Late stage combination drug development for improved portfolio-level decision-making

Emily Zara Graham

MMath (Hons.), Durham University, 2015

MRes, Lancaster University, 2016

Submitted for the degree of

Doctor of Philosophy

at Lancaster University

October 2019

STOR-i
excellence with impact

Lancaster University

Abstract

Combination therapies are becoming increasingly used in drug development for a range of therapeutic areas such as oncology and infectious diseases, providing potential benefits such as minimising drug resistance and toxicity. Typically, a pharmaceutical company will have multiple treatments in different stages of development in their portfolio and the problem of portfolio decision-making will include decisions such as which studies to initiate and how to prioritise studies.

This problem is more complex for portfolios of combinations since sets of combination studies may be related, for example if they have at least one treatment in common and are used in the same indication. However, in this setting, value can be gained by sharing information between related combination studies in terms of improving the treatment effect estimates and improving the portfolio-level decisions. We discuss the challenges of portfolio decision-making for a portfolio of combinations and present methodology to assist with this.

One of the key estimates that is used in decision-making regarding a clinical study is the probability of study success. We present a framework that allows the study success probabilities of a set of related combination therapies to be updated based on the outcome of a single combination study. This allows us to incorporate both direct and indirect data on a combination therapy in the decision-making process for future studies.

Existing methods for portfolio decision-making do not account for the differences between single agent and combination drug development. We extend the existing methodology to consider the relationship between combinations and the effect that observing certain outcomes may have on the portfolio decisions we make.

This is achieved by updating the study success probabilities throughout the decision-making process whenever a relevant outcome is observed.

This thesis is dedicated to James.

Acknowledgements

I would like to acknowledge the STOR-i Centre for Doctoral Training for this opportunity and for the support of the leadership team and support staff. I gratefully acknowledge the EPSRC and Roche for their financial support. I would like to thank my supervisors, Jack and Chris, for their support. I am so grateful to have had the opportunity to work with you both. I would also like to thank the Biostatistics team at Roche; my extended visits were one of the highlights of my PhD and were invaluable to my learning.

My second set of acknowledgements is to my family. I would like to thank Mum for giving me the drive that led me through the PhD and Dad for telling me that girls can do anything. I am grateful for Gran and Grandpa nurturing my inquisitive mind and for the unwavering support that John continues to provide. I am so grateful to have Philippa, Imogen and Isabella as my sisters; with them by my side, I know I can achieve anything. The support of The Houghtons has been hugely appreciated, especially the evening cups of tea and their willingness to listen.

I am grateful for the friendships that have been fundamental to my happiness and growth during the PhD. The support of Kirsten, Katie and Sophie has been paramount to my successes over the last eight years. One of the best, unexpected things to come out of the PhD was my friendship with Emma and I am so grateful for this. I am also grateful for the support, advice and laughter that Lucy has provided. I would like to thank Kim, Pearl and Alice for their unconditional love and support. Finally, I am grateful to have shared this experience with my cohort at STOR-i and wish them the best in the future.

Declaration

I declare that the work in this thesis has been done by myself and has not been submitted elsewhere for the award of any other degree.

The work presented in Chapter 2 has been accepted for publication as Graham, E., Jaki, T. and Harbron, C. (2019) A comparison of stochastic programming methods for portfolio level decision making. *Journal of Biopharmaceutical Statistics.*

The work presented in Chapter 3 has been submitted for publication as Graham, E., Harbron, C. and Jaki, T. (2019) Updating the probability of study success for combination therapies using related combination study data.

The word count for this thesis is approximately 45000 words.

Emily Zara Graham

Detailed Contents

Figures

Tables

Symbols

$i \in I$	index for the programme in the set of all programmes in the portfolio
$j \in J_i$	index for the study in the set of all studies for programme i
(i, j)	notation for the jth ordered study in programme i
$s \in S$	index for the scenario in the set of all possible trial outcome scenarios
$t \in T$	index for the time point in the set of all time points
c_{ij}	the cost of study (i, j)
ϕ_{ij}	the probability of success of study (i, j)
τ_{ij}	the duration of study (i, j)
γ_i^D	the loss per time period due to shorter active patent life of drug i
γ_i^L	the loss per time period due to smaller market share of drug i
$r \in R$	index for the resource type in the set of all resource types
λ_{ijr}	the resource requirement of study (i, j) of resource type r
$\lambda_r^{\max}$	the availability of resource type r per time period
θ_i	the treatment effect of drug i
μ_i	the prior mean of θ_i
σ_i	the prior standard deviation of θ_i
$\rho_{ii'}$	the prior correlation between θ_i and $\theta_{i'}$
Z	the score statistic of a study
V	the Fisher information of a study
α	the significance level of a test
$1 - \beta$	the power of a test

Chapter 1

Chapter 1

Introduction

1.1 Introduction

There has been a recent rise in popularity of **combination therapies**, with reports of over 10000 clinical trials being registered as ongoing in 2017 in the US alone [1]. Combination therapies combine **new molecular entities** and existing drugs with potential benefits including achieving synergy, minimising drug resistance and reducing side effects.

While this is an exciting advancement in terms of the development of new treatments, it also brings with it many other benefits, one of which being a new potential for learning. Due to the nature of combination therapies, it may be realistic to assume that the performance of similar combinations of drugs in similar indications is related in some way. This can be seen as a potential for learning across different combinations.

These benefits, however, do come at the cost of new challenges relating to combination drug development and the optimal way to conduct studies in a portfolio of combinations. Typically, a pharmaceutical company will have multiple drugs undergoing development in their **pharmaceutical portfolio**, which will be in different stages of development, and other potential studies that they may consider including in the portfolio.

In this thesis, we present a framework for sharing information across related combination studies and build this into a procedure for **portfolio-level decision-making**.

1.2 Drug development

A **drug development programme** is the sequence of tasks that a new drug must undertake in order to be launched onto the market for use by the target population. This process can be summarised as a sequence of five parts, as depicted in Figure 1.1.

The first stage of drug development is drug discovery. This takes place in a laboratory and involves identifying targets of interest within the body that are believed to be associated with the indication of interest. Once these targets have been identified, existing drugs and new molecular entities that are believed to have beneficial interactions with the targets are tested in in vitro studies to see if they show potential for benefit [2]. These studies are conducted outside of living organisms, for example on tissue samples.

The next stage is preclinical development, which includes both in vivo and in vitro studies. In vivo studies refer to tests that are carried out in living organisms. The main aims of preclinical development are to assess the toxicity of the treatment and to find an appropriate dosage level. At the end of preclinical development, a decision is made based on the collected evidence as to whether or not the treatment should continue to **clinical trials**, where it will be tested in humans, and what dosage is safe to consider in these trials [2].

Clinical development is typically made up of three phases that each have a different aim. In Phase I, the treatment is typically given to a small sample of healthy volunteers, or patients if the intervention is expected to have adverse

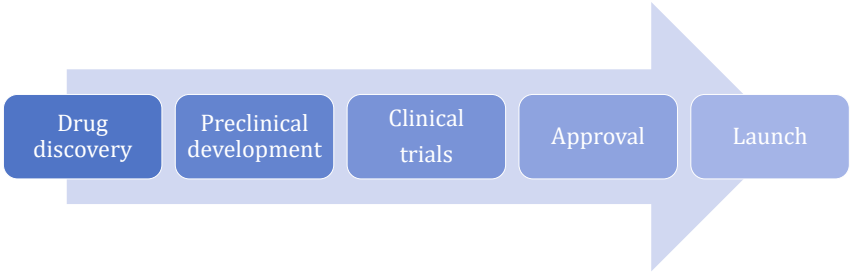

Figure 1.1: The different stages of a drug development programme.

effects, such as in oncology. The main aim of Phase I is to assess safety and to find the safe doses to take forwards to the later phases. This phase often only lasts several months. Providing that none of the Phase I studies in a development programme fail, the treatment is taken to Phase II where it will tested in a larger sample of patients with an aim to look at both safety and efficacy. If there are no safety concerns and there is evidence of benefit, the drug may proceed to Phase III. Phase III is much larger and much longer than the other phases, typically recruiting hundreds, if not thousands, of patients and lasting between one and four years [2]. The aim of this phase is to confirm efficacy and also to monitor safety. Phase III trials are often referred to as confirmatory trials. It should be noted that a particular development programme might not complete all of the phases of clinical development sequentially, for example a company might skip Phase II, run a Phase II/III trial or run two Phase III trials simultaneously.

If a drug is successful in all associated clinical trials, then it will be taken to regulatory bodies, such as the European Medicines Agency (EMA) or the US Food and Drug Administration (FDA), for approval. If a drug is granted regulatory approval, then it can be marketed and launched for use by the target population. After a new treatment is launched, further studies may be run in order to, for example, assess its performance in different populations or indications. There will also be a long-term follow-up with the people who were enrolled in the

clinical trials to monitor long-term effects that might not have been captured during the trials and also to monitor any rare side effects of the treatment.

The drug development process takes approximately 10-15 years [2] and costs hundreds of millions of pounds [3]. Not only is it long and expensive, but it also involves a high level of uncertainty since a new treatment can fail at any stage, exceed the allocated budget or take longer than expected. Furthermore, even if the drug is successfully developed, it might not have the expected impact on the market due to either a lower benefit than was expected or competition in the marketplace.

Throughout a drug development programme, there are many decisions to be made from whether development should continue to how each stage should be conducted. Decisions that might be considered in a clinical trial include the most appropriate choice of primary endpoint, for example tumour shrinkage or progression free survival time, and whether interim analyses should be included. Another decision that a company might make during drug development is the decision to submit the drug for accelerated approval [4]. The accelerated approval programme allows drugs to be approved earlier when they are for use in indications with a high unmet medical need or when they will provide a step change in the treatment of the disease, such as the introduction of immunotherapies in cancer treatments. Typically, for a treatment to be considered for accelerated approval, a significant result will have been observed in the surrogate endpoint, which is believed to be an indication of overall benefit in the primary endpoint.

The decisions made throughout a drug development programme are by no means trivial and the consequences of making suboptimal decisions can include increasing the likelihood of failure. The success rates and the most common reasons for clinical trial failure are summarised in Table 1.1. Grignolo and Pretorius [5] also discuss the reasons for failure at Phase III and the trends that

	Success rate	Reason for failure			
		Safety	Efficacy	Strategic	Other
Phase I	63.2%	-	-	-	-
Phase II	30.7%	22%	59%	16%	3%
Phase III	58.1%	35%	52%	0%	13%

Table 1.1: Success rates in the different phases of drug development [6] and the reasons for drugs failing to continue to the next stage of development between 2011-2012 [7].

are seen in Phase III failures. Some of the reasons that they identify for failure include flawed study design, such as insufficient sample size, and flawed data collection and analysis, such as incorrect assumptions on treatment effect. They also presented several approaches to help prevent these failures in the future.

In this thesis, we present methodology that can be used to assist decision-making in a pharmaceutical portfolio since improving the decision-making process can help to reduce the failure rates observed in clinical trials, especially in Phase III.

1.3 Combination therapies

One of the motivations for our research is the recent rise in popularity of combination therapies. Combination therapies refer to combinations of drugs that can contain both existing drugs and new molecular entities (NMEs). Typically, a combination therapy consists of a **backbone treatment** and one or more **add-on treatments**. There will often be several potential add-on treatments for each backbone (see Figure 1.2 for an illustration of this). The backbone treatment is often a well-established drug for which there is already a reasonable amount of information. The potential add-on treatments, however, will often include NMEs and external drugs, for which there will be less information available than for the backbone. In oncology, for example, an existing chemotherapy might be used as a backbone treatment and NMEs and immunotherapies might be used as potential add-on treatments to this chemotherapy.

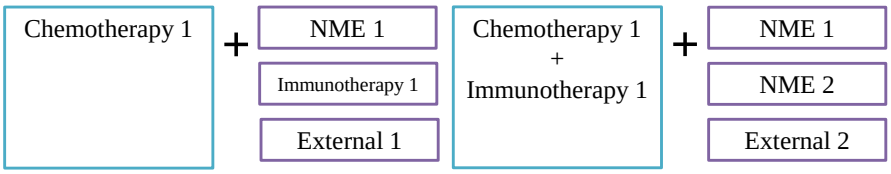

Figure 1.2: Diagram showing the backbone (blue boxes) and add-on (purple boxes) structure of six potential combination therapies in an oncology setting.

Combination therapies are becoming increasingly used in a range of therapeutic areas, but the areas that they are most commonly used in include oncology and infectious diseases. Glickman and Sawyers [8] compare these two areas and discuss the way that experiences of combination drug development in one can help inform decisions in the other.

There can be many different reasons for using a combination therapy over a single agent. Chou [9] describes the main aims of combination drug development to be achieving a synergistic effect, minimising drug resistance and reducing dose and toxicity. This means that combination therapies are often able to achieve efficacy whilst reducing side effects. This is achieved by combining drugs that are similar in terms of therapeutic effect but different in terms of toxicity. Doing this allows the doses of the individual drugs to be lower than if they were to be used separately [10]. Furthermore, Podolsky and Greene [11] discuss the fact that combinations might increase effectiveness when they are not able to increase efficacy due to their potential to make treatment adherence easier. They use Combivir as an example to highlight this, which reduced the patient pill burden by a half [11].

Synergy is defined by Chou [9] to be the realisation of an effect that is more than additive between drugs. Antagonism is the opposite of synergy and is defined to be a less than additive effect. However, even if the effect of the combination therapy is greater than the effect of each drug alone, it does not necessarily mean that it is a synergistic effect as this could just be an additive effect [9]. Therefore, methods are needed to quantify the synergy.

Two main approaches for exploring the effect of a combination of inhibitors are discussed in [10]: Loewe additivity [12] and Bliss independence [13]. Inhibitors aim to reduce the symptoms of an indication by reducing the rate at which the reactions that cause the disease occur within the body. Loewe additivity assumes that the inhibitors act similarly whereas Bliss independence assumes that they act independently. These methods are most commonly used in a preclinical setting where we are interested in understanding the mode of action, and can run detailed experiments at many dose levels. In a clinical setting, we generally ask the much simpler question of whether the combination is much better than the individual components on their own, which could still technically be additive but would give the patient a sizeable benefit.

The Loewe additivity method [12] works as follows. Let c'_A and c'_B denote the concentrations of inhibitors A and B respectively that achieve α% target inhibition separately. Also, let c_A and c_B denote the concentrations of A and B necessary when combined to achieve α% target inhibition. We can calculate the combination index to find the level of interaction between the inhibitors, which is given in [10] by

$$I = \frac{c_A}{c'_A} + \frac{c_B}{c'_B}.$$

Using this equation, antagonism is represented by $I > 1$, synergy is represented by $I < 1$ and if $I = 1$ then the effect is additive and there is no interaction between the inhibitors [10].

Using the Bliss independence method [13], we calculate the effect, θ_C, of the inhibitor combination given the effects for inhibitors A and B, θ_A and θ_B, by

$$\theta_C = \theta_A \times \theta_B.$$

This takes into account the fact that if inhibitor A has already acted upon a certain amount of the target then B has less to inhibit [10].

Foucquier and Guedj [14] provide an overview of the methodology that can be used to analyse combination effects and split these approaches into effect-based and dose-effect-based approaches. They define effect-based approaches to be those that consider the effect of the combination against the effect of the components and dose-effect-based approaches to be those that also consider the dose. Under this definition, the Bliss independence method can be classed as an effect-based approach whereas the combination index arising from the Loewe additivity method is a dose-effect-based approach. Other effect-based approaches that are discussed in [14] include combination subthresholding, highest single agent and response additivity.

Although there is a rich literature regarding synergy, there is little available literature on decision-making regarding combination therapies. Existing methodology for single agents can be applied in this setting, but the decisions that these lead to might not be optimal. This is because of the relationships that are likely to exist between similar combinations and the way that these relationships would affect our decision-making. We believe that there is potential to be gained from sharing information between trials with similar combinations of treatments and using this to assist decision-making. For example, if a study is run for combination $A + B$ in a particular indication, the outcome of this trial might inform our beliefs about how combination $A + C$ might perform in that same indication, or in an indication that is linked to C in the same way that B was linked to the original indication.

Furthermore, combination and single agent drug development have different logistical constraints and considerations that must be accounted for in any decision-making process. Woodcock et al. [15] discuss novel combination drug development and the complexities of co-development. They note that co-development increases uncertainty regarding the separate components thus should only be used in specific situations such as when there are no alternative therapies avail-

able for the indication. This means that a company will typically need to develop the individual agents separately to the combination.

In this thesis, we present methodology that allows us to use the **relationship** between combination studies to increase learning and assist decision-making. We use the presented methodology to update the **treatment effect estimates** whenever a relevant outcome is observed, which can lead to an improvement in the accuracy of the treatment effect estimates and improved decision-making in the portfolio. This will include decisions such as which drugs in the large set of possible combinations to add to the portfolio and how to organise the studies within these programmes.

1.4 Clinical trials

In this thesis, our main focus will be on clinical development. Thus, in this section, we provide background to the design and analysis of clinical trials with a focus on the later phases.

1.4.1 Outline

In Section 1.2, it was noted that there are three clinical trial phases, each with a different aim. Studies in each phase will be designed and analysed differently to accommodate these different aims.

Phase I trials are often referred to as "first-in-human" trials as they are the first time that the new intervention is given to humans. At this point, the aim is to assess the safety and find the appropriate dosage of the drug; the design of a Phase I study will be dependent upon its main aim [16]. Phase I studies usually recruit a small number of healthy volunteers, unless the treatment is believed

to be toxic. For example, for oncology drugs, these studies would always use cancer patients. Dose escalation studies are used to identify the range of safe doses at Phase I and find the maximum tolerated dose. These studies typically add individuals to the study sequentially, often in small cohorts, and wait to observe any side effects before adding the next cohort to the study and deciding what dose this cohort should receive, or concluding the study. Dose escalation studies can take on a variety of forms [17][18]. Efficacy might also be considered at this point but it will not be the primary aim of the study. The sample size is not usually related to the power of the study unlike in the later phases of clinical development.

The aims of a Phase II study include assessing the safety and benefit of the drug in patients and these studies are often referred to as "exploratory" or "proof-of-concept" studies. Early Phase II studies focus more on dose exploration and proof of concept whereas late Phase II studies focus more on efficacy [16]. The design of Phase II studies can vary and ranges from single-arm studies [19] to multi-arm multi-stage studies [20].

The largest and most expensive studies are in Phase III. These studies aim to show efficacy and are often called "confirmatory" studies. It is at this stage that the drug will typically be compared to the current standard of care and the aim will be to show superiority, but some studies might aim to show non-inferiority alongside different benefits such as a reduction in side effects [16]. As with the previous phases, studies in Phase III can take on a range of different designs but will typically have similar qualities.

1.4.2 Design and analysis

We will now summarise some of the key elements of clinical trial design that will be relevant in the later chapters of this thesis. Typically, a study will contain two

arms - an **experimental** arm and a **control** arm. Patients will be **randomised** to the experimental treatment with allocation ratio $R:1$, where R is often equal to 1. This means that when a patient enters the trial, they have a probability of $R/(R+1)$ of being allocated to the experimental group and a probability of $1/(R+1)$ of being allocated to the control group. Studies are often double-blind, which means that both the patient and their doctor do not know which treatment group they are in. This is done to reduce the various types of bias that can otherwise be introduced.

A clinical trial will monitor many different things, but the main effect relating to efficacy that the team is interested in will be referred to as the **primary endpoint**. This must be selected at the design stage and the choice will feed into decisions regarding the sample size and the planned duration of the study. The endpoint can take on many different forms in terms of the type of measure. For example, progression free survival time is often considered in oncology whereas a binary outcome of survival is more appropriate in Ebola. Regardless of the type of response, we need to be able to capture this information in a measure of **treatment effect**, which captures the difference between the experimental and control treatment, for which we will denote the true value by θ. This true value of the treatment difference is unknown and can only ever be estimated using relevant data. Examples of measures that can be used for binary data and survival data are the log-odds ratio and the log-hazard ratio, which are given respectively by

$$\theta = \log\left(\frac{p_E\,(1-p_C)}{p_C\,(1-p_E)}\right) \text{ and } \theta = -\log\left(\frac{h_E\,(t)}{h_C\,(t)}\right),$$

where p_x is the success probability in group x and $h_x\,(t)$ is the hazard function in group x [21]. We use survival data as an illustration here, and in future illustrations, as our work is motivated by the development of combination therapies in oncology and survival data are commonly used in this setting. We

will also consider a real world example in Chapter 3 that uses survival data.

The type of test that will be used in the analysis of a clinical trial must also be specified during the design stage. Parameters that are needed at this point will include the **significance level**, α, the **power**, $1-\beta$, and the **minimally important difference**, δ. The type of test will also need to be specified at this point i.e. one-sided or two-sided. The null hypothesis for superiority testing is usually chosen to be $H_0 : \theta = 0$.

Two values that can be used to analyse the study data are the **efficient score** of the study, Z, and the **Fisher information**, V. The efficient score is a measure of benefit of the experimental treatment over control and the Fisher information is used to quantify the information contained in Z relating to θ [21]. It can also be shown that $Z \sim N(\theta V, V)$ approximately, when θ is small and the sample size is large [21]. Note that this is not the same measure as the Z-score, which follows the standard normal distribution.

In a two-sided test, the null hypothesis will be rejected when $|Z| \geq d$, where d is the **critical value** of the test. If $Z \geq d$, then this would signify evidence that the experimental treatment is superior to the control treatment and the converse would be true if $Z \leq -d$. This means that we require

$$P(Z \geq d|\theta = 0) = \alpha/2 \text{ and } P(Z \geq d|\theta = \delta) = 1-\beta.$$

From this, we can deduce the critical value of the test and the required Fisher information to be

$$d = \frac{Z_{\alpha/2}\left(Z_{\alpha/2} + Z_{\beta}\right)}{\delta} \text{ and } V = \left(\frac{Z_{\alpha/2} + Z_{\beta}}{\delta}\right)^2 \tag{1.1}$$

respectively, where $Z_x = \Phi^{-1}(1-x)$ and $\Phi(\cdot)$ is the standard normal cumulative distribution function [21]. The value of V here can be used to find the appropriate

sample size for the study based on approximations of V for different responses. When θ is the log-hazard ratio, the approximation of these variables for survival data is given by

$$V \approx \frac{R}{(R+1)^2} e,$$

where e is the number of survival events observed and the allocation ratio is $R:1$, as defined previously. [21].

The calculation of Z and V given the study data will depend on the type of response. An overview of the different types of responses and the associated calculations for Z and V are provided in [21]. For survival data, these are given by

$$Z = e_C - \sum_{i=1}^{k} \frac{o_i r_{iC}}{r_i} \text{ and } V = \sum_{i=1}^{k} \frac{o_i (r_i - o_i) r_{iE} r_{iC}}{(r_i - 1) r_i^2}$$

where $t_1, \ldots t_k$ represent the survival times, o_i is the number of survival times equal to t_i, r_i is the number of survival times greater than or equal to t_i and e_C is the number of events observed in the control arm.

Upon the conclusion of the trial, we are able to calculate Z and V and draw conclusions about the treatments based on the critical value d and the previously outlined hypothesis test. Note that we use the actual calculation of V when study data is available rather than the planned value that was given in Equation (1.1).

1.5 Probability of success

The **probability of success** (PoS) is one of the measures that is commonly used to assist decision-making regarding a study or a set of studies. In this section, we provide details of how the PoS can be calculated and discuss additional aspects such as updating this estimate and its use in decision-making.

1.5.1 Methods for calculation

One of the most common approaches for calculating the PoS of a study is linked to the idea of the **assurance** or the expected power. This is also referred to in the literature as the average success probability and the Bayesian predictive power. Where the power is the probability that a trial will be successful given a particular value of the measure of treatment difference, the assurance is the unconditional probability that the trial will be successful. Chuang-Stein [22] and O'Hagan et al. [23] both discussed this value and defined the PoS to be

$$\int_{-\infty}^{\infty} P\left(\text{study success}|\theta\right) P\left(\theta|\text{data}\right) \mathrm{d}\theta$$

where study success will typically be linked to rejecting the null hypothesis. In this expression, $P\left(\theta|\text{data}\right)$ is the posterior distribution of the treatment effect, θ, given some historical data but it could also be set to an informative or uninformative prior.

In the calculation of the assurance, we are combining frequentist and Bayesian ideas. The specification of a prior for the treatment effect and the use of this in calculating the PoS is Bayesian, but it is assumed that the analysis of the trial will be frequentist and is linked to testing a hypothesis. This is a typical feature of methods for calculating the probability of study success.

O'Hagan et al. [23] provide several closed form solutions for the assurance. As a simple example of this, in the case of a two-sided superiority trial with normally distributed outcomes with known variance,

$$1 - \Phi\left(\frac{V^{-0.5} Z_{\alpha/2} - \mu}{\sqrt{V^{-1} + \sigma^2}}\right) \tag{1.2}$$

is the closed form solution of the probability of rejecting the null hypothesis and concluding that the experimental treatment is superior to the control. It is

also noted that, when a closed form solution for the assurance is not available, Bayesian clinical trial simulation can be used. The outline of this process is given as follows [23]:

1. Set counter, n, equal to 0.
2. Draw θ^* from the prior distribution for θ.
3. Draw the test statistic from the distribution based on θ^*.
4. If the test statistic is significant, increase n by 1.
5. Repeat steps 2-4 N times.
6. Find PoS = n/N.

