2DT (1)* and *EffToxU (2)* methods.
- Table 5.5 compares the the R2DT stopping rule as part of EffTox in *EffToxU (5)* and *EffToxU (2)*
- Tables 5.6 and 5.7 compare different stopping rules for R2DT.
- Table 5.8 compares *EffToxU (7)* and *EffTox (6)*, specified as a sensitivity for the specification of *EffToxU (2)*.

The *R2DT (1)* and *EffToxU (2)* designs are simulated and contrasted in 10 scenarios (Table 5.4). To define which dose is the most desirable in any given scenario doses are first excluded by the stopping rule, i.e any dose that has greater than 40% toxicity or less than 50% efficacy cannot be the optimum dose. The optimum dose is then defined by the maximum utility value. Scenarios 1 & 2 have minimal toxicity and relatively steep efficacy with the $50mg/kg$ dose optimal. Both methods have very similar percentage of selection and numbers of patients treated at each dose. Scenarios 3 and 4 mirror the efficacy of scenarios 1 and 2 but increase the toxicity with both methods indicating the $40mg/kg$ dose as optimal. *R2DT (1)* recommends the optimum dose more often particularly in scenario 3.

Scenario 5 is steeply increasing with efficacy but is also very toxic, relative to the reference point, with the $20mg/kg$ dose optimal according to *R2DT (1)* and the $50mg/kg$ dose according to *EffToxU (2)*. In this scenario *EffToxU (2)* has an equal utility with rounding to 2 decimal places at the $40mg/kg$ dose, without the rounding the $50mg/kg$ has a marginally higher utility. In practice this means that in this scenario the doses are considered practically equivalent. But when combining with the toxicity admissibility rule however all but the $20mg/kg$ for *EffToxU(2)*. *R2DT (1)* chooses the lower two doses more often under

this scenario. Scenario 6 is flat for efficacy with the $20mg/kg$ dose optimal according to both utility functions, *EffToxU (2)* out performs *R2DT (1)* in terms of correct selection. Scenario 7 has an efficacy plateau at the $40mg/kg$ dose; *R2DT (1)* strongly outperforms in this scenario. Scenario 8 has a steep increase in toxicity with the $30mg/kg$ optimal. *R2DT (1)* out performs *EffToxU (2)*. Scenario 9 is specified with all doses overly toxic. The two designs perform similarly despite *EffToxU (2)* suggesting the $40mg/kg$ is optimal according to the utility function. The decision making process in this scenario is dominated by the admissibility rules which are identical between the designs. Scenario 10 is minimally efficacious for all doses with similar interpretation to the previous scenario.

The probability of selection of each dose for Scenarios 2 - 7 with the two methods is contrasted with sample size in Figure 5.10. In all scenarios after 12 patients the *EffToxU (2)* has a greater proportion of simulated trials selecting the 50mg/kg dose as optimal suggesting that the *R2DT (1)* method is initially more conservative in escalation. The probability of correct selection of both methods increase with sample size. The choice of 45 patients was deemed appropriate in this setting based upon the slower rate of improvement in accuracy after 45 patients and is seen as a clinically realistic sample size for this number of doses and setting.

Applying the R2DT stopping rule to EffTox in *EffToxU (5)*, Tables 5.5, makes little difference to scenarios 1, 2, 3, 4, 6, 7 in comparison to *EffToxU (2)*. In these scenarios some of the lower dose levels may be unacceptable but the main driver of design performance is the utility function, which is the same between the two designs. In Scenario 5 the contour stipulates that all doses are acceptable and maximises more frequently to the $50mg/kg$ dose which has a toxicity of 51%. Whether this is acceptable will need clinical judgement, if it isn't this would suggest that an inappropriate stopping rule has been specified. In scenarios 8, 9 and 10 the *EffToxU (5)* suggest that higher doses have acceptable toxicity given high efficacy. Take scenario 9 for example it is only the $20mg/kg$ dose that has unacceptable utility in contrast to all doses in *EffToxU (2)*. This results in the alternative stopping rule more frequently recommending higher doses and a lower proportion of trials stopping early without selecting a dose. Similarly in scenario 10 the alternative stopping rule recommends the $50mg/kg$ a high proportion of the time. Here the $50mg/kg$ is acceptable according to the specified stopping rule. It is unlikely that a contour could be specified that accommo-

dates a threshold for unacceptability. This set of simulations highlights the dependence of the EffTox method on the stopping rules to restrict treating and recommending doses that are overly toxic or (and) not efficacious enough.

Specification of different stopping rules for R2DT makes minimal difference in scenarios 1, 3 and 7, Table 5.6 and Table 5.7. In general in the other scenarios the admissibility stopping rules (*EffToxU (3)*) are more likely to exclude doses and more likely to recommend stopping the trial in contrast to the trial stopping rule in (*EffToxU (4)*). This is a comparison between designs where the contour is the same, denoted by the same Roman numeral. The difference is slightly larger in designs requiring the highest utility, (*iii*), but still less than 5%. This result is not unexpected as the only time that the decisions will differ is if a dose maximises the expected utility but also meets the threshold to be classed as inadmissible. In most instances the dose with the maximum expected utility will also be admissible. The two designs will recommend stopping at the same point but clearly there are some different decisions at earlier decision points.

In scenario 2 (*R2DT (3iii)*) recommends stopping with no dose selected in 18% of simulations. This is because the utility of 0.73 is close to the stopping contour with utility of 0.69. In scenario 4 the trial is more likely to stop and select no dose with the alternative stopping rules this is predominantly a reduced number of times selecting the $50mg/kg$ dose. There appears to be a noticeable difference for Scenario 5, where the stopping rule based upon utility for (*R2DT (3i)*) suggests that all doses are acceptable (Utility at each dose is greater than the reference utility values) while the stopping rule based upon the individual probabilities would exclude all but the $20mg/kg$ dose. In scenarios 6, 8, 9 and 10 the contour stopping rules are more likely to end the trial without recommending a dose. This is proportional to how strict the stopping rule is with the designs needing a higher utility stopping more often. The rules are not directly comparable, and in practice which one is sensible would need clinical input for the given situation. If the utility approach cannot be considered acceptable this would indicate that the contour needs changing.

EffToxU (7) and *EffTox (6)*, specified as a sensitivity for the specification of *EffToxU (2)*, make little difference (Table 5.8). There is a difference in scenario 3 with *EffToxU (7)* suggesting dose level 4 is optimal and selecting this dose level more often.

Subsequent pages contain all of the tables and figures associated with the operating charac-

teristics of the simulation study. A summary of the key findings from each table and figure is as follows:

- Table 5.4; the R2DT method can lead to considerable improvement in operating characteristics in comparison to EffToxU.
- Figure 5.10; the R2DT method is initially more conservative in escalation than the EffToxU method as configured.
- Table 5.5; the EffTox method has a high dependence on the stopping rules to restrict treating and recommending doses that are overly toxic or (and) not efficacious enough. Without them the method tends to recommend overly toxic doses.
- Tables 5.6 and 5.7; the R2DT method the new stopping rules are relatively consistent wt.
- Table 5.8; the main conclusion for the R2DT method is consistent when comparing with two alternative specifications of the comparator.

Table 5.4: Comparison between RT2D and EffToxU: data of form: [utility at scenerio probability (π_E, π_T)] percentage selection (average number of patients treated). Percentage of trials with no dose selected abbreviated to NDS

	Dose (mg/kg)				
Method	20	30	40	50	NDS
		Scenario 1 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.08)	(0.75, 0.12)	(0.85, 0.15)	
R2DT (1)	[0.41] 0.9 (5.1)	[0.76] 4 (4.9)	[0.85] 8.6 (5.8)	[**0.88**] 86.2 (29.1)	0.3
EffToxU (2)	[0.39] 1.5 (4.8)	[0.60] 4.1 (5.3)	[0.72] 3.8 (4.1)	[**0.77**] 90.4 (30.8)	0.2
		Scenario 2 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.08)	(0.51, 0.12)	(0.55, 0.15)	
R2DT (1)	[0.49] 14.1 (11.6)	[0.58] 7 (6.2)	[0.70] 8 (5.1)	[**0.73**] 63.3 (20.6)	7.6
EffToxU (2)	[0.45] 15.3 (11.6)	[0.50] 5.8 (5.9)	[0.53] 6.5 (4.2)	[**0.55**] 65 (21.8)	7.4
		Scenario 3 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.13)	(0.75, 0.23)	(0.85, 0.35)	
R2DT (1)	[0.41] 0.9 (4.9)	[0.75] 11.7 (7.4)	[**0.80**] 54.9 (16.4)	[0.76] 32.2 (16.2)	0.4
EffToxU (2)	[0.39] 1.2 (4.8)	[0.57] 9 (6.7)	[**0.65**] 29.2 (10.2)	[0.64] 60.1 (23.1)	0.5
		Scenario 4 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.13)	(0.51, 0.23)	(0.55, 0.35)	
R2DT (1)	[0.49] 14.7 (11.7)	[0.57] 10.4 (7.1)	[**0.66**] 28 (9.7)	[0.63] 38.6 (14.6)	8.3
EffToxU (2)	[0.45] 16.2 (11.8)	[0.48] 13.3 (7.5)	[**0.48**] 21.3 (7.4)	[0.45] 40.5 (16.4)	8.7
		Scenario 5 (π_E, π_T)			
	(0.55, 0.35)	(0.75, 0.42)	(0.85, 0.47)	(0.9, 0.51)	
R2DT (1)	[**0.63**] 19.7 (8.5)	[0.62] 33.2 (13.2)	[0.60] 8.9 (6.5)	[0.58] 29.8 (15.3)	8.3
EffToxU (2)	[0.45] 14.8 (7.6)	[0.54] 26.9 (11.4)	[0.56] 15.6 (7.3)	[**0.56**] 34.2 (17.2)	8.5
		Scenario 6 (π_E, π_T)			
	(0.6, 0.26)	(0.62, 0.35)	(0.63, 0.42)	(0.64, 0.48)	
R2DT (1)	[**0.72**] 31 (13.1)	[0.67] 35.2 (16)	[0.57] 13.5 (6.8)	[0.52] 18.1 (8.6)	2.1
EffToxU (2)	[**0.53**] 39.1 (14.8)	[0.49] 24.6 (12.8)	[0.46] 9.8 (5.8)	[0.44] 24.3 (11.2)	2.1
		Scenario 7 (π_E, π_T)			
	(0.26, 0.05)	(0.6, 0.13)	(0.7, 0.23)	(0.7, 0.35)	
R2DT (1)	[0.37] 0.3 (4.4)	[0.77] 15.2 (8)	[**0.78**] 46.9 (14.9)	[0.70] 36.9 (17.4)	0.8
EffToxU (2)	[0.36] 0.9 (4.6)	[0.59] 11.9 (7.4)	[**0.61**] 27.9 (9.4)	[0.54] 58.6 (23.4)	0.8
		Scenario 8 (π_E, π_T)			
	(0.26, 0.18)	(0.6, 0.35)	(0.7, 0.5)	(0.7, 0.62)	
R2DT (1)	[0.35] 3.4 (5.6)	[**0.66**] 61.4 (18.2)	[0.53] 22.2 (11)	[0.44] 6.2 (8.8)	6.8
EffToxU (2)	[0.32] 3.9 (6.3)	[**0.48**] 50.8 (14.4)	[0.46] 26.5 (10.6)	[0.39] 11.8 (12.3)	7
		Scenario 9 (π_E, π_T)			
	(0.55, 0.45)	(0.75, 0.57)	(0.85, 0.64)	(0.9, 0.7)	
R2DT (1)	[**0.51**] 32.4 (13.1)	[0.49] 10.8 (8.5)	[0.46] 0.9 (4)	[0.43] 2.9 (8.3)	52.9
EffToxU (2)	[0.40] 29.4 (12.3)	[0.45] 13.2 (8.9)	[**0.45**] 2 (4.3)	[0.43] 2.9 (8.6)	52.5
		Scenario 10 (π_E, π_T)			
	(0.2, 0.05)	(0.3, 0.08)	(0.38, 0.12)	(0.45, 0.15)	
R2DT (1)	[0.31] 1.6 (6.2)	[0.40] 0.6 (3.7)	[0.48] 0.9 (3.7)	[**0.57**] 51.1 (21.3)	45.7
EffToxU (2)	[0.31] 1.5 (6.1)	[0.38] 0.9 (3.7)	[0.43] 1.5 (3.6)	[**0.47**] 51.1 (21.7)	44.9

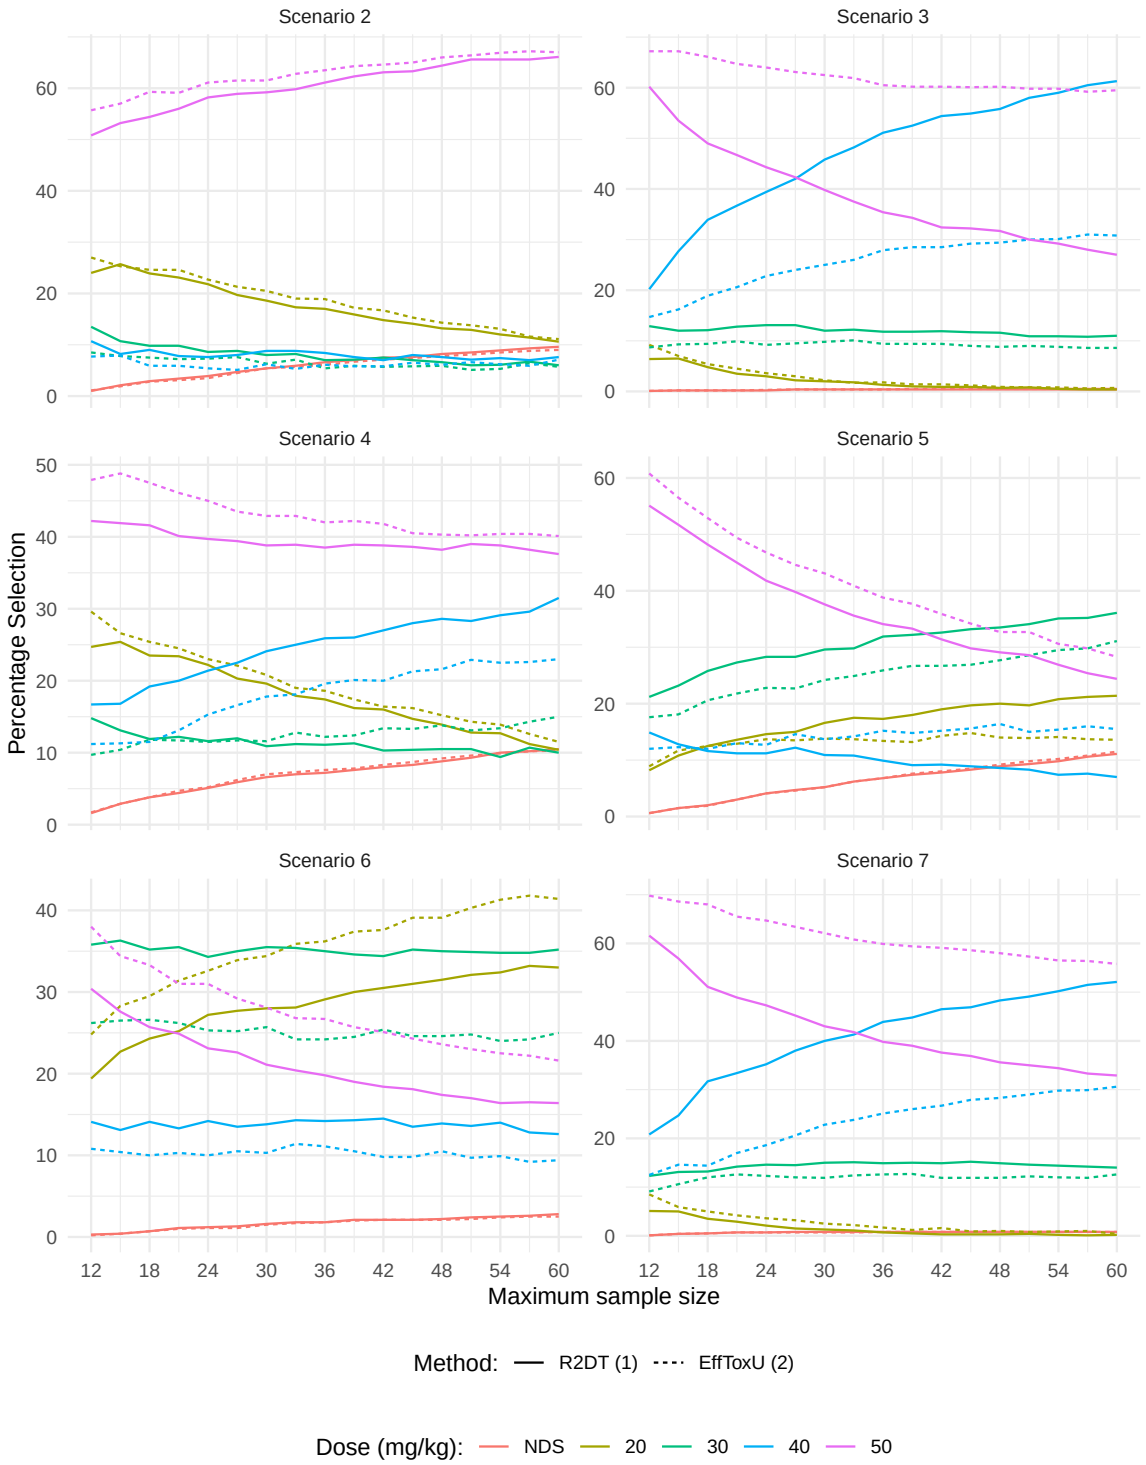

Figure 5.10: Percentage selection by sample size: Each scenario is plotted separately with solid lines representing *R2DT (1)* method, dashed lines *EfftoxU (2)* and the recommended dose at the end of the trial by colour.