Su [24] presented a method for calculating the success probability of a Phase III study based on Phase II data that considers both Bayesian and likelihood approaches. The calculation of the success probability is similar to the assurance, but the consideration of the component $P(\theta|\text{data})$ is different. They include a user-specified parameter that defines the trade-off between the Bayesian and likelihood approaches and are able to consider non-parametric and semi-parametric settings. Liu [25] noted that previous methods for calculating the assurance consider the uncertainty in θ, but not the uncertainty in the variance. They present the extended Bayesian expected power to consider this uncertainty in the success probability calculation, which considers the joint distribution of the treatment effect and the variance parameter in the assurance calculation. A fully Bayesian approach to calculating the probability of success was presented by Ibrahim et al. [26]. This method is more general than the calculation of the assurance, which can be found as a special case of the approach, and is able to consider covariates and patient characteristics. Wang et al. [27] present a method for calculating the probability of programme success and consider the fact that trials within the same programme are correlated in this calculation.

1.5.2 Updating

As data accumulates during a trial, or outside of a trial, we may want to update our estimate of the probability of success. This is discussed by Rufibach et al. [28], who present an approach that allows both internal and external information to be used to update the probability of success for a time-to-event endpoint. External information can include other supporting studies that complete during the trial of interest and internal information includes the information gained from an interim analysis. Wang et al. [27] provide a method for updating the probability of success at an interim analysis for a trial with binary outcomes. They also discuss how the times at which the interim analyses are planned can affect the estimate of the probability of success and the uncertainty in this value.

In this thesis, we will build on the existing methodology to consider how we can use information from related programmes to update the probability of success of a study.

1.5.3 Combination therapies

Although there is a rich literature on calculating the probability of success of a study, there is little discussion of doing this for combination studies. Wang et al. [29], however, do consider a combination study example when describing the ways that the assurance can be used in the real world.

They present an example where there was data available for the two separate components of the combination, but not of the combination itself, prior to the Phase III study. A Bayesian hierarchical model was used separately for the two different components to collate the historical data and find a posterior distribution for each of the parameters. Then, given the lack of information on the combination, they defined two different scenarios for how the combination

would perform - an optimistic one and a pessimistic one. Based on these definitions, they were able to simulate study outcomes under the two scenarios and find the associated success probabilities.

1.5.4 Decision-making

It was previously mentioned that the probability of success is one of the estimates that is frequently used to assist decision-making regarding clinical studies. In the real world, this might simply involve a study team calculating the probability of success and deciding whether or not the study is worth running. However, there are also several approaches that specifically consider decision-making and use the probability of success in this process.

Stallard et al. [30] discuss the use of Phase II data in decision-making regarding a Phase III study. It is assumed that, if the PoS of Phase III exceeds some predefined threshold, the Phase III study will be run. They also discuss the way that it can be implemented for a Phase II study with interim analyses where the study will be stopped if the PoS falls higher or lower than two critical values. The use of the probability of success in interim monitoring is also discussed in [31]. A comparison of interim stopping rules is given for Bayesian predictive, predictive power and conditional power approaches.

A more general approach for decision-making regarding the running of a study is provided in [32]. Emphasis is placed upon the importance of using new data when it emerges and a discussion of the probability of making a correct decision is provided alongside the probability of making a go decision.

Sabin et al. [33] presented a two-stage method for enhancing end of Phase II decisions. The first stage involves collating evidence linked to both the treatment and the indication and using these to address a set of key questions. The second

stage involves finding the probability of success of Phase III and consists of a four-step process.

1.5.5 Prior elicitation

One of the perceived difficulties in calculating the assurance is finding an appropriate prior distribution since the value of the PoS will rely upon these values. However, there are methods for prior elicitation that can be used to assist with this process, such as the SHELF framework [34]. Prior elicitation involves taking the opinions of experts in the indication or therapy and combining these into a single probability distribution. Typically, this will be performed with more than one expert, but it can also be used for a single opinion.

Rufibach et al. [35] provide a discussion of the choice of prior when calculating the PoS and review some of the recommendations in the literature. They provide an exploration of different types of prior and provide their own recommendations for using the PoS. These include considering the density of the power and considering non-normal priors.

1.6 Portfolio decision-making

Typically, a pharmaceutical company will have several new products in different stages of development. We call this collection a **pharmaceutical portfolio** and within the industry, this is often referred to as the "pipeline." In this thesis, our main focus is on clinical development, therefore when we refer to the portfolio, we are referring to the drugs that are currently undergoing clinical development. In this section, we outline some of the methods that are used for portfolio decision-making and provide background to the techniques that are used in these procedures.

1.6.1 Existing literature

Portfolio decision-making can cover a range of different decisions from the programmes to include within the portfolio [36] to the design that these studies should use [37] to out-licensing [38].

The decisions that are made within one drug development programme will often have an impact on the other programmes in the portfolio because a company will have finite resources for which the drugs in the portfolio will be competing. Therefore, when making decisions regarding the programmes within the portfolio it is important to take this into account. This is one of the complexities of portfolio decision-making as the comparison of different programmes and their value is not always clear. The value of a programme will be linked to many different things such as the expected revenue, the probability of programme success and the competition in the marketplace. Consequently, when designing a method for portfolio decision-making, we need to be able to consider these in order to make useful comparisons and appropriate decisions.

Existing methods for portfolio decision-making often draw upon optimisation techniques and simulation models. The two main methods that we will consider in this thesis both use **stochastic programming** approaches for portfolio decision-making. The first method was presented by Rogers et al. [36] and considers the uncertainty in the value of the programmes in the portfolio. The value of a programme is calculated using a real options approach and the value is tracked through time. The decisions that should be made after each study, assuming it is successful, are given based on the value at the time that the study is concluded. The method also provides the optimal set of programmes to include subject to budget constraints. The second method was presented by Colvin and Maravelias [39][40] and the focus of this approach is on the trial outcomes, rather than the programme value. The result of the approach is a set of schedules for each trial

outcome scenario, which can also be presented as a decision tree. These methods will be discussed in more detail in Chapter 2.

There have been a wide range of methods for pharmaceutical portfolio decision-making presented in the literature, given the variety of decisions that can be considered and the different definitions of optimality than can be specified for each of these decisions. Patel et al. [37] presented an optimisation procedure that aims to maximise the **expected net present value** (ENPV) of the portfolio by finding the optimal designs of each of the Phase III trials in the portfolio. They also extend the approach to consider the uncertainty in the availability of each of the drugs at Phase III. This is done using integer programming and stochastic integer programming. Unlike most of the other approaches in the literature, they do not consider the probability of study success to be fixed, but dependent on the design of the study. Blau et al. [41] take a different approach to the portfolio management problem and use discrete event stochastic simulation alongside genetic algorithms to find the optimal portfolio and sequence of projects. Unlike other approaches, they consider the expected positive net present value rather than the ENPV and use the negative part of the distribution to quantify risk. Further literature in this area will be discussed in Chapter 2.

1.6.2 Linear programming

The portfolio decision-making methods that we consider in this thesis use mathematical programming to find the set of **optimal decisions**. In this section, we provide details of **linear programming** before extending this to integer and non-linear programming in the next section. For further details, we refer the reader to [42].

Linear programming is an optimisation tool for finding the set of decisions that optimise a particular outcome. There are three main components to any

linear programme (LP): the **objective function**, the **constraints** and the **decision variables**.

The decision variables are used to represent the different decisions that can be made and, in a linear programme, these variables are continuous. We will use $\mathbf{x}$ to denote the vector of decision variables. The standard form of a linear programme includes non-negativity constraints on the decision variables, but this does not mean that the original variables should be defined such that they are non-negative as we can convert any linear programming problem into one of standard form [42]. We do not specify the value of the decision variables, as the optimal values of these variables will be found as part of the solution to the linear programme.

The constraint set is used to represent the constraints on the decisions that we can make and they are expressed in the standard form of a linear programme as equalities. Note that, again, this does not mean that the original set of constraints must be equalities as we can convert inequality constraints to equality constraints using slack variables [42]. We summarise the set of constraints using a matrix, $\mathbf{A}$, and a vector, $\mathbf{b}$. Then, the constraint set is given by $\mathbf{Ax} = \mathbf{b}$. Here, the matrix $\mathbf{A}$ cannot include any of the decision variables, so that the resulting constraints will be linear. The values of $\mathbf{A}$ and $\mathbf{b}$ are fixed input parameters.

The objective function links values to the decisions that we make and we can summarise these values in a vector, $\mathbf{c}$, such that the objective function is given by $\mathbf{c}^T\mathbf{x}$, which is also linear. The value of $\mathbf{c}$ is also a fixed input parameter. The standard form of a linear programme maximises this objective function, but if our original problem was one of minimisation, we could simply multiply the objective function by minus one.

The standard form of a linear programme is given by

$$\begin{aligned} &\text{maximise } \mathbf{c}^T\mathbf{x} \\ &\text{subject to } \mathbf{A}\mathbf{x} = \mathbf{b} \\ &\qquad\qquad \mathbf{x} \geq \mathbf{0}. \end{aligned}$$

Solving a linear programme, for example using the Simplex algorithm, returns the value of the decision variables in the **optimal solution** along with the associated value of the objective function [42]. Optimisation software, such as JuMP [43] in Julia [44], can also be used to model and solve linear programmes.

1.6.3 Non-linear and integer programming

A **non-linear programme** (NLP) is one that does not require the objective function or the constraints to be linear, an **integer programme** (IP) contains only integer decision variables and a **mixed integer linear programme** (MILP) contains both continuous and integer decision variables. The latter two types of programme will be considered in the later chapters of this thesis.

We cannot solve these different types of programmes using the same techniques as can be used for linear programmes. In a linear programme, the set of values for the decision variables that obey the constraint set is called the feasible region and this will be a convex set [42]. The optimal solution, if it exists, will then be at one of the corner points of the feasible region [42]. However, for a NLP the feasible region might not be convex and, even if it is, the above property of the optimal solution will not necessarily hold [42]. In an IP or a MILP, if we found the feasible region in the same way as for a LP, this would assume that all of the variables are continuous, which is not the case. Thus, if we tried to find the solution of an IP or MILP using the same approach as for a LP, this might result

in non-integer values of the decision variables at the optimal solution. Hence, different techniques will be needed to find the optimal solution for these types of programmes. It should be noted that, as for linear programmes, these types of programmes can be modelled and solved using optimisation software such as JuMP [43] in Julia [44].

1.6.4 Knapsack problem

The **knapsack problem** is one of the most famous types of integer programme. This problem contains only binary decision variables and a single constraint [42]. This type of programme is referred to as the knapsack problem as it can be thought of as the problem of packing a knapsack such that the value of the items packed is maximised. In this problem, we have n different items that we would like to pack in the knapsack and we must decide which ones to pack, assuming that we cannot pack them all. Hence, the decision variables, $\mathbf{x}$, will be binary variables where $x_i = 1$ will correspond to the decision to pack item i in the knapsack.

Since our aim is to maximise the total value of the items packed, each item must be assigned a value in advance. We will denote the vector of values by $\mathbf{c}$. We assume that the knapsack has a weight capacity, b, and that each item has an associated weight. We will use $\mathbf{a}$ to denote the vector of weights.

The integer programme for the knapsack problem is then given below.

$$\begin{aligned} &\text{maximise } \mathbf{c}^T\mathbf{x} \\ &\text{subject to } \mathbf{a}^T\mathbf{x} \leq b \\ &\qquad x_i \in \{0,1\} \ \forall i \in 1,...,n \end{aligned}$$

The knapsack problem has been extended in many different ways, for example it has been extended to consider multiple dimensions [45] and quadratic objective functions [46]. This makes it a very useful problem and it can be applied to many different settings such as online advertising [47] and the cutting stock problem [48], which aims to find the cutting pattern of a material that reduces waste. In Chapter 4, we will provide an example of how the knapsack problem can be used in the setting of scheduling clinical studies in a portfolio, which was originally presented in [49].

1.6.5 Stochastic programming

Stochastic programming is the main focus of Chapter 2 and so we will limit this section to a basic outline of the topic.

In the types of mathematical programme discussed previously, there was no uncertainty considered. However, in the real world, and indeed in the problem of portfolio decision-making, uncertainty plays an important role. In order to account for the uncertainty that we may have regarding outcomes or parameters in a mathematical programme, we can use stochastic programming.

A **multi-stage stochastic programme** (MSSP) considers the fact that uncertainty might be revealed at different points, or stages, in the planning horizon and takes this into account when finding the set of optimal decisions. At each stage, decisions are made based on the current state of events, whilst also accounting for the possible future outcomes. Then, dependent on the uncertainty that is revealed at the next stage, we are able to take **recourse action** to account for the effect of the uncertain event.

Scenario based stochastic programming breaks down the uncertainty into a set of scenarios that summarise all of the different potential outcomes. These

scenarios and their associated probabilities are taken into account when finding the optimal decisions.

Further details on stochastic programming, specifically in the context of pharmaceutical portfolio decision-making, will be given in Chapter 2, where we provide a detailed comparison of the two main methods of interest.

1.7 Thesis outline

This thesis will discuss methods for portfolio decision-making and extend existing methodology to consider combination therapies. In Chapter 2, we will provide a critical discussion of two existing methods for portfolio management that draw upon stochastic programming techniques. We conclude the chapter by making a recommendation of the most appropriate method for a portfolio of combination therapies and briefly discuss how this method will be extended.

The probability of study success and its use in decision-making is discussed in Chapter 3. We discuss the potential to learn across related combination studies and present a method that allows the probability of success of a study to be updated based on related study data.

In Chapter 4, this method is incorporated into a heuristic for one of the portfolio management approaches from Chapter 2. This allows the results of combination studies to be accounted for in the decision-making process for related studies, which previous methods were not able to achieve. This novel approach for decision-making in portfolios of combinations is compared to existing portfolio management techniques and further extensions are presented.

The thesis is concluded with a summary of the main contributions along with the limitations of the work and potential areas for further research.

CHAPTER 2

Chapter 2

A comparison of stochastic programming methods for portfolio-level decision-making

2.1 Introduction

It was noted in Chapter 1 that the drug development process is long, expensive and contains a high level of uncertainty. The biggest source of uncertainty comes from the fact that the outcomes of the clinical studies are unknown. Supposing a novel treatment does perform well in all of its associated studies and is approved and then launched, it will also encounter uncertainty in the revenue that it will generate and the impact that it will have due to the many external factors that affect these outcomes, such as competitors in the marketplace and demand [50].

Typically, a pharmaceutical company will have several products undergoing clinical development within their portfolio including new products and existing products that are being tested in a different indication. One challenge that arises when we consider portfolios rather than individual drug development programmes comes from the fact that the decisions made within one drug development programme are likely to have an effect on the other programmes in

the portfolio. This is because a pharmaceutical company will have a finite level of resources and budget for which the drugs in the portfolio compete. Consequently, within a pharmaceutical portfolio the decisions will include selecting the studies that should be prioritised, considering when to either abandon or defer development for certain drugs given the current state of the portfolio and other logistic decisions such as scheduling and budget/resource allocation.

One of the biggest challenges within drug development and portfolio management is the previously mentioned stochasticity of the underlying process. While a pharmaceutical company will make sure that care is taken in ensuring that the drug they are developing will be beneficial to both the company and the target patient population, the true performance of a drug is an unknown parameter and can only ever be estimated given relevant data. Therefore, ensuring that portfolio decisions are well informed and consider the associated uncertainties in each stage of development is of high importance. This will allow resources to be allocated to where they will have the highest potential for benefit, especially when this is considered alongside the expected costs and potential revenue in comparison with the other drugs in the portfolio.

Several methods have been presented in the literature for the management of a pharmaceutical portfolio in terms of selecting which studies should be conducted. These methods typically employ optimisation techniques but draw upon a range of different areas and often have different focuses.

Schmidt and Grossman [51] presented a model for scheduling non-sequential testing tasks with an aim to maximise the expected net present value and provided several reformulations that assist in solving the model and focus on different aspects of the problem. Rotstein et al. [52] presented a two-stage stochastic programme that also aims to maximise the expected net present value of the portfolio. This model considers decisions such as capacity investments, product selection and resource allocation in the first stage and capacity adjustments in

the second stage. Blau et al. [41] model drug development programmes using probabilistic network models. Bubble charts are used to find a prioritisation scheme for the drugs that is then used as a starting point for the genetic algorithm based search. The search takes into account interdependencies between drugs and resource constraints. The aim of the method is to find the solution that maximises the expected net present value for a given level of risk. Varma et al. [53] presented a method that combines stochastic simulation and mixed-integer linear programming in order to maximise the expected net present value of the portfolio while also evaluating different strategies on the pipeline. Sundaramoorthy et al. [54] presented a multi-scenario multi-period mixed-integer linear programme that aims to maximise the expected net present value of the portfolio. The decisions considered in the formulation include things such as building/expanding facilities, capacity decisions, production levels and storage decisions but the approach considers no resource constraints during product development.

In this chapter, we will focus on two approaches that are of particular interest because they provide models that capture the core aspects of pharmaceutical portfolio management. These approaches focus on the process of decision-making when we are considering multiple drug development programmes and consider how we might compare different programmes that are in different stages of development rather than focusing on different aspects of production. They also provide clear results and information on the decisions that should be made to help the decision maker achieve their goal. These are the reasons why we will limit our attention to these approaches.

Both of the methods that we will discuss draw upon **stochastic programming** techniques and formulate the problem using mixed integer linear programmes. Stochastic programming is beneficial in the setting of pharmaceutical portfolio management due to the inherent stochastic nature of the process. Stochastic

programming allows us to model the uncertainty of the process and let this contribute to the decisions that are made. Furthermore, it allows us to consider what the optimal decisions might be based on different outcomes of the uncertain process. This will then allow decision makers to consider the impact of certain decisions and to compare different sets of decisions in terms of the costs incurred and potential benefits.

In Section 2.2, we will describe the two approaches that are the main interest of this chapter. We will highlight that, while both use similar methodology, the focus of the approaches is actually quite different. The first approach [36] draws upon **real option valuation** from the financial setting and the focus of this approach is the stochasticity in the value of the drug development programmes. The second approach [39][40] is similar to the formulation of the **resource constrained project scheduling problem** and the stochasticity considered here is in the uncertainty of the trial outcomes.

2.2 Stochastic programming methods

In this section, we will review two portfolio management approaches that are based on stochastic programming and formulate the decision-making process as a mixed integer linear programme (MILP). We will then provide a comparison and critical discussion of the implementation of these approaches in the next section.

The methods that we will review in this section were presented by Rogers et al. [36] and Colvin and Maravelias [39]. Both methods model the decision-making process as **scenario-based multi-stage stochastic programmes**. In a scenario-based multi-stage stochastic programme, the **scenarios** correspond to realisations of the vector of random variables and the **stages** correspond to the times at

which some uncertainty is resolved. The two approaches that we will discuss consider different types of uncertainty; hence, the scenarios and the stages will be different in each approach. Both methods consider the pharmaceutical planning horizon and the potential decisions that can be made in the planning horizon (e.g. continue or abandon development at each stage) and when the programme is solved, the optimal set of decisions is returned along with the value of the optimal solution.

Stochastic programming is beneficial when we are modelling a process that involves randomness as it allows us to take into account the uncertainty in the underlying process and considers the **recourse action** that should be taken given different observations of the random variable. The recourse action is the decision that should be made in order to compensate for the effect of what has just been observed. The recourse action relating to each uncertain observation will be contained in the solution to the stochastic programme. For example, if the observation was that the value of the drug dropped significantly over the most recent stage then the recourse action might be to abandon development. Alternatively, if the observation was that a study failed in the sense that the drug was shown to lack efficacy or be harmful then the recourse action might be to allocate resources to the development of a different drug.

Stochastic programming is able to consider the different potential outcomes of the uncertain component when finding the best set of decisions overall. For example, let us consider the problem of scheduling studies with uncertain outcomes. A stochastic programme would consider all of the potential combinations of study outcomes and return the set of initial studies to run along with those that should be run in the event of the different potential study outcomes. We could represent this using a set of Gantt charts, one for each trial outcome scenario, or using a decision tree.

If we do not want to use a stochastic programme, we could instead consider

the problem in a **deterministic** setting using expectations. For example, in the scheduling problem, we could calculate the expected resource requirements and expected revenues using the probability of trial success. We could use these to build a deterministic programme that does not consider the study outcomes. This deterministic model will assume that we are able to run all studies, regardless of the outcome of a preceding study. Solving the programme would return a single schedule that maximises the expected revenue subject to constraints on the expected resource requirements under the assumption that we are able to run everything. If we do this, then we would need to build a new programme after we observe a study outcome to reflect the new information gained and solve this programme for the next set of decisions.

A deterministic model is not able to consider the different possible outcomes of the stochastic element, or the effect that these will have on the optimal decisions. Therefore, the recourse action that is considered in a stochastic programme is neglected when we model a stochastic problem deterministically. It is likely that the solution found using this approach would be suboptimal compared to the solution that is found when it is modelled using stochastic programming [39]. When we consider the problem deterministically, it cannot take into account the fact that it may be beneficial to wait to observe certain outcomes before making some decisions. A stochastic programme, however, is able to account for future outcomes and decisions when finding the optimal decisions.

Colvin and Maravelias [39] provide a comparison between using stochastic and deterministic models for a small portfolio management problem. The same information is used in both of the models, but the deterministic model uses the study success probabilities, resource requirements and revenues to find the expected resource requirements and revenues. When solved, the deterministic approach selects to run two studies simultaneously for which the future studies cannot be run together due to resource constraints. The stochastic method only

Nomenclature	
I	the set of all drugs under consideration
i	the index used to denote a specific drug in the set I
J_i	the set of all studies remaining for drug i
j	the index used to denote a specific study
(i, j)	the notation denoting the jth ordered study for drug i
N_{ij}	the number of value scenarios at the beginning of (i, j)
k_{ij}	the index used to denote a particular value scenario at the beginning of study j
S	the set of all possible trial outcome scenarios
s	the index used to denote a particular trial outcome scenario in the set S
T	the set of all time points considered in the planning horizon
t	the index to denote a specific time point in the planning horizon
V_{0i}	the present value of drug i at $t = 0$
c_{ij}	the cost of study (i, j)
ϕ_{ij}	the probability of success of study (i, j)
τ_{ij}	the duration of study (i, j)
ξ_i	the market volatility of drug i
γ_i^D	the loss per time period due to shorter active patent life of drug i
γ_i^L	the loss per time period due to smaller market share of drug i
B_t	the available budget at time t
r	the risk free interest rate

Table 2.1: Notation used in the ROV and PS approaches.

chooses to run one of these studies and then waits until the study has been completed in order to make the decision for whether to run the second study or continue development of the first. This leads to a higher expected net present value of the solution.

2.2.1 Real option valuation (ROV) approach

Rogers et al. [36] noted that the sequential nature of the investments made for each study in a drug development programme are comparable to a series of call options. A call option is the right but not the obligation to buy an asset by a given future date for a specified price [55].

Real options are similar to financial options but, instead of the asset of interest

being a financial asset, it is a real, non-financial asset. This means that, rather than having the right to buy the underlying asset by a future date, we instead have the right, but not the obligation, to take an action on the asset by a future date where the action could be to continue development, for example [55].

When a pharmaceutical company invests in the current stage of a drug development programme this in turn gives them the option to invest in later stages, should the current stage be successful. The asset in this setting is the present value of the future cash flows of the product should it be commercialised. The cost of buying the real option is the cost of the current study and the predetermined price of the asset is the cost of future studies [36]. For example, after investing in a Phase I study, we have the chance to invest in a Phase II study and potentially a Phase III study if Phase I was successful enough to be carried forwards.

Drawing these parallels between real options and drug development programmes allows us to use methods for real option valuation to assign values to each of the drug development programmes [36]. The real option value of each of the drug development programmes, as presented in [36], takes into account many different aspects that affect the value of a development programme including the uncertainty in the trial outcomes, the potential market movements throughout the development process and the potential to abandon development.

After each of the drugs within the portfolio have been assigned a **real option value** (ROV), we will have an ordering of the most attractive programmes to run where the programme with the highest ROV is the most attractive and the programme with the lowest ROV is the least attractive. However, the pharmaceutical portfolio management problem is a real world problem with finite resources. Therefore resource constraints must also be included in the decision-making process. Hence, the selected portfolio will not always contain the most attractive programmes if they do not satisfy the constraints of the model.

In order to find the real option value of the drug development programmes, Rogers et al. [36] use a quadranomial pricing approach. This approach considers the **market movements** at discrete time intervals and assigns probabilities to the movement being either upward or downward. These movements correspond to the value of the drug increasing or decreasing, respectively. The market movements are represented by a multiplier that is calculated using the standard deviation of the value of drug i, ξ_i, which represents the beliefs that the team has about the **volatility** of the value of the drug in the marketplace. This value can be predicted by looking at historical data for similar products. The **multipliers** for upward and downward market movements are then given in [36] to be

$$u_i = \exp\left(\xi_i \sqrt{\Delta T}\right) \quad \text{and} \quad d_i = 1/u_i \tag{2.1}$$

respectively, where ΔT is the discrete time interval that the market movements are considered over. These potential movements over each time step, u_i and d_i, will be treated as constant through time. That is, the multiplier used to calculate the value after the market movement will not depend on where we are in the planning horizon.

At the end of each study, (i, j), there will be a set of possible values for the drug based on the different combinations of upward and downward market movements over the length of the study. For example, if we consider a study that lasts for two time intervals then there will be three possible values of the drug at the end of the study corresponding to: two downward market movements, one downward and one upward market movement or two upward market movements.

We will refer to these possible final values as the **value scenarios** and these are the scenarios that make up the scenario-based multi-stage stochastic programme presented in [36]. We will denote the value scenario at the end of study (i, j) by k_{ij+1} where the set of all value scenarios at the end of study (i, j), or equivalently

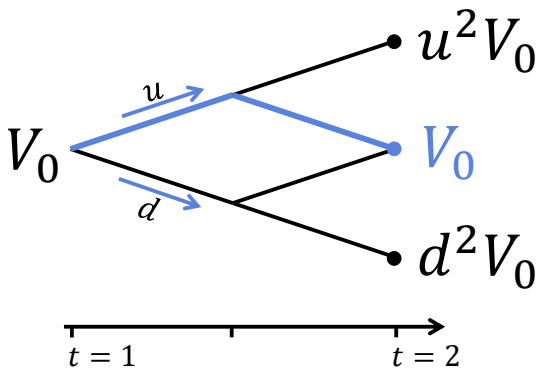

Figure 2.1: Diagram showing a potential path of market movements as considered in the discrete pricing approach [36].

at the beginning of study $(i, j+1)$, is given by $\{1, 2, ..., N_{ij+1}\}$. Using this notation, $k_{ij} = 1$ will correspond to the worst value scenario (the scenario with the lowest value) and $k_{ij} = N_{ij}$ will correspond to the best value scenario (the scenario with the highest value) at the beginning of study (i, j). At the beginning of the first study for drug i in the planning horizon, we will only have one value scenario, $N_{i1} = 1$, as we have a fixed estimate of the present value of the product at the initial time point.

An example of potential market movements can be seen in Figure 2.1. In the diagram in Figure 2.1, we have $N_{i2} = 3$ where $k_{i2} = 1$ represents the scenario with two downward market movements over the course of the study and final value d^2V_0 and $k_{i2} = 3$ represents the scenario with two upward movements over the study and final value u^2V_0. A possible path is highlighted, which leads to scenario $k_{i2} = 2$ and consists of an upward market movement followed by a downward market movement resulting in a final value of V_0.

The ROV approach does not only consider the market movements but also the probability of study success, ϕ_{ij}, and the potential to either continue or abandon development dependent on what is observed. Note that, in this setting, when we refer to study success we are referring to the situation where the drug may continue to further stages of development after the study in question is con-

cluded. Also, this success probability is conditional on the previous studies being successful. There are many different approaches for calculating the probability of study success, which were discussed in Chapter 1, and the ROV approach does not require a particular definition or method of calculation for this input parameter, but we do recommend that the same method of calculation is used across the different drug development programmes for comparability. This success probability is included in the calculation of the ROV of the drug, which is given below.

Under the discrete pricing approach for real options, the value of drug i at the beginning of study j in value scenario k_{ij} is given by

$$M_{ijk_{ij}} = \max\left\{-c_{ij} + \frac{\phi_{ij}\sum_{k_{ij+1}=1}^{N_{i,j+1}} p_{ik_{ij}k_{ij+1}} M_{i,j+1,k_{ij+1}}}{(1+r\Delta T)^{\tau_{ij}/\Delta T}}, 0\right\} \tag{2.2}$$

where, for drug i and study j, $p_{ik_{ij}k_{ij+1}}$ is the probability of moving from scenario k_{ij} to k_{ij+1} during study j and ΔT is the length of the time interval that we consider the market movements over [36]. Then, as given in Table 2.1, τ_{ij} is the study duration, c_{ij} is the study cost and ϕ_{ij} is the study success probability. We could rescale this by setting $\Delta T = 1$ and adjusting the study durations accordingly. However, we will not rescale in order to provide commentary on the original approach and compare it to the PS approach.

This is a recursive formula for which we begin at the expected reward received during launch. The reward received during launch in each scenario is given by

$$M_{i,|J_i|+1,k_{i|J_i|+1}} = u_i^{k_{i|J_i|+1}-1} d_i^{N_{i,|J_i|+1}-k_{i|J_i|+1}} V_{0i}. \tag{2.3}$$

We can then iteratively work backwards from this reward to find the values at the beginning of each study in each value scenario for drug development programme i.

Let us consider a simple example of a single drug with a single study with $V_{01} = 50$, $c_{11} = 10$, $\phi_{11} = 0.8$, $\xi_1 = 0.6$, $\tau_{11} = 1$ and $r = 0.05$. We will consider two market movements per time step, $\Delta T = 1/2$. The number of value scenarios at the end of the study will be given by $N_2 = 3$, and at the beginning we have $N_1 = 1$. This is the same setting as was considered in Figure 2.1. We use the values associated with the final scenarios to find the values $M_{12k_{12}}$, as in Equation (2.3). This gives

$$\begin{aligned} k_{12} = 1: \;\; M_{121} &= u^{1-1}d^{3-1}V_0 = 21.4 \\ k_{12} = 2: \;\; M_{122} &= u^{2-1}d^{3-2}V_0 = 50 \\ k_{12} = 3: \;\; M_{123} &= u^{3-1}d^{3-3}V_0 = 116.8 \end{aligned}$$

where $u = 1.53$ and $d = 0.65$, to two decimal places, using Equation (2.1). The transition probabilities in this example are given by $p_{111} = 0.33$, $p_{112} = 0.49$ and $p_{113} = 0.18$. We will not go through the details of the transition probabilities here, however, and refer the reader to [36] for the full details. We can find the value of M_{111} using Equation (2.2).

$$M_{111} = \max\left\{-10 + \frac{0.8\,[0.33 \times 21.4 + 0.49 \times 50 + 0.18 \times 0.18]}{(1 + 0.05 \times 0.5)^{1/0.5}}, 0\right\} = 30.02$$

The form of $M_{ijk_{ij}}$ given in Equation (2.2) does not take into account the fact that we may choose to abandon development due to limited resources, for example, and then the value of the drug would be equal to zero, as it cannot add any value to the portfolio if it is not a part of the portfolio. Therefore the calculation of $M_{ijk_{ij}}$ must be reformulated to include the continue/abandon decision variable $Y_{ijk_{ij}}$ that is equal to one when study (i, j) is continued in scenario k_{ij} and zero otherwise.

The objective function of this programme, which will be defined later in the section, is to maximise the overall value of the portfolio. Hence, if $M_{ijk_{ij}}$ dropped

below zero then $Y_{ijk_{ij}}$ would be set equal to zero, as it would not be profitable to continue with the study. Therefore, we can write the reformulation as

$$M_{ijk_{ij}} = \left[-c_{ij} + \frac{\phi_{ij} \sum_{k_{ij+1}=1}^{N_{i,j+1}} p_{ik_{ij}k_{ij+1}} M_{i,j+1,k_{ij+1}}}{(1 + r\Delta T)^{\tau_{ij}/\Delta T}} \right] \times Y_{ijk_{ij}}. \tag{2.4}$$

This reformulation satisfies Equation (2.2) whilst also allowing for abandon decisions, $Y_{ijk_{ij}} = 0$, which can be related to limited resources.

In our simple example, the value of the drug at the beginning of the study was given by $M_{111} = 30.02$. Since this value is positive and we only considered one drug in the example, we would expect to select to run this study, which would correspond to $Y_{111} = 1$. However, if the available budget was less than the cost of the study, $c_{11} = 10$, we would not be able to run the study. This would lead to $Y_{111} = 0$ and, from the reformulation given in Equation (2.4), we would have $M_{111} = 0$ as we have not been able to run the study and therefore it has not added any value to the portfolio.

When $M_{i11} > 0$, the value of M_{i11} is equal to the ROV of drug i as there is only one value scenario, $k_{i1} = 1$, at the starting point in the planning horizon, $j = 1$. When $M_{i11} = 0$ in the optimal solution this means that drug i has not been selected as part of the optimal portfolio.

The calculation of the values of $M_{ijk_{ij}}$ is core to both the decision-making process and the model formulation of the ROV approach. The uncertainty modelled in this approach is in the value of the drug and the values calculated above will be the values that we want to maximise in the decision-making process. For example, if $M_{111} > M_{211} > 0$ this means that drug $i = 1$ is preferable to drug $i = 2$ as drug $i = 1$ has a higher ROV, which is given by M_{i11}. Furthermore, the objective value of the optimisation will be to maximise the sum of M_{i11}. This is because the drugs with the highest values of M_{i11} are deemed the most attractive

under this approach and hence our aim will be to select the drugs that lead to the maximal value of $\sum_i M_{i11}$ subject to practical constraints.

The formulation of the decision-making process is presented in [36] as a mixed integer linear programme, which means that the objective function and the constraints of the model are linear in terms of the variables. The variables for this formulation are the values, $M_{ijk_{ij}}$, and the continue/abandon binary decision variables, $Y_{ijk_{ij}}$; the values of these variables in the optimal solution are found by solving the MILP.

This means that the above formula for $M_{ijk_{ij}}$ should be added as a constraint to the MILP so that the values can be found by solving the MILP. However, the form of this constraint as given above is not linear in terms of the variables, $M_{ijk_{ij}}$ and $Y_{ijk_{ij}}$. Therefore, in order to include this in the MILP, a reformulation is required to linearise this constraint. This reformulation requires upper bounds on the values of $M_{ijk_{ij}}$ to be found by solving a separate LP and adding these upper bounds as inputs to the final MILP in order to find the optimal solution. For full details of the linearisation, see Appendix A.1.

The objective function for this formulation, as mentioned previously, is given by maximising the real option value (ROV) of the portfolio and can be written as

$$\text{maximise ROV} = \sum_i M_{i1k_{i1}}.$$

This is subject to constraints including: the calculation of the values $M_{ijk_{ij}}$; drug precedence constraints to ensure that future studies are not selected when a previous study was abandoned; value monotonicity constraints to represent the fact that if an abandon decision is made in a particular value scenario then the same decision must be made for all worse value scenarios; investment constraints to ensure that the expected budget required at a particular time point, t, does

not exceed the available budget, B_t. For a full model formulation, see Appendix A.1.

The investment constraints consider the expected budget required at each time point, t, and assume that each study begins as soon as possible and the cost of a study is incurred at the commencement of the study. These constraints are given by

$$\sum_{i,j} \sum_{k_{ij}=1}^{N_{ij}} p_{ik_{ij-1}k_{ij}} c_{ij} Y_{ijk_{ij}} w_{ijt} \leq B_t \ \forall t$$

where w_{ijt} is an indicator for if study (i, j) starts at time t. Here, the expectation is taken over the different value scenarios since the market value of the drug is the uncertainty that is modelled in this approach. The probability of study success, ϕ_{ij}, is not used directly in calculating the expected budget required, but it is included in the calculation of the values $M_{ijk_{ij}}$. These values affect the continue/abandon decisions, $Y_{ijk_{ij}}$, which are included in the budget constraint. If the value of $M_{ijk_{ij}}$ dropped below zero, for example due to low success probabilities, then $Y_{ijk_{ij}}$ would be set to zero and this would reduce the expected budget required at the time that study (i, j) is initiated.

Solving the resulting mixed integer linear programme returns the set of values, $M_{ijk_{ij}}$, and the continue/abandon decisions, $Y_{ijk_{ij}}$, in the optimal solution. In the ROV approach, each drug development programme can be thought of as a series of continue/abandon decisions that are dependent on the value scenario of the programme at a particular time point. This can be represented by a diagram such as the one seen in Figure 2.2.

The diagram in Figure 2.2 shows the decisions that should be made after each phase in a single drug development programme, assuming that the phase was successful, based on the value of the programme at that point. For example, at the end of Phase I, if the trial was successful, the decision should be to continue no matter what the value of the programme is at this point. At the end of Phase II,

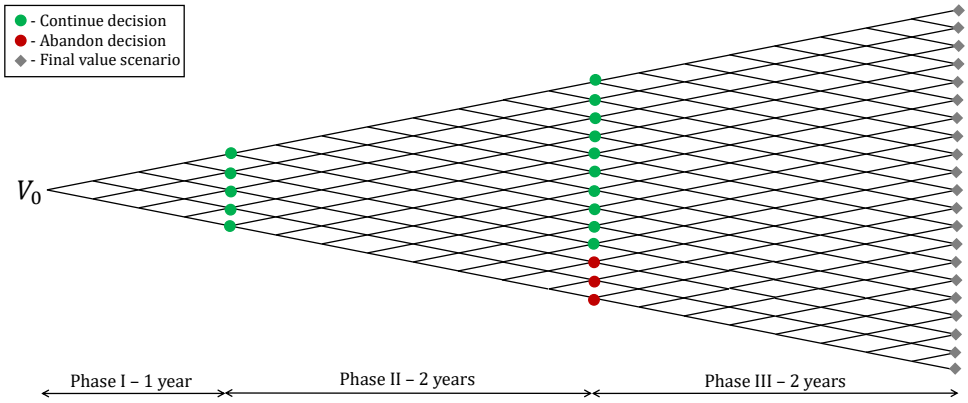

Figure 2.2: Diagram showing the continue/abandon decisions dependent on the value scenario at the end of Phases I and II for an example drug development programme.

if it is observed that the value of the programme is less than or equal to d^8V_0 then the team should abandon development, whereas if the value of the programme at this point is greater than d^8V_0 then the team should continue development if Phase II is successful.

Note that no scheduling is considered in this approach; the decision points are at the completion of each of the studies hence the next study is either started as soon as possible or not at all. This is because the method does not model the study outcomes; it models the uncertainty in the market value. In addition, if scheduling was considered and studies were allowed to be delayed then this would affect the ROV of the drugs in the portfolio. The focus of this approach is on portfolio selection rather than portfolio scheduling. The fact that a study may not be successful is considered by including the study success probabilities in the calculation of the values of $M_{ijk_{ij}}$, which also affects the expected budget required. The full details of this formulation can be found in Appendix A.1.

An alternative formulation for this approach was presented by Lo Nigro et al. [56], which uses a continuous pricing approach to find the ROV of each of the drug development programmes rather than the binomial pricing approach. This reduces the size of the formulation in terms of the number of variables and number of constraints and decreases the computational time required to find an

optimal solution.

Lo Nigro et al. [56] also presented two extensions to the model: reinvestment of attained profits in the future and joint development with another partner company. However, this approach does not use stochastic programming, which means that it does not consider the recourse action that should be taken under different value scenarios; it simply selects the programmes to include at the initial time point.

2.2.2 Project scheduling (PS) approach

A second stochastic programming method was presented by Colvin and Maravelias [39][40] that compares the problem to the stochastic version of the resource constrained project scheduling problem. Hence, this approach considers the scheduling of the tasks, unlike the previous approach. In this approach, we consider a set of projects that correspond to drug development programmes. Each of these projects is made up of a series of tasks, which correspond to the clinical studies within the programme. In order to complete a project, all associated tasks must be completed successfully. In the setting of a pharmaceutical portfolio, this corresponds to a set of drug development programmes containing studies that must be successfully completed in order to complete the programme. Unlike the ROV approach, which considers the value of the programmes to be stochastic, the stochasticity that is considered here is in the **trial outcomes**. The value of the programme, if successful, is not considered to be stochastic and is a linearly decreasing function of time in this approach.

Colvin and Maravelias [40] note that previous stochastic programming methods for resource allocation problems treated the timing of the tasks as fixed, as was seen in Section 2.2.1 when each trial was considered to start as early as possible and no later. The case that is considered in the PS approach treats both the

outcomes of the developmental tasks and the timing at which these outcomes are observed as uncertain. This is because the scheduling of the tasks is dependent on what is observed and the timings of the observations are dependent on the schedule. It is further noted in [40] that this type of stochastic programme is difficult to solve and hence most approaches use the repeated solution of deterministic models rather than solving the full stochastic programme. Since this does not allow for the consideration of the "wait and see" approach this may lead to optimal deterministic solutions that are in fact suboptimal compared to the stochastic solution.

The problem is modelled in [39] as a scenario-based multi-stage stochastic programme and formulated using a mixed integer linear programme, which requires similar information to that in the ROV approach. In this setting, however, the scenarios correspond to the different combinations of trial outcomes rather than the value of the drug development programmes.

Let us consider the simple case where there is a series of studies that must be completed sequentially and failure in a study means that no further studies can be run. Then we can represent the outcome of each drug development programme by the number of successful studies and hence each scenario will correspond to the number of successful studies in each drug development programme. Hence the number of scenarios will be given by

$$|S| = \prod_{i \in I} (|J_i| + 1) \tag{2.5}$$

where J_i is the set of studies of drug i. For example, if we consider a portfolio containing three drugs, $|I| = 3$, with the number of remaining studies being given by $|J_1| = 3$ and $|J_2| = |J_3| = 2$ then the number of scenarios will be given by $|S| = 4 \times 3 \times 3 = 36$ and the set of scenarios can be denoted by

$$S = \{(0,0,0), (0,0,1), (0,0,2), \cdots, (3,2,1), (3,2,2)\}$$

where the ith element of each scenario, $s \in S$, gives the number of successful studies for drug i in scenario s.

We can find the probability of each scenario, $p(s)$, using the study success probabilities, ϕ_{ij}. If we consider the scenario $s = (0, 0, 2)$ from the above example, then the probability of this scenario would be given by

$$p(s) = (1 - \phi_{11}) \times (1 - \phi_{21}) \times \phi_{31}\phi_{32}.$$

The study success probabilities were also a required input parameter for the ROV approach and, as in the ROV approach, there are no assumptions on the modelling technique used to calculate these parameters but we would again recommend using the same approach across different drug development programmes to avoid misleading results.

The decision variable in this approach is given by X_{ijts}, which is equal to one when study (i, j) is chosen to start at time t, where $t = 1, \ldots, |T|$, in scenario s, where $s \in S$. Note that, in this approach, the decision variable is also dependent on time due to the fact that this formulation considers scheduling. The model also requires variables that track when a study is completed and when a study is able to start. These variables will be equal to either zero or one, but they do not need to be specified as binary variables since they are calculated using the binary decision variable, X_{ijts}, thus are naturally constrained to be zero or one. This is what makes the programme mixed integer rather than integer.

The constraints of this formulation are largely similar to those in the ROV approach: each trial can be performed at most once; resource requirements at each time point must not exceed limits (note that budget can be considered as a resource in this approach if required); studies must be completed in the correct order; a study must not be performed if a previous study has failed. The resource/budget constraints in this approach are quite different, however, to

Number of drugs in portfolio	3	4	5
Non-anticipativity constraints	41472	294912	1843200
Other constraints	15024	79104	392448
Variables	14080	74752	372736

Table 2.2: Number of constraints and variables in the PS approach for example portfolios containing 3, 4 and 5 drugs each with $|J_i| = 3$ and $|T| = 8$.

their alternatives in the ROV approach. Firstly, multiple resources, $r \in R$, can be considered in this approach, rather than just the budget and $\lambda_r^{\max}$ represents the total resource available for type r at each time point. The exact resources required are used, as opposed to the expected resources required, since the constraints are considered in each individual trial outcome scenario. In addition, this approach typically considers resources to be required throughout the study, rather than being incurred at the beginning of the study. This, however, can easily be modified to incur these costs at the beginning of the study, as in the ROV approach, which we will do in the next section to make a fair comparison of the approaches.

The main difference in the constraint set from the ROV approach is the addition of **non-anticipativity constraints** (NACs). In this approach, the scenarios correspond to the study outcomes. Since the decision variables are dependent on the scenarios, we must ensure that the stochastic programme does not exploit the scenario information before the corresponding uncertainty has been resolved - it must not anticipate future outcomes. Including non-anticipativity constraints vastly increases the scale of the problem, as can be seen in Table 2.2. Therefore, some possible model reductions were presented in [39] and [57] that consider, for example, the structure of the problem and the fact that expressing certain subsets of NACs ensures that all NACs are satisfied.

The objective function for the PS approach maximises the **expected net present value** (ENPV) and is given by

$$\text{maximise} \quad \text{ENPV} = \sum_s p(s) \{Rv_s + FRv_s - Cst_s\}$$

where: Rv_s is the overall total **revenue** generated in scenario s; FRv_s is the **future revenue** in scenario s supposing ongoing drug development programmes are completed; Cst_s is the total development **cost** in scenario s. These values are functions that are dependent on the decisions, X_{ijts}, made in each of the scenarios during the planning horizon. Both the revenue, Rv_s, and future revenue, FRv_s, are linearly decreasing in time, which encourages trials to be run earlier rather than being postponed until the end of the time frame. The components included in these functions that make them linearly decreasing include a reduction for reduced active patent life, γ_i^D, and a reduction for late completion hence reduced market share, γ_i^L. This means that it is better to run a programme earlier in the time frame and without delay.

The revenue, Rv_s, is used to capture the total value of the drug development programmes that have had all studies initiated during the planning horizon in scenario s. The future revenue, FRv_s, however, aims to capture the value that has been gained by running additional studies for drug development programmes that have not yet been completed in scenario s. This encourages the selection of studies, even if the programme cannot be completed during the planning horizon, which allows the model to consider further into the future and reflects the way that decisions would be made in the real world. The calculation of FRv_s assumes that the first study that has not been initiated in each drug development programme during the planning horizon is either initiated at the end of the planning horizon or upon the completion of the preceding study in the programme, whichever is later. It then assumes that all subsequent studies are initiated upon completion of their preceding study. For a full formulation, see Appendix A.2.

Solving the MILP returns not only the optimal portfolio but also the optimal schedule under each scenario, which can be represented using Gantt charts, as seen in Figure 2.3. This way, after each trial outcome is observed the decision makers can discard the set of schedules that have scenarios that do not match the

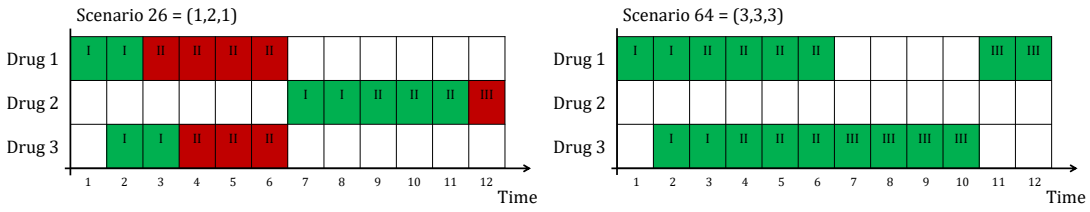

Figure 2.3: Gantt charts representing the optimal schedule for two scenarios in a simple three drug example. The shaded regions correspond to ongoing tasks and the numbers in the boxes correspond to the values of $j \in J_i$ that are ongoing at each time point.

observed outcome and use the schedules that do correspond to what has been observed so far.

In Figure 2.3, the effect of the NACs can be seen as (D1, PII) is selected to be run in both scenarios despite the fact that it will be unsuccessful in Scenario 26. Until the point where we observe the difference in these scenarios at the end of (D1, PII), the schedules are the same. We also see that the development of Drug 3 is delayed such that (D1, PII) and (D3, PII) complete at the same time. This is due to there not being enough available resources to run (D1, PII) and (D3, PIII) at the same time in this formulation that considers resources to be required throughout a study hence the method has chosen a schedule that means that (D3, PIII) would not need to be delayed should (D3, PII) be successful. As mentioned previously, delaying a study in the middle of the development programme results in a smaller revenue due to a reduced active patent life in this method. Furthermore, in this example the revenue reduction, or "penalty", incurred due to a shorter active patent life was larger than that which was incurred due to a reduced market share. This means that it is preferable in some circumstances to run a full drug development programme later in the planning horizon without delays between studies than it is to start it early and delay certain studies within the programme. This is because the method assumes that the patent is filed at the beginning of the first study in the development programme.

	ROV approach	PS approach
Stochasticity modelled	Programme value	Trial outcome
Method of valuation	Real option valuation	Expected net present value
Scheduling	No	Yes
Recourse action	Continue/abandon decisions based on value	Selection of next study to run (over all drugs)

Table 2.3: Summary of the main differences between the ROV approach [36] and the PS approach [39].

2.3 Comparison

A summary of the key differences that will be discussed in this section can be found in Table 2.3. The examples discussed in this section were generated and solved using the JuMP package [43] in Julia [44] with the Cbc solver.