Table 5.5: Evaluation of proposed stopping rules: data of form: [utility at scenerio probability (π_E, π_T)] percentage selection (average number of patients treated). Percentage of trials with no dose selected abbreviated to NDS

	Dose (mg/kg)				
Method	20	30	40	50	NDS
		Scenario 1 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.08)	(0.75, 0.12)	(0.85, 0.15)	
EffToxU (2)	[0.39] 1.5 (4.8)	[0.60] 4.1 (5.3)	[0.72] 3.8 (4.1)	[**0.77**] 90.4 (30.8)	0.2
EffToxU (5)	[0.39] 3.2 (5.3)	[0.60] 4 (5.2)	[0.72] 3.6 (4.1)	[**0.77**] 89.1 (30.4)	0
		Scenario 2 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.08)	(0.51, 0.12)	(0.55, 0.15)	
EffToxU (2)	[0.45] 15.3 (11.6)	[0.50] 5.8 (5.9)	[0.53] 6.5 (4.2)	[**0.55**] 65 (21.8)	7.4
EffToxU (5)	[0.45] 24.9 (13.2)	[0.50] 6.2 (6)	[0.53] 5 (3.9)	[**0.55**] 63.5 (21.6)	0.5
		Scenario 3 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.13)	(0.75, 0.23)	(0.85, 0.35)	
EffToxU (2)	[0.39] 1.2 (4.8)	[0.57] 9 (6.7)	[**0.65**] 29.2 (10.2)	[0.64] 60.1 (23.1)	0.5
EffToxU (5)	[0.39] 2.9 (5.2)	[0.57] 8.8 (6.6)	[**0.65**] 28.3 (10.1)	[0.64] 60 (23)	0.1
		Scenario 4 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.13)	(0.51, 0.23)	(0.55, 0.35)	
EffToxU (2)	[0.45] 16.2 (11.8)	[0.48] 13.3 (7.5)	[**0.48**] 21.3 (7.4)	[0.45] 40.5 (16.4)	8.7
EffToxU (5)	[0.45] 28.2 (14.1)	[0.48] 13.5 (7.5)	[**0.48**] 17.6 (6.7)	[0.45] 38.4 (16)	2.2
		Scenario 5 (π_E, π_T)			
	(0.55, 0.35)	(0.75, 0.42)	(0.85, 0.47)	(0.9, 0.51)	
EffToxU (2)	[0.45] 14.8 (7.6)	[0.54] 26.9 (11.4)	[0.56] 15.6 (7.3)	[**0.56**] 34.2 (17.2)	8.5
EffToxU (5)	[0.45] 8.1 (6.4)	[0.54] 20.6 (10.4)	[0.56] 16.6 (7.8)	[**0.56**] 54.6 (20.4)	0
		Scenario 6 (π_E, π_T)			
	(0.6, 0.26)	(0.62, 0.35)	(0.63, 0.42)	(0.64, 0.48)	
EffToxU (2)	[**0.53**] 39.1 (14.8)	[0.49] 24.6 (12.8)	[0.46] 9.8 (5.8)	[0.44] 24.3 (11.2)	2.1
EffToxU (5)	[**0.53**] 38 (14.6)	[0.49] 22.8 (12.6)	[0.46] 8.6 (5.6)	[0.44] 29.8 (12)	0.8
		Scenario 7 (π_E, π_T)			
	(0.26, 0.05)	(0.6, 0.13)	(0.7, 0.23)	(0.7, 0.35)	
EffToxU (2)	[0.36] 0.9 (4.6)	[0.59] 11.9 (7.4)	[**0.61**] 27.9 (9.4)	[0.54] 58.6 (23.4)	0.8
EffToxU (5)	[0.36] 2.2 (4.9)	[0.59] 11.8 (7.4)	[**0.61**] 27.1 (9.3)	[0.54] 58.6 (23.3)	0.4
		Scenario 8 (π_E, π_T)			
	(0.26, 0.18)	(0.6, 0.35)	(0.7, 0.5)	(0.7, 0.62)	
EffToxU (2)	[0.32] 3.9 (6.3)	[**0.48**] 50.8 (14.4)	[0.46] 26.5 (10.6)	[0.39] 11.8 (12.3)	7
EffToxU (5)	[0.32] 5.9 (6.5)	[**0.48**] 28.3 (10.8)	[0.46] 24.1 (10.2)	[0.39] 35.7 (16.5)	5.9
		Scenario 9 (π_E, π_T)			
	(0.55, 0.45)	(0.75, 0.57)	(0.85, 0.64)	(0.9, 0.7)	
EffToxU (2)	[0.40] 29.4 (12.3)	[0.45] 13.2 (8.9)	[**0.45**] 2 (4.3)	[0.43] 2.9 (8.6)	52.5
EffToxU (5)	[0.40] 13.7 (8.6)	[0.45] 20.8 (9.3)	[**0.45**] 13.4 (6.6)	[0.43] 48.9 (19.9)	3.2
		Scenario 10 (π_E, π_T)			
	(0.2, 0.05)	(0.3, 0.08)	(0.38, 0.12)	(0.45, 0.15)	
EffToxU (2)	[0.31] 1.5 (6.1)	[0.38] 0.9 (3.7)	[0.43] 1.5 (3.6)	[**0.47**] 51.1 (21.7)	44.9
EffToxU (5)	[0.31] 6.1 (8)	[0.38] 1.9 (3.9)	[0.43] 2.4 (3.5)	[**0.47**] 77.8 (26.4)	11.8

Table 5.6: Stopping rules of R2DT: data of form: [utility at scenerio probability (π_E, π_T)] percentage selection (average number of patients treated). Percentage of trials with no dose selected abbreviated to NDS

	Dose (mg/kg)				
Method	20	30	40	50	NDS
	Scenario 1 (π_E, π_T)				
	(0.3, 0.05)	(0.57, 0.08)	(0.75, 0.12)	(0.85, 0.15)	
R2DT (1)	[0.41] 0.9 (5.3)	[0.76] 4.1 (5)	[0.85] 8.7 (5.8)	[**0.88**] 86.2 (29)	0.1
R2DT (3i)	[0.41] 1.7 (5.5)	[0.76] 4.3 (5)	[0.85] 8.1 (5.8)	[**0.88**] 85.7 (28.6)	0.3
R2DT (3ii)	[0.41] 1.5 (5.3)	[0.76] 4 (5)	[0.85] 8.2 (5.8)	[**0.88**] 85.9 (28.8)	0.4
R2DT (3iii)	[0.41] 0.9 (5.1)	[0.76] 4 (4.9)	[0.85] 8.1 (5.8)	[**0.88**] 83.9 (28.1)	3
R2DT (4i)	[0.41] 3.4 (5.8)	[0.76] 4 (4.9)	[0.85] 8 (5.8)	[**0.88**] 84.4 (28.5)	0.2
R2DT (4ii)	[0.41] 3.2 (5.7)	[0.76] 4 (4.9)	[0.85] 8 (5.8)	[**0.88**] 84.4 (28.4)	0.4
R2DT (4iii)	[0.41] 3.2 (5.7)	[0.76] 4 (4.9)	[0.85] 8 (5.7)	[**0.88**] 82 (27.6)	2.9
	Scenario 2 (π_E, π_T)				
	(0.37, 0.05)	(0.45, 0.08)	(0.51, 0.12)	(0.55, 0.15)	
R2DT (1)	[0.49] 16 (12.2)	[0.58] 7.2 (6.3)	[0.70] 7.8 (5)	[**0.73**] 64.5 (20.6)	4.5
R2DT (3i)	[0.49] 18 (12.6)	[0.58] 7 (6.2)	[0.70] 7 (5)	[**0.73**] 62.6 (20.1)	5.3
R2DT (3ii)	[0.49] 15.3 (11.8)	[0.58] 6.7 (6.1)	[0.70] 7 (5)	[**0.73**] 62.4 (20.1)	8.6
R2DT (3iii)	[0.49] 11.7 (10.8)	[0.58] 6.6 (6)	[0.70] 7.2 (4.9)	[**0.73**] 56.5 (18.6)	18
R2DT (4i)	[0.49] 21.4 (13.1)	[0.58] 6.8 (6.1)	[0.70] 6.8 (4.9)	[**0.73**] 60.7 (19.8)	4.4
R2DT (4ii)	[0.49] 19.9 (12.8)	[0.58] 6.7 (6.1)	[0.70] 6.7 (4.9)	[**0.73**] 59.4 (19.5)	7.3
R2DT (4iii)	[0.49] 18.1 (12.4)	[0.58] 6.6 (6.1)	[0.70] 6.7 (4.7)	[**0.73**] 53 (17.6)	15.6
	Scenario 3 (π_E, π_T)				
	(0.3, 0.05)	(0.57, 0.13)	(0.75, 0.23)	(0.85, 0.35)	
R2DT (1)	[0.41] 0.9 (5)	[0.75] 11.7 (7.4)	[**0.80**] 54.9 (16.3)	[0.76] 32.4 (16.2)	0.1
R2DT (3i)	[0.41] 1 (5.2)	[0.75] 12.5 (7.4)	[**0.80**] 54.9 (16.3)	[0.76] 30.9 (16)	0.7
R2DT (3ii)	[0.41] 0.9 (5)	[0.75] 12 (7.4)	[**0.80**] 55.2 (16.3)	[0.76] 30.6 (15.9)	1.2
R2DT (3iii)	[0.41] 0.6 (4.7)	[0.75] 11.8 (7.4)	[**0.80**] 53.9 (16)	[0.76] 28.5 (15.2)	5.1
R2DT (4i)	[0.41] 2.5 (5.4)	[0.75] 12 (7.4)	[**0.80**] 54.1 (16.1)	[0.76] 30.6 (15.9)	0.7
R2DT (4ii)	[0.41] 2.4 (5.4)	[0.75] 12 (7.4)	[**0.80**] 53.9 (16.1)	[0.76] 30.4 (15.8)	1.2
R2DT (4iii)	[0.41] 2.1 (5.3)	[0.75] 11.9 (7.4)	[**0.80**] 52.6 (15.8)	[0.76] 28.4 (14.9)	4.9
	Scenario 4 (π_E, π_T)				
	(0.37, 0.05)	(0.45, 0.13)	(0.51, 0.23)	(0.55, 0.35)	
R2DT (1)	[0.49] 17.8 (12.6)	[0.57] 10.9 (7.3)	[**0.66**] 27.4 (9.4)	[0.63] 39 (14.8)	5
R2DT (3i)	[0.49] 20.2 (12.8)	[0.57] 9.7 (7)	[**0.66**] 25.1 (9)	[0.63] 34.6 (13.9)	10.4
R2DT (3ii)	[0.49] 17.4 (12)	[0.57] 9.2 (6.9)	[**0.66**] 24.6 (9)	[0.63] 33.1 (13.6)	15.8
R2DT (3iii)	[0.49] 14 (10.9)	[0.57] 8.7 (6.8)	[**0.66**] 21.4 (8.4)	[0.63] 25.7 (11.9)	30.2
R2DT (4i)	[0.49] 24.5 (13.6)	[0.57] 9.4 (7)	[**0.66**] 23.8 (8.8)	[0.63] 32.8 (13.4)	9.5
R2DT (4ii)	[0.49] 23.3 (13.3)	[0.57] 9.2 (6.9)	[**0.66**] 23 (8.7)	[0.63] 31 (13)	13.6
R2DT (4iii)	[0.49] 20.8 (12.8)	[0.57] 8.8 (6.8)	[**0.66**] 21.2 (8.2)	[0.63] 23.8 (11)	25.5
	Scenario 5 (π_E, π_T)				
	(0.55, 0.35)	(0.75, 0.42)	(0.85, 0.47)	(0.9, 0.51)	
R2DT (1)	[**0.63**] 19.2 (8.4)	[0.62] 33.8 (13.2)	[0.60] 9.1 (6.6)	[0.58] 33.1 (16)	4.8
R2DT (3i)	[**0.63**] 15.3 (7.8)	[0.62] 29 (12.4)	[0.60] 9.6 (6.8)	[0.58] 40.6 (16.9)	5.6
R2DT (3ii)	[**0.63**] 15.2 (7.7)	[0.62] 28.4 (12.1)	[0.60] 9 (6.7)	[0.58] 36.2 (16.2)	11.1
R2DT (3iii)	[**0.63**] 13.2 (7.4)	[0.62] 23.6 (11.3)	[0.60] 7.7 (6.4)	[0.58] 26.1 (14.3)	29.4
R2DT (4i)	[**0.63**] 14.2 (7.5)	[0.62] 30.1 (12.6)	[0.60] 9.8 (6.8)	[0.58] 40.7 (16.9)	5.2
R2DT (4ii)	[**0.63**] 13.9 (7.5)	[0.62] 29.5 (12.5)	[0.60] 9.4 (6.8)	[0.58] 36.9 (16.1)	10.2
R2DT (4iii)	[**0.63**] 11.7 (7)	[0.62] 26 (11.9)	[0.60] 8.4 (6.5)	[0.58] 27.5 (14.3)	26.4

Table 5.7: Stopping rules of R2DT: data of form: [utility at scenerio probability (π_E, π_T)] percentage selection (average number of patients treated). Percentage of trials with no dose selected abbreviated to NDS