2.3.1 Comparison of approaches

The main difference between the two approaches is the choice of which element is modelled as stochastic i.e. which uncertain element we observe and track over the decision-making process. The ROV approach [36] considers the value of the drug development programmes to be stochastic and tracks the potential market movements over time. The standard deviation of the value of the drug is used in calculating the possible value scenarios. The potential failure of a study is only included via the probability of success, which is used in the calculation of both the expected cost/resource requirement at each stage and the values associated with the drug development programmes. The PS approach [39] considers the trial outcomes to be stochastic and does not consider potential market movements or the variance of the value of the drug. The trial outcomes are tracked and future decisions depend on the trial outcomes. However, there is potential to include the variance in the value of the drug due to external competition into the PS

approach via the specification of the objective value and different penalties for drugs that have higher competition in the market place.

The stochasticity in the two approaches is considered through the generation of scenarios. In the ROV approach, these were value scenarios that captured the value of each development programme after each study, or stage of development. The number of value scenarios for each programme was dependent on the duration of the different stages and the discrete time interval, ΔT, that these were considered over. We considered a very simple example in Figure 2.1, where the duration of the study was one unit and the market movements were applied every 0.5 units. This resulted in three value scenarios at the end of the study, with the associated values being dependent on the number of upward and downward market movements. These scenarios can be determined in advance of solving the programme.

In the PS approach, the scenarios represented the different potential combinations of trial outcomes and so the total number of scenarios was equal to the total number of combinations of study outcomes ranging from all programmes being unsuccessful in their first study to all programmes being successful in all of their studies. Equation (2.5) gave the calculation for the number of scenarios in the PS approach, which can be specified in advance of solving the resulting programme.

The difference in the stochastic component modelled leads to three main differences in the approaches. These differences are the type of **recourse action** considered, the **values** assigned to the drug development programmes and the **scheduling** decisions made.

The recourse action is the action that should be taken to compensate for the effect of an observed outcome and is one of the decisions that is returned upon solving the mixed integer linear programmes. In the ROV approach, the recourse

action relates to the value of the programme. If the value drops below zero then the recourse action would be to abandon development whereas, for all values above zero, development should continue and the next study should be invested in, unless a study results in a failure. If a study does end in failure, the ROV approach does not offer any alternative solutions in the sense that it does not recommend investing in alternative drug development programmes to replace the one that has just been abandoned. Therefore, if we wanted to find which other studies should be run given the study failure, the model would need to be updated given the new information and solved again. This can lead to suboptimal decisions compared to the stochastic version of the problem as the ROV approach cannot plan ahead and consider what might happen if a study is unsuccessful. The PS approach, however, is able to do this. A limitation of the ROV approach is that it assumes that the utility of the company is directly related to the ROV of the product, which may not necessarily be true. For example, if a company only had a single product undergoing development then, even if the ROV was negative due to low success probabilities, they may still wish to continue development should the clinical studies be successful. Hence, the recourse action in the ROV approach might not always be realistic.

In the PS approach, the recourse action is more complex as it can involve other drug development programmes. If a study is successful, the recourse action could be to continue development immediately or to delay development if, for example, the available resources are limited at that time point. If a study is unsuccessful, the recourse action will consist of selecting which study should be run next for another drug development programme. This allows the commencement of studies for other drug development programmes to compensate for a given study failure.

The values assigned to the drug development programmes are different in the two approaches. In the ROV approach, the values are calculated using real

option theory and they incorporate the market volatility of the drug, which was denoted by ξ_i for drug i. Values are assigned to each drug development programme at the beginning of each stage and in each value scenario and they were denoted by $M_{ijk_{ij}}$. M_{i11} corresponds to the ROV of drug i, when it is non-zero, and these values are found recursively by starting at the possible values at the conclusion of the drug development programme, should all associated studies be successful, and working backwards. This results in a clear ordering in terms of the attractiveness of the drugs with the most attractive having the highest value of M_{i11} and the least attractive having the lowest value of M_{i11}. The value of ξ_i is used in the ROV approach to capture the key beliefs about the stochasticity in the value of a drug development programme and is described by Rogers et al. [36] as the estimated annual standard deviation in the value of the drug after commercialisation. This parameter has a significant impact on the final value of the drug in the ROV approach; increasing the market volatility increases the value of the drug making it more likely to be included in the optimal portfolio. In the current model, it is assumed that the decision points are at study completion but we could also consider the decisions at an interim analysis in this framework. Considering the market information at additional points in the drug development process could add value and improve the decision-making process.

However, in the PS approach the market volatility is not considered at all. The only parameters that affect the value of the drug over time in the PS approach are the reductions in revenue for shorter active patent life and smaller market share and a time discounting factor. This can lead to discrepancies in the attractiveness of the drugs under the two methods. Therefore, a direct comparison for which approach selects the best portfolio is not simple as the portfolio that is deemed the best is dependent on how the values are assigned. In the PS approach, the values that are calculated for the revenue, future revenue and costs are dependent upon the scenario, which, in the PS approach, refers to the trial outcomes.

Furthermore, the ROV approach considers discounting of the value through time but the PS approach does not, it only considers the previously mentioned reductions for shorter active patent life and market share. This could be included in the objective function if required to ensure that promising programmes are not delayed unnecessarily.

Solving the mixed integer linear programme in the PS approach will provide the optimal order in which the studies should be run, but this order will not necessarily correspond to the value order of the individual drug development programmes. For example, the first study to be run will not necessarily belong to the most valuable individual programme. This is because the PS approach considers multiple drug development programmes that are competing for resources and the optimal ordering of the studies will depend on more than the individual programme values. Therefore, it may be preferable, for example, to run two programmes simultaneously that have a higher combined value than a third highly valuable programme with resource requirements equal to the total requirements of the other two studies. If we want to learn about the value of the individual programmes, we can consider each drug separately and assume that all studies begin as soon as possible in order to find the revenue and associated costs. Then the drug with the highest profit could be considered as the most valuable and the drug with the lowest profit could be considered as the least valuable, if we assume that utility is directly related to profit, providing an ordering in terms of the value of the individual drug development programmes. This value order may be different to the optimal running order due to the reasons discussed above.

The final main difference due to the uncertainty considered is scheduling. Since the ROV approach focuses on the value uncertainty rather than the trial uncertainty, scheduling is not considered. It is assumed that each study begins as soon as possible or not at all; the choice to postpone a study is not available in the

model formulation. One of the downfalls of this is that the ROV approach cannot take into account the potential benefit of waiting to observe certain outcomes in advance of making decisions. Instead, when a trial is concluded the model may be adjusted to reflect the current state of the portfolio in order to select which studies should be run. This is because the ROV approach focuses on the uncertainty in the market value rather than the trial outcomes.

Colvin and Maravelias [39] provide a simple example that illustrates the differences between a model that considers the different potential study outcomes and one that uses expectations as an attempt to capture the uncertainty in trial outcomes. That is, they compare a stochastic programme that models trial outcomes with its deterministic alternative. They show how this may lead to suboptimal decisions compared to when we allow studies to be postponed until certain outcomes are observed. The reason that the stochastic version finds a solution with a higher expected revenue in this example is that the deterministic model selects two drugs to run at the initial time point for which the second studies cannot be run at the same time due to limited resources. The stochastic model, however, considers this and selects a different pair of studies to be run at the same time that can be completed simultaneously. Since a penalty is incurred when studies are delayed this leads to the stochastic model achieving a higher expected revenue.

Furthermore, the resource constraints in the ROV approach are all calculated in the expected sense because the scenarios in the ROV approach do not contain information on trial success. This means that if the expected resource usage for a portfolio at a time point exceeds the resource limit then the development of some of the drugs included in this portfolio will need to be abandoned since they cannot be postponed. Using the expected resource utilisation in the constraints may even allow the ROV approach to select studies that have costs exceeding the available budget; we will see an example of this in the next section.

Conversely, modelling the trial outcomes and scheduling is one of the key features of the PS approach. The PS approach takes into account the fact that it may be beneficial to wait and observe certain outcomes before making some decisions. It also allows studies to be postponed if there are not enough resources/budget available to run the study immediately. This means that the PS approach is often able to find schedules that facilitate more development programmes being run than in the ROV solution. In addition, the PS approach is able to calculate resource requirements exactly, rather than the expected requirements, since the information about trial outcomes is included in the scenarios of the PS approach. This means that, under the decisions suggested by the PS approach, there will never be a case where the resources required exceed the resources available, which can happen when we only calculate the expected resources required.

It should also be noted that, in both of the approaches, the resources/budget are treated as fixed at each time point and any resources that are not used in one time point do not carry over to be used in the next time point. This, however, could be added as an extension to the methods if it was required by a company to make the budget allocation more realistic. Another potential modification could involve neglecting to include budget constraints entirely and instead including constraints on different resource types, e.g. staffing resources, and considering the trial costs within the objective value alone. Note that this is actually how the formulation of the PS approach is set up. The PS formulation includes resource constraints, which we have taken to be the budget in order to compare it to the ROV approach, and considers the trial costs within the objective function. The ROV approach assumes that the cost of a study is incurred at the initiation of the study, whereas the PS approach is able to incur these costs, or resource requirements, either at the beginning of the study or throughout the study. In the comparison that follows, we have modified the resource constraint in the PS approach so that the study costs are incurred at the beginning of the study in order to provide a fair comparison to the ROV approach.

In the PS approach, some drug development programmes might not be completed in the time frame that is considered in the model. This is because studies are allowed to begin at any time, subject to the previous study having been completed and resulting in a success. In order to compensate for the fact that some programmes may not complete in the planning horizon, the future revenue is considered, which is calculated by assuming that the ongoing studies are completed as quickly as possible. Considering the future revenue is beneficial as it encourages studies to be run where possible, even if the revenue will not be realised in the planning horizon, which is the type of forward planning that we would expect to see in real life decisions. However, it should be noted that the future revenue is not able to capture the information in the same way as extending the planning horizon would. In fact, if we change the length of the planning horizon in the PS approach, even by a single time step, the set of optimal schedules may change. Hence, when using this method a sensitivity analysis might be required to study the effect of different planning horizon lengths on the set of optimal schedules for a particular portfolio.

In terms of the model formulation, there are three main things to discuss: flexibility, complexity and size.

Flexibility will often be desirable so that the model can be adjusted to accurately represent the portfolio in question. A company may also wish to add constraints that reflect their decision-making process e.g. there might be two drugs that they would only want to develop at most one of. Due to the fact that the ROV approach is based on the way that the values of the drugs are calculated, there is little flexibility in the choice of the objective function. In the PS approach, however, there is a lot more flexibility in terms of the objective function. If we chose, we could modify the PS approach to include an objective function that maximises the ROV of the portfolio. Adding further constraints is relatively straightforward in both approaches, provided that some consideration is given

to what the scenarios refer to in each of the approaches and what this means in terms of the constraints.

In terms of the complexity of the models, the difficulty arises in different areas. For the ROV approach, the complexity arises in the calculation of the value of drug i at the beginning of study j in value scenario k_{ij}, which was denoted by $M_{ijk_{ij}}$. In this approach, a separate linear programme must be solved in order to find upper bounds on the value of $M_{ijk_{ij}}$, which are then used as some of the input parameters of the mixed integer linear programme that returns the optimal portfolio. Essentially, the upper bounds that we find by solving the first programme correspond to the values of the drugs if we did not include any resource constraints. As was mentioned in Section 2.2.1, including the decision to continue/abandon development in the calculation of $M_{ijk_{ij}}$ results in a non-linear constraint. The linearisation of this constraint also adds to the complexity of the model. In the PS approach, the complexity arises from the fact that we require non-anticipativity constraints for this formulation. While the interpretation of these constraints is relatively straightforward, the formulation of them is less so.

This leads us to the final consideration to make, which is the size of the model formulation. The non-anticipativity constraints in the PS approach vastly increase the number of constraints in the formulation of the PS approach compared to the ROV approach. Also, the number of variables required for the PS approach typically exceeds the number in the ROV approach. This is due to the number of scenarios typically being larger for the PS approach. This will be illustrated in the next section.

(a)

i	$\lvert J_i \rvert$	ϕ_{i1}	ϕ_{i2}	ϕ_{i3}	τ_{i1}	τ_{i2}	τ_{i3}	c_{i1}	c_{i2}	c_{i3}	V_{0i}
1	3	0.75	0.7	0.85	1	2	2	20	55	80	180
2	3	0.6	0.8	0.95	1	2	2	30	55	120	380
3	2	0.8	0.9	-	2	2	-	30	60	-	100
4	2	0.8	0.9	-	2	2	-	75	180	-	400
5	1	0.75	-	-	2	-	-	180	-	-	350

(b)

ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ΔT	r
0.55	0.35	0.8	0.3	0.6	1/6	0.05

(c)

γ^L	γ^D	n_t
10	20	0.1

Table 2.4: Parameter values used for the comparison in (a) both approaches, (b) the ROV approach and (c) the PS approach. A planning horizon of $|T| = 6$ was used in both approaches.

2.3.2 Results

We implemented both of the approaches using JuMP [43] in Julia [44] for an illustrative example portfolio with parameters given in Table 2.4. This example is used to highlight our main findings regarding the use of the two methods and the differences between them. The optimal portfolios were considered for three different budgets: 200, 250 and 300 per time point. That is, $B_t = 200$, $B_t = 250$ and $B_t = 300$ for all time points, t, in the planning horizon. A summary of the comparison is provided in Table 2.5. Gantt charts showing some of the schedules selected under the PS approach with a budget of 200 per time point, $B_t = 200$, are shown in Figure 2.4. The aspects for which comparative results are provided include: speed to obtain the solution, size of the problem in terms of the variables and constraints, selected portfolio and ordering of the most attractive drug development programmes.

Note that we are not assuming an underlying truth for the illustrative example. In this section, our aim is to illustrate the way the methods work and to highlight any differences between the methods and the results that they may lead to through the use of our illustrative example. Therefore, we will not draw conclusions on which method has performed better. It should be further noted

		ROV approach	PS approach
	Variables	208 (a) + 2438 (b)	58176
	Constraints	208 (a) + 6760 (b)	288624
	Value order	4,5,2,3,1	2,4,5,1,3
$B_t = 200\ \forall t$	ROV	158.61	-
	ENPV	-	125.15
	Selection	1,2,3,4	1,2,3,4,5
	Time (CPUs)	58.85	4128.45
$B_t = 250\ \forall t$	ROV	158.61	-
	ENPV	-	129.65
	Selection	1,2,3,4	1,2,3,4,5
	Time (CPUs)	18.62	3104.95
$B_t = 300\ \forall t$	ROV	229.36	-
	ENPV	-	129.65
	Selection	2,4,5	1,2,3,4,5
	Time (CPUs)	24.23	5247.81

Table 2.5: Comparison of the ROV approach with the PS approach. Note that for the ROV approach (a) refers to the LP used to find the upper bounds of $M_{ijk_{ij}}$ and (b) refers to the MILP used to find the optimal portfolio.

that, even if an underlying truth was assumed, we would still not necessarily be able to conclude which method performs better. This is because the methods assign values to drug development programmes differently and so the most valuable programme under the ROV approach may be different to the most valuable programme under the PS approach.

One of the first things that is apparent in Table 2.5 is the difference in the size of the formulations. Despite the fact that the ROV approach requires two mathematical programmes to be solved (the first to find upper bounds on the values in the second and the second to find the optimal portfolio), it still has far less variables and constraints than the PS approach. This is because the scheduling in the PS approach comes at a high computational burden due to the inclusion of the non-anticipativity constraints, which ensure that the optimisation does not use information regarding trial outcomes before they have been revealed. The pattern observed here will be the same for most sets of input parameters. The main thing that could increase the size of the ROV formulation past the size of the PS formulation would be if we applied the market movements over

much smaller periods as this would increase the number of value scenarios to consider. However, it is unlikely that a user would require a level of granularity that would be small enough to cause the ROV formulation size to exceed the PS formulation size.

We see a similar ordering in the time taken as we saw in the numbers of variables and constraints and this is typical of what we would observe for most example portfolios. For this example containing five drugs, all of the mixed integer linear programmes were solved within reasonable time. This will not necessarily scale as we increase the portfolio size, unfortunately, as it is noted in [58] that without model reductions the PS approach cannot solve the problem for portfolios containing more than six drugs in reasonable time. There have been heuristics presented for the PS approach that aim to tackle this problem [49]. The most promising heuristic was presented by Christian and Cremaschi [49] and it decomposes the problem into a series of smaller knapsack problems. While the solution of the knapsack decomposition algorithm will not always match the optimal solution of the PS approach, the reported results were within 3% of the optimal solution and the solution is found much more quickly than in the full formulation of the PS approach and can be found for much larger portfolios.

In Table 2.5, we see that, although the two methods assign values to drugs differently, the ordering of the drugs in terms of the value of the associated drug development programmes is similar in both. The biggest difference is that Drug 2 is deemed the most attractive in the PS approach but only the third most attractive in the ROV approach. This highlights the fact that the different methods of assigning values to the drug development programmes can lead to different decisions being made. Although Drug 2 has the second highest present value, it has the second lowest volatility causing it to be attractive under the PS approach but not as attractive in the ROV approach. This is because increasing the volatility of a drug also increases the ROV of the drug.

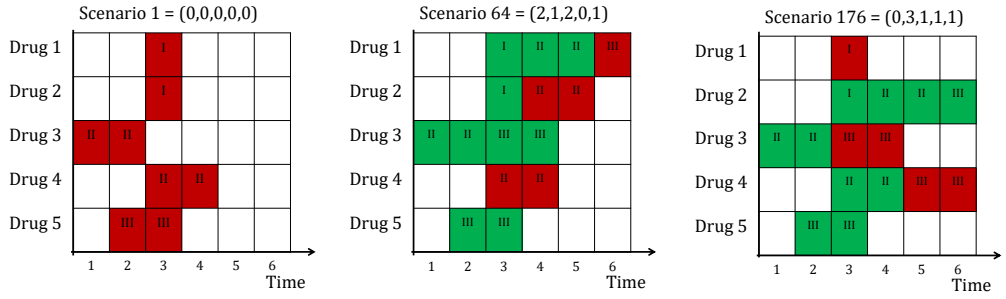

Figure 2.4: Gantt charts representing the schedules found by the PS approach in the comparison with $B_t = 200$.

It should be noted that even if the portfolios selected in each approach were similar, or even the same, the objective value under each method would still be quite different due to the different methods of valuation. Also, we are not able to calculate the objective function of one model in a meaningful way for the decisions made under the other model. The objective function in the ROV approach is calculated under the assumption that trials start as soon as possible, which will not necessarily be the case with the decisions made in the PS approach, as we can see in Figure 2.4. The ROV objective function does not consider any reductions in revenue given the late completion of a drug development programme, and the market movements are considered over the minimum time taken to complete the programme. Therefore, the ROV objective function cannot be calculated appropriately for the decisions made under the PS approach. Similarly, looking at the decisions made in the ROV approach in the PS framework will not lead to a meaningful objective value. This is because scheduling was not an option in the ROV framework and there will be "better" scheduling decisions available. Hence, the non-scheduled decisions will be suboptimal in terms of the PS objective function.

Rather than comparing the objective values across the approaches, it is more interesting to look at how the objective values change in each approach given additional budget/resources. For example, when we increase the budget from 200 to 250, the objective value of the ROV approach remains the same, but the

objective value of the PS approach increases due to the fact that it is able to find a more profitable schedule. This leads us on to the next part of the comparison.

One of the most significant differences in terms of the portfolio that we see in Table 2.5 is that the PS approach is able to select more drugs than the ROV approach. This is because it is able to schedule when the studies should be run and hence allows delays in the commencement of a study, which often facilitates the selection of more drugs as it has done here for our example portfolio. In Table 2.5, we see that when we have budgets of 200 and 250 for each time period, the ROV approach is unable to select Drug 5 despite the fact that it has the second highest ROV. The PS approach, however, is able to include this drug in its portfolio under all three budgets. The PS approach is in fact able to include all five drugs in its optimal portfolio under all three budgets. We see that for a budget of 300 both methods are able to select their most attractive drugs to include in their optimal portfolio.

If we consider a much smaller budget than those considered in the table, $B_t = 100 \ \forall t$, the ROV approach selects Drugs 1 and 4. However, Drug 4 has a cost of 180 for its second study. This highlights the way that taking the expected value for resource constraints may not always be realistic. Here, the cost of a study is almost double the available budget yet it is still selected. While this might not always be an issue, it could certainly be a problem in the real world if a company were to run into the situation where the extra budget could not be found and hence the investment in previous studies would be wasted.

In Figure 2.4, we see that, under the assumption that the costs are incurred at the commencement of a trial, (D5, PIII) is only selected to be run in the second time period due to budget constraints not allowing it to start alongside any of the other studies. This also applies to (D4, PIII). This leads to a sparse amount of studies being run in the first two time steps and a much denser schedule later on. This may not be realistic in terms of what would be preferred by a company.

If this situation were to arise, a company may add constraints to the model that reflect their preferences and explore the effect of these preferences on the overall scheduling and optimal value. For example, these constraints might be to run certain studies straight away or to ensure that the pipeline is not left idle when there are available studies to be run.

We also see that the two most valuable drug development programmes under the PS approach, Drugs 2 and 4, only have their first study initiated at the third time point. This shows how, as we discussed in the previous section, the optimal order to initiate the studies may not reflect the most valuable development programmes due to the complexities added by the consideration of multiple programmes, resource constraints and the effect of early decisions on later decisions. Since (D4, PIII) can only be initiated at the same time as (D1, PI) due to budget constraints, its preceding study, (D4, PII) is only initiated later in the planning horizon to allow for the initiation of other studies. If we assume that the optimal time to initiate (D4, PIII) is at $t = 5$ then there is no reason to start (D4, PII) until $t = 3$ as the PS approach discounts study costs throughout the planning horizon to encourage studies to be selected. This means that leaving (D4, PII) to start as late as possible without affecting the initiation of (D4, PIII) will increase the ENPV.

Another component that will affect the optimal decisions are the penalties incurred due to shorter active patent life and smaller market share. It is assumed that the patent life of a drug starts at the initiation of Phase I, therefore the penalties associated to shorter active patent life are only incurred after Phase I is initiated. It is therefore beneficial to reduce the delays between phases. We see the effect of this in Figure 2.4 for Drugs 1 and 2 as both programmes are completed without any delays, despite the fact that the earlier phases could be initiated sooner but with delays between the later phases as this would lead to a reduction in the ENPV.

2.4 Discussion

In this chapter, we have provided a comparison and discussion of two stochastic programming approaches for pharmaceutical portfolio management. We used an illustrative example to highlight our findings and the differences between the two approaches.

The first approach [36] uses real option theory from the financial setting and focuses on the uncertainty in the value of a drug. While this approach has a reasonably sized formulation and is quick to solve, it lacks flexibility and there is some discussion to be had in terms of the relevance of the market volatility, which drives this approach in the pharmaceutical setting. Also, this approach assumed that every study would start as soon as possible and it is not able to schedule tasks, which in turn leads to drugs being omitted from the optimal portfolio that could in fact be included if scheduling were considered.

The second approach [39][40] provides a modification of the stochastic version of the resource constrained project scheduling problem in order to find the optimal portfolio and the optimal schedule under each scenario. While this approach was preferable in terms of scheduling and flexibility, this came at a high computational burden.

Extensions have been presented for both of these methods. For the ROV approach, Rogers et al. [38] presented an extension that considers partnership opportunities. This extension considers both the optimal timing of the partnership and the best investment policy. Lo Nigro et al. [56] noted the perceived complexity of the method presented by Rogers et al. [36] and presented a more user-friendly simplification that uses a different option pricing approach and considers partnership opportunities. However, this approach does not track the continue/abandon decisions over time.

Colvin and Maravelias [57] presented methods to tackle the problem of the size of the model formulation for the original PS approach presented in [39]. This was then extended further to consider solution methods and a branch and cut algorithm was presented in [58]. Then, Colvin and Maravelias [50] presented several extensions that include: resource planning decisions such as expansion and outsourcing; task interdependencies in terms of uncertainty, resources and revenues; risk management approaches.

Although the two approaches use different methods for valuation, neither of them capture the effect that treatment efficacy and safety can have on the revenue of a new drug. Treatment safety and efficacy estimates will often be driving factors in our estimate of the revenue and these safety and efficacy estimates will evolve throughout the drug development programme. The consideration of these estimates in the valuation method could make it more appropriate and realistic in the setting of drug development. Also, both methods are reliant upon point estimates for the revenue generated upon successful programme completion and they are not able to consider our prior uncertainty in this estimate. We think that the ability to do so would add further benefit to the methods as the prior estimates of the revenue are not always representative of what is actually seen.

Both of these methods have different benefits and focuses and therefore the best approach will depend on what the decision maker deems more important - modelling trial uncertainty or modelling value uncertainty. We believe that the flexibility offered in the project scheduling approach offers an advantage over the real option approach as, if one chose to, the valuation of the drugs in this approach could be modified to consider the market volatility. We also believe that, in the setting of clinical trials, the impact of the trial outcomes is more significant than the market volatility. The pharmaceutical industry has high development failure rates compared to many other industries. Therefore, the stochastic modelling of the trial outcome may be more important than market volatility in this

setting. Furthermore, the ROV approach is centred on the estimates of the net present value of the future cash flows of the drug and the estimates of the market volatilities. If these estimates are not accurate or representative of the drug then the focus of this approach is wasted.

In conclusion, we believe that the advantages of the PS approach for tackling the true portfolio management problem are many and, with the flexibility in this approach, further adjustments to make the approach match the individual requirements of a company can be added easily making this approach very useful and applicable for real world problems.

Furthermore, our main interest in this thesis lies in combination drug development and we believe that the PS approach is best suited to this setting. This is because the recourse action in the PS approach is able to consider different development programmes. This would be beneficial in the context of combination therapies as the results of one combination study might affect our beliefs about similar combination studies, thus the decisions that we would make. The way that these beliefs regarding performance might change can be captured via changes in the probability of success of a study. Therefore, in the next chapter we present a method for updating the probability of success of a study based on related combination study data. This will then be incorporated into a heuristic [49] for the PS approach in Chapter 4.

CHAPTER 3

Chapter 3

Updating the probability of study success for combination therapies using related combination study data

3.1 Introduction

The recent rise in popularity of combination therapies has brought with it several new questions and challenges. The question that we will aim to answer in this chapter is associated with the potential relationships between combinations. We may expect a relationship between two combination studies when, for example, they have a particular treatment in common. We look at how we can use the information from **related combination studies** to inform the probability of success of a particular combination study of interest. It was noted in Chapter 1 that combination therapies often consist of a **backbone therapy**, such as a chemotherapy in oncology, and one or more different **add-on treatments**. Hence, in this scenario, there are clear groups of associated combinations, which correspond to the combination therapies that share a backbone treatment.

Therefore, there is much to be gained by considering related combinations. This gain is even more significant when there is little available information on the combination of interest, but a much larger amount of available information on a related combination, such as the outcome of a Phase III study. This is because of the potential for strong **correlations** between the outcomes of related combination studies. Using the additional information from related combinations appropriately may improve the accuracy of the treatment effect estimates, which in turn may lead to improved decision-making in the planning of combination studies through the calculation of the study success probabilities. Improved decision-making may help to reduce the failure rates in the later clinical trial phases or optimise the portfolio.

One of the key estimates used in order to assist decision-making regarding a potential study is the **probability of success** (PoS), as discussed in Chapter 1. Existing methods for calculating the PoS are often based upon the expected power, (Bayesian) predictive power or **assurance**. These terms are often used interchangeably in the literature. O'Hagan and Stevens [59] presented the concept of the assurance and detailed how it can be used and interpreted. O'Hagan et al. [23] then provided further discussion of the assurance and how this can be used instead of the power in calculating the required sample size of a study. Rufibach et al. [35] provided discussion of the choice of prior when calculating the assurance and also provided some recommendations.

The literature on the PoS also covers how this can be used to assist decision-making. Stallard et al. [30] present an approach that combines Bayesian and frequentist ideas. The decision-making process uses Bayesian methodology in the calculation of the PoS but it is assumed that the study design and analysis in Phase III will be frequentist. This approach can be used both at the end of Phase II and at any interim analyses. Thus, it may help with the decision to begin planning Phase III rather than just whether to run a Phase III study. Sabin et al.

[33] further discussed the use of the PoS in decision-making and presented a two-stage method that starts before Phase II and takes the user through to the end of Phase II decision.

Another area in the literature relevant to the problem that we are interested in relates to the planning of sequences of trials. Whitehead [60] discusses the problem of designing a series of Phase II studies when the aim is to identify the treatment that should be taken to Phase III. Rather than considering the treatments separately, they are considered together, which allows sample sizes to be reduced. The methodology presented also provides the optimal number of treatments to be tested in Phase II. Existing literature regarding the planning of sequences of trials also includes platform trials [61] and multi-arm multi-stage trials [20]. In our setting, however, the related combination studies might not share the same target population or the same indication, as is typical in the literature for planning sequences of studies, and the studies might not be available to begin simultaneously.

In this chapter, we present a framework that allows us to update the probability distributions of the effect sizes of a group of related combination therapies based on the outcome of a single combination study. This will allow us to update the PoS of related combination studies. This procedure allows emerging information on related combination therapies to feed into the decision-making process for other potential combinations and assist in the planning of these studies. In line with existing literature, we assume that the design and analysis of the studies are conducted using frequentist methods while the calculation of the PoS will use a Bayesian framework.

In order to provide further motivation and context to this problem, we will consider two historic Phase III trials that are related and use them to illustrate the methodology throughout the chapter. We will consider the CLEOPATRA (NCT00567190) [62] and MARIANNE (NCT01120184) [63] trials, which both

considered similar combination therapies in the treatment of patients with HER2-positive breast cancer and used **progression-free survival** (PFS) as the primary endpoint. The purpose of our method is to capture the relationship between study outcomes, rather than considering the differences between the studies themselves, and to use this to help inform decision-making. Therefore, while these two studies do have several differences, they serve as an example of the type of situation in which the proposed method can be applied.

The CLEOPATRA study is a double-blind study that compared trastuzumab plus docetaxel, a type of taxane, plus placebo with trastuzumab plus docetaxel plus pertuzumab with a 1:1 allocation ratio [62]. The MARIANNE study is a multi-arm study that compared trastuzumab plus taxane to trastuzumab emtansine plus placebo and trastuzumab plus pertuzumab [63]. The control arm in the MARIANNE study was open label, whereas the two experimental arms were blinded with respect to pertuzumab or placebo and the allocation ratio was 1:1:1.

Both studies have a control arm that contains trastuzumab and a taxane (the CLEOPATRA study also contains placebo) and an experimental arm that includes pertuzumab. We will therefore aim to use the CLEOPATRA study to draw inference upon the outcome of a modified two-arm version of the MARIANNE study that compares trastuzumab plus taxane to trastuzumab emtansine plus pertuzumab; we will not consider the trastuzumab emtansine plus placebo arm. From now on, we will refer to this study as the modified MARIANNE (mod-MARIANNE) study.

It is clear that there are several differences between the two studies, such as the blinding and the number of arms. Irrespective of this, we believe that the outcome of one study is informative for the other and hence these studies will be used for illustration of the methodology.

In Section 3.2, we present the framework and methodology for updating the PoS of related combination studies. We also provide an extension that allows us to account for the fact that the treatment effects, hence study success probabilities, of the "related" combinations might not be correlated. In Section 3.3, we present the results of a simulation study. We provide a discussion of the approach in Section 3.4.

3.2 Methods

In this section, we build the framework that allows us to update the PoS of a combination study based on the outcome of a related combination study in order to assist decision-making.

First, we update the distributions of related combination therapies based on the outcome of a single combination study using **Gaussian Markov Random Field** theory. Then, using the updated marginal distributions, we can find the PoS for all remaining studies. We also provide an extension that allows us to consider the fact that the assumption of "related" combinations being positively correlated might not always hold and account for this in our PoS calculations.

3.2.1 Framework

For illustrative purposes, let us first consider a pair of related combinations, for example $A + B$ and $A + C$, which we might be comparing to a similar control treatment, before extending the problem to a set of n related combinations. We will refer to combinations as "related" when they have at least one monotherapy in common and there is reason to believe that the performance of the combinations will be related. An example of this might be in oncology where A is a

backbone treatment, such as a chemotherapy, and B and C are potential add-on treatments with different modes of action.

We are interested in calculating the PoS for one combination study based on the study results of a related combination. In order to calculate the PoS we will follow the method presented by O'Hagan et al. [23] to calculate the assurance, which is defined by

$$\text{PoS} = \int P\left(\text{study success} \mid \theta\right) P\left(\theta \mid \text{data}\right) \mathrm{d}\theta \tag{3.1}$$

where θ represents the treatment effect. We can often find a closed form solution for the assurance. For example, in the case of a two-sided superiority trial with normally distributed outcomes and known variance, the assurance for rejecting the null hypothesis of no treatment difference in favour of the experimental treatment is given in [23] by

$$\text{PoS} = 1 - \Phi\left(\frac{V^{-0.5} Z_{\alpha/2} - \mu}{\sqrt{V^{-1} + \sigma^2}}\right). \tag{3.2}$$

Here, V^{-1} is the sampling variance of the planned study, α is the significance level and μ and σ^2 are the mean and variance of the distribution representing our beliefs on θ, respectively. When a closed form solution for the assurance is not available, we can use Bayesian clinical trial simulation to estimate it, as presented in [23] and discussed in Chapter 1. Alternative distribution-based definitions of the PoS could also be used in our presented framework, such as the Bayesian probability of success presented in [26].

In order to calculate the PoS and update it using related combination study data, we need to consider the treatment effects of the combinations of interest. For our simple example of combinations $A + B$ and $A + C$, we will use θ_1 and θ_2 to represent the treatment effects of $A + B$ and $A + C$, respectively. We will assume that θ_1 and θ_2 are measured on the same scale and are therefore directly

comparable.

Before a clinical trial is run, the study team will have some idea as to how the therapy might perform based on historical data and expert opinion. In order to capture these beliefs we can specify a prior distribution on the parameter of interest. This prior distribution is able to capture the expected value of the treatment effect and the level of uncertainty in this value. There is an extensive literature on prior elicitation in the setting of a clinical trial, with one of the most commonly discussed methods being the SHELF framework [34].

We will represent the prior beliefs for the treatment effects of our two combination therapies, $\boldsymbol{\theta} = (\theta_1, \theta_2)^T$, by the multivariate normal (MVN) distribution. We can write this as $\boldsymbol{\theta} \sim \text{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ or, alternatively for the two combination example,

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 \\ \rho_{12}\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right)$$

where μ_i and σ_i^2 represent the prior expectation and prior variance for θ_i, respectively, and $\rho_{ii'}$ is the prior correlation between θ_i and $\theta_{i'}$. In this model, the parameter $\rho_{ii'}$ will be used to define the level of borrowing across the two combinations. Therefore, the user does not need to be concerned with calculating an accurate estimate of the true underlying correlation for this model to be appropriate or useful. The reasons why we can use the interpretation of $\rho_{ii'}$ as the degree of borrowing will be discussed in Section 3.2.4.

When determining an appropriate value for ρ_{12}, one could consider a thought experiment using studies relating to θ_1 and θ_2. For example, if these relate to the combinations of $A + B$ and $A + C$ then one might consider either the outcome of these combinations in different indications or alternatively the outcomes of B and C when paired with different backbone treatments. If these outcomes are typically positive or negative simultaneously then a higher value of ρ_{12} may be

appropriate. If there was little or no pattern between the pairs then a lower value of ρ_{12} would be more appropriate.

Further note that this model does not aim to capture synergism or antagonism within the components of the combinations, instead it aims to capture similarities across the combinations, which will allow us to learn across the combinations.

In the case where we have n related combinations, we would specify the prior beliefs using

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_n \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 & \cdots & \rho_{1n}\sigma_1\sigma_n \\ \rho_{12}\sigma_1\sigma_2 & \sigma_2^2 & \cdots & \rho_{2n}\sigma_2\sigma_n \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{1n}\sigma_1\sigma_n & \rho_{2n}\sigma_2\sigma_n & \cdots & \sigma_n^2 \end{pmatrix}\right).$$

This is the distribution that we will update based on the outcome of a combination study relating to one of the θ_i variables. We will then use the updated distribution to calculate the PoS for future combination studies.

In order to specify the prior distribution for our illustrative example, we will let θ_M be the treatment effect for the mod-MARIANNE study comparing trastuzumab plus taxane with trastuzumab emtansine plus pertuzumab. We will further let θ_C be the treatment difference for the CLEOPATRA study comparing trastuzumab plus docetaxel plus placebo with trastuzumab plus docetaxel plus pertuzumab.

We will specify the prior means of both θ_M and θ_C to be equivalent to a hazard ratio of 0.75, $\mu_M = \mu_C = -\log(0.75)$, which is equal to the reference value that was used to power both of the studies. We will specify a prior correlation of $\rho_{MC} = 0.6$ to reflect the belief that the study outcomes are related along with our interest in using the outcome of one of the studies to inform our beliefs about the other. If the two studies only differed in one aspect but were otherwise identical, we could use a higher correlation. However, since the studies differ in several ways,

we have decided to use a lower correlation to reflect the uncertainty caused by the differences. Finally, we will specify a prior variance of 0.08 on both treatment effects. This is equivalent to the posterior variance after observing approximately 50 PFS events, a common size of a Phase II study, given an uninformative prior variance. This will give a bivariate prior of

$$\begin{pmatrix} \theta_M \\ \theta_C \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} 0.288 \\ 0.288 \end{pmatrix}, \begin{pmatrix} 0.08 & 0.05 \\ 0.05 & 0.08 \end{pmatrix}\right).$$

3.2.2 Score statistics

As we observe further clinical studies on the combinations, we will want to update this distribution to reflect the information gained from these new studies. We assume that these studies will be designed and analysed using frequentist methodology. Therefore, to summarise the outcome of study i, we will use the efficient **score statistic**, Z_i, and the **Fisher information**, V_i, of the test with null hypothesis $\theta_i = 0$ [21]. This is not the same as the Z-score, which follows the standard normal distribution. Note that, in this chapter, we are only considering one study per programme and so when we refer to study i, we are referring to the study of interest in programme i. Thus, we will not use the study identifier, j, in the notation in this chapter, but the programme identifier, i, alone.

The score statistic can be considered as a measure of benefit of the experimental treatment over the control treatment, based on what was observed in the study. The Fisher information, on the other hand, is a measure of how much information on θ_i is contained in Z_i. When the study sample size is large and θ_i is small, that is, close to zero, the score statistic is approximately normally distributed with mean given by $V_i\theta_i$ and variance given by V_i where θ_i is the true value of the treatment effect [21]. Note that this normal approximation holds for many endpoints, which

is one of the main reasons that we consider the score statistic in our framework.

$$Z_i \dot{\sim} N(V_i\theta_i, V_i)$$

If we only consider the marginal prior distribution of θ_i, then, since the normal distribution is a conjugate prior for normally distributed data, we could find the posterior distribution of $\theta_i|Z_i$ via a typical Bayesian update, and this would also be normally distributed.

In our setting, however, we consider these parameters in a vector represented by $\boldsymbol{\theta}$ and we do not observe all dimensions of $\boldsymbol{\theta}$ simultaneously, but observe the outcome of one combination study at a time. The distribution of the score statistic, Z_i, will remain one-dimensional, but we will still want to update the distribution of $\boldsymbol{\theta}$ each time we observe new data. To allow updating regardless of the inconsistency in dimensions, we will utilize the properties of Gaussian Markov Random Fields.

In our illustrative example, the CLEOPATRA study was the first of the two studies to be conducted therefore we will use the information from the CLEOPATRA study to update our beliefs about the PoS of the mod-MARIANNE study. The CLEOPATRA study observed 604 PFS events, 320 in the control arm and 284 in the experimental arm, and the observed hazard ratio (HR) was 0.68 [64]. We are able to find Z_C and V_C for the study using

$$V \approx e \times R/(R+1)^2 \quad \text{and} \quad Z = -V\log(\text{HR})$$

where e is the number of PFS events and R is the allocation ratio [21]. This gives $V_C = 151$ and $Z_C = 58.235$.

3.2.3 Gaussian Markov Random Fields

In order to update the distribution of $\boldsymbol{\theta}$ based on a single combination study, we will formulate the problem using Gaussian Markov Random Fields [65].

A Gaussian Markov Random Field (GMRF) is finite dimensional random vector $\mathbf{x} = (x_1, ..., x_{n+1})^T$ that follows the multivariate normal distribution and has some additional Markovian properties to be defined. First, let us write

$$\mathbf{x} \sim \text{MVN}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$$

where $\boldsymbol{\mu}$ is the mean vector and $\mathbf{Q}$ is the precision matrix.

Recall that two variables, x_a and x_b, are conditionally independent given another variable, x_c, if and only if

$$\pi(x_a, x_b|x_c) = \pi(x_a|x_c)\pi(x_b|x_c)$$

where $\pi(\cdot)$ represents a probability density function. We can represent the **conditional independence** structure of $\mathbf{x}$ using a graph $G = (N, E)$. In this graph, the nodes, N, represent the random variables, $x_1, ..., x_n$, and the edges, E, provide us with information regarding the conditional independence structure. We can say that $\mathbf{x}$ is a GMRF with respect to G if and only if

$$\pi(\mathbf{x}) = (2\pi)^{-(n+1)/2}|\mathbf{Q}|^{1/2} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T\mathbf{Q}(\mathbf{x} - \boldsymbol{\mu})\right)$$

and

$$Q_{ab} \neq 0 \quad \Leftrightarrow \quad (a, b) \in E \quad \forall\ a \neq b.$$

The properties of GMRFs that are of particular use for our problem are the conditional properties. First, let us partition the random vector into two sets, $\mathbf{x} = (\mathbf{x}_A, \mathbf{x}_B)^T$. We can similarly partition the mean vector, $\boldsymbol{\mu} = (\boldsymbol{\mu}_A, \boldsymbol{\mu}_B)^T$, and the

precision matrix,

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_{AA} & \mathbf{Q}_{AB} \\ \mathbf{Q}_{BA} & \mathbf{Q}_{BB} \end{pmatrix}.$$

Then, it can be shown that the conditional distribution of $\mathbf{x}_A \mid \mathbf{x}_B$ is also a GMRF with respect to the subgraph $G^A = (N^A, E^A)$ [65]. N^A is a subset of N containing the nodes that represent the variables in $\mathbf{x}_A$. E^A is a subset of E containing the edges between the nodes in N^A. The mean and precision matrix of this conditional distribution are given, respectively, by

$$\boldsymbol{\mu}_{A|B} = \boldsymbol{\mu}_A - \mathbf{Q}_{AA}^{-1}\mathbf{Q}_{AB}\left(\mathbf{x}_B - \boldsymbol{\mu}_B\right) \text{ and } \mathbf{Q}_{A|B} = \mathbf{Q}_{AA}. \tag{3.3}$$

Hence, we can write

$$\mathbf{x}_A \mid \mathbf{x}_B \sim \text{MVN}\left(\boldsymbol{\mu}_{A|B}, \mathbf{Q}_{A|B}^{-1}\right), \tag{3.4}$$

the conditional distribution of $\mathbf{x}_A$ given $\mathbf{x}_B$ [65].

3.2.4 Method

Let us now formulate our problem in the setting of GMRFs. We will illustrate the method using our two drug example, but will also provide the results for n combinations. In our setting, the random variables are the treatment effects of the different combinations alongside the score statistic for the study that is going to be run. We can write this as $(\boldsymbol{\theta}, Z_i)^T$ where we will be observing the outcome of a study on combination i.

We will also need to consider the conditional independence structure of this vector of random variables. For our simple example of combinations $A + B$ and $A + C$, θ_1 and Z_2 will be conditionally independent given θ_2. This represents the

beliefs that θ_1 and θ_2 are "related" and the score statistic, Z_2, from a study on combination 2 is dependent only on the value of θ_2.

Suppose that we are going to observe the outcome of a study on $A + C$, which we will summarise using Z_2 and V_2. Under the assumption that Z_2 is approximately normally distributed, $Z_2 \sim N(V_2\theta_2, V_2)$, we can write

$$Z_2 = V_2\theta_2 + \epsilon_2 \quad \text{where} \quad \epsilon_2 \sim N(0, V_2).$$

Using this formulation, we are able to find the expectation and variance of Z_2 in this framework and the covariance between Z_2 and $\boldsymbol{\theta}$. We can then write our GMRF as

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ Z_2 \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ V_2\mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 & V_2\rho_{12}\sigma_1\sigma_2 \\ \rho_{12}\sigma_1\sigma_2 & \sigma_2^2 & V_2\sigma_2^2 \\ V_2\rho_{12}\sigma_1\sigma_2 & V_2\sigma_2^2 & V_2^2\sigma_2^2 + V_2 \end{pmatrix}\right).$$

The precision matrix, $\mathbf{Q}$, of this GMRF is given by

$$\begin{aligned} \mathbf{Q} &= \begin{pmatrix} \mathbf{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}} & \mathbf{Q}_{\boldsymbol{\theta}Z_2} \\ \mathbf{Q}_{Z_2\boldsymbol{\theta}} & \mathbf{Q}_{Z_2Z_2} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{\sigma_1^2(1-\rho_{12}^2)} & \frac{-\rho_{12}}{\sigma_1\sigma_2(1-\rho_{12}^2)} & 0 \\ \frac{-\rho_{12}}{\sigma_1\sigma_2(1-\rho_{12}^2)} & \frac{1}{\sigma_2^2(1-\rho_{12}^2)} + V_2 & -1 \\ 0 & -1 & \frac{1}{V_2} \end{pmatrix}. \end{aligned}$$

Note that if we have n combination programmes and we observe a study in programme i then we can find the mean vector and covariance matrix in the

same way and the GMRF for this problem would be given by

$$\begin{pmatrix} \theta_1 \\ \vdots \\ \theta_n \\ Z_i \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \\ V_i\mu_i \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \cdots & \rho_{1n}\sigma_1\sigma_n & V_i\rho_{1i}\sigma_1\sigma_i \\ \vdots & \ddots & \vdots & \vdots \\ \rho_{1n}\sigma_1\sigma_n & \cdots & \sigma_n^2 & V_i\rho_{in}\sigma_i\sigma_n \\ V_i\rho_{1i}\sigma_1\sigma_i & \cdots & V_i\rho_{in}\sigma_i\sigma_n & V_i^2\sigma_i^2 + V_i \end{pmatrix}\right).$$

Using the conditional properties of GMRFs, we are able to find the conditional distribution of $\boldsymbol{\theta}|Z_i$. Applying Equations (3.3) and (3.4) as given in [65], we find that

$$\boldsymbol{\theta}|Z_i = z_i \sim \text{MVN}\left(\boldsymbol{\mu}_{\text{post}}, \boldsymbol{\Sigma}_{\text{post}}\right)$$

where

$$\boldsymbol{\mu}_{\text{post}} = \boldsymbol{\mu} - \mathbf{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}}^{-1}\mathbf{Q}_{\boldsymbol{\theta}Z_i}(z_i - V_i\mu_i) \text{ and } \boldsymbol{\Sigma}_{\text{post}} = \mathbf{Q}_{\boldsymbol{\theta}\boldsymbol{\theta}}^{-1}.$$

For our simple example containing combinations $A + B$ and $A + C$, this gives

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \mid Z_2 \sim \text{MVN}\left(\begin{pmatrix} \mu_1 - \frac{\rho_{12}\sigma_1\sigma_2V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}z_2 \\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}z_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2} & \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} \\ \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} & \frac{\sigma_2^2}{1+V_2\sigma_2^2} \end{pmatrix}\right).$$

Here we see that the parameter ρ_{12} defines how far and in which direction the mean of θ_1 shifts from its prior mean. If ρ_{12} is positive and $Z_2/V_2 > \mu_2$ then the posterior mean for θ_1 will be greater than μ_1. This represents the assumption that if θ_1 and θ_2 are correlated, then our prior beliefs will also be correlated. Therefore, if we observe an outcome on θ_2 that suggests that our prior mean was an underestimate of the truth, $Z_2/V_2 > \mu_2$, then we might believe that μ_1 is also an underestimate of the truth hence we should increase the mean of θ_1. The amount by which we shift the mean will naturally be dependent on our prior variance and observed variance, but it will also depend on the value of ρ_{12} that we specify in advance. Consequently, when specifying ρ_{12}, we should consider this as a measure of how far we would want our unobserved treatment effect

mean to shift based on indirect data.

Note that interestingly we can also use the Kalman filter [66] or the methods presented in [67] to tackle the problem presented here, which leads to the same posterior distribution. Details of this can be found in Appendix B.1.

We can then find the updated PoS for a study on $A + B$ using Equation (3.1), where $P(\theta|\text{data})$ will correspond to the marginal distribution for θ_1, or using Bayesian clinical trial simulation to estimate this expression. Note that, in order to calculate this value, we will also require the definition of study success for the study of $A + B$.

Following the above approach, the posterior distribution for our illustrative example is given by

$$\begin{pmatrix} \theta_M \\ \theta_C \end{pmatrix} \mid Z_C = 58.235 \sim \text{MVN}\left(\begin{pmatrix} 0.342 \\ 0.378 \end{pmatrix}, \begin{pmatrix} 0.053 & 0.004 \\ 0.004 & 0.006 \end{pmatrix}\right).$$

We can find the PoS of the mod-MARIANNE study using Equation (3.1) along with this posterior distribution and information on the study design. We will use a significance level of $\alpha = 0.05$ for the mod-MARIANNE study, which was also used in the MARIANNE study although it was split between the two comparisons. A power of 80% and a target HR of 0.75 will be used, as in the MARIANNE study design. This results in

$$\begin{aligned} V_M &= \left(\frac{Z_{0.05/2} + Z_{1-0.8}}{-\log(0.75)}\right)^2 \\ &= 94.838 \end{aligned}$$

following the method presented by Whitehead [21]. Hence, the PoS of the mod-MARIANNE study, based on the results of the CLEOPATRA study, is 0.711. If we had not included the information from the CLEOPATRA study, the PoS based

on the marginal prior distribution would have been 0.613.