	Dose (mg/kg)				
Method	20	30	40	50	NDS
	Scenario 6 (π_E, π_T)				
	(0.6, 0.26)	(0.62, 0.35)	(0.63, 0.42)	(0.64, 0.48)	
R2DT (1)	[**0.72**] 30.6 (13.2)	[0.67] 35.1 (16)	[0.57] 13.8 (6.8)	[0.52] 19.4 (8.9)	1
R2DT (3i)	[**0.72**] 30 (13)	[0.67] 33.1 (15.6)	[0.57] 12.3 (6.5)	[0.52] 17.8 (8.4)	6.7
R2DT (3ii)	[**0.72**] 30 (12.8)	[0.67] 32.3 (15.3)	[0.57] 11 (6.3)	[0.52] 15.7 (8)	11.1
R2DT (3iii)	[**0.72**] 27.6 (12.1)	[0.67] 27.8 (14.3)	[0.57] 9.2 (6)	[0.52] 11.3 (7.1)	24.1
R2DT (4i)	[**0.72**] 29.3 (12.9)	[0.67] 32.9 (15.6)	[0.57] 12.4 (6.5)	[0.52] 18.6 (8.4)	6.8
R2DT (4ii)	[**0.72**] 28.1 (12.5)	[0.67] 32.7 (15.6)	[0.57] 12.1 (6.5)	[0.52] 16.4 (7.9)	10.6
R2DT (4iii)	[**0.72**] 24.6 (11.6)	[0.67] 31.6 (15.3)	[0.57] 10.5 (6.1)	[0.52] 10.9 (6.7)	22.4
	Scenario 7 (π_E, π_T)				
	(0.26, 0.05)	(0.6, 0.13)	(0.7, 0.23)	(0.7, 0.35)	
R2DT (1)	[0.37] 0.4 (4.5)	[0.77] 15 (8)	[**0.78**] 46.8 (14.9)	[0.70] 37.4 (17.5)	0.3
R2DT (3i)	[0.37] 0.6 (4.6)	[0.77] 14.4 (7.8)	[**0.78**] 46.2 (14.9)	[0.70] 37.5 (17.3)	1.4
R2DT (3ii)	[0.37] 0.5 (4.5)	[0.77] 14.2 (7.8)	[**0.78**] 46.1 (14.9)	[0.70] 36.8 (17.2)	2.4
R2DT (3iii)	[0.37] 0.5 (4.4)	[0.77] 14.5 (7.8)	[**0.78**] 44.9 (14.7)	[0.70] 33.2 (16.1)	6.9
R2DT (4i)	[0.37] 1.3 (4.8)	[0.77] 14.4 (7.8)	[**0.78**] 45.9 (14.8)	[0.70] 37.1 (17.2)	1.4
R2DT (4ii)	[0.37] 1.1 (4.8)	[0.77] 14.3 (7.8)	[**0.78**] 45.8 (14.8)	[0.70] 36.6 (17.1)	2.2
R2DT (4iii)	[0.37] 1.1 (4.7)	[0.77] 14.2 (7.7)	[**0.78**] 44.9 (14.5)	[0.70] 32.8 (15.9)	7
	Scenario 8 (π_E, π_T)				
	(0.26, 0.18)	(0.6, 0.35)	(0.7, 0.5)	(0.7, 0.62)	
R2DT (1)	[0.35] 3.7 (5.7)	[**0.66**] 59.4 (17.8)	[0.53] 25 (11.3)	[0.44] 7.9 (9.5)	4.1
R2DT (3i)	[0.35] 2.5 (5.3)	[**0.66**] 43.8 (15.4)	[0.53] 18.6 (10)	[0.44] 7.6 (8.6)	27.6
R2DT (3ii)	[0.35] 1.9 (5.3)	[**0.66**] 41.3 (14.9)	[0.53] 14.1 (9.1)	[0.44] 4.4 (7.5)	38.2
R2DT (3iii)	[0.35] 2.1 (5)	[**0.66**] 28.3 (12.4)	[0.53] 7.6 (7.6)	[0.44] 2.5 (6.1)	59.7
R2DT (4i)	[0.35] 2.8 (5.3)	[**0.66**] 41.3 (15.1)	[0.53] 21.9 (10.5)	[0.44] 9.2 (8.8)	24.8
R2DT (4ii)	[0.35] 2.2 (5.2)	[**0.66**] 39.4 (14.7)	[0.53] 18 (9.8)	[0.44] 5.7 (7.7)	34.8
R2DT (4iii)	[0.35] 1.2 (4.8)	[**0.66**] 32 (13.2)	[0.53] 10.8 (8.1)	[0.44] 2 (5.9)	54
	Scenario 9 (π_E, π_T)				
	(0.55, 0.45)	(0.75, 0.57)	(0.85, 0.64)	(0.9, 0.7)	
R2DT (1)	[**0.51**] 39.6 (14.2)	[0.49] 13.6 (9.1)	[0.46] 1.3 (4.1)	[0.43] 4.2 (9.7)	41.2
R2DT (3i)	[**0.51**] 21.1 (10.4)	[0.49] 12.6 (8.4)	[0.46] 1.8 (4.3)	[0.43] 10.3 (11.5)	54.2
R2DT (3ii)	[**0.51**] 16.7 (9.6)	[0.49] 8.6 (7.4)	[0.46] 1.1 (4.1)	[0.43] 5.8 (9.5)	67.8
R2DT (3iii)	[**0.51**] 8.6 (7.8)	[0.49] 3.4 (6)	[0.46] 0.7 (3.8)	[0.43] 1.8 (7.2)	85.6
R2DT (4i)	[**0.51**] 17.2 (9.5)	[0.49] 18.2 (9.4)	[0.46] 2.5 (4.4)	[0.43] 13.6 (11.9)	48.5
R2DT (4ii)	[**0.51**] 13.8 (8.6)	[0.49] 14.3 (8.8)	[0.46] 1.8 (4.2)	[0.43] 7.6 (9.9)	62.5
R2DT (4iii)	[**0.51**] 7.8 (7.2)	[0.49] 7.8 (7.3)	[0.46] 0.7 (3.8)	[0.43] 2.5 (7.3)	81.2
	Scenario 10 (π_E, π_T)				
	(0.2, 0.05)	(0.3, 0.08)	(0.38, 0.12)	(0.45, 0.15)	
R2DT (1)	[0.31] 2.3 (6.6)	[0.40] 1 (3.8)	[0.48] 0.9 (3.7)	[**0.57**] 60.5 (23.3)	35.3
R2DT (3i)	[0.31] 2 (6.6)	[0.40] 0.5 (3.7)	[0.48] 1 (3.6)	[**0.57**] 59.4 (22.4)	37.2
R2DT (3ii)	[0.31] 1.3 (6.2)	[0.40] 0.6 (3.6)	[0.48] 0.9 (3.6)	[**0.57**] 50.2 (20.5)	46.9
R2DT (3iii)	[0.31] 1.2 (5.5)	[0.40] 0.2 (3.6)	[0.48] 0.9 (3)	[**0.57**] 32.7 (14.8)	65
R2DT (4i)	[0.31] 2.9 (7.1)	[0.40] 0.6 (3.6)	[0.48] 1.2 (3.6)	[**0.57**] 60.5 (22.4)	34.8
R2DT (4ii)	[0.31] 2.4 (6.8)	[0.40] 0.5 (3.6)	[0.48] 1.2 (3.6)	[**0.57**] 51.6 (20.3)	44.2
R2DT (4iii)	[0.31] 2.1 (6.5)	[0.40] 0.5 (3.6)	[0.48] 1.2 (3.1)	[**0.57**] 34.6 (14.6)	61.5

Table 5.8: Sensitivity of EffToxU: data of form: [utility at scenario probability (π_E, π_T)] percentage selection (average number of patients treated). Percentage of trials with no dose selected abbreviated to NDS

	Dose (mg/kg)				
Method	20	30	40	50	NDS
		Scenario 1 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.08)	(0.75, 0.12)	(0.85, 0.15)	
EffToxU (2)	[0.39] 1.5 (4.8)	[0.60] 4.1 (5.3)	[0.72] 3.8 (4.1)	[**0.77**] 90.4 (30.8)	0.2
EffTox (6)	[0.55] 1.4 (4.8)	[0.71] 4.2 (5.4)	[0.81] 3.5 (4.1)	[**0.85**] 90.6 (30.6)	0.2
EffToxU (7)	[0.49] 1.5 (4.8)	[0.67] 4.4 (5.3)	[0.77] 2.2 (3.7)	[**0.82**] 91.6 (31.2)	0.2
		Scenario 2 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.08)	(0.51, 0.12)	(0.55, 0.15)	
EffToxU (2)	[0.45] 15.3 (11.6)	[0.50] 5.8 (5.9)	[0.53] 6.5 (4.2)	[**0.55**] 65 (21.8)	7.4
EffTox (6)	[0.59] 15.4 (11.7)	[0.64] 5.9 (6.1)	[0.66] 5.9 (4)	[**0.68**] 65.6 (21.8)	7.2
EffToxU (7)	[0.54] 15.3 (11.4)	[0.58] 6 (6)	[0.61] 6.6 (4.2)	[**0.62**] 65 (21.9)	7.1
		Scenario 3 (π_E, π_T)			
	(0.3, 0.05)	(0.57, 0.13)	(0.75, 0.23)	(0.85, 0.35)	
EffToxU (2)	[0.39] 1.2 (4.8)	[0.57] 9 (6.7)	[**0.65**] 29.2 (10.2)	[0.64] 60.1 (23.1)	0.5
EffTox (6)	[0.55] 1 (4.8)	[0.69] 8.4 (6.6)	[**0.76**] 31 (10.6)	[0.75] 59.2 (22.9)	0.4
EffToxU (7)	[0.49] 1.1 (4.8)	[0.64] 7.3 (6.2)	[0.72] 19.9 (8)	[**0.73**] 71.2 (25.9)	0.5
		Scenario 4 (π_E, π_T)			
	(0.37, 0.05)	(0.45, 0.13)	(0.51, 0.23)	(0.55, 0.35)	
EffToxU (2)	[0.45] 16.2 (11.8)	[0.48] 13.3 (7.5)	[**0.48**] 21.3 (7.4)	[0.45] 40.5 (16.4)	8.7
EffTox (6)	[0.59] 16.2 (11.8)	[**0.62**] 14.1 (7.7)	[0.62] 21.3 (7.5)	[0.60] 39.7 (16.1)	8.6
EffToxU (7)	[0.54] 17.2 (11.9)	[0.56] 13.2 (7.4)	[**0.56**] 18.4 (6.8)	[0.54] 43 (17)	8.2
		Scenario 5 (π_E, π_T)			
	(0.55, 0.35)	(0.75, 0.42)	(0.85, 0.47)	(0.9, 0.51)	
EffToxU (2)	[0.45] 14.8 (7.6)	[0.54] 26.9 (11.4)	[0.56] 15.6 (7.3)	[**0.56**] 34.2 (17.2)	8.5
EffTox (6)	[0.60] 14.1 (7.6)	[0.67] 27.7 (11.3)	[**0.69**] 15.3 (7.5)	[0.69] 34.4 (17.1)	8.5
EffToxU (7)	[0.54] 12.4 (7.2)	[0.64] 26.2 (10.7)	[0.67] 14.5 (7.1)	[**0.69**] 38.6 (18.6)	8.3
		Scenario 6 (π_E, π_T)			
	(0.6, 0.26)	(0.62, 0.35)	(0.63, 0.42)	(0.64, 0.48)	
EffToxU (2)	[**0.53**] 39.1 (14.8)	[0.49] 24.6 (12.8)	[0.46] 9.8 (5.8)	[0.44] 24.3 (11.2)	2.1
EffTox (6)	[**0.66**] 39.1 (14.8)	[0.63] 24.8 (12.8)	[0.61] 9.2 (5.8)	[0.58] 25 (11.3)	2
EffToxU (7)	[**0.61**] 40.6 (15.7)	[0.59] 21.4 (11.8)	[0.56] 9.7 (5.5)	[0.54] 26 (11.6)	2.1
		Scenario 7 (π_E, π_T)			
	(0.26, 0.05)	(0.6, 0.13)	(0.7, 0.23)	(0.7, 0.35)	
EffToxU (2)	[0.36] 0.9 (4.6)	[0.59] 11.9 (7.4)	[**0.61**] 27.9 (9.4)	[0.54] 58.6 (23.4)	0.8
EffTox (6)	[0.52] 0.9 (4.6)	[0.71] 12.7 (7.5)	[**0.73**] 27.4 (9.4)	[0.68] 58.4 (23.4)	0.7
EffToxU (7)	[0.46] 0.9 (4.7)	[0.66] 10.4 (6.8)	[**0.69**] 20.2 (7.8)	[0.64] 67.7 (25.5)	0.8
		Scenario 8 (π_E, π_T)			
	(0.26, 0.18)	(0.6, 0.35)	(0.7, 0.5)	(0.7, 0.62)	
EffToxU (2)	[0.32] 3.9 (6.3)	[**0.48**] 50.8 (14.4)	[0.46] 26.5 (10.6)	[0.39] 11.8 (12.3)	7
EffTox (6)	[0.49] 4.3 (6.6)	[**0.62**] 50.3 (14.4)	[0.60] 26 (10.4)	[0.54] 12 (12.4)	7.2
EffToxU (7)	[0.42] 4.7 (6.4)	[**0.57**] 42.8 (12.9)	[0.57] 30.8 (10.3)	[0.52] 14.6 (14.2)	7.2
		Scenario 9 (π_E, π_T)			
	(0.55, 0.45)	(0.75, 0.57)	(0.85, 0.64)	(0.9, 0.7)	
EffToxU (2)	[0.40] 29.4 (12.3)	[0.45] 13.2 (8.9)	[**0.45**] 2 (4.3)	[0.43] 2.9 (8.6)	52.5
EffTox (6)	[0.55] 30 (12.5)	[0.59] 13.2 (8.8)	[**0.60**] 1.6 (4.3)	[0.58] 3 (8.6)	52.1
EffToxU (7)	[0.50] 28.9 (11.9)	[0.57] 15 (9.1)	[**0.59**] 2 (4.5)	[0.59] 3.1 (8.8)	51
		Scenario 10 (π_E, π_T)			
	(0.2, 0.05)	(0.3, 0.08)	(0.38, 0.12)	(0.45, 0.15)	
EffToxU (2)	[0.31] 1.5 (6.1)	[0.38] 0.9 (3.7)	[0.43] 1.5 (3.6)	[**0.47**] 51.1 (21.7)	44.9
EffTox (6)	[0.49] 1.7 (6.1)	[0.54] 0.7 (3.7)	[0.58] 1.8 (3.6)	[**0.61**] 51.1 (21.7)	44.8
EffToxU (7)	[0.42] 1.4 (6.1)	[0.48] 0.8 (3.7)	[0.52] 1.8 (3.6)	[**0.56**] 51.2 (21.7)	44.9

5.7 Discussion

The R2DT method applies novel utility functions based upon reference dependence and attitudes to risk for both efficacy and toxicity attributes. Efficacy and toxicity utility are combined by consideration of payoffs and interaction effects to give a joint utility used to determine dosing at each stage. A method to elicit the values of the utility function by the consideration of lotteries has been proposed. The design has been compared to *EffToxU*, an established method, which has been shown to be a special case of R2DT, and found some initial evidence that the method could lead to considerable improvement in operating characteristics. A stopping rule based upon a contour from the utility function has been proposed which may closer resemble stopping rule preferences that consider both efficacy and toxicity simultaneously.

The R2DT design splits the specification of the utility function into separate attributes and whether levels of the attribute are a gain or a loss upon a reference. The approach is an example of a reference dependent utility function similar to prospect theory; a theory of behavioral economics and finance [148]. The two attributes are combined by consideration of how they interact according to utility independence axioms. The use of references in the marginal utility functions allows us to consider the joint utility from four distinct regions defined from the combinations of gain and loss for each of the attributes. In the region of a gain in both toxicity and efficacy the utility function takes a similar form to previously proposed designs. When one of the attributes is a loss, this attribute dominates the utility function and further improvement in the other attribute adds only very marginally to the overall utility function. The interaction component of the joint utility and the loss aversion parameters are predominantly responsible for this effect. The change in effect of one attribute conditional on the level of the other captures the motivating clinical example in Chapter 2. Where a steep contour was specified with the issue that doses with high efficacy and unacceptable toxicity were considered equivalent to doses with acceptable efficacy and toxicity profiles.

Admissibility criteria are typically necessary components of trial design in this setting to prevent unethical choices for patients. Step functions that require a level of evidence that one of the attributes is below or above a fixed probability before a dose is excluded are used to meet this aim. This is an important part of the design process, whereby the design will

push for higher doses until there is strong level of evidence for the dose to be excluded [41]. The EffTox design doesn't perform well when the admissibility rules are poorly specified. This was seen when applying the R2DT stopping rule to the EffTox utility design. The design process is a balance between exploring untried doses by defining steep contours and preventing unethical choices for patients with admissibility rules. The most extreme attitude to risk in each domain for R2DT ($\alpha_{GE} = \alpha_{LE} = \alpha_{GT} = \alpha_{LT} = 0$) creates a step function at the reference point. The functional form of the utility is now the same as the admissibility rules. Sigmoidal utilities are a more coherent and less ad-hoc way of expressing the basic type of preferences that people try to express when combining admissibility rules with risk and loss neutral utilities. The intention of the R2DT is to shift the focus toward specification of the utility function that reflects unethical choices for patients.

The contour stopping method appeared to be similarly effective to the more conventional stopping rules. The major difference was that the conventional stopping rules and the variations of the contours were declaring different things as acceptable or not. A difference with the contour is efficacy and toxicity are being considered at the same time. A threshold of acceptable toxicity depends upon the level of efficacy with more toxicity being acceptable for higher efficacy. We can see with the *R2DT 3(iii)* design particularly that the specified utility contour can still be quite flat, allowing approximately an additional 10% toxicity for a jump between 60% and 100% efficacy. The use of the single stopping rule is desirable as it simplifies the design process to specifying the utility function that reflects clinical preferences and one additional parameter to be tuned. There was a small difference between the trial stopping rule and the admissibility rule which excluded individual doses from the decisions. This highlights for the R2DT design that it is the utility function that is driving the decisions. From a theoretic point of view it is preferable to allow the utility function to dictate decisions, in practice however a safety review committee would be provided with all information in order to make a decision that does not harm patients.