If we use a prior correlation of 0.4, instead of 0.6 as above, then we would have a posterior distribution of

$$\begin{pmatrix} \theta_M \\ \theta_C \end{pmatrix} \mid Z_C = 58.235 \sim \text{MVN}\left(\begin{pmatrix} 0.324 \\ 0.378 \end{pmatrix}, \begin{pmatrix} 0.068 & 0.002 \\ 0.002 & 0.006 \end{pmatrix}\right),$$

which would lead to a PoS of 0.669 for the mod-MARIANNE study. These results are more conservative as we are choosing to borrow less information from the CLEOPATRA study, but the posterior PoS is still increased compared to the prior PoS. Alternatively, the posterior distribution based on a prior correlation of 0.8 results in a PoS of 0.777 for the mod-MARIANNE study. This illustrates the effect that the prior correlation has on the inference we make based on the output of this approach.

The three-arm MARIANNE study [63] was completed with study parameters as described previously and $\alpha = 0.05$ split between the two comparisons of the experimental treatments with control. The results of the study showed both experimental arms to be non-inferior, but not superior, to the control arm. The stratified hazard ratio for PFS for trastuzumab emtansine plus pertuzumab vs trastuzumab plus taxane was 0.87 [63]. It is noted by Perez et al. [63] that the median PFS of the control arm that was assumed when designing the study was shorter than what was observed. The median PFS of the control arm was assumed to be 11 months, which was based on information that was available at the time. The median PFS observed in the study control group was 13.7 months, which is similar to the estimate from more recent studies [63].

Note that, in this illustrative example, there were several differences between the two studies, yet our method is still able to add benefit in this case. This is because our method allows the user to consider how the beliefs regarding a treatment

effect change based on related study outcomes and the effect that this has on the probability of study success. The method does not require a high level of correlation between the treatment effects, nor does it require specific information on the similarities between the studies, it simply requires a parameter for the level of borrowing across the studies. This means that it is applicable in a wide range of settings and can be used to help inform and assist decision-making.

If there is doubt regarding the relationship between the study outcomes, the user might prefer the amount of borrowing to be dependent on the observed data. This would allow for a small amount of borrowing when the observed data suggests little correlation between study outcomes and a higher level of borrowing when the data suggests a relationship between outcomes. We present a robustification in the next section that aims to capture this requirement.

3.2.5 Robustification

In Section 3.2.4, we outlined a method that can be used to update the distribution of the treatment effects for a set of related combination therapies based on the outcome of a single combination study. Updating a distribution given relevant observations will always improve the accuracy of our estimates. However, so far, we have assumed that all of the therapies in our set of "related" combinations are truly correlated hence there is something to be gained from sharing information across the different combinations, but this might not always be the case.

In this section, we will consider an extension to the method that allows us to take into account the fact that two combinations might not be correlated and **robustify** our procedure against this. Since we are only observing one combination study at a time, we do not have the opportunity to learn from pairs of outcomes. Therefore, we cannot learn about the correlation and update our model using

this. Instead, we will consider how emerging data aligns with our prior beliefs, which is similar to recent work on extrapolation [68].

If we observe the outcome of a study in combination programme i, which we summarise using Z_i and V_i, we will naturally want to update our beliefs about θ_i using the study data. However, we may not necessarily want to update our beliefs about $\theta_{i'}$, for $i \neq i'$. When the posterior expectation of $\theta_{i'}$ given Z_i is similar to our prior expectation of $\theta_{i'}$, we would likely wish to include this additional information, as it does not seem too controversial given what we believed initially. However, if the marginal posterior of $\theta_{i'}$ is shifted "too far" in location given Z_i, this might cause some concern as to whether or not we should be including this indirect information. Therefore, our extension will allow us to include less of the indirect information when the jump size is large.

First, we will consider a **mixture prior** on $\boldsymbol{\theta}$ made up of two distributions. In the first distribution, the correlation between combinations will be set equal to zero, which implies no borrowing across combinations, and in the second distribution, the correlation will be set to the level that we would choose if we knew that they were in fact correlated. This value, as before, can be thought of as the amount that we would like to borrow across the combinations. We will write this mixture prior as

$$\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} \sim \omega_0^0 \times \text{MVN}\left(\begin{pmatrix}\mu_1\\ \mu_2\end{pmatrix}, \begin{pmatrix}\sigma_1^2 & 0\\ 0 & \sigma_2^2\end{pmatrix}\right) + \omega_1^0 \times \text{MVN}\left(\begin{pmatrix}\mu_1\\ \mu_2\end{pmatrix}, \begin{pmatrix}\sigma_1^2 & \rho_{12}\sigma_1\sigma_2\\ \rho_{12}\sigma_1\sigma_2 & \sigma_2^2\end{pmatrix}\right)$$

where the weights ω_0^0 and ω_1^0 can be thought of as the prior probabilities that θ_1 and θ_2 are uncorrelated and correlated, respectively, and $\omega_0^0 + \omega_1^0 = 1$.

If we were to update this mixture in the standard way then, as shown in Appendix B.2, the weights would remain unchanged despite the gain in information. Therefore, we will develop some further methodology in order to update the

weights and use the methodology from Section 3.2.4 to update the separate components.

Let us first consider the properties that we will want this procedure to have. Firstly, we want it to consider the amount that the distribution has shifted and to assign a higher weight to the uncorrelated distribution if this shift is too large, that is, it moves "too far" from what we initially thought was realistic. Conversely, if the shift in the marginal posterior mean is small and the study size is large, we would want to assign a higher weight to the correlated distribution. Finally, if the observed study is small, then we only want the weights to shift by a small amount compared to how much they would have shifted given equivalent results from a large study.

We want to update the weights by combining the prior weights, ω_0^0 and ω_1^0, with some new information that we will contain in a yet to be defined measure, p. This value will be used to quantify how much of the new information we want to borrow. The posterior weights will be given by

$$\omega_0^1 = \frac{(1-p)\omega_0^0}{(1-p)\omega_0^0 + p\omega_1^0} \quad \text{and} \quad \omega_1^1 = \frac{p\omega_1^0}{(1-p)\omega_0^0 + p\omega_1^0}. \tag{3.5}$$

We will require the value of p to take on the properties outlined above.

We will consider two ways of specifying p: a **hypothetical** posterior approach and a **limiting** posterior approach. Both of these approaches have desirable properties that align with the requirements outlined above for our weighting procedure.

For the hypothetical posterior approach, we construct a hypothetical normally distributed posterior for θ_1 given Z_2 that has posterior mean equal to the prior mean, μ_1, and posterior variance equal to the posterior variance found doing the

usual update given V_2. Hence, the hypothetical posterior is given by

$$N\left(\mu_1, \sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2}\right),$$

since $\mu_1^{\text{post}} = \mu_1$ and $\hat{\mu}_2 = Z_2/V_2$.

For the limiting posterior approach, we construct the limiting posterior distribution for θ_1 given Z_2 as $V_2 \to \infty$ that has posterior mean equal to the prior mean. Hence, here the posterior mean will be given, as before, by μ_1, and the posterior variance will be given by $\sigma_1^2\left(1-\rho_{12}^2\right)$ so that the limiting posterior is given by

$$N\left(\mu_1, \sigma_1^2\left(1-\rho_{12}^2\right)\right).$$

Our interest, however, lies in the location of the mean. Therefore, we will consider the lower and upper quartiles of these distributions, which we will denote by $[\theta_{1,l}^H, \theta_{1,u}^H]$ and $[\theta_{1,l}^L, \theta_{1,u}^L]$ for the hypothetical and limiting posterior distributions, respectively.

We then want to compare these quartiles with the posterior that we find using the standard GMRF procedure given the observed value of Z_2 and V_2. In order to do this we will truncate the posterior at its upper and lower quartiles, $\theta_{1,l}$ and $\theta_{1,u}$.

$$\theta_1|Z_2 = z_2 \sim \text{TN}\left(\mu_1 - \frac{\rho_{12}\sigma_1\sigma_2V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}z_2, \sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2}; \theta_{1,l}, \theta_{1,u}\right)$$

Then, we will take the value of p, the value that we use to update the weights, to be

$$p = P\left(\theta_1 \in \left[\theta_{1,l}^q, \theta_{1,u}^q\right] \mid Z_2 = z_2\right)$$

using the truncated posterior distribution of θ_1 where $q = H, L$ represents the hypothetical or limiting posterior distributions. This value is the probability of

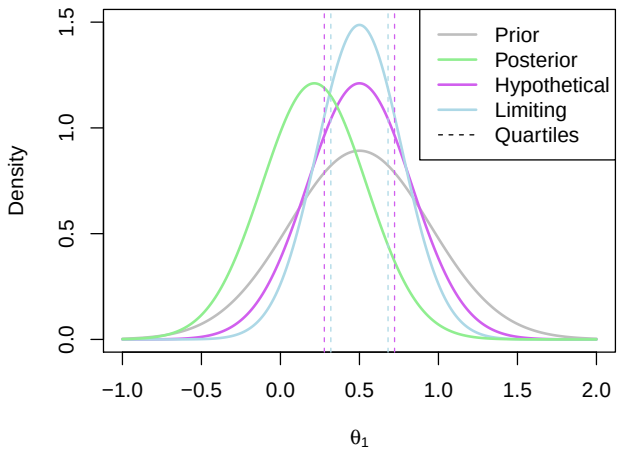

Figure 3.1: The hypothetical and limiting posterior distributions for an illustrative example.

the truncated posterior distribution lying within the lower and upper quartiles of the hypothetical/limiting posterior distributions. Hence, when the jump size is small, this probability will be large as there will be a large overlap between the distributions. On the other hand, when the jump size is large, this probability will be small, especially since we are taking the posterior truncated at the lower and upper quartiles. Note that if the posterior is perfectly aligned with the hypothetical or limiting posterior distribution then p will take a value of 1.

An example of what the hypothetical and limiting posterior distributions may look like can be found in Figure 3.1. In this example, our posterior beliefs do not align with our prior beliefs as we see a shift in mean. However, there still seems to be quite a large amount of overlap between the posterior distribution and the hypothetical posterior distribution, while there is less so with the limiting posterior distribution. This figure also serves to illustrate why we consider the truncated posterior rather than the original posterior. Recall that the posterior distribution here is based only on "related" data, and not on direct data. Therefore, considering the truncated distribution allows us to reduce the overlap in cases such as this one where the means are far enough apart for us to consider it to be a reason not to borrow from the "related" combination.

Once we have found our chosen value of p, we are able to find the updated weights using Equation (3.5) and then our mixture posterior will be given by

$$\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} \mid Z_2 = z_2 \sim \omega_0^1 \times \text{MVN}\left(\begin{pmatrix}\mu_1\\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}z_2\end{pmatrix}, \begin{pmatrix}\sigma_1^2 & 0\\ 0 & \frac{\sigma_2^2}{1+V_2\sigma_2^2}\end{pmatrix}\right) + \omega_1^1 \times \text{MVN}\left(\begin{pmatrix}\mu_1 - \frac{\rho_{12}\sigma_1\sigma_2 V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}z_2\\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}z_2\end{pmatrix}, \begin{pmatrix}\sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2} & \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}\\ \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} & \frac{\sigma_2^2}{1+V_2\sigma_2^2}\end{pmatrix}\right).$$

Following the approach outlined in Section 3.2.4, this posterior can be used to calculate the success probability of a combination study of interest by using the assurance as presented in [23].

If we applied this approach to our illustrative example then we would have a posterior distribution of

$$\begin{pmatrix}\theta_M\\ \theta_C\end{pmatrix} \mid Z_C = 58.235 \sim \omega_0^1 \times \text{MVN}\left(\begin{pmatrix}0.288\\ 0.378\end{pmatrix}, \begin{pmatrix}0.08 & 0\\ 0 & 0.006\end{pmatrix}\right) + \omega_1^1 \times \text{MVN}\left(\begin{pmatrix}0.342\\ 0.378\end{pmatrix}, \begin{pmatrix}0.053 & 0.004\\ 0.004 & 0.006\end{pmatrix}\right)$$

where the values of Z_C, V_C, μ and Σ were given in Section 3.2.4.

If we set $\omega_0^0 = 0.5$ and $\omega_1^0 = 0.5$, then the hypothetical posterior approach would lead to $\omega_0^1 = 0.16$ and $\omega_1^1 = 0.84$, which would give a PoS of the mod-MARIANNE study of 0.689. The limiting posterior approach yields similar results with $\omega_0^1 = 0.17$ and $\omega_1^1 = 0.83$ and a PoS of 0.688. As we would expect, the PoS under the robustified approach is between the PoS from the marginal prior of θ_M, 0.613, and the standard GMRF procedure, 0.711. They are also higher than the PoS when the prior correlation was set to 0.4, but this would not necessarily be the case if the observed data were further away from our prior beliefs. Figure

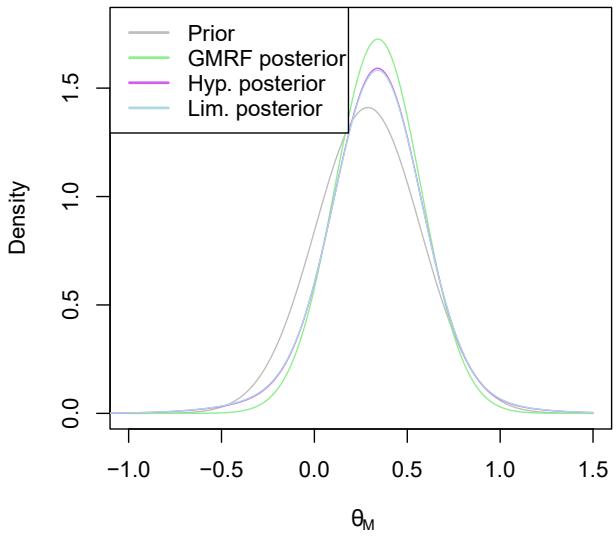

Figure 3.2: The marginal posterior distributions of θ_M for the illustrative example.

3.2 shows the different marginal posterior distributions of θ_M for this illustrative example.

If we wished to use this extension for more than two combinations, we would simply need to split our vector of random variables, $\boldsymbol{\theta}$, into pairs of random variables, $(\theta_i, \theta_{i'})\ \forall\, i \neq i'$. Each pair would need to contain θ_i, the treatment effect that we will observe some data for, alongside one of the correlated treatment effects. This would allow us to find the values of p in the same way presented here and would account for the fact that some pairs of θ_i and $\theta_{i'}$ might be strongly correlated, which would lead to a high weight on the correlated component, whereas other pairs might be uncorrelated, which would lead to a high weight on the uncorrelated component of the mixture. Thus, splitting the full n-dimensional problem into $n - 1$ two-dimensional problems could be the most appropriate approach in this setting, even though it is a heuristic.

3.3 Results

In this section, we will illustrate the performance of these methods by looking at the posterior distributions and the success probabilities that these methods lead to in a simulation study. We will compare the results of the proposed multivariate methods to the results of only marginal updating i.e. the univariate alternative. We will use the assurance to calculate the study success probabilities in the simulation study, as in previous sections, but it should be noted that other methods for calculating the PoS could be used instead.

In order to provide a complete picture of the way these multivariate methods perform compared to the univariate alternative, we will consider different sets of prior distributions that may have arisen from historical data such as the results of a small study. We will take the true value of θ_1 and θ_2 to be equal to 0.5.

We will assume that the prior information on both of these parameters is equivalent to having a prior variance of 0.2. This is approximately equal to having an uninformative "pre-prior" and updating based on the outcome of a study involving 20 patients with normally distributed responses.

We will assume that we observe the outcome of a study on $A + C$ and want to update the distributions for both $A + B$ and $A + C$ based on this. If we do not consider borrowing information across the combinations, the prior distribution will represent all of the information, or beliefs, that we have regarding $A + B$ and we will make our decisions based on this distribution in the univariate setting.

3.3.1 Effect of the sample size on the PoS

In order to illustrate what might happen to the PoS for different sample sizes, we can consider a fixed study outcome and find the PoS using this outcome

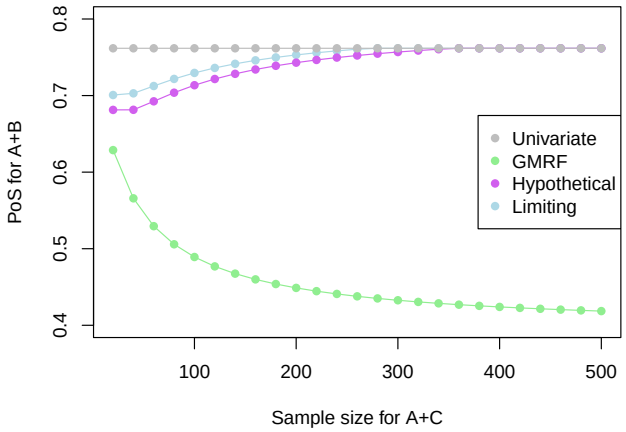

Figure 3.3: Plot showing the PoS for a study on θ_1 as the sample size of a study on θ_2 increases, which has an outcome of $Z_2/V_2 = 0$.

with different sample sizes. Figure 3.3 shows how the PoS for a study on $A + B$ changes as the sample size changes in the observed study on $A + C$ for a fixed outcome of no effect, $Z_2/V_2 = 0$, and a prior mean of $\boldsymbol{\mu} = (0.5, 0.5)^T$. We define the PoS to be the assurance for a future two-sided superiority study on $A + B$ that has a planned sample size of 500 and a significance level of 0.05.

We see that as the sample size increases, the PoS of the GMRF approach decreases as there is more evidence to suggest that $\theta_2 = 0$, which would suggest that, if θ_1 and θ_2 are correlated, then our prior mean for θ_1 is also an overestimate. However, the hypothetical and limiting posterior approaches originally have a lower PoS than the univariate method, but as the sample size increases, so will the jump from the prior to the posterior mean of θ_1. Since these approaches will assign a higher weight to the univariate approach when the jump size increases, the success probabilities from the hypothetical and limiting posterior approaches tend towards the PoS of the univariate approach as the sample size increases. Note that no method here is performing better than the other, as we have not defined what the truth is and our prior mean for θ_1 might be an overestimate, or θ_1 and θ_2 might be uncorrelated. This figure simply serves as an illustration as to how the different methods assign the PoS.

3.3.2 Simulation set-up

We will consider the sample size of the study on A + C to be equal to 500 as we would be most interested in borrowing information and using this methodology when we observe the outcome of a large (e.g. Phase III) study.

In order to account for the variability in the treatment effect estimate that we might have based on such a small study, we will consider three different prior means for θ_2. (Results for different prior means of θ_1 can be found in Appendix B.3.) We will consider prior means of 0.2, 0.5 and 0.8 for θ_2. These values correspond to the posterior means we would find given an uninformative "pre-prior" and an update based on the quartiles of the distribution of the score statistic when the true value of θ_i is equal to 0.5 and the value of V_i is equivalent to a study size of 20 patients with normally distributed responses.

We set up the different prior distributions and we simulate 10000 replications of Z_2 from $Z_2|\theta_2 = 0.5 \sim N(0.5 \times 125, 125)$ where $V_2 = 125$ corresponds to approximately 500 patients with normally distributed responses. We then update each of the different prior distributions to find the set of 10000 posterior distributions for each prior using a correlation of $\rho_{12} = 0.8$.

As before, we considered the definition of the PoS for a future study on $A + B$ to be equal to the assurance for a two-sided superiority study with a planned sample size of 500 and a significance level of 0.05. We further assumed that, in order to run a study on combination $A + B$, we would need to observe a PoS of at least 0.6. The selection of an appropriate decision criterion on the PoS is discussed in [69]. We recorded the PoS of each replication along with whether or not this would lead to a "go" decision.

3.3.3 Results

The results of the simulation study are given in Table 3.1. For the univariate approach, we do not need to consider multiple replications of a study on combination $A + C$, as we would only consider direct information on combination $A + B$ in this approach. Therefore, the mean PoS and the proportion of "go" decisions for the univariate approach actually correspond to the PoS and the "go" decision based on the prior distribution, since no direct information on combination $A + B$ is observed in the simulation study. The mixture (no updates) approach in the table is the standard Bayesian approach as mentioned in Section 3.2.5 and presented in Appendix B. This approach has the same distributional components as the hypothetical and limiting posterior approaches but the weights are not updated.

	$\mu_1 = 0.2$	$\mu_2 = 0.2$	$\mu_2 = 0.5$	$\mu_2 = 0.8$
Univariate	Mean PoS	0.520	0.520	0.520
	% "Go" (PoS > 0.6)	0	0	0
GMRF	Mean PoS	0.802	0.530	0.242
	% "Go" (PoS > 0.6)	99.6	22.7	0
Hypothetical	Mean ω_0^1	0.623	0.137	0.625
	Mean PoS	0.615	0.529	0.429
	% "Go" (PoS > 0.6)	74.9	18.6	0
Limiting	Mean ω_0^1	0.639	0.153	0.642
	Mean PoS	0.610	0.529	0.434
	% "Go" (PoS > 0.6)	70.9	17.9	0
Mixture	Mean PoS	0.661	0.525	0.381
(no updates)	% "Go" (PoS > 0.6)	95.9	4.9	0

Table 3.1: Table showing the results for combination $A + B$ of the simulation study where the true values of θ_1 and θ_2 are given by 0.5 and $\mu_1 = 0.2$ and μ_2 represent the prior means for each combination. Note that the univariate approach does not update the distribution of combination $A + B$ based on the results of combination $A + C$ and the mixture (no updates) approach is the mixture approach using standard Bayesian updating, which does not update the weights in this setting.

In Figure 3.4, fifty posterior distributions under the GMRF approach are plotted for fifty replications with the prior means given by $\mu_1 = 0.2$ and (a) $\mu_2 = 0.2$, (b) $\mu_2 = 0.5$ and (c) $\mu_2 = 0.8$.

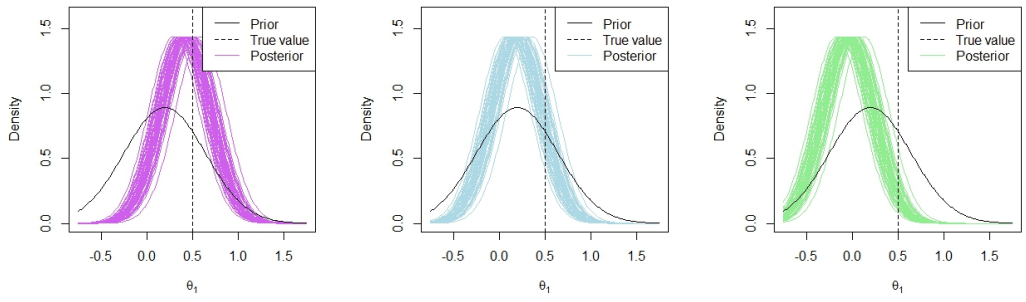

Figure 3.4: Marginal posterior distributions of θ_1 for 50 of the 10000 replications using the GMRF approach. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a) $\mu_2 = 0.2$, (b) $\mu_2 = 0.5$ and (c) $\mu_2 = 0.8$.

We see that when we have a prior mean of $\mu_1 = 0.2$, this leads to a PoS of 0.520 in the univariate approach, which does not exceed the required threshold to make a decision to run the next study. Therefore, under the "go" rule of the PoS exceeding 0.6, if we do not use any indirect data, we will never run a study based on this univariate prior, despite the true value of θ_1 being equal to 0.5.

However, when we do include the indirect data, we make many more "go" decisions. This, however, is also dependent on what the prior mean for combination $A + C$ was. When $\mu_2 = 0.2$, this is underestimating the true value of θ_2, therefore many of the observed studies will result in an estimate that exceeds the prior mean. This means that the posterior mean of combination $A + C$ will be increased in the majority of cases and, since we have set $\rho_{12} = 0.8$, the posterior mean of combination $A + B$ will also increase from a prior mean of $\mu_1 = 0.2$. This will cause an overall increase in the PoS compared to when we did not include indirect data, hence we will choose to "go" in the majority of cases. This is what we observe in Table 3.1, with the mean PoS being equal to 0.802 and the majority of the PoS values exceeding 0.6 resulting in 99.6% of decisions being "go" decisions. This is also reflected in Figure 3.5(a), which provides a histogram of the success probabilities for this set of prior means.

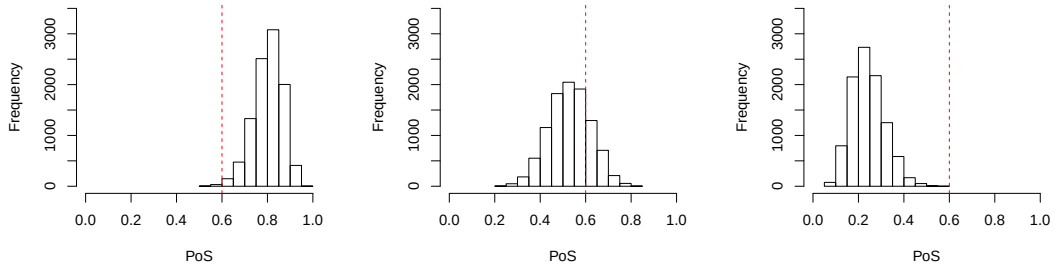

Figure 3.5: Histograms of the PoS for a study on θ_1 in the simulation study using the GMRF approach. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a) $\mu_2 = 0.2$, (b) $\mu_2 = 0.5$ and (c) $\mu_2 = 0.8$.

When $\mu_1 = 0.2$ and $\mu_2 = 0.5$, the majority of the replications will lead to little change in the posterior mean from the prior mean of combination $A + C$ as the true treatment effect is given by $\theta_2 = 0.5$. Therefore, this will cause the posterior mean of combination $A + B$ to also remain similar to its prior mean as there is little difference between the data and our prior beliefs. However, since we will still be borrowing information, our posterior variance for combination $A + B$ will decrease. This will cause the PoS to increase slightly compared to the univariate PoS. We observe a mean PoS of 0.530 and we would make the decision to run a study on combination $A + B$ in 22.2% of cases, which is reasonably higher than had we not included the indirect information, despite the fact that our posterior mean for θ_1 will still actually be an underestimate.

However, when our prior means for θ_1 and θ_2 underestimate and overestimate the truth, respectively, the multivariate method performs worse than the univariate method. This is because the data will cause the posterior mean for θ_2 to reduce from $\mu_2 = 0.8$ to be closer to 0.5. This, in turn, will also cause the posterior mean of θ_1 to decrease from its already low prior mean of $\mu_1 = 0.2$. This leads to a mean PoS of 0.242 and zero "go" decisions in all 10000 replications. Consequently, if there is a chance that the prior estimates of the effects modelled may be incorrect in opposite directions, this methodology may not be appropriate.

However, one might assume that there will be some correlation between the prior estimates of related correlations and so we might expect that they will often be wrong in the same direction, given the nature of the problem. This is, in fact, what is assumed by the methodology presented here.

These results are also highlighted in the plots provided in Figure 3.4. In Figure 3.4(a), we see that when θ_2 is underestimated by its prior mean, this leads to a posterior mean for θ_2 that is on average higher than the prior mean. When θ_2 is equal to its prior mean, on average the posterior mean is equal to the prior mean for θ_1, as seen in Figure 3.4(b), and when the prior mean for θ_2 is an overestimate, the posterior mean for θ_1 is on average lower than the prior mean, as seen in Figure 3.4(c).

These patterns will hold for other values of μ_1, μ_2 and θ_2 when the prior correlation is positive. When μ_2 is an overestimate of the true value of θ_2, the mean of θ_1 will decrease and cause a lower PoS than when μ_2 is equal to the true value of θ_2. Similarly, when μ_2 underestimates the true value of θ_2, the mean of θ_1 will increase and cause a higher PoS than when μ_2 is equal to the true value of θ_2. Results for our example with $\mu_1 = 0.5$ and $\mu_1 = 0.8$ can be found in Appendix B.3.

In this simulation study, we considered both θ_1 and θ_2 to be equal. If this was not the case, the patterns observed here would still be the same. That is, if we observe a value of Z_2/V_2 that is greater than the prior mean, μ_2, this will cause the posterior mean of θ_1 to be greater than μ_1, assuming a positive prior correlation. Similarly, if we observe $Z_2/V_2 < \mu_2$, this would result in a posterior mean of θ_1 that is less that μ_1. The size of the jump from the prior mean to the posterior mean of θ_1 is related to the size of the difference between Z_2/V_2 and μ_2, the prior correlation, the observed study size and the prior variances.

We also considered the performance of the two possible mixture approaches that

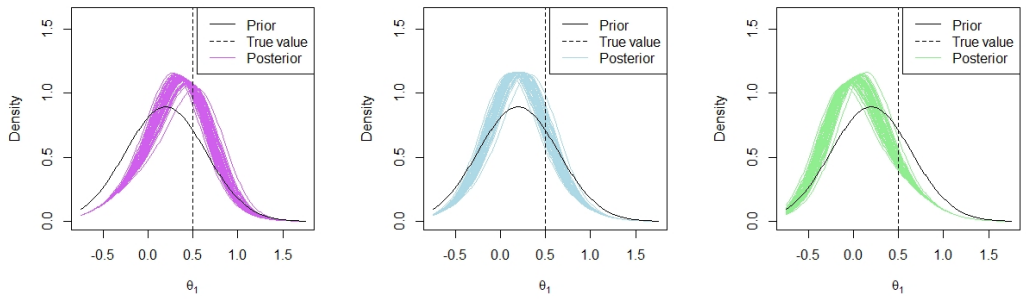

Figure 3.6: Marginal posterior distributions of θ_1 for 50 of the 10000 replications for the mixture approach with no weight updates. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a) $\mu_2 = 0.2$, (b) $\mu_2 = 0.5$ and (c) $\mu_2 = 0.8$.

were presented in Section 3.2.5 along with the standard Bayesian approach in the simulation study and the results are presented in Table 3.1. The posterior distributions for fifty replications are given in Figures 3.6 and 3.7 and the histograms of the success probabilities are given in Figures 3.8 and 3.9 for the mixture approaches. We used prior weights of 0.5 for the correlated and uncorrelated components in each of the mixture approaches.

Since the mixture prior approach was introduced to account for the fact that two combinations may not be correlated and to borrow less when this is the case, the results that we see for all three mixture approaches are not as extreme as in the GMRF approach in most cases. That is, the mean PoS and the proportion of "go" decisions are lower in the mixture approach than in the GMRF approach when the indirect data causes an increase in mean, i.e. $\mu_2 = 0.2$, and the mean PoS and the proportion of "go" decisions are higher than in the GMRF approach when the indirect data causes a decrease in mean, i.e. $\mu_2 = 0.8$. This is what we would hope to see given that when there is a large jump in means, a higher weight is assigned to the uncorrelated component of the model. In addition, the values of the mean PoS and the proportion of "go" decisions for all three mixture approaches in Table 3.1 lie between the values of the univariate and GMRF approaches. This is intuitive given that the mixture approaches are weighted mixtures of these two

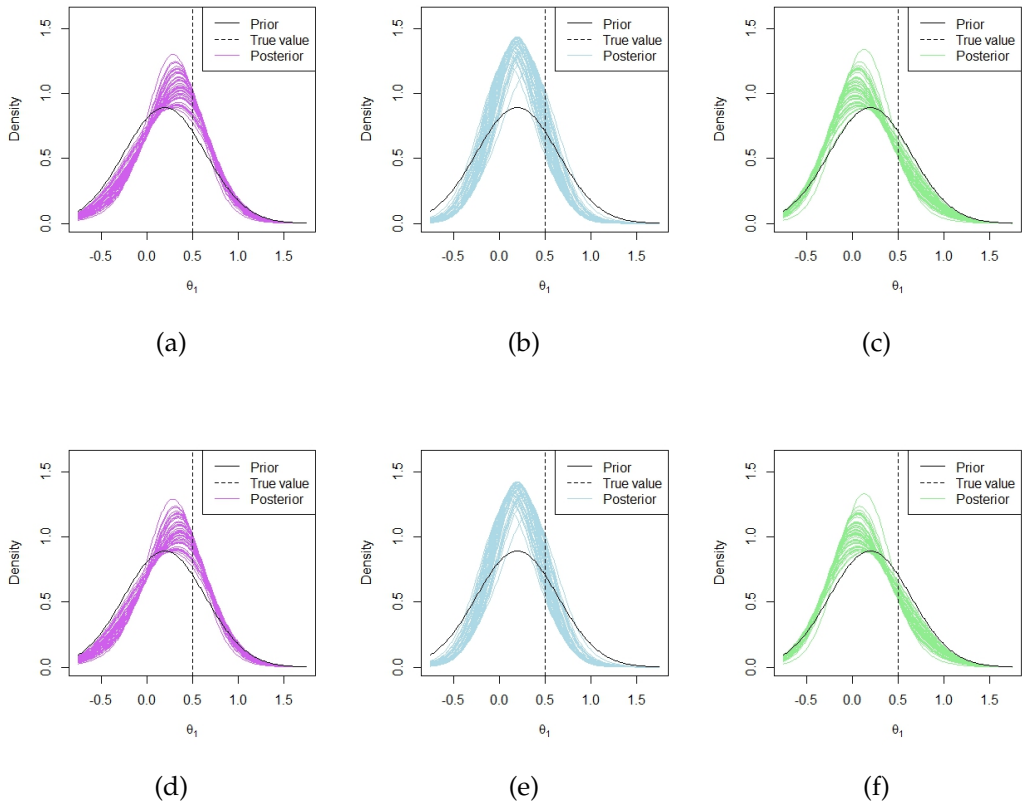

Figure 3.7: Marginal posterior distributions of θ_1 for 50 of the 10000 replications using (a) - (c) the hypothetical posterior approach and (d) - (f) the limiting posterior approach. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a,d) $\mu_2 = 0.2$, (b,e) $\mu_2 = 0.5$ and (c,f) $\mu_2 = 0.8$.

models hence the PoS of each replication will be bound by the univariate PoS and the GMRF PoS for that replication.

In Figure 3.7, we see very similar patterns in terms of the posterior distributions under the two mixture approaches. One of the key things that we notice from these posteriors is that the peak of the distributions are much closer to the prior means than in Figure 3.4, which showed the posteriors under the GMRF approach. This is due to the way in which we specified the weightings in Section 3.2.5.

One of the places where the effect of the mixture approach is the most apparent is when $\mu_1 = 0.2$ and $\mu_2 = 0.8$. Here, the GMRF method observed, in most cases, evidence of a lower value of θ_2 than was predicted by μ_2, causing the posterior

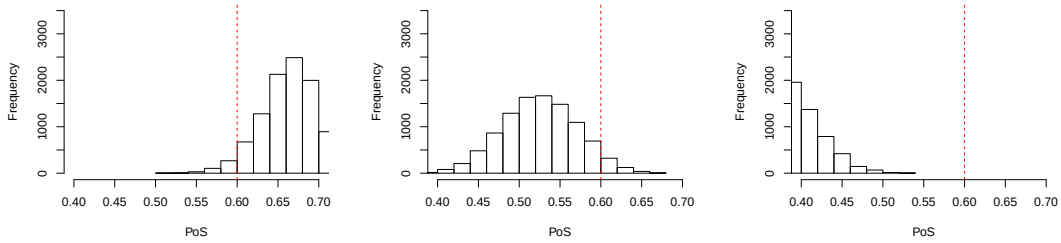

(a) (b) (c)

Figure 3.8: Histograms of the PoS for a study on θ_1 in the simulation study using the mixture approach with no weight updates. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a) $\mu_2 = 0.2$, (b) $\mu_2 = 0.5$ and (c) $\mu_2 = 0.8$.

means of both θ_1 and θ_2 to decrease compared to their prior means resulting in a low mean PoS and no "go" decisions in all 10000 replications. In both of the mixture approaches, however, this jump from the prior to the posterior mean of θ_2 caused the method to assign a higher posterior weight to the uncorrelated component of the mixture. This meant that the posterior mean did not drop as low as in the GMRF case, hence the mean PoS is much higher in the mixture approaches than in the GMRF approach. This is exactly the situation where the mixture approach provides a benefit over the GMRF approach.

The difference in the performance of the hypothetical posterior approach as compared to the limiting posterior approach is less obvious in Table 3.1. To learn more about the differences, we look at the histograms presented in Figure 3.9. In the histograms presented, we see that the limiting posterior approach is less likely to assign more extreme values of the PoS than the hypothetical posterior approach. In Figure 3.9(d), the final bar on the histogram is much smaller than those that precede it, despite an overall upwards trend until that point. This contrasts with what we see in Figure 3.9(a). Similarly, in Figure 3.9(c), we see that the first bar in the histogram is quite an amount higher than the first bar in the histogram in Figure 3.9(f). This is also reflected in Table 3.1 where we see that, in the simulation study, the limiting posterior approach assigns a higher posterior weight on average to the uncorrelated component of the model than

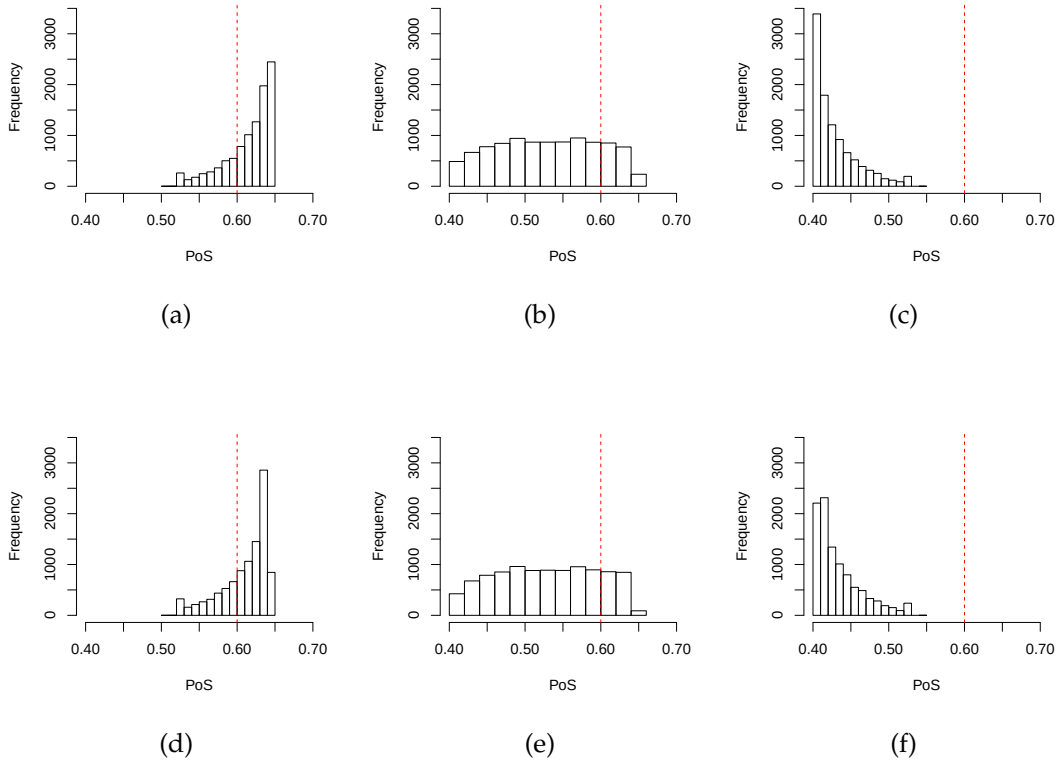

Figure 3.9: Histograms of the PoS for a study on θ_1 in the simulation study using (a) - (c) the hypothetical posterior approach and (d) - (f) the limiting posterior approach. The prior mean for θ_1 was set to $\mu_1 = 0.2$ and the prior mean for θ_2 was set to (a,d) $\mu_2 = 0.2$, (b,e) $\mu_2 = 0.5$ and (c,f) $\mu_2 = 0.8$.

the hypothetical posterior approach in all three cases.

The reason that the limiting posterior approach is less likely to assign these more extreme values is related to the way that the weights are assigned. In both mixture approaches, a value of $p = 1$ means that the posterior under the GMRF approach is perfectly aligned with either the hypothetical or limiting posterior. In the simulation study, there is potential for values of $p = 1$ in the hypothetical posterior approach. In fact, values of $p = 1$ will always be possible no matter the size of the study. Even a study with only 20 patients could cause a weight of 1 to be assigned to the correlated component. However, the limiting posterior approach is more cautious in how it assigns the values of p. To have a value of $p = 1$, one would need a study that is large enough to result in a posterior variance equal to the limiting posterior variance. This means that the

hypothetical posterior mixture approach has potential to assign more extreme weights based on less information than the limiting posterior mixture approach would need to assign an equally extreme weighting. This is why we see the differences in the histograms for the success probabilities.

3.4 Discussion

In this chapter, we have presented a method that allows us to update the estimates of the treatment effects for a set of related combination studies based on a single observation. This allows us to include both direct and indirect data in the treatment effect estimates, which reduces the variance, hence improves the accuracy of these estimates. The probability distributions representing our beliefs about a particular therapy can often be used to gain insight into the expected performance of a new therapy, but they are also often used to calculate the probability of success of an upcoming study through the calculation of the assurance [23] or other distribution based definitions of the probability of study success. This success probability is then used to assist decision-making regarding the study. Improving the accuracy of treatment effect estimates will allow decision-making to improve by providing the decision makers with the ability to recognise beneficial, or ineffective, treatments sooner.

The CLEOPATRA and MARIANNE studies were used to build an example that illustrates how the proposed methodology can be used in the real world. Identical marginal priors were used for both studies and the results of the CLEOPATRA study were used to update the joint distribution.

The methods presented provide an overall advantage over traditional univariate approaches due to the fact that they are able to use the available data appropriately. The time when these methods might not perform as well as the traditional

methods are when the prior means of the treatment effects are incorrect by quite a significant distance in opposite directions. However, we presented an extension to our method in Section 3.2.5 that allows this to be accounted for and limits the jump size from the prior mean to the posterior mean when indirect data has caused the jump.

In order to use the presented methods to calculate the PoS of an upcoming study based on the results of a related study, three types of information are required. The first is the significance level and the planned sample size of the upcoming study, both of which should be readily available if we are considering whether or not to run the study. The second is the score statistic and the Fisher information of the completed study, which should also be available at the conclusion of the study. The final type of information relates to the prior parameters for the distribution of treatment effects. The prior mean and variance for a treatment effect is a standard requirement when calculating the PoS, whereas the prior correlation is an additional requirement of our approach over standard approaches. Rather than trying to quantify the correlation between treatments, one may instead consider this parameter as the amount of indirect information they would like to use (i.e. the strength of borrowing) when calculating the updated PoS as shown in Section 3.2.4. We also presented a robustification in Section 3.2.5 that allows the alignment of our prior beliefs with the data to guide the degree of borrowing across combinations. The value of the PoS will always be dependent upon the prior parameters when we use methods such as the assurance [23] to calculate it. Therefore, our method will naturally have some sensitivity towards the choices of these prior parameters and users should explore this when specifying prior parameters. However, since our method allows the user to incorporate relevant study data in this calculation, the PoS calculated under this approach will be less reliant upon the prior mean and variance than traditional univariate approaches.

	Benefits	Limitations
Univariate	Well established/accepted No bias from indirect data	Cannot capture relationships
GMRF	Captures relationships Indirect data used	Relies on prior assumptions
Robustification	Captures relationships Mixture of methods	Relies on prior assumptions

Table 3.2: Summary of the conclusions and recommendations of when to use each approach.

We highlighted the performance of the multivariate method in Section 3.3 and showed that, compared to the univariate approach for calculating for the PoS, it leads to improved decision-making regarding whether or not a particular combination study should be run.

A summary table of the approaches presented in this chapter can be found in Table 3.2.

While most of the examples discussed, and the results in Section 3.3, were presented for a pair of combinations, it should be noted that the method can be used for any number of combinations. The method could also be used to assist internal decision-making based on external data. For example, multiple companies have developed PD-1/PD-L1 inhibitors, which are often combined with chemotherapies to treat different cancers. Companies could use the results of an external study to update the PoS of a study of their PD-1/PD-L1 inhibitor combination in the same indication. Furthermore, these methods could be applied not only in the setting of related combination studies but also in different settings where there is potential to share information across studies. Some potential settings that could benefit from these methods include the same combination but also in different indications and programmes in different regions.

Although the methodology presented in this chapter is clearly useful in itself, our main interest lies in decision-making for portfolios of combinations. Therefore, in Chapter 4 we will look at decision-making for portfolios of combinations. The

methodology presented in this chapter will be added to an existing heuristic [49] for the PS approach [39], which was discussed in Chapter 2. Incorporating this methodology into portfolio management techniques allows the relationships between studies to be considered at the planning stage and for study outcomes to guide the decisions made for future related combination studies, which was not accounted for previously.

Chapter 4

Chapter 4

Improving decision-making for portfolios of combination therapies

4.1 Introduction

Current methods for pharmaceutical portfolio decision-making do not take into account the differences between single agent and combination drug development. That is, they consider each development programme as independent of each other while in many instances some relationships can be expected. In this chapter, we present a method that extends the work discussed in Chapter 2 to consider relationships between different drug development programmes using the approach detailed in Chapter 3.

Existing methods for pharmaceutical portfolio management address decisions such as the scheduling of studies [39], out-licensing [38] and study design [37]. These methods generally use mathematical programming techniques, often stochastic programming, and aim to find the set of decisions that maximise the value of the portfolio. Portfolio management techniques help to ensure that the most valuable programmes are selected and conducted optimally. This in turn can lead to reductions in the failure rates of different phases, and programmes overall, alongside helping to ensure that beneficial treatments reach the patient

population as soon as possible and that investments are made in the most promising areas.

However, none of the existing methods for portfolio management are able to capture certain aspects of combination drug development such as the relationships between similar combinations or the additional logistical decisions that are required in combination drug development. Taking into account the relationships between combination therapies allows us to make better informed decisions and improve portfolio outcomes.

In this chapter, we present a method that extends the work of Colvin and Maravelias [39], which was presented and discussed in detail in Chapter 2, to consider combination drug development. Motivation and context for the approach is provided in Appendix C through the use of an illustrative example that is based on a real world pharmaceutical portfolio, the Roche neuroscience pipeline [70] as of January 2019.

In Section 4.2, we provide the setting for the problem of scheduling studies in a pharmaceutical portfolio. We also provide details of a heuristic for the project scheduling approach [39] that was presented by Christian and Cremaschi [49] and extended in [71]. This heuristic overcomes the issue of the full multi-stage stochastic programme (MSSP) of the project scheduling approach [39] not being solvable for portfolios containing more than six drugs in reasonable time. It achieves this by decomposing the programme into a series of smaller knapsack subproblems.

In Section 4.3, we highlight the relevance of updating the study success probabilities in the portfolio management problem. We then detail how this procedure can be added to the **"after every realisation"** (AER) version [71] of the **knapsack decomposition algorithm** (KDA) heuristic [49] and provide the full details of our approach, the **adaptive knapsack decomposition algorithm** (aKDA), in

Section 4.4. The implementation and performance of the aKDA is compared to the KDA and the MSSP for the project scheduling approach in Section 4.5.

Several extensions to the aKDA are presented in Section 4.6, such as the inclusion of more than two study outcomes in the decision tree along with additional constraint sets that address different aspects of combination drug development. This chapter is concluded with a discussion of the aKDA and how it can be implemented in the real world in Section 4.7.

4.2 Portfolio-level decision-making

Portfolio management can cover many different types of decisions regarding the portfolio, from which programmes to include in the portfolio [36] to how the programmes should be designed [37]. We are interested in the problem of which studies to select and how to schedule the selected studies. In this section, we will formalise this problem, briefly discuss the approach of Colvin and Maravelias [39] that was detailed in Chapter 2 and present a heuristic for this approach [49][71].

4.2.1 Problem formulation

We assume that there is a finite number of programmes, $|I|$, in the portfolio, and that each programme, $i \in I$, has a known, finite number of studies remaining, $|J_i|$. We assume that studies within a programme must be completed sequentially and a study, (i, j), can only be initiated if all previous studies have been completed successfully.

We further assume that there is information available that will allow us to estimate the probability of success (PoS) of each of the studies, ϕ_{ij}; study durations, τ_{ij}; study costs, c_{ij}; study resource requirements, λ_{ijr}; resource availability, $\lambda_r^{\max}$; the revenue that would be realised upon successful completion of the programme, $\text{rev}_i^{\max}$.

We are interested in finding the set of programmes that should be initiated, along with when each of the studies within the programmes should be initiated. Given that the outcome of a study is uncertain, we require that the scheduling takes this uncertainty into account and provides decisions based on the different potential study outcomes. The optimal decisions will be dependent on the objective, which we will assume is linked to the value of the portfolio. These decisions will also be subject to resource constraints.

4.2.2 Project scheduling approach

Colvin and Maravelias [39] presented a multi-stage stochastic programme (MSSP) for the problem of scheduling clinical studies within a portfolio. This approach finds the set of decisions that should be made in order to maximise the expected net present value (ENPV) of the portfolio, whilst accounting for the uncertainty in the outcomes of the clinical studies. This approach was discussed in detail in Chapter 2, so here we will not go into detail and will simply provide a summary of the key points.

The input parameters required for this model include those listed in Section 4.2.1. The method also requires penalties to be specified that should be incurred per time period for late completion and reduced active patent life. These penalties encourage studies to be completed as soon as possible and without delays between sequential studies. A discounting factor is also required that is applied

to the cost of a study and encourages a study to be run later in the planning horizon, had it not been initiated earlier.

The scenario-based MSSP of the project scheduling approach considers all of the different possible trial outcome scenarios and, upon solving, it returns the optimal set of decisions, which can be represented using a decision tree. For more details of this approach, see Chapter 2 and [39].