There was some variability considered when trying to elicit the reference points. The elicitation procedure was designed to minimise any bias associated with determining a single value. The Bayesian decision theoretic way of thinking about uncertainty is to consider the parameter as something that is unknown. The alternative efficacy rate for example could be considered an unknown state of nature with an associated distribution. The parameter

could be updated through Bayes theorem as a control arm within the trial or be given a distribution as a prior that reflects the uncertainty. The R2DT utility function could have the same functional form but the reference parameter for the marginal efficacy rate would now be the uncertain parameter. The design has good operating characteristics and the formal elicitation process is not strictly necessary if we consider the design to have a set of components that need tuning through simulation and less formal clinical consultation; this is a similar approach to a practical implementation of the *EffTox* method [41]. An application of the method is needed for the R2DT design to establish whether the functional form of R2DT sufficiently captures clinical preferences. It does appear to capture the admissibility rules that are routinely specified.

The pilot study assessed the feasibility of the R2DT elicitation protocol, demonstrating that while the process is more involved than existing methods, it is achievable within a reasonable time frame and was well-received by the clinician involved. Training and practice are crucial for understanding, and further studies with multiple clinicians are recommended to capture the broader variability and ensure the utility function reflects the wider clinical community. Challenges in the loss domain highlight the need for careful framing of questions and avoiding terminology related to risk attitudes. Elicitation of the R2DT method needs a far greater understanding of utility theory and is more time consuming at the design stage than the EffToxU patient utility. The same utility function is also used by model assisted designs such as BOIN12 and U-BOIN detailed in Chapter 2.

The method of indifference used to elicit the joint utility function for R2DT in Section 5.4.3, is similar to what is elicited as part of the EffTox method. There are two key differences between what was described in the R2DT method. The first is that the suggested points were chosen to mimic choices that are similar to those made in clinical practice, with modest improvements in one of the attributes to reduce elicitation bias. This avoids more challenging ideas for the DM of considering treatments that work perfectly, or treatments that don't have any associated toxicity. The second is about the approach to this method as a whole; the elicitation of R2DT should reflect clinical beliefs and these beliefs should not be shifted to give good operating characteristics. The reason for this is that EffTox is a simplification and compromises in some regions in the domain of consequence are made so that contours are sufficiently steep when efficacy and toxicity are both acceptable. The more complex

structure of R2DT aims to be a far closer approximation of preference under uncertainty in all areas of the joint consequence space.

Chapter 6

Discussion

This thesis has proposed a novel Bayesian decision theoretic approach to dose finding based on joint toxicity and efficacy outcomes, and this has been investigated in 3 chapters. Chapter 3 looked at a component of the joint probability model associated with dose finding, Chapter 4, the statistical theory necessary to achieve the aim of the trial (to identify an optimal dose whilst treating patients optimally at each stage) and a novel Bayesian decision theoretic design, R2DT, was introduced in Chapter 5. An initial literature review of dose finding designs in oncology found that the definition of an optimal dose when using efficacy and toxicity endpoints differed between designs. A motivating example found that there can be conflict between clinical preferences and the ability of a statistical design to achieve the trial objectives. It was hoped that by using a scientifically more robust method in contrast to a more ad-hoc procedure that the objectives of a dose finding trial could be more closely met. A particular motivation for this was to better capture clinical preferences in order to make more informed choices for each cohort of patients. This discussion chapter starts by providing a summary of the conclusions from each of the previous chapters and how they contributed to the wider objectives of the thesis. The scope for the conclusions as well as limitations of the thesis are discussed in addition to recommendations for further work.

6.1 Thesis summary

There are two separate components necessary for a decision analysis; the probability model and a decision component. In the setting of joint outcomes to inform decision making, the

correlation component of the probability model was investigated through copula modeling. A copula is a multivariate distribution function that describes the dependence structure between efficacy and toxicity outcomes. A key recommendation from this work is that when investigating and contrasting correlation models in the setting, a consistent correlation measure is needed. It was demonstrated that the correlation parameter used as part of a copula model was specific to the model itself; it was proposed that Kendall's tau provides consistency across multiple models to enable comparison. The calculation of Kendall's tau is possible from different models and depends upon the probability of efficacy and the probability of toxicity at each dose. With a low number of patients treated in a dose finding trial the precision in estimating marginal probability distributions is limited, adding to the lack of precision in estimating correlation. The Farlie-Gumbel-Morgenstern copula that has previously been used in dose finding trial design [77] has a limited ability to measure stronger correlation owing to the structural form and conditions necessary to make it a copula.

Decision making in the setting of joint outcomes is primarily made with respect to the population parameters for the chance of efficacy and toxicity at each dose. This is achieved through the calculation of an expectation of a function where the two population parameters are the arguments. As such any differences between operating characteristics for a dose finding trial specifying an independent model and a copula model will be down to differences in the posterior distributions. The simulation study performed in Chapter 3 demonstrated that the inclusion of a copula model had minimal effect on each of the marginal posterior distributions. The independent model was therefore concluded to be a more parsimonious model in the setting. A similar conclusion has previously been proposed via a small simulation study [111]. This thesis adds to the existing literature by offering a theoretical exploration of copulas and their application to binary data; providing a broader analytical framework to offer insights into the wider applicability of copula models in the setting.

The review of the statistical methods for decision making in Chapter 4 defined a value function as a numerical scale to denote an ordering of preferences of an attribute. Considering the attributes of efficacy and toxicity independently in dose finding, there is a coherent ordering from the raw attribute but the difference between two levels of an attribute is not comparable. For example, the clinical interpretation of a 10% improvement in the probability of efficacy will depend upon the probability from which the 10% is an improvement.

A further issue is that decision making in a situation where there is uncertainty about the possible states of nature leads to different choices than when there is certainty. Optimizing the expected consequence function using an unsuitable scale can result in poor recommendations that wouldn't be taken in practice. To mitigate these challenges, employing a consistent scale, known as utility, resolves the issue by framing the problem in terms of lotteries between possible outcomes. The Von Neumann-Morgenstern axioms provide a foundation for rational decision-making under uncertainty. When following these axioms the expected utility can be used to decide between possible courses of action in a decision analysis.

In dose finding the attributes of efficacy and toxicity are described by the population parameters describing the chance of an event for a patient occurring. Both attributes are value functions on an ordinal scale. This feature is important when trying to combine the two attributes into an order of preference; ideas about payoffs for one attribute with respect to another are unlikely to be represented by a constant or a simple function. This is because the interpretation of an incremental increase in an attribute changes depending on the attribute level, as described above. This feature was apparent in the motivating example in Chapter 2.

The novel R2DT method proposed in Chapter 5 sought to account for decision making under uncertainty, that is most prominent early on in a dose finding trial, by using a consistent scale - utility. The method of defining a utility function with two attributes by direct-assessment doesn't exploit any features of the attributes (such as monotonicity). The requisite information is difficult to assess and the result is difficult to work with in calculating expected utility and further sensitivity analyses [110]. The approach described in Chapter 4 to overcome these difficulties was to inspect utility independence axioms and to split the utility function into more easily assessed separate univariate utility functions. A simple joint utility function was shown to result from following independence axioms [142], with two parameters to be set according to how each utility function interacts. The interaction component was considered to be positive interaction in the dose finding setting. The main component of the marginal utility functions in R2DT is reference dependence defining the merit of efficacy and toxicity attributes in relation to a reference point. The attitude to risk is determined by whether attributes are above or below this reference point. The R2DT

utility function was used to define a novel trial stopping rule that jointly accounted for levels of efficacy and toxicity. The proposed stopping rule appeared to be similarly effective to the more conventional stopping rules.

There was a detailed elicitation procedure proposed to capture the extensions to the proposed novel utility that were detailed in Chapter 5. The set of questions proposed in order to obtain the parameters of the R2DT utility function were designed to minimise biases to ensure that preferences were as accurate as possible and unintentional biases were not introduced. Questions posed were for simple lotteries that were designed to be reflective of choices made routinely in clinical practice (i.e around the reference). In doing so the the bias in the elicitation process is reduced in contrast to hypothetical choices such as perfect efficacy or zero toxicity associated with the EffTox method [25]. This procedure was demonstrated to be feasible in a small pilot study.

6.2 Limitations and future work

The stages of development for a new method have been likened to the four conventional stages of clinical trials [156]. This benchmark is helpful in understanding what evidence has been produced in support of a new method and what are the next stages necessary for a new method to be widely adopted. Methods in bivariate dose finding are not commonly used [41]. It is hoped that following this process can increase the potential uptake of R2DT. The R2DT trial design has been introduced from a theoretical perspective (phase I) and then applied to a small simulation study to provide proof of concept (phase II). Further work is needed to better characterise the design in a wider set of scenarios and settings and for the method to be contrasted against further designs to understand its limitations better.

6.2.1 Probability model

A major component of further work is to better understand the role of the probability model. The R2DT design was specified using the EffTox probability model. The EffTox probability model was initially calibrated using the accompanying software to give good operating characteristics for the simulation study, assuming the EffTox decision function. There is some evidence (Table 5.4, Scenario 1 & Scenario 7) that both the R2DT and EffToxU methods have a tendency to favour higher doses and that this is influenced by the

probability model. This particular feature would need further simulations to understand if R2DT was to be implemented and justified in a trial protocol. The EffTox approach justifies and calibrates the model by the presence of good operating characteristics. A probability model may need to be mildly informative or skewed towards preferences for higher doses to ensure that it doesn't get stuck at lower doses, whilst not producing undesirable safety characteristics [78].

The probability model in a Bayesian decision theoretic approach should capture the uncertainty of the situation and be based upon plausible facts, science, expert judgement and available data [25]. A model that is used as an algorithm is different to one that tries to capture the inherent uncertainty in any given situation. This point was made in Chapter 2 when contrasting the choice of one parameter or two parameter logistic regression models for toxicity only designs. Given the increased complexity of the R2DT utility function it is possible that a probability model that closer resembles the situation, or is more flexible, could yield improved operating characteristics. There are two components to this: the functional form of the probability model, and the specification of priors. Better understanding of the role of the probability model and its impact upon operating characteristics of R2DT could be the topic of further work.

There is an increased use of model assisted designs in the literature not only due to the relative simplicity of implementation but also because of superior operating characteristics [83]. Designs such as U-BOIN and BOIN12, detailed in Chapter 2, specify the 4-outcome utility function to make decisions. This utility function was used as the main comparator in this thesis, but the more complex EffTox logistic regression probability model was specified. Further work would be needed to see if the advantages of using model assisted approaches are further extended by the use of R2DT utility function. The motivation for using a utility function that better captures the clinical situation, as is the aim of R2DT, remains pertinent.

The R2DT utility function is based upon outcome probabilities which are frequently used in phase I-II designs. The Bayesian decision theoretic approach however separates the probability and decision components [25], making R2DT method applicable in a wide variety of settings. Changing the component of the probability model, such as a plateau ([35]), for example, to better model the modern drug paradigm requires further evaluation, without a change to the utility function. R2DT uses utility independence to join two separate marginal

utility functions. The individual utility functions could be adapted to accommodate different endpoints. Marginal, continuous and time to event outcomes could be transformed into the utility scale [157]. Ordinal outcomes for toxicity or efficacy would first need combining into a single measure or utility function. The methods to do achieve this can be adapted from the methods presented in Chapter 4. This further work is important because the effect of a targeted treatment is less likely to be accurately reflected over a short period of time [15].

6.2.2 Search strategy

The literature review in Chapter 2 was not a systematic review with pre-specified search terms that would have enabled reproducible results. As such it is not an exhaustive review of trial designs and did not consider every design in the setting. A systematic review can give more reassurance in the conclusions in that the literature is considered in its entirety rather than a potentially biased search procedure conducted to support preconceived conclusions. A "snowballing" [158] (alternatively called "pearl growing" [159]) approach was adopted to to identify the key work in the field. Specifically the Von Neumann–Morgenstern utility approach inspected the references within the decision theoretic designs as a starting point [56, 85, 86, 82, 78, 88, 89]. The copula approach inspected the references for the copula simulation study initially identified through a search of Google Scholar. Given the similarity in the sigmoidal shaped value function in prospect theory and its prominence in the economic literature, a similar technique found no references for the application of prospect theory relating to dosing in clinical trials.

6.2.3 Copulas when correlation is an attribute

The simulation study in the thesis for copulas didn't investigate the effect on operating characteristics of the copula when correlation was a component of the decision model. Instead, conclusions were restricted to designs where efficacy and toxicity were the sole attributes in decision making. The EffTox utility design is a design that gives a utility to each of the separate possible outcomes for a patient. The probabilities responsible for each of these events are defined using the marginal probabilities and a parameter for the correlation. It was shown that the calculation for expectation of the utility function depended upon the correlation parameter. Graphically however when specifying an FGM copula with the EffTox utility design, the marginal probabilities of efficacy and toxicity tend to dominate

the correlation component in the consequence function. Correlation was also shown to be challenging to estimate in the setting to diminish any effect of correlation in decision making further. The role of correlation as an attribute, and incorporation into the R2DT utility function hasn't been considered. Copulas are also able to measure the correlation between binary endpoints and a continuous endpoint for example. The copula work presented in this thesis could be adapted to such a setting.

6.2.4 Exploration-exploitation in decision making

Following a strategy of choosing the optimal dose at each stage introduces a possibility of getting stuck at a sub-optimal dose and not finding the global optimal [160, 161]. This is known as the exploration-exploitation dilemma in decision making. Changing the balance away from exclusively optimising at each stage, exploration, may lead to finding the OD more often at the end of the trial. This needs careful consideration at the trial design stage as to whether it is ethical. A solution is to introduce some randomisation into the process. A number of authors have suggested Thompson sampling [162] as a solution [82, 35]. When there is little information available at the start of the trial in-particular a randomised approach could be seen as unethical. An approach taken by some authors is to define two stages, with an initial period of staged dose finding (with or without an efficacy component) before continuing with randomisation [95]. Linking this to different endpoints, if the endpoint isn't measurable over a shorter period of time, there won't be information to adapt at each stage and a randomised approach may be acceptable. There are also designs that account for partial information that are highlighted in Chapter 2.

R2DT assumes the same utility function at each of the multiple decision points. The decision maker reference point(s) may change over time due to adaptation and shifting expectations. This could be linked to the slight shift in prioritising the objectives of the trial. An alternative approach would be to incorporate gain of information to change the balance between the objectives of determining an optimal dose and treating patients optimally at each decision point. In some contexts, the dose finding trial will be one trial within a larger clinical development pathway. The merit of continuing the trial to determine more information or deciding to initiate the next stage of clinical development isn't specified as one of the potential actions. The benefit of using the utility scale is that this action could be given a utility and considered as an add on to R2DT. An alternative approach could be a

dynamic-programming method which tries to make each decision optimally, considering all the subsequent outcomes and decisions that will come afterwards [163]. There are increasing calls for the next study in the clinical development pathway of an oncology agent to be a separate randomised dose ranging study [164, 165]. The objective of the dose finding study in this paradigm is to suggest a range of doses suitable for further study. The proposed stopping rule given in Equation 5.20 could be used as part of the R2DT design to achieve this aim. Individual doses studied as part of the dose finding trial are tested at the end of the trial to see if they are acceptable for further study.

6.2.5 Computer simulation studies

The evaluation of R2DT and the copula chapter featured simulation studies that were used to inform conclusions. It is important that a simulation study is well conducted in order to give confidence in results. There are a number of recent guidance publications around good practices for simulation studies [166, 39]. The simulation study did not explicitly follow the guidance presented in these publications; rather an ad hoc approach based upon good programming practices, reproducibility and scientific method. Retrospectively assessing what was implemented against the guidance would suggest that the simulation study was of high quality. The aims of each simulation study and metrics to assess performance were stated. The data generated as part of each simulation study was created in a separate program with a pre-specified seed. Additionally the assessment of different designs used the same large dataset of generated data. The main metrics to assess design operating characteristics were pre-specified and presented both in tables and figures. A seed for the Bayesian model fitting was not specified. This was an oversight that has been corrected in the code (Appendix B) but the simulation study has not been rerun in its entirety. Given the high number of draws from the posterior distribution and that different methods have the same model fitting and draws (Appendix A) the effect is likely to be negligible if rerun, with the issue being reproducibility. Further work in the setting could be to better define Monte Carlo errors associated with the performance measures and better understand the number of simulated trial replicates. The number of specified repetitions, two thousand in the case of the R2DT design,was specified to be a compromise between computational efficiency and accuracy. A factor surrounding the uptake of novel designs is the availability of user friendly software [47]. The R code associated with the design has been submitted

as part of this thesis (Appendix B). Future work to put the code into a R package with accompanying documentation would increase the chance of the work being implemented.

6.2.6 Performance measures

There is a potential need for further work in assessing the performance metrics of a phase I-II design. The performance of the R2DT design was assessed by inspecting the operating characteristics in contrast to an alternative design. Assessing the performance of the design at each decision point for patient benefit doesn't have an easy metric by which to compare different designs. This was assessed in this thesis by studying operating characteristics longitudinally. It was possible to observe that one design had a tendency to make different decisions earlier in the trial. This was seen in the escalation with overdose control design with a toxicity endpoint reviewed as part of the literature review. The design defined a parameter that controlled how conservative escalation was. Defining exactly what is an acceptable decision in the presence of uncertainty at a particular decision point is linked to a utility function that is capable of reflecting clinical attitudes. This highlights the intention of the decision theoretic method. This relies upon accurate elicitation that is expanded upon in the next section.

The probability of correct selection is one of the main metrics used to assess whether a design meets its objectives. What constitutes the optimal dose is used as part of this metric to define what is "correct". When comparing two different designs with different ways of defining the optimal dose it is possible to create a scenario where the definition of which dose is "correct" and the optimal dose differs between the designs. In this case, which design describes the correct dose? This judgement is clinical and highlights the importance of using a consequence function that captures the situation. The overall assessment across multiple different scenarios for the R2DT was subjective. This brings up the question as to how to compare performance across different scenarios. For some scenarios it may be more difficult for a design to find the OD because two doses are similarly optimal in contrast to another scenario where there is a clear OD. In the phase I setting with a single toxicity attribute it has been proposed that the toxicity odds ratio could be used as a measure of the difficulty of a scenario [54]. Here the odds of toxicity at a dose are contrasted with the target toxicity level that defines the optimal dose. If two doses have very similar toxicity close to the target toxicity level then the odds ratio would be close to 1 and the design would

struggle to differentiate between two doses in terms of correct selection. This idea doesn't naturally extend to the setting with bivariate outcomes because there isn't an agreed upon benchmark against which to contrast.

Utility is specified on a scale that measures strength of preference, further work could be to use this standard, defined at part of R2DT method, to compare operating characteristics. A difficult scenario would be one where the highest utility (at a dose) is similar to another and an easier scenario is one where there is a larger jump in utility between the two doses with highest utility. If the utility function now represents a proportional measure of clinical preference between doses this could be adapted to generate and assess scenarios. The proposal depends upon the ability to accurately capture a utility function, which stresses the importance of the elicitation methods.

6.2.7 Elicitation

Part of the motivation for the R2DT method is to more closely capture clinical preferences. The stage of method development for the elicitation component of this thesis similarly has been introduced from a theoretical perspective with a set of questions to elicit each parameter of R2DT (phase I) in addition to a pilot study to demonstrate proof of principle and feasibility (phase II). The functional form of the utility is multifaceted and has been proposed from principles rather than specific clinical input. While the form is flexible to accommodate different preferences, further work is needed for the R2DT design to establish whether it fully captures clinical preferences. Particular features of this are the reference dependence and the piece-wise power model.

Reducing bias in eliciting the parameters in the utility function and ensuring that what is elicited reflects preferences is key. There are a number of factors that are known to induce measurement error from the behavioural sciences including how the question is asked [138]. Elicitation methods for this aim using probability lottery equivalencies are well established [139]. The impact of the reference point in the elicitation of utilities in the healthcare setting has also been investigated [137]. Elicitation in the more established setting of probabilistic judgements such as the Sheffield elicitation framework provides a far more comprehensive package in order to obtain and document multiple expert opinions [167]. Further work could be writing and understanding an elicitation protocol in the setting. This thesis referred to

a single unitary decision maker, which in practice means that there is one utility function that describes preferences. In dose finding there are multiple stakeholders interested in any given decision. A rational impartial observer (RIO) was suggested as a possibility to overcome difficulties in obtaining preferences for multiple experts [127]. This idea wasn't fully developed and further work is needed for an increased understanding of how each of the different stake holders are impacted and any differences in opinion.

6.2.8 Patient engagement

A particularly important stakeholder is the patient. There is a strong ethical argument in favour of involving patients in research, with the expectation that their engagement will improve research by better addressing their concerns [168]. With complex clinical trial design it is important that patient perspectives are not just sought but add value [169]. One of the aims of dose finding that the novel design aims to better account for is dosing optimally at each stage. Each patient entering the trial should receive a dose believed to be optimal and this should account for the lack of available evidence so that the risk-benefit is acceptable. Demonstrating with patient input that this is a key objective that needs further consideration can only enhance the importance of the R2DT design. Concepts such as patient risk-benefit are challenging to articulate and quantify; it is hoped that the work to elicit preference in this thesis could be adapted to help to better collaborate with patients in clinical trial design.

6.3 Conclusion

My research has looked at a decision theoretic approach to dose finding based on joint outcomes of toxicity and efficacy, and demonstrated its potential use in the oncology setting. Copula models had previously been proposed in the setting to account for correlation between the binary outcomes of efficacy and toxicity without exploration as to whether they achieved this aim and to what extent. I demonstrated using an analytical approach that an independent model is more parsimonious in this setting. The analytical theory that I presented explained why a consistent correlation measure was needed in order to compare different copula models and that this parameter couldn't be estimated accurately in small sample sizes associated with dose finding. My work has the potential to simplify the de-

sign process by reassuring researchers that an independent model sufficiently captures the situation. My proposed R2DT method is a novel method of defining the utility function as part of a decision theoretic approach to dose finding in oncology. My approach offers a design for researchers aiming to identify an optimal dose of treatment where there is a need to optimise and control for the risk-benefit ratio of patients entering the study. The key to this was specifying a decision process that closer reflected preferences for the situation, in particular using a measure that accounted for how preferences may change when faced with an uncertain prospect. I have achieved this with respect to Von Neumann–Morgenstern utility theory defined in terms of preferences for lotteries. Expressing utilities in this manner is a more coherent and less ad-hoc way of expressing the basic type of preferences that researchers try to express when combining admissibility rules with simpler objective functions to define the optimal dose. My design has been compared to an established method (EffToxU), which has been shown to be a special case of R2DT. I have demonstrated that my novel method could lead to considerable improvement in operating characteristics in the setting of dose finding with joint outcomes. I have also demonstrated that my novel method can be used to generate a trial stopping rule that incorporates both efficacy and toxicity measures and that this could be more reflective of clinical preferences. The work applies a broad framework to give insight to existing methods and potential to adapt to different endpoints and trial features. I proposed an elicitation process in order to bridge the gap between a theoretical model and considered process to captures these preferences. The elicitation process was demonstrated to be feasible in a small pilot study. My method aligns with some of the goals of the wider FDA dose finding and dose optimisation initiative of project Optimus [16] by using a strategy of dose selection that not only limits the toxicity but contrasts this with the efficacy of a treatment. In the context of the clinical trial development pathway, my method has the potential to better identify optimal doses of treatment in the initial stages of the pathway, while meeting the treatment goals of patients entering the study; ultimately enabling more reliable dose finding to take forward into later stages of definitive evidence generation.

Bibliography

[1] Elisabeth Mahase. Cancer overtakes CVD to become leading cause of death in high income countries. *BMJ*, page l5368, sep 2019.

[2] Dejene Tolossa Debela, Seke GY Muzazu, Kidist Digamo Heraro, Maureen Tayamika Ndalama, Betelhiem Woldemedhin Mesele, Dagimawi Chilot Haile, Sophia Khalayi Kitui, and Tsegahun Manyazewal. New approaches and procedures for cancer treatment: Current perspectives. *SAGE Open Medicine*, 9:205031212110343, jan 2021.

[3] Michael J Grayling, Munyaradzi Dimairo, Adrian P Mander, and Thomas F Jaki. A review of perspectives on the use of randomization in phase ii oncology trials. *JNCI: Journal of the National Cancer Institute*, 111(12):1255–1262, June 2019.

[4] ICH Harmonised Tripartite Guideline. General considerations for clinical trials e8. In *International Conference on Harmonization of Technical Requirements for Registration of Pharmaceuticals for Human Use*, 1997.

[5] J. A. DiMasi, H. G. Grabowski, and R. W. Hansen. Innovation in the pharmaceutical industry: New estimates of r&d costs. *J Health Econ*, 47:20–33, 2016.

[6] Olivier J. Wouters, Martin McKee, and Jeroen Luyten. Estimated research and development investment needed to bring a new medicine to market, 2009-2018. *JAMA*, 323(9):844, March 2020.

[7] V. Prasad and S. Mailankody. Research and development spending to bring a single cancer drug to market and revenues after approval. *JAMA Internal Medicine*, 177(11):1569–1575, 2017.

[8] FDA. Fast track, breakthrough therapy, accelerated approval, priority review. `https://www.fda.gov/about-fda/oncology-center-excellence/project-optimus`,

2023. [Accessed 24-Jul-2023].

[9] Yojana Gadiya, Philip Gribbon, Martin Hofmann-Apitius, and Andrea Zaliani. Pharmaceutical patent landscaping: A novel approach to understand patents from the drug discovery perspective. *Artificial Intelligence in the Life Sciences*, 3:100069, dec 2023.

[10] Asher Mullard. The high, and redundant, cost of failure in cancer drug development. *Nature Reviews Drug Discovery*, 22(9):688–688, August 2023.

[11] Chi Heem Wong, Kien Wei Siah, and Andrew W Lo. Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2):273–286, January 2018.

[12] T. J. Hwang, D. Carpenter, J. C. Lauffenburger, B. Wang, J. M. Franklin, and A. S. Kesselheim. Failure of investigational drugs in late-stage clinical development and publication of trial results. *JAMA Internal Medicine*, 176(12):1826–1833, 2016.

[13] Richard K. Harrison. Phase ii and phase iii failures: 2013–2015. *Nature Reviews Drug Discovery*, 15(12):817–818, November 2016.

[14] Khanh Do, Chris H. Takimoto, and Shivaani Kummar. Changing landscape of early phase clinical trials. In *Novel Designs of Early Phase Trials for Cancer Therapeutics*, pages 1–6. Elsevier, 2018.

[15] Jeanne Fourie Zirkelbach, Mirat Shah, Jonathon Vallejo, Joyce Cheng, Amal Ayyoub, Jiang Liu, Rachel Hudson, Rajeshwari Sridhara, Gwynn Ison, Laleh Amiri-Kordestani, Shenghui Tang, Thomas Gwise, Atiqur Rahman, Richard Pazdur, and Marc R. Theoret. Improving dose-optimization processes used in oncology drug development to minimize toxicity and maximize benefit to patients. *Journal of Clinical Oncology*, 40(30):3489–3500, October 2022.

[16] FDA. Project Optimus — fda.gov. `https://www.fda.gov/about-fda/oncology-center-excellence/project-optimus`, 2023. [Accessed 24-Jul-2023].

[17] M. J. R. Healy and Richard Simon. New methodology in clinical trials. *Biometrics*, 34(4):709, dec 1978.

[18] Elizabeth A. Eisenhauer, Christopher Twelves, and Marc Buyse, editors. *Phase I Cancer Clinical Trials*. Oxford University Press, mar 2015.

[19] Philip Pallmann, Alun W. Bedding, Babak Choodari-Oskooei, Munyaradzi Dimairo, Laura Flight, Lisa V. Hampson, Jane Holmes, Adrian P. Mander, Lang'o Odondi, Matthew R. Sydes, Sofía S. Villar, James M. S. Wason, Christopher J. Weir, Graham M. Wheeler, Christina Yap, and Thomas Jaki. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC Medicine*, 16(1), feb 2018.

[20] Gail A. Van Norman. Phase ii trials in drug development and adaptive trial design. *JACC: Basic to Translational Science*, 4(3):428–437, June 2019.

[21] Branimir K. Hackenberger. Bayes or not bayes, is this the question? *Croatian Medical Journal*, 60(1):50–52, February 2019.

[22] Deborah Ashby. Bayesian statistics in medicine: a 25 year review: Bayesian statistics in medicine. *Statistics in Medicine*, 25(21):3589–3631, August 2006.

[23] K.R.A.J.P.M. David J. Spiegelhalter, D.J. Spiegelhalter, K.R. Abrams, and J.P. Myles. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. Statistics in Practice. Wiley, 2004.

[24] Scott M. Berry. *Bayesian adaptive methods for clinical trials*. Chapman & Hall/CRC, 2011.

[25] Jim Q Smith. *Bayesian decision analysis: principles and practice*. Cambridge University Press, 2010.

[26] Tim Bedford and Roger Cooke. *Probabilistic Risk Analysis*. Cambridge University Press, apr 2001.

[27] S. French and D.R. Insua. *Statistical Decision Theory: Kendall's Library of Statistics 9*. Wiley, 2000.

[28] C. L. Tourneau, J. J. Lee, and L. L. Siu. Dose escalation methods in phase i cancer clinical trials. *J Natl Cancer I*, 101, 2009.

[29] National Cancer Insttitute. Common terminology criteria for adverse events (ctcae). `https://ctep.cancer.gov/protocoldevelopment/electronic_applications/ctc.htm`, 2023. [Accessed 25-Jul-2023].

[30] H E SKIPPER, F M SCHABEL, and W S WILCOX. Experimental evaluation of potential anticancer agents. xiii. on the criteria and kinetics associated with "curability" of experimental leukemia. *Cancer chemotherapy reports*, 35:1–111, February 1964.

[31] W M Hryniuk. More is better. *Journal of Clinical Oncology*, 6(9):1365–1367, sep 1988.

[32] Siddhartha Mukherjee. *The Emperor of All Maladies*. Harper Collins Publ. UK, 2011.

[33] Pattanaik Smita, Patil Amol Narayan, Kumaravel J, and Prakash Gaurav. Therapeutic drug monitoring for cytotoxic anticancer drugs: Principles and evidence-based practices. *Frontiers in Oncology*, 12, dec 2022.

[34] Kristian Brock, Victoria Homer, Gurjinder Soul, Claire Potter, Cody Chiuzan, and Shing Lee. Is more better? an analysis of toxicity and response outcomes from dose-finding clinical trials in cancer. *BMC Cancer*, August 2020.

[35] Marie-Karelle Riviere, Ying Yuan, Jacques-Henri Jourdan, Frédéric Dubois, and Sarah Zohar. Phase i/II dose-finding design for molecularly targeted agent: Plateau determination using adaptive randomization. *Statistical Methods in Medical Research*, 27(2):466–479, mar 2016.

[36] Ahmedin Jemal, Freddie Bray, Melissa M. Center, Jacques Ferlay, Elizabeth Ward, and David Forman. Global cancer statistics. *CA: A Cancer Journal for Clinicians*, 61(2):69–90, feb 2011.

[37] S.D. Durham, N. Flournoy, and W. Li. A sequential design for maximizing the probability of a favourable response. *Canadian Journal of Statistics*, 26(3):479–495, sep 1998.

[38] J. O'Quigley, M. Pepe, and L. Fisher. Continual reassessment method: a practical design for phase i clinical trials in cancer. *Biometrics*, 46, 1990.

[39] Tim P. Morris, Ian R. White, and Michael J. Crowther. Using simulation studies to evaluate statistical methods. *Statistics in Medicine*, 38(11):2074–2102, January 2019.

[40] Andrew P. Grieve. Idle thoughts of a 'well-calibrated' bayesian in clinical drug development. *Pharmaceutical Statistics*, 15(2):96–108, January 2016.

[41] Kristian Brock, Lucinda Billingham, Mhairi Copland, Shamyla Siddique, Mirjana Sirovica, and Christina Yap. Implementing the EffTox dose-finding design in the matchpoint trial. *BMC Medical Research Methodology*, 17(1), July 2017.

[42] Nolan A. Wages, Bethany Jablonski Horton, Mark R. Conaway, and Gina R. Petroni. Operating characteristics are needed to properly evaluate the scientific validity of phase i protocols. *Contemporary Clinical Trials*, 108:106517, September 2021.

[43] Oleksandr Sverdlov, Weng Kee Wong, and Yevgen Ryeznik. Adaptive clinical trial designs for phase i cancer studies. *Statistics Surveys*, 8(0):2–44, 2014.

[44] Stephen K Carter. The phase i study. *Fundamentals of Cancer Chemotherapy. McGraw-Hill: New York, NY, USA, pp xv*, 527, 1987.

[45] Zhengjia Chen, Mark D. Krailo, Junfeng Sun, and Stanley P. Azen. Range and trend of expected toxicity level (etl) in standard a+b designs: A report from the children's oncology group. *Contemporary Clinical Trials*, 30(2):123–128, mar 2009.

[46] Y. Lin. Statistical properties of the traditional algorithm-based designs for phase i cancer clinical trials. *Biostatistics*, 2(2):203–215, jun 2001.

[47] Sharon B Love, Sarah Brown, Christopher J Weir, Chris Harbron, Christina Yap, Birgit Gaschler-Markefski, James Matcham, Louise Caffrey, Christopher McKevitt, Sally Clive, Charlie Craddock, James Spicer, and Victoria Cornelius. Embracing model-based designs for dose-finding trials. *British Journal of Cancer*, 117(3):332–339, jun 2017.

[48] Douglas Faries. Practical modifications of the continual reassessment method for phase i cancer clinical trials. *Journal of Biopharmaceutical Statistics*, 4(2):147–164, jan 1994.

[49] Steven N. Goodman, Marianna L. Zahurak, and Steven Piantadosi. Some practical improvements in the continual reassessment method for phase i studies. *Statistics in Medicine*, 14(11):1149–1161, jun 1995.

[50] J OQuigley. A stopping rule for the continual reassessment method. *Biometrika*, 85(3):741–748, sep 1998.

[51] S Zohar and S Chevret. The continual reassessment method: comparison of bayesian stopping rules for dose-ranging studies. *Statistics in medicine*, 20:2827–2843, October 2001.

[52] J. O'Quigley. Continual reassessment designs with early termination. *Biostatistics*, 3(1):87–99, mar 2002.

[53] John O'Quigley and Larry Z. Shen. Continual reassessment method: A likelihood approach. *Biometrics*, 52(2):673, jun 1996.

[54] Ying Kuen Cheung. *Dose Finding by the Continual Reassessment Method (Chapman & Hall/CRC Biostatistics Series Book 41)*. Chapman and Hall/CRC, 2011.

[55] Alexia Iasonos, Andrew S Wilton, Elyn R Riedel, Venkatraman E Seshan, and David R Spriggs. A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in phase i dose-finding studies. *Clinical Trials*, 5(5):465–477, September 2008.

[56] John Whitehead and David Williamson. Bayesian decision procedures based on logistic regression models for dose-finding studies. *Journal of Biopharmaceutical Statistics*, 8(3):445–467, jan 1998.

[57] John Whitehead and Hazel Brunier. BAYESIAN DECISION PROCEDURES FOR DOSE DETERMINING EXPERIMENTS. *Statistics in Medicine*, 14(9):885–893, may 1995.

[58] Beat Neuenschwander, Michael Branson, and Thomas Gsponer. Critical aspects of the bayesian approach to phase i cancer trials. *Statistics in Medicine*, 27(13):2420–2439, 2008.

[59] James Babb, André Rogatko, and Shelemyahu Zacks. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine*, 17(10):1103–1120, may 1998.

[60] X. Paoletti and A. Kramar. A comparison of model choices for the continual reassessment method in phase i cancer trials. *Statistics in Medicine*, 28(24):3012–3028, August 2009.

[61] Ruitao Lin and Ying Yuan. On the relative efficiency of model-assisted designs: a conditional approach. *Journal of Biopharmaceutical Statistics*, 29(4):648–662, jun 2019.

[62] Wentian Guo, Sue-Jane Wang, Shengjie Yang, Henry Lynn, and Yuan Ji. A bayesian interval dose-finding design addressingOckham's razor: mTPI-2. *Contemporary Clinical Trials*, 58:23–33, jul 2017.

[63] Fangrong Yan, Sumithra J. Mandrekar, and Ying Yuan. Keyboard: A novel bayesian toxicity probability interval design for phase i clinical trials. *Clinical Cancer Research*, 23(15):3994–4003, aug 2017.

[64] Denis Heng-Yan Leung and You-Gan Wang. Isotonic designs for phase i trials. *Controlled Clinical Trials*, 22(2):126–138, feb 2001.

[65] Suyu Liu and Ying Yuan. Bayesian optimal interval designs for phase i clinical trials. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(3):507–523, 2015.

[66] Ying Yuan, J. Jack Lee, and Susan G. Hilsenbeck. Model-assisted designs for early-phase clinical trials: Simplicity meets superiority. *JCO Precision Oncology*, pages 1–12, December 2019.

[67] Revathi Ananthakrishnan, Ruitao Lin, Chunsheng He, Yanping Chen, Daniel Li, and Michael LaValley. An overview of the BOIN design and its current extensions for novel early-phase oncology trials. *Contemporary Clinical Trials Communications*, 28:100943, aug 2022.

[68] Y. Yuan, K. R. Hess, S. G. Hilsenbeck, and M. R. Gilbert. Bayesian optimal interval design: A simple and well-performing design for phase i oncology trials. *Clinical Cancer Research*, 22(17):4291–4301, jul 2016.

[69] Bethany Jablonski Horton, Nolan A. Wages, and Mark R. Conaway. Performance of toxicity probability interval based designs in contrast to the continual reassessment method. *Statistics in Medicine*, 36(2):291–300, jul 2016.

[70] Mirat Shah, Atiqur Rahman, Marc R. Theoret, and Richard Pazdur. The drug-dosing conundrum in oncology — when less is more. *New England Journal of Medicine*,

385(16):1445–1447, oct 2021.

[71] Ted A. Gooley, Paul J. Martin, Lloyd D. Fisher, and Mary Pettinger. Simulation as a design tool for phase i/II clinical trials: An example from bone marrow transplantation. *Controlled Clinical Trials*, 15(6):450–462, dec 1994.

[72] S. Hunsberger, L. V. Rubinstein, J. Dancey, and E. L. Korn. Dose escalation trial designs based on a molecularly targeted endpoint. *Stat Med*, 24(14):2171–81, 2005.

[73] Anastasia Ivanova. A new dose-finding design for bivariate outcomes. *Biometrics*, 59:1001–1007, December 2003.

[74] Janis Hardwick, Mary C. Meyer, and Quentin F. Stout. Directed walk designs for dose-response problems with competing failure modes. *Biometrics*, 59(2):229–236, jun 2003.

[75] J O'Quigley, M D Hughes, and T Fenton. Dose-finding designs for hiv studies. *Biometrics*, 57(4):1018–1029, December 2001.

[76] Thomas M Braun. The bivariate continual reassessment method. extending the CRM to phase I trials of two competing outcomes. *Controlled clinical trials*, 23:240–256, June 2002.

[77] Peter F. Thall and John D. Cook. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*, 60(3):684–693, aug 2004.

[78] Ying Yuan, Hoang Q. Nguyen, and Peter F. Thall. *Bayesian Designs for Phase I-II Clinical Trials (Chapman & Hall/CRC Biostatistics Series)*. Chapman and Hall/CRC, 2016.

[79] D. MORGENSTERN. Einfache beispiele zweidimensionaler verteilungen. *Mitt, Math, Statist.*, 8:234–235, 1956.

[80] Roger B. Nelsen. *An Introduction to Copulas*. Springer New York, June 2007.

[81] John D Cook. Efficacy-toxicity trade-offs based on lp norms: Technical report utmdabtr-003-06. Technical report, Technical report, 2006.

[82] Peter F. Thall and Hoang Q. Nguyen. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics*, 22(4):785–801, may 2012. PMID: 22651115.

[83] Jingyi Zhang, Nolan A. Wages, and Ruitao Lin. Sfu: Surface-free utility-based design for dose optimization in cancer drug combination trials. *Statistics in Biosciences*, April 2024.

[84] Suyu Liu and Valen E. Johnson. A robust bayesian dose-finding design for phase i/II clinical trials. *Biostatistics*, 17(2):249–263, oct 2015.

[85] John Whitehead, Yinghui Zhou, John Stevens, and Graham Blakey. An evaluation of a bayesian method of dose escalation based on bivariate binary responses. *Journal of Biopharmaceutical Statistics*, 14(4):969–983, dec 2004.

[86] John Whitehead, Yinghui Zhou, John Stevens, Graham Blakey, Jim Price, and Joanna Leadbetter. Bayesian decision procedures for dose-escalation based on evidence of undesirable events and therapeutic benefit. *Statistics in Medicine*, 25(1):37–53, 2005.

[87] Shenghua K. Fan and You-Gan Wang. Decision-theoretic designs for dose-finding clinical trials with multiple outcomes. *Statistics in Medicine*, 25(10):1699–1714, 2006.

[88] Yee-Chong Loke, Say-Beng Tan, YiYu Cai, and David Machin. A bayesian dose finding design for dual endpoint phase i trials. *Statistics in Medicine*, 25(1):3–22, 2005.

[89] Meihua Wang and Roger Day. Adaptive bayesian design for phase i dose-finding trials using a joint model of response and toxicity. *Journal of Biopharmaceutical Statistics*, 20(1):125–144, dec 2009.

[90] Daniel H. Li, James B. Whitmore, Wentian Guo, and Yuan Ji. Toxicity and efficacy probability interval design for phase i adoptive cell therapy dose-finding clinical trials. *Clinical Cancer Research*, 23(1):13–20, jan 2017.

[91] Pin Li, Rachael Liu, Jianchang Lin, and Yuan Ji. Tepi-2 and ubi: designs for optimal immuno-oncology and cell therapy dose finding with toxicity and efficacy. *Journal of Biopharmaceutical Statistics*, 30(6):979–992, September 2020.

[92] Haolun Shi, Jiguo Cao, Ying Yuan, and Ruitao Lin. utpi: A utility-based toxicity probability interval design for phase i/ii dose-finding trials. *Statistics in Medicine*, 40(11):2626–2649, March 2021.

[93] Kentaro Takeda, Masataka Taguri, and Satoshi Morita. BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharmaceutical Statistics*, 17(4):383–395, apr 2018.

[94] Ruitao Lin and Guosheng Yin. STEIN: A simple toxicity and efficacy interval design for seamless phase i/II clinical trials. *Statistics in Medicine*, 36(26):4106–4120, aug 2017.

[95] Yanhong Zhou, J. Jack Lee, and Ying Yuan. A utility-based bayesian optimal interval (u-BOIN) phase i/II design to identify the optimal biological dose for targeted and immune therapies. *Statistics in Medicine*, 38(28), oct 2019.

[96] Ruitao Lin, Yanhong Zhou, Fangrong Yan, Daniel Li, and Ying Yuan. BOIN12: Bayesian optimal interval phase i/II trial design for utility-based dose finding in immunotherapy and targeted therapies, nov 2020.

[97] World Medical Association. World medical association declaration of helsinki: ethical principles for medical research involving human subjects. *JAMA*, 310:2191–2194, November 2013.

[98] Vladimir Dragalin and Valerii Fedorov. Adaptive designs for dose-finding based on efficacy–toxicity response. *Journal of Statistical Planning and Inference*, 136(6):1800–1823, jun 2006.

[99] B. Nebiyou Bekele and Yu Shen. A bayesian approach to jointly modeling toxicity and biomarker expression in a phase i/II dose-finding trial. *Biometrics*, 61(2):343–354, jun 2005.

[100] Peter F. Thall and Kathy E. Russell. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase i/II clinical trials. *Biometrics*, 54(1):251, mar 1998.

[101] Ying Kuen Cheung and Rick Chappell. Sequential designs for phase i clinical trials with late-onset toxicities. *Biometrics*, 56(4):1177–1182, dec 2000.

[102] Ick Hoon Jin, Suyu Liu, Peter F. Thall, and Ying Yuan. Using data augmentation to facilitate conduct of phase i–II clinical trials with delayed outcomes. *Journal of the American Statistical Association*, 109(506):525–536, apr 2014.

[103] Peter F. Thall, Hoang Q. Nguyen, and Elihu H. Estey. Patient-specific dose finding based on bivariate outcomes and covariates. *Biometrics*, 64(4):1126–1136, mar 2008.

[104] Beat Neuenschwander, Simon Wandel, Satrajit Roychoudhury, and Stuart Bailey. Robust exchangeability designs for early phase clinical trials with multiple strata. *Pharmaceutical Statistics*, 15(2):123–134, dec 2015.

[105] Claire C. Villette, David Orrell, Jim Millen, and Christophe Chassagnole. Should personalised dosing have a role in cancer treatment? *Frontiers in Oncology*, 13, may 2023.

[106] Razelle Kurzrock, Chia-Chi Lin, Tsung-Che Wu, Brian P. Hobbs, Roberto Carmagnani Pestana, and David S. Hong. Moving beyond 3+3: The future of clinical trial design, jun 2021.

[107] Yusuke Yamaguchi, Kentaro Takeda, Satoshi Yoshida, and Kazushi Maruo. Optimal biological dose selection in dose-finding trials with model-assisted designs based on efficacy and toxicity: a simulation study. *Journal of Biopharmaceutical Statistics*, 34(3):379–393, April 2023.

[108] L W Huson. Phase i oncology trials incorporating patient choice of dose. *British Journal of Cancer*, 107(7):1022–1024, August 2012.

[109] Graham M. Wheeler, Adrian P. Mander, Alun Bedding, Kristian Brock, Victoria Cornelius, Andrew P. Grieve, Thomas Jaki, Sharon B. Love, Lang'o Odondi, Christopher J. Weir, Christina Yap, and Simon J. Bond. How to design a dose-finding study using the continual reassessment method. *BMC Medical Research Methodology*, 19(1), jan 2019.

[110] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives*. Cambridge University Press, jul 1993.

[111] Kristen Cunanan and Joseph S Koopmeiners. Evaluating the performance of copula models in phase i-II clinical trials under model misspecification. *BMC Medical*

Research Methodology, 14(1), apr 2014.

[112] Hadi Safari-Katesari, S. Yaser Samadi, and Samira Zaroudi. Modelling count data via copulas. *Statistics*, 54(6):1329–1355, nov 2020.

[113] M. G. KENDALL. A NEW MEASURE OF RANK CORRELATION. *Biometrika*, 30(1-2):81–93, jun 1938.

[114] Abe Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231, 1959.

[115] Christian Genest and Johanna Nešlehová. A primer on copulas for count data. *ASTIN Bulletin*, 37(2):475–515, 2007.

[116] R. L. Plackett. A class of bivariate distributions. *Journal of the American Statistical Association*, 60(310):516–522, 1965.

[117] Olivier P. Faugeras. Inference for copula modeling of discrete data: a cautionary tale and some facts. *Dependence Modeling*, 5(1):121–132, 2017.

[118] Christian Genest, Aristidis K. Nikoloulopoulos, Louis-Paul Rivest, and Mathieu Fortin. Predicting dependent binary outcomes through logistic regressions and meta-elliptical copulas. *Braz. J. Probab. Stat.*, 27(3):265–284, 08 2013.

[119] Aristidis K. Nikoloulopoulos and Dimitris Karlis. Modeling multivariate count data using copulas. *Communications in Statistics - Simulation and Computation*, 39(1):172–187, dec 2009.

[120] Hakim Bekrizadeh, Gholam Ali Parham, and Mohamad Reza Zadkarmi. The new generalization of farlie–gumbel–morgenstern copulas. *Applied Mathematical Sciences*, 6(71):3527–3533, 2012.

[121] Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.

[122] Yuxi Tao, Junlin Liu, Zhihui Li, Jinguan Lin, Tao Lu, and Fangrong Yan. Dose-finding based on bivariate efficacy-toxicity outcome using archimedean copula. *PLoS ONE*, 8(11):e78805, nov 2013.

[123] Guosheng Yin and Ying Yuan. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 58(2):211–224, jan 2009.

[124] Peter F. Thall, Hoang Q. Nguyen, Thomas M. Braun, and Muzaffar H. Qazilbash. Using joint utilities of the times to response and toxicity to adaptively optimize schedule–dose regimes. *Biometrics*, 69(3):673–682, aug 2013.

[125] W. Scott Richardson, Mark C. Wilson, Jim Nishikawa, and Robert S.A. Hayward. The well-built clinical question: a key to evidence-based decisions. *ACP Journal Club*, 123(3):A12, nov 1995.

[126] Kenneth J Arrow. *Social choice and individual values*. Cowles Foundation Monographs Series. Yale University Press, New Haven, CT, 3 edition, July 2012.

[127] Anthony O'Hagan. Expert knowledge elicitation: Subjective but scientific. *The American Statistician*, 73(sup1):69–81, mar 2019.

[128] Massimo Di Maio, Ciro Gallo, Natasha B. Leighl, Maria Carmela Piccirillo, Gennaro Daniele, Francesco Nuzzo, Cesare Gridelli, Vittorio Gebbia, Fortunato Ciardiello, Sabino De Placido, Anna Ceribelli, Adolfo G. Favaretto, Andrea de Matteis, Ronald Feld, Charles Butts, Jane Bryce, Simona Signoriello, Alessandro Morabito, Gaetano Rocco, and Francesco Perrone. Symptomatic toxicities experienced during anticancer treatment: Agreement between patient and physician reporting in three randomized trials. *Journal of Clinical Oncology*, 33(8):910–915, mar 2015.

[129] Amanda Ribbands, Natalie Boytsov, Abigail Bailey, Boris Gorsh, Emily Luke, and Annabel Lambert. Real-world patient-reported outcomes and concordance between patient and physician reporting of side effects across lines of therapy in multiple myeloma within the USA. *Supportive Care in Cancer*, 31(6), jun 2023.

[130] Shing M. Lee, Xiaoqi Lu, and Bin Cheng. Incorporating patient-reported outcomes in dose-finding clinical trials. *Statistics in Medicine*, 39(3):310–325, dec 2019.

[131] Anaïs Andrillon, Lucie Biard, and Shing M. Lee. Incorporating patient-reported outcomes in dose-finding clinical trials with continuous patient enrollment. *Journal of Biopharmaceutical Statistics*, pages 1–12, July 2023.

[132] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23, jan 1954.

[133] Martin Peterson. The St. Petersburg Paradox. `https://plato.stanford.edu/archives/fall2023/entries/paradox-stpetersburg/`, 2023.

[134] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, , 1947.

[135] Lindsay J Mangham, Kara Hanson, and Barbara McPake. How to do (or not to do) ... designing a discrete choice experiment for application in a low-income country. *Health Policy and Planning*, 24(2):151–158, dec 2008.

[136] John Brazier, Jennifer Roberts, and Mark Deverill. The estimation of a preference-based measure of health from the sf-36. *Journal of Health Economics*, 21(2):271–292, March 2002.

[137] Eva Rodríguez-Míguez, José Luis Pinto-Prades, and Jacinto Mosquera-Nogueira. Eliciting health state utilities using paired-gamble methods: The role of the starting point. *Value in Health*, 22(4):446–452, apr 2019.

[138] Thomas Gilovich and Daniel Kahneman. *Heuristics and Biases*. Cambridge University Press, 2002.

[139] Peter H. Farquhar. State of the art—utility assessment methods. *Management Science*, 30(11):1283–1300, nov 1984.

[140] Kenneth Arrow. Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, NBER Chapters, pages 609–626. National Bureau of Economic Research, Inc, 1962.

[141] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122, jan 1964.

[142] Ralph L. Keeney. Utility independence and preferences for multiattributed consequences. *Operations Research*, 19(4):875–893, 1971.

[143] John O'Quigley and Mark Conaway. Continual reassessment and related dose-finding designs. *Statistical Science*, 25(2), may 2010.

[144] John K. Kruschke. *Doing Bayesian Data Analysis*. Academic Press, 2015.

[145] Satoshi Morita, Peter F. Thall, and Peter Müller. Determining the effective sample size of a parametric prior. *Biometrics*, 64(2):595–602, jun 2008.

[146] Peter F Thall, Richard C Herrick, Hoang Q Nguyen, John J Venier, and J Clift Norris. Effective sample size for computing prior hyperparameters in bayesian phase i–II dose-finding. *Clinical Trials*, 11(6):657–666, sep 2014.

[147] Kristian Brock. *trialr: Clinical Trial Designs in 'rstan'*, 2023. R package version 0.1.6.

[148] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263, mar 1979.

[149] John M Miyamoto and Stephen A Eraker. Parametric models of the utility of survival duration: Tests of axioms in a generic utility framework. *Organizational Behavior and Human Decision Processes*, 44(2):166–202, oct 1989.

[150] Lia C.G. Verhoef, Anton F.J. De Haan, and Willem A.J. Van Daal. Risk attitude in gambles with years of life. *Medical Decision Making*, 14(2):194–200, apr 1994.

[151] Han Bleichrodt and Jose Luis Pinto. A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science*, 46(11):1485–1496, nov 2000.

[152] Arthur E. Attema, Werner B.F. Brouwer, and Olivier l'Haridon. Prospect theory in the health domain: A quantitative assessment. *Journal of Health Economics*, 32(6):1057–1065, dec 2013.

[153] Nigel Stallard, Peter F. Thall, and John Whitehead. Decision theoretic designs for phase II clinical trials with multiple outcomes. *Biometrics*, 55(3):971–977, sep 1999.

[154] Daniel Kahneman, Amos Tversky, and Paul Slovic. *Judgment under uncertainty*. Cambridge University Press, 1982.

[155] Peter F. Thall and John D. Cook. Efftox biostatistics software, Jul 2021.

[156] Georg Heinze, Anne-Laure Boulesteix, Michael Kammer, Tim P. Morris, and Ian R. White and. Phases of methodological research in biostatistics—building the evidence base for new methods. *Biometrical Journal*, page 2200222, feb 2023.

[157] Ying Yuan and Guosheng Yin. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(5):719–736, dec 2009.

[158] Claes Wohlin. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*, EASE '14. ACM, May 2014.

[159] Ralf W. Schlosser, Oliver Wendt, Suresh Bhavnani, and Barbara Nail-Chiwetalu. Use of information-seeking strategies for developing systematic reviews and engaging in evidence-based practice: the application of traditional and comprehensive pearl growing. a review. *International Journal of Language & Communication Disorders*, 41(5):567–582, January 2006.

[160] Jay Bartroff and Tze Leung Lai. Approximate dynamic programming and its applications to the design of phase i cancer trials. *Statistical Science*, 25(2), may 2010.

[161] J. C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164, jan 1979.

[162] William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285, dec 1933.

[163] L. Pitt. Optimising first in human trials. `https://researchportal.bath.ac.uk/en/studentTheses/optimising-first-in-human-trials`, November 2021.

[164] D. Araujo, A. Greystoke, S. Bates, A. Bayle, E. Calvo, L. Castelo-Branco, J. de Bono, A. Drilon, E. Garralda, P. Ivy, O. Kholmanskikh, I. Melero, G. Pentheroudakis, J. Petrie, R. Plummer, S. Ponce, S. Postel-Vinay, L. Siu, A. Spreafico, A. Stathis, N. Steeghs, C. Yap, T.A. Yap, M. Ratain, and L. Seymour. Oncology phase i trial design and conduct: time for a change - mdict guidelines 2022. *Annals of Oncology*, 34(1):48–60, January 2023.

[165] Bahru A. Habtemariam, Lian Ma, Brian Booth, Olanrewaju O. Okusanya, and Nitin Mehrotra. Dose selection of targeted oncology drugs in early development. In *Novel Designs of Early Phase Trials for Cancer Therapeutics*, pages 73–94. Elsevier, 2018.

[166] Ian R White, Tra My Pham, Matteo Quartagno, and Tim P Morris. How to check a simulation study. *International Journal of Epidemiology*, October 2023.

[167] Anthony O'Hagan, Caitlin E. Buck, Alireza Daneshkhah, J. Richard Eiser, Paul H. Garthwaite, David J. Jenkinson, Jeremy E. Oakley, and Tim Rakow. *Uncertain Judgements*. Wiley, 2006.

[168] Juan Pablo Domecq, Gabriela Prutsky, Tarig Elraiyah, Zhen Wang, Mohammed Nabhan, Nathan Shippee, Juan Pablo Brito, Kasey Boehmer, Rim Hasan, Belal Firwana, Patricia Erwin, David Eton, Jeff Sloan, Victor Montori, Noor Asi, Abd Moain Abu Dabrh, and Mohammad Hassan Murad. Patient engagement in research: a systematic review. *BMC Health Services Research*, 14(1), February 2014.

[169] Stuart D. Faulkner, Fabian Somers, Mathieu Boudes, Begõna Nafria, and Paul Robinson. Using patient perspectives to inform better clinical trial design and conduct: Current trends and future directions. *Pharmaceutical Medicine*, 37(2):129–138, January 2023.

[170] Colin Gillespie and Robin Lovelace. *Efficient R Programming: A Practical Guide to Smarter Programming*. O'Reilly Media, 2017.

Appendix A

Additional material

A.1 Chapter 2 supplementary material

A.1.1 Multiple dose example

Scenario 1 from the CSS paper [111] has in part been recreated as further evidence that the mechanism for estimating the marginal distributions do not affect the findings of the simulation study in Chapter 3. The data generating mechanism for the 4 doses had fixed probability of efficacy at of $[0.38, 0.55, 0.71, 0.83]$, and toxicity $[0.05, 0.12, 0.27, 0.5]$. The correlation followed the FGM copula with $\theta = 1$. The joint probability model followed an FGM copula (Equation 3.13) with a uniform prior $\theta \sim U(-1, 1)$. The dose level covariates were defined as $x = [0, 1, 2, 3]$ and $x^2 = [0, 1, 4, 9]$. The efficacy marginal distribution assumed logistic regression with parameters for the intercept, slope and a quadratic term. Priors were assumed to all follow independent normal distributions with means $[-1, 1, 0]$ and standard deviations $[3, 2, 0.25]$ for the intercept, slope and quadratic parameters respectively. The toxicity probability model was logistic with intercept and slope parameters. Priors were specified to be normal with means, $[-3, 1]$ and standard deviations $[3, 2]$ for the intercept and slope respectively.

The priors for the slope parameters from both models in the CSS were specified to be from a Gamma distribution with the same means and standard deviations. When attempting to replicate Gamma priors, Stan produced a number of warnings that the posterior distribution was difficult to evaluate. This may have been a symptom of the method of MCMC with

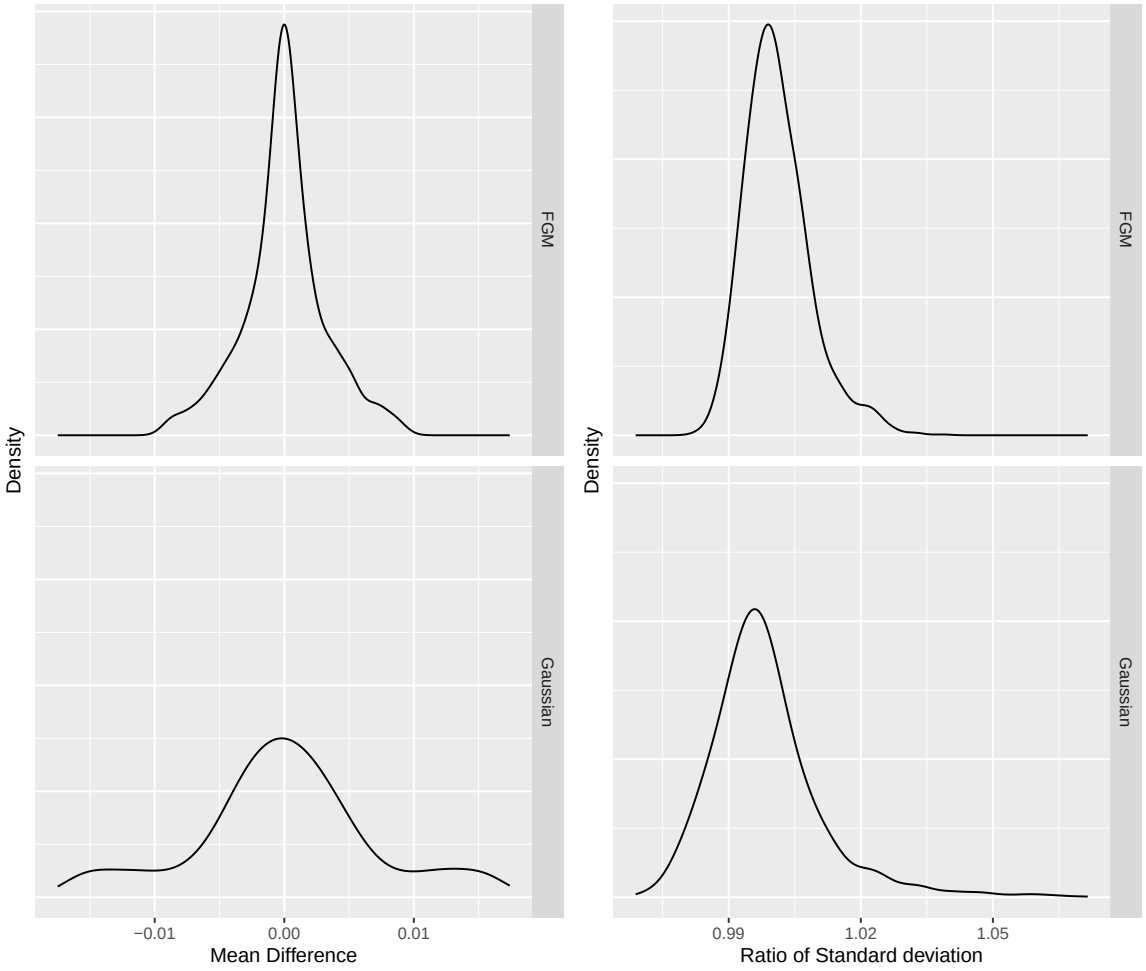

Figure A.1.1: Difference in toxicity marginal distributions for all possible combination of data for 20 patients recruited at a single dose between copula models and independent models. First row of plots is from the FGM copula model and the second row from the Gaussian Copula. First column is the difference in means between the copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

the prior placing a lot of mass just above zero and no mass below. For the purpose of this example, Normal priors were used for the two slope parameters with the same mean and standard deviation. Stan software was used to conduct the simulation with 3 chains, 3000 replicates and a 1000 replicate warm up.

The simulation study wasn't a staged study example with 30 patients, 6 treated at each dose levels 1 and 2 and 9 patients at dose levels 3 and 4. 5000 draws were made from the data generating mechanism with the data then fit to the FGM copula and an independent model. The posterior plots for difference between intercept and slope parameters are given in Figure A.1.4 for the efficacy model. Equivalent plots for toxicity intercept and slope parameters is given in Figure A.1.5. Note that in both figures the difference is on the odds

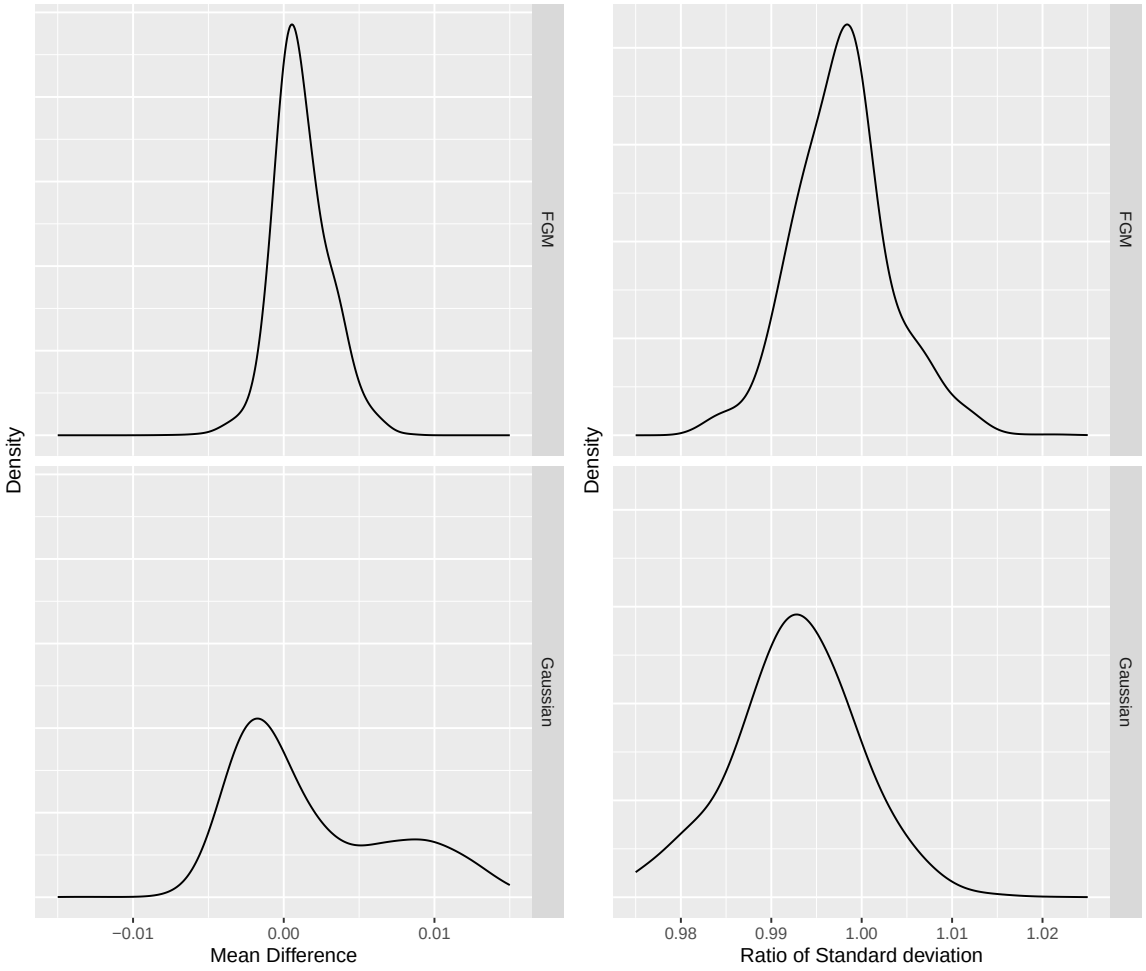

Figure A.1.2: Difference in efficacy marginal distributions of data for 20 patients recruited at a single dose between copula models and independent models. Data is generated from a dose with $\pi_E = 0.7$, $\pi_T = 0.3$ and $\tau_b = 0.168$. First row of plots is from the FGM copula model and the second row from the Gaussian copula. First column is the difference in means between the copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

scale, 0.05 on the odds scale representing 0.012 on a probability scale. Given the values for the covariates, the intercept is the odds at dose level 1. The maximum magnitude of difference for the slope parameter occurs at dose level 4 where the difference is multiplied by 3. Overall there are only very minor differences in the means and standard deviations with the magnitude consistent with the values obtained in the single dose simulation study.

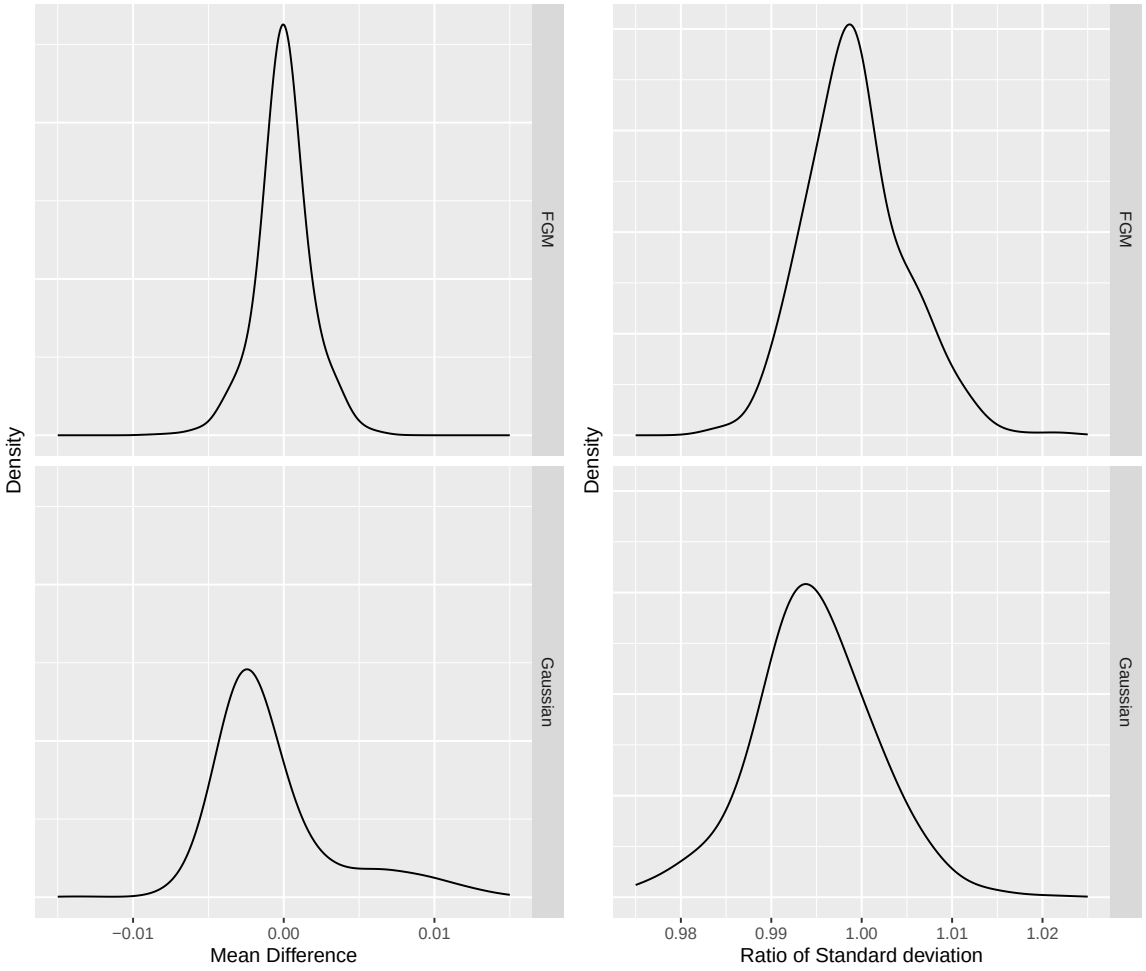

Figure A.1.3: Difference in efficacy marginal distributions of data for 20 patients recruited at a single dose between copula models and independent models. Data is generated from a dose with $\pi_E = 0.7$, $\pi_T = 0.3$ and $\tau_b = 0$ or an independent data generating process. First row of plots is from the FGM copula model and the second row from the Gaussian Copula. First column is the difference in means between the Copula and independent models. Second column in the ratio of standard deviations between copula model and independent model fit to the same data.

A.2 Chapter 4

A.2.1 Efficient simulation

Programming efficiency is the speed at which a computer can complete a task. A simple approach to conducting a simulation study in dose finding is as follows, simulate outcome data for the first cohort, fit this data to a probabilistic model to obtain draws from the posterior, before using these to decide the mean utility and dose for the next cohort. This iterative procedure would continue until the maximum sample size has been reached or the trial stopped early. The process is repeated a high number of times for a single scenario to evaluate the performance of a particular design (or decision function). A loop statement

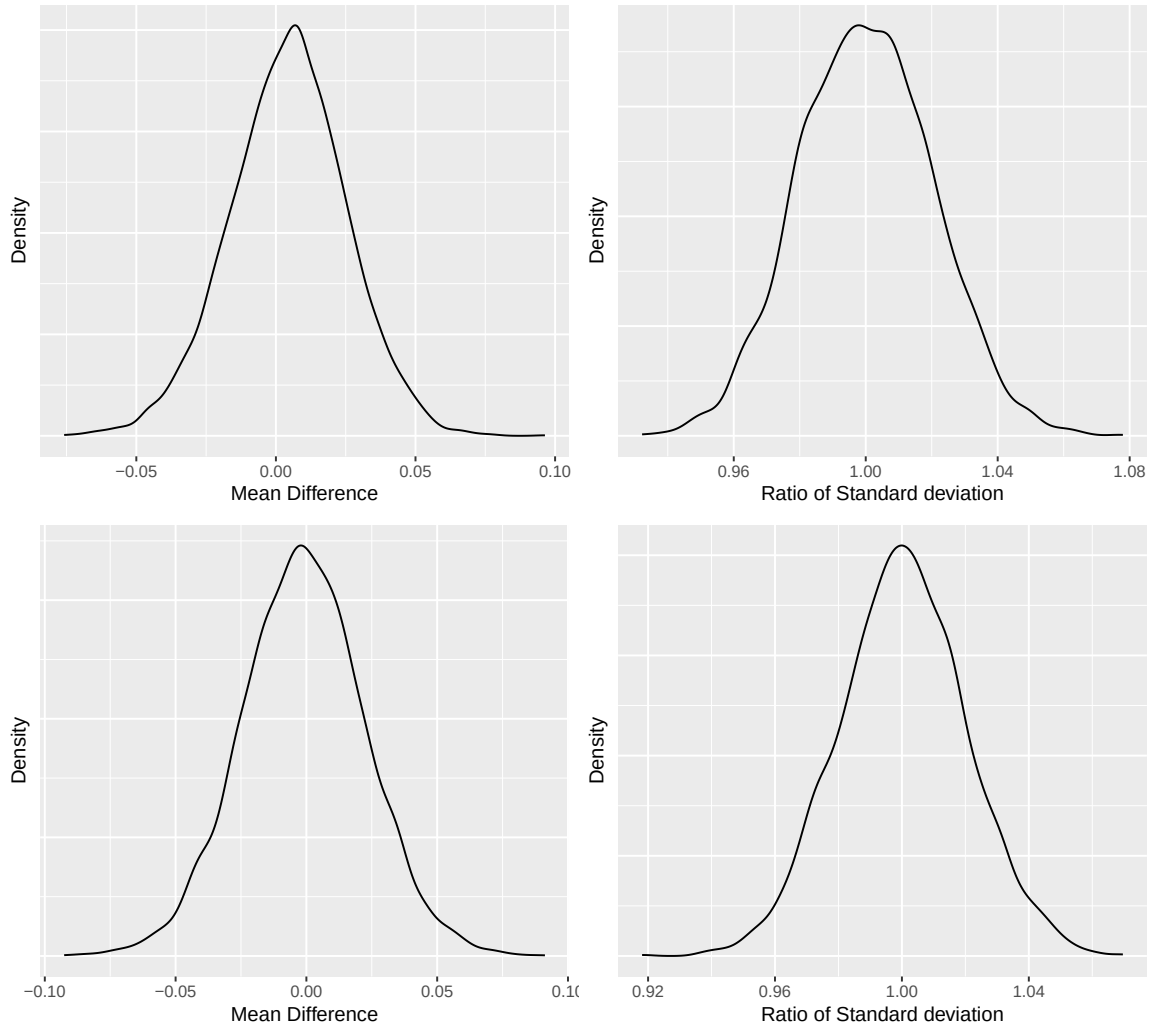

Figure A.1.4: Difference in efficacy marginal parameters of data for 30 patients recruited at 4 doses between copula models and independent models. First row of the plot is the difference in the intercept parameter between FGM model and an independent model. Second row is for the slope parameter. The x axis is on the odds scale with an increment of 0.05 on the odds scale representing 0.012 on the probability scale

allows us to write the code for a single trial and then cycle back through for the given number of simulation replicates.

There is nothing inherently wrong with this approach, however it is computationally inefficient. If each model fit takes approximately a single second, a simulation study looking at a trial with 20 cohorts and 2000 simulations, would take approximately 12 hours to complete. The purpose here isn't to detail the miniature of code optimisation but how this problem can be tackled differently in order to achieve large gains in efficiency.

Removing loops is typically a sensible first step in optimising R code [170]. In dose finding however each trial needs to iteratively grow the data with each stage dependent on the last,

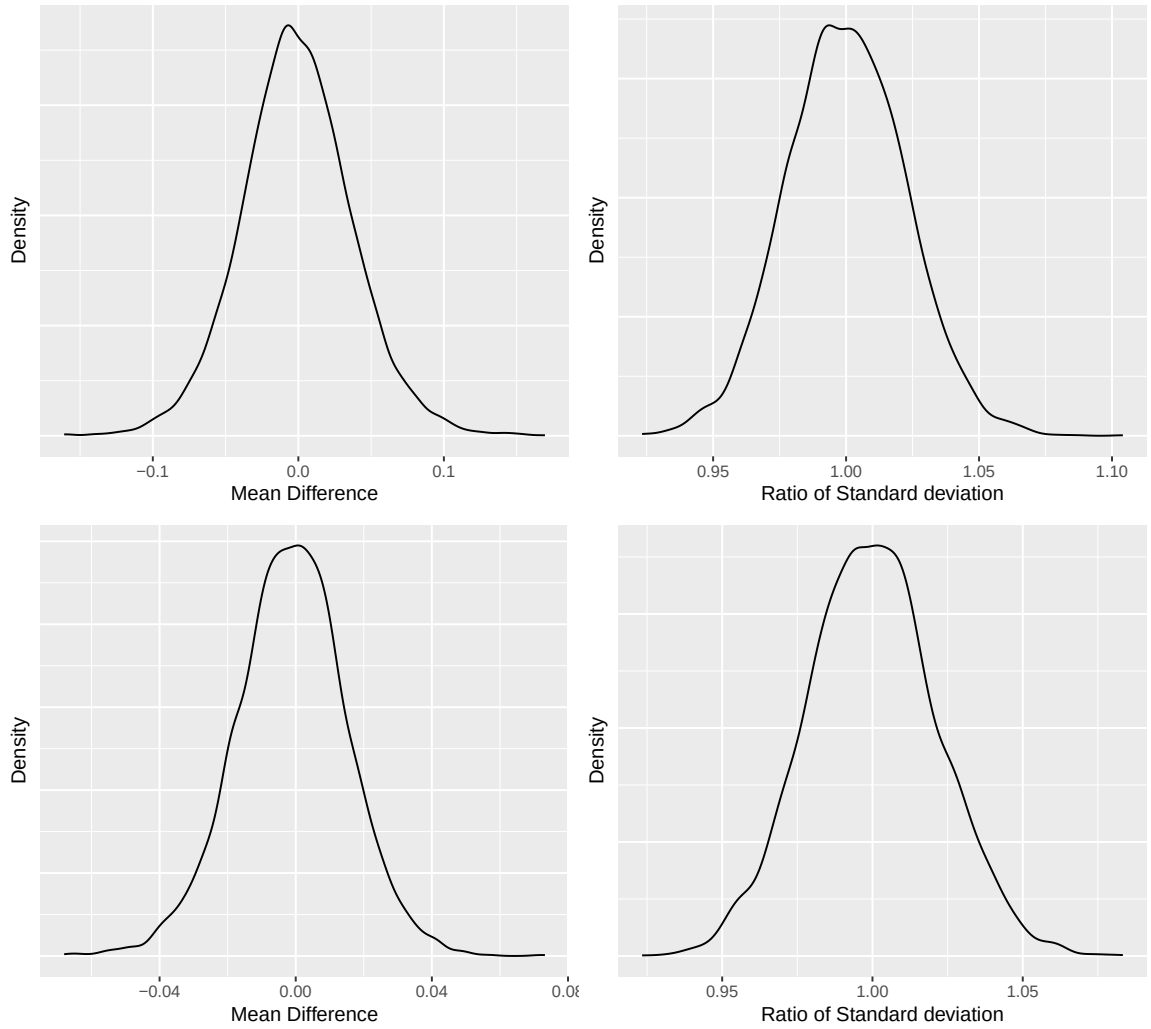

Figure A.1.5: Difference in toxicity marginal parameters of data for 30 patients recruited at 4 doses between copula models and independent models. First row of the plot is the difference in the intercept parameter between FGM model and an independent model. Second row is for the slope parameter. The x axis is on the odds scale with an increment of 0.05 on the odds scale representing 0.012 on the probability scale

so this isn't easily possible. The main computational burden of simulation in this setting is in the model fitting function. In this simulation study Stan was used for model fitting [121]. There a number of improvements, such as writing the likelihood as a vector, that can be done to optimise Stan code that results in small improvements in the time to achieve each model fit.

When considering a simulation study in its entirety there is a lot of repetition in fitting a model to the same data. For example, if the cohort size is one, after the first cohort there are only 4 possible data outcomes. The looping approach would however fit the model the number of times that the simulation is to be replicated (2000 times). To improve on efficiency

the approach taken for the wider simulations study is to simulate data for all simulation repetitions, then to subset all repetitions so that the model fits once to each unique data combination. Decisions are then made and merged back to the original simulated data set.

Efficacy and toxicity outcomes are defined in the R2DT method to be independent. Individual patient data for efficacy and toxicity at each dose are modelled as Bernoulli random variables but due to exchangeabilty and sufficency uniqueness of data for the model is defined by the cumulative number of events and total number of patients treated at each dose (See Chapter 3). It is the uniqueness of the efficacy and toxicity data vectors separately that define a unique dataset and model fit. For example consider 6 patients treated at a single dose, there are 7 possible efficacy data sets around the total number of efficacy events. Similarly for the toxicity data there are 7 possible unique data sets to fit the toxicity model. Considering these as combined models would result in 49 possible combinations of the data. The strategy in separating efficacy and toxicity allows a further gain in efficiency.

A simulation study can incorporate multiple different scenarios and multiple decision functions (or arguments to the same function) that all have the same probability model. This allows further savings in efficiency by running everything at the same time and only feeding unique data sets to the model fitting. Consider the simulation study given in Chapter 5, each trial has 20 cohorts simulated 2000 replicates, repeated for 10 scenarios and 21 different decision functions. After the first cohort of three patients there are just 4 toxicity and 4 different efficacy models possible. Contrast this to fitting everything separately, the probability model would be fit 410000 times ($10 \times 2000 \times 21$) in the simple approach.

If each decision function and scenario from earlier takes 12 hours to run, this approximates to 105 days of computing time to run the entire simulation study. This becomes only feasible with high performance computing. Running the entire simulation with the approach described above by only fitting unique data sets allowed the code to be run in 27 hours upon a modest desktop computer. The bespoke software available to fit *EffTox* is quicker still [155]. The software doesn't have the ability to edit the code to be suitable to fit R2DT however.

Appendix B

Programming code

The R computer code used to conduct computer simulation in Chapters 3 and 5 are saved on a Github repository here: `https://github.com/medahala/PhD.git`