Although this approach is able to capture many of the different, crucial aspects of the process of portfolio management, it generally cannot be solved in reasonable time for portfolios containing more than six drugs. This is due to the size of the MSSP, which increases quickly as we increase the number of programmes due to the requirement of the non-anticipativity constraints. Given that our interest lies in portfolios containing combination therapies, for which there may be a large number of potential combinations of interest due to the nature of the problem, this makes the MSSP formulation unsuitable.

4.2.3 Knapsack decomposition algorithm

Christian and Cremaschi [49] presented multiple heuristics to tackle the issue of the MSSP formulation [39] being unsolvable for portfolios containing more than six drugs. The most promising heuristic was a knapsack decomposition algorithm (KDA) that breaks the full MSSP down into a series of smaller subproblems that are solved iteratively at each node in the decision tree. The subproblems are similar in construction to the traditional knapsack problem discussed in Chapter 1, where the items correspond to the studies and the capacity of the knapsack corresponds to the resource constraints. In this section, we present the "after every realisation" (AER) version of the KDA that was presented in [71] as it is noted that this is an improvement on the original version of the KDA that was presented in [49]. The differences between these versions will be discussed at

the end of this section, along with some of the other modifications presented by Christian and Cremaschi [71].

The algorithm starts at the beginning of the planning horizon, at which point there is only one subproblem corresponding to the current state of the portfolio. This is represented in the resulting decision tree by the first node. The knapsack subproblem at this point, which will be defined below, must then be solved. The solution to this subproblem will contain the studies that should be run at this time point. Then, based on the studies selected, further subproblems are generated that correspond to the time at which we will observe the next set of study outcomes and the potential outcomes that will be observed i.e. study success or failure. We denote a particular subproblem by $[t, k]$, where $t \in T$ is the time of the subproblem and $k \in K_t$ is the particular subproblem at this time point. Here, K_t is the set of subproblems at time t and this set will be built during the KDA. The different subproblems at each time point will be represented by different nodes in the decision tree at that time point.

In each subproblem, the first step is to find the set of eligible studies. A study is eligible if all prerequisite studies have been completed successfully, the study is not ongoing and it has not already been completed. This is represented by an eligibility indicator, E_{ijtk}, which is equal to 1 when study (i, j) is eligible to be run in the kth subproblem at time t, $[t, k]$.

Then, the values of the studies must be calculated. The value of study (i, j) at time t is given in [49] to be

$$R_{ijt} = \left[R_{ij}^{p} - \gamma_i^{L} \left(t + \sum_{j' \geq j} \tau_{ij'} \right) \right] \times \prod_{j' \geq j} \phi_{ij'} \tag{4.1}$$

This value is calculated by considering the potential revenue, R_{ij}^{p}, that is associated with study (i, j) and applying the deductions in revenue due to a shorter active patent life, γ_i^{L}, that were applied in the original PS approach. Then, we

multiply this by the probability of the remaining trials being successful to return the expected revenue.

R_{ij}^{p} is the potential revenue generated by study (i, j) and is given by

$$R_{ij}^{p} = \text{rev}_{i}^{\max} - \sum_{j' \geq j} c_{ij'} \left[1 - 0.025 \sum_{j''=j+1}^{j'} \tau_{ij''-1} \right],$$

where $\text{rev}_{i}^{\max}$ is the total maximum possible revenue generated by drug i, c_{ij} is the cost of study (i, j), τ_{ij} is the duration of study (i, j), ϕ_{ij} is the PoS of study (i, j) and γ_{i}^{L} is the penalty due to a smaller market share.

The costs here are linearly depreciated by a factor of 0.025 to make it attractive to run trials later in the planning horizon if they have not already been initiated. This depreciation was also performed in the objective function of the PS approach, as can be seen in Appendix A.2.

Note that, unlike the full MSSP, the penalty for reduced active patent life, γ_{i}^{D}, is not considered in the value calculation given in Equation (4.1). This is because the KDA assumes that each trial is initiated as soon as the previous one is concluded when calculating R_{ijt}. Hence, a penalty for a programme being idle is not incurred.

Each study, (i, j), will have a corresponding weight, which corresponds to the resource requirement of the study and is given by $\lambda_{ijr} \ \forall r \in R$.

The objective function of the knapsack subproblem at $[t, k]$ is given by

$$\text{maximise} \sum_{i,j} R_{ijt} X_{ij} \tag{4.2}$$

where X_{ij} is a binary decision variable that is equal to 1 when study (i, j) is selected to be run in subproblem $[t, k]$. The constraints of this knapsack subproblem

ensure that any studies, (i, j), that we select are eligible to be run,

$$X_{ij} \leq E_{ijtk} \ \forall i, j \tag{4.3}$$

and that the resource requirements do not exceed the resource availability,

$$\sum_{i,j} \lambda_{ijr} \left(X_{ij} + y_{ij} \right) \leq \lambda_r^{\max} \ \forall r, \tag{4.4}$$

where y_{ij} is equal to 1 when study (i, j) was initiated prior to the current subproblem and is currently ongoing in $[t, k]$.

Overscheduling constraints are included to ensure that there will be enough resources available in the future to complete all of the selected programmes.

$$\sum_{i,j' \geq j} X_{ij} \lambda_{ij'r} \tau_{ij'} \leq \lambda_r^{\max} \max \left\{ \sum_{j'>j} \tau_{ij'} + 1 \forall i, j : E_{ijtk} = 1 \right\} \ \forall r \tag{4.5}$$

This constraint aims to capture the ability of the MSSP to consider future outcomes and find the decision that results in the best expected outcome based on these future outcomes.

Finally, there are constraints on the form of the decision variables.

$$X_{ij} \in \{0, 1\} \ \forall i, j \tag{4.6}$$

The algorithm for the AER version of the KDA [71] is given below.

```
1  t := 1;
2  while t ≤ |T| do
3      k := 1 ;
4      while k ≤ |K_t| do
5          Calculate R_ijt, E_ijtk ∀ i, j;
6          Solve knapsack problem to find solution X_ij ∀i, j;
7          Find the time t' until the next observation given X_ij and y_ij;
8          Generate set S of knapsack problems given observations at time t + t';
9          K_{t+t'} = K_{t+t'} ∪ S;
10         k := k + 1;
11     t := t + 1;
```

The KDA significantly reduces the size of the formulation and the time taken to find the optimal solution. Christian and Cremaschi [49] noted that, in their analysis of the original KDA, most objective values were within 3% of the true optimal value of the MSSP. This suboptimality can arise for different reasons, which we will now briefly discuss.

The "do nothing" solutions of the MSSP were found to be the setting where the KDA was found to perform the worst. However, it should be noted that, in reality, running no studies when there are studies available to run is not an option a company would consider. Therefore, this drawback of the KDA may not be as much of a drawback in the real world as it appears to be in the analyses. Christian and Cremaschi [49] also found that the optimal set of decisions found by the KDA is not sensitive to changes in the parameters.

One drawback is that the KDA has a tendency to underutilise resources due to the inclusion of the overscheduling constraint. The overscheduling constraint considers the case where all selected studies are successful hence resources must

be available for all future studies to be completed simultaneously, which is an unlikely situation hence leads to underutilisation of resources. This underutilisation limits future investments in additional products hence leading to suboptimal decisions compared to the full MSSP discussed in detail in Chapter 2.

At the beginning of the section it was noted that the KDA presented here is the "after every realisation" (AER) version of the algorithm [71]. The original KDA [49] only generated subproblems upon the conclusion of all selected studies, "after all realisations". This means that it did not need to consider ongoing studies when generating subproblems and it led to a sparser decision tree. The AER version [71] was then presented as a modification and it was shown to provide solutions that are closer to those of the full MSSP presented in [39]. Another alternative generation rule that was presented in [71] generated subproblems at "each time period" (ETP). In this version of the KDA, subproblems are generated iteratively at every time period in the planning horizon regardless of whether or not any studies complete at that point. This was shown to provide improved results to the AER version in a small number of cases, but the authors concluded that the added benefit was not worth the increase in computational burden [71]. For this reason, we only consider the AER version of the KDA in this chapter and so wherever we refer to the KDA, we are referring to the AER version of the KDA.

Another set of modifications to the KDA that was presented by Christian and Cremaschi [71] is related to the overscheduling constraint, given in Equation (4.5). It was noted in [49], and mentioned previously, that the overscheduling constraint often leads to conservative solutions compared to the solutions of the MSSP. Therefore, Christian and Cremaschi [71] presented an objective function penalty approach and a probabilistic constraint approach as alternatives to the original constraint. However, these two approaches were shown to be inferior to the original constraint in most cases. Therefore we will only consider the original

overscheduling constraint in the work that follows.

4.3 Updating study success probabilities

It was previously noted that the PoS of a study is one of the measures that is most commonly used to assist decision-making regarding the running of a study. It can also be used, amongst other measures, to compare different studies that are competing for resources. The full MSSP formulation of the PS approach [39] uses the PoS to calculate the probability of the different trial outcome scenarios, which are then averaged over to calculate the expected net present value (ENPV) of the portfolio. The KDA [49][71], however, does not consider the full set of trial outcome scenarios and instead uses the PoS to calculate the value of a study in each subproblem, leading to a direct comparison of the value of the different studies and how the PoS affects this value.

Although the MSSP and KDA require the study success probabilities as input parameters that are fixed throughout the optimisation procedure, they do not require a specific definition of the PoS. One of the most common methods for calculating the PoS uses the concept of the assurance [23] and was discussed in Chapter 1. Using this approach, the PoS of a study is given by

$$\text{PoS} = \int P\left(\text{study success} \mid \theta\right) P\left(\theta \mid \text{data}\right) \mathrm{d}\theta \tag{4.7}$$

where θ represents the treatment effect. There is often a closed form solution available for this calculation and when there is not, it can be estimated using Bayesian clinical trial simulation [23].

It is clear that, using this formulation of the PoS, the inclusion of new data regarding θ can be easily incorporated via Bayesian updating of $P\left(\theta \mid \text{data}\right)$. The inclusion of indirect data, for example data from a related combination study,

is less straightforward. One of the main differences between single agent and combination drug development is the potential to learn across related combination studies, where we deem two studies to be related if they have at least one treatment in common and are used in the same indication, for example. In Chapter 3, we discussed the way in which we can use these relationships to update the PoS for a combination study based on the outcome of related studies. This was achieved by updating $P(\theta \mid \text{data})$ given the outcome of a related study and is recalled here.

Let θ_1 and θ_2 represent the treatment effects of two related combination studies and further assume that we can represent our prior beliefs for these using a bivariate normal distribution. Then, if the outcome of the study on θ_2 has a score statistic of Z_2 and a Fisher information of V_2, we can find

$$\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} \mid Z_2 \sim \text{MVN}\left(\begin{pmatrix}\mu_1 - \frac{\rho_{12}\sigma_1\sigma_2 V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}z_2\\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}z_2\end{pmatrix}, \begin{pmatrix}\sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2} & \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}\\ \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} & \frac{\sigma_2^2}{1+V_2\sigma_2^2}\end{pmatrix}\right) \tag{4.8}$$

where μ_i and σ_i^2 give the prior mean and variance of θ_i, respectively, and $\rho_{ii'}$ defines the prior correlation between θ_i and $\theta_{i'}$, which may be considered as the level of borrowing.

As discussed in Chapter 3, we can find the updated success probability for the study of θ_1 given the outcome of the study of θ_2 by taking the marginal distribution of θ_1 in Equation (4.8) and using this to calculate the assurance via Equation (4.7).

4.4 Method

In this section, we present a framework that extends the work of Christian and Cremaschi [71] to consider relationships between the different programmes in

the portfolio. We do this using the methods presented in Chapter 3 to update the PoS throughout the decision-making procedure whenever a relevant outcome is observed.

This method requires the same input parameters as the KDA in addition to a set of parameters relating to updating the PoS. For each programme, i, we need to consider the main treatment effect of interest, which we will denote by θ_i. This is the treatment effect that we consider in the assurance calculation given in Equation (4.7) and so it might be chosen to be the treatment effect relating to the primary endpoint in Phase III, for example. Then, once the effect has been chosen, we require a binary indicator for whether study (i, j) being "successful" is related to the main treatment effect of interest in programme i, θ_i. We will denote this indicator by θ_{ij}^{ind}. This is because the method that we use to update the study success probabilities was designed for use in the later phases of drug development, where a study success is related to the main treatment effect of interest. For example, a Phase I study being deemed successful is not typically related to the main treatment effect of efficacy of the drug whereas the success of a Phase III study typically is related to the efficacy endpoint. We require this indicator so that we only update the study success probabilities after relevant outcomes are observed.

For each study, (i, j), with $\theta_{ij}^{\text{ind}} = 1$, we require the significance level, α_{ij}, the power, $1 - \beta_{ij}$, and the minimally important difference, δ_{ij}. For studies with $\theta_{ij}^{\text{ind}} = 0$, we require the PoS, ϕ_{ij}, to be specified as an input parameter. Then for each programme, i, in the portfolio we require a prior mean, $\mu_i^{[1,1]}$, for θ_i, which could be set equal to δ_{ij}, and a prior standard deviation, $\sigma_i^{[1,1]}$. The final requirement is the correlation between the different drugs in the portfolio, $\rho_{ii'}$, which can be considered as the level of borrowing of information across the drugs, as noted in Chapter 3.

Using these parameters, we are able to find the critical value of the hypothesis

test of each of the studies, d_{ij}, and the planned Fisher information of the study, V_{ij}. We will require both of these during the decision-making procedure and they are given by

$$d = \frac{Z_{\alpha/2}\left(Z_{\alpha/2} + Z_{\beta}\right)}{\delta} \qquad V = \left(\frac{Z_{\alpha/2} + Z_{\beta}}{\delta}\right)^2 \tag{4.9}$$

where $Z_x = \Phi^{-1}(1-x)$ and $\Phi(\cdot)$ is the standard normal cumulative distribution function.

The aKDA begins at the first time point, $t = 1$, in the first and only subproblem at this time point, $k = 1$. This node in the decision tree is referred to as subproblem $[1, 1]$, as was used in the KDA in Section 4.2.3. The first step at this node, and at any node that has eligible studies available, is to build the distribution of $\boldsymbol{\theta}$. This is required as the aKDA calculates the PoS based on the current distribution of $\boldsymbol{\theta}$ for the eligible studies in each subproblem rather than using the same PoS across all subproblems as in the KDA.

The distribution of $\boldsymbol{\theta}$ is built in each subproblem, $[t, k]$, using the most recent parameter estimates of the mean and standard deviation of θ_i, which we denote by $\mu_i^{[t,k]}$ and $\sigma_i^{[t,k]}$. We use the prior correlation, $\rho_{ii'}$, rather than using an updated version, as this is treated as fixed throughout the decision-making process due to the fact that it controls the level of borrowing across studies.

The multivariate normal distribution is used to model the distribution of $\boldsymbol{\theta}$ at each node. In subproblem $[t, k]$, this is given by

$$\begin{pmatrix} \theta_1 \\ \vdots \\ \theta_{|I|} \end{pmatrix} \sim \text{MVN}\left(\begin{pmatrix} \mu_1^{[t,k]} \\ \vdots \\ \mu_{|I|}^{[t,k]} \end{pmatrix}, \begin{pmatrix} \sigma_1^{[t,k]^2} & \cdots & \rho_{1|I|}\sigma_1^{[t,k]}\sigma_{|I|}^{[t,k]} \\ \vdots & \ddots & \vdots \\ \rho_{1|I|}\sigma_1^{[t,k]}\sigma_{|I|}^{[t,k]} & \cdots & \sigma_{|I|}^{[t,k]^2} \end{pmatrix}\right).$$

After the distribution of $\boldsymbol{\theta}$ has been built, the next step in the aKDA is to find the

eligible studies along with their values and success probabilities. The eligible studies are defined in the same way as in the KDA. A study is eligible if it has not yet been initiated and if all prerequisite studies have been completed successfully. We use the marginal distribution of θ_i along with α_{ij}, β_{ij} and δ_{ij} to find the PoS of study (i, j) in subproblem $[t, k]$, $\phi_{ij}^{[t,k]}$, using the calculation of the assurance as given in Equation (4.7) for all studies with $\theta_{ij}^{\text{ind}} = 1$. For studies with $\theta_{ij}^{\text{ind}} = 0$, the PoS should be specified as one of the input parameters. The value of each of the studies, R_{ijt}, can then be calculated using the same formula as in the KDA approach, which is given in Equation (4.1).

At this point, everything necessary has been specified and so we can solve the knapsack problem at $[t, k]$, which is defined by Equations (4.2) - (4.6). Solving this knapsack problem will return the values of the decision variables, X_{ij}, that maximise the value of the objective function in this node. When $X_{ij} = 1$ this means that study (i, j) has been selected to run in subproblem $[t, k]$.

After the knapsack subproblem has been solved and the studies to run in $[t, k]$ have been identified, we must find which studies will be the next to complete out of those that have just been selected and those that were already ongoing in $[t, k]$. The time until this observation will be made is denoted by t'. Therefore, the time that we must generate the next set of subproblems for is given by $t + t'$.

The generation of subproblems requires two sets of information. The first set of information is the same as in the KDA and relates to the studies that have been completed by this time point and whether they were successful. The second set of information relates to the distribution of $\boldsymbol{\theta}$. The distribution of $\boldsymbol{\theta}$ will be updated if any of the studies due to complete at $t + t'$ have $\theta_{ij}^{\text{ind}} = 1$. If all of the completing studies have $\theta_{ij}^{\text{ind}} = 0$, then the current distribution will be sent to the generated subproblems.

In order to update the distribution of $\boldsymbol{\theta}$ given relevant outcomes, we will follow

the method presented in Chapter 3. This method requires the score statistic, Z_{ij}, and the Fisher information, V_{ij}, of the two-sided test with null hypothesis given by $H_0 : \theta_i = 0$. Since we do not have this information at the planning stage, we will need to select appropriate values to use in place of the true study outcomes.

For each study that is next to complete with $\theta_{ij}^{\text{ind}} = 1$, we will consider two outcomes, success and failure. This is the same as in the KDA and the subproblems are generated based on this information. This means that we will require a value of Z_{ij} to represent a study success, which we will denote by Z_s, and study failure, which we will denote by Z_f. We will use the same value for the Fisher information in both of these cases, which we will set equal to the planned Fisher information given in Equation (4.9).

Let us consider a study, (i, j), that is next to complete and has $\theta_{ij}^{\text{ind}} = 1$. For the values of Z_s and Z_f, we will consider the normal distribution truncated at the critical value for the study, d_{ij}, as given in Equation (4.9). We know that Z_{ij} is approximately normally distributed with mean $\theta_i V_{ij}$ and variance V_{ij}, but we do not know θ_i as this is the true treatment effect, which can only ever be estimated. Therefore, we will use our most recent estimate of θ_i, which is given by $\mu_i^{[t,k]}$. This defines the parent distribution for our truncated normal distributions. The distributions of Z_s and Z_f are given by

$$Z_s \sim \text{TN}\left(\mu_i^{[t,k]} V_{ij}, V_{ij}; d_{ij}, \infty\right) \text{ and } Z_f \sim \text{TN}\left(\mu_i^{[t,k]} V_{ij}, V_{ij}; -\infty, d_{ij}\right), \tag{4.10}$$

respectively. However, we still require a single value for Z_s and Z_f. Therefore, we will use the expectations of these two distributions in our updating.

$$E\left(Z_s\right) = \mu_i^{[t,k]} V_{ij} + \sqrt{V_{ij}} \left(\frac{\phi\left(\frac{d_{ij} - \mu_i^{[t,k]} V_{ij}}{\sqrt{V_{ij}}}\right)}{1 - \Phi\left(\frac{d_{ij} - \mu_i^{[t,k]} V_{ij}}{\sqrt{V_{ij}}}\right)} \right)$$

$$E\left(Z_f\right) = \mu_i^{[t,k]} V_{ij} - \sqrt{V_{ij}} \left(\frac{\phi\left(\frac{d_{ij} - \mu_i^{[t,k]} V_{ij}}{\sqrt{V_{ij}}}\right)}{\Phi\left(\frac{d_{ij} - \mu_i^{[t,k]} V_{ij}}{\sqrt{V_{ij}}}\right)} \right)$$

We are then able to use the method discussed in Section 4.3 and presented in Chapter 3 to update the distribution of $\boldsymbol{\theta}$ given study success, $E\left(Z_s\right)$, or study failure, $E\left(Z_f\right)$.

If more than one study is due to complete next after $[t,k]$ with $\theta_{ij}^{\text{ind}} = 1$, then, rather than considering the distribution of $\left(\boldsymbol{\theta}^T, Z_{ij}\right)$ as in Chapter 3, we will consider all relevant Z_{ij} in this distribution, along with the correlations between them. For example, in a three drug example where we decide to run the first studies for drugs 2 and 3 simultaneously and $\theta_{21}^{\text{ind}} = \theta_{31}^{\text{ind}} = 1$, the distribution that we consider would be given by

$$(\boldsymbol{\theta}, \mathbf{Z})^T \sim \text{MVN}\left(\text{E}((\boldsymbol{\theta}, \mathbf{Z})^T), \text{Var}((\boldsymbol{\theta}, \mathbf{Z})^T)\right),$$

where $(\boldsymbol{\theta}, \mathbf{Z})^T = (\theta_1, \theta_2, \theta_3, Z_{21}, Z_{31})^T$, $\text{E}((\boldsymbol{\theta}, \mathbf{Z})^T) = \left(\mu_1^{[t,k]}, \mu_2^{[t,k]}, \mu_3^{[t,k]}, V_{21}\mu_2^{[t,k]}, V_{31}\mu_3^{[t,k]}\right)^T$ and the covariance matrix, $\text{Var}((\boldsymbol{\theta}, \mathbf{Z})^T)$, is given by

$$\begin{pmatrix}
\sigma_1^{[t,k]^2} & \rho_{12}\sigma_1^{[t,k]}\sigma_2^{[t,k]} & \rho_{13}\sigma_1^{[t,k]}\sigma_3^{[t,k]} & V_{21}\rho_{12}\sigma_1^{[t,k]}\sigma_2^{[t,k]} & V_{31}\rho_{13}\sigma_1^{[t,k]}\sigma_3^{[t,k]} \\
\rho_{12}\sigma_1^{[t,k]}\sigma_2^{[t,k]} & \sigma_2^{[t,k]^2} & \rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & V_{21}\sigma_2^{[t,k]^2} & V_{31}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} \\
\rho_{13}\sigma_1^{[t,k]}\sigma_3^{[t,k]} & \rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & \sigma_3^{[t,k]^2} & V_{21}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & V_{31}\sigma_3^{[t,k]^2} \\
V_{21}\rho_{12}\sigma_1^{[t,k]}\sigma_2^{[t,k]} & V_{21}\sigma_2^{[t,k]^2} & V_{21}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & V_{21}^2\sigma_2^{[t,k]^2} + V_{21} & V_{21}V_{31}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} \\
V_{31}\rho_{13}\sigma_1^{[t,k]}\sigma_3^{[t,k]} & V_{31}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & V_{31}\sigma_3^{[t,k]^2} & V_{21}V_{31}\rho_{23}\sigma_2^{[t,k]}\sigma_3^{[t,k]} & V_{31}^2\sigma_3^{[t,k]^2} + V_{31}
\end{pmatrix}.$$

The values in this covariance matrix are found using the same approach as in Chapter 3. We can then find the distribution of $\boldsymbol{\theta} \mid \mathbf{Z}$ using the approach in Chapter 3.

Once the updated distribution of $\boldsymbol{\theta}$ has been found, the individual updated means, $\mu_i^{[t+t',k']}$, and the individual updated standard deviations, $\sigma_i^{[t+t',k']}$, should

be used in the generated subproblem. The updated covariances should not be used in the generated subproblem. This is because the correlation between drugs shrinks after performing the update since no information is observed in the update relating to this quantity. Furthermore, in this setting the correlation is considered as the level of borrowing, which we require to be constant through time. Thus, we calculate the covariances using the updated standard deviations and the prior correlations.

The set of generated subproblems are then fully defined and we are able to move onto the next subproblem in the algorithm. The algorithm concludes when there are no eligible studies to be run in any subproblem or when it reaches the predefined end of the planning horizon, $|T|$. The full details of the aKDA are given below with additions to the original KDA shown in italics.

$t := 1$;
while $t \leq |T|$ **do**
  $k := 1$;
  while $k \leq |K_t|$ **do**
    Build distribution for $\boldsymbol{\theta}$ *given* $\mu_i^{[t,k]}$, $\sigma_i^{[t,k]}$ *and* $\rho_{ii'}$;
    Calculate $\phi_{ij}^{[t,k]}$ *for all studies with* $\theta_{ij}^{ind} = 1$;
    Calculate $R_{ijt}, E_{ijtk}\ \forall\, i, j$;
    Solve knapsack problem to find solution $X_{ij}\ \forall i, j$;
    Find the time t' until the next observation given X_{ij} and y_{ij};
    Generate set S of knapsack problems given observations at time $t + t'$;
    for $s \in S$ **do**
      Find the updated parameters for each θ_i;
    $K_{t+t'} = K_{t+t'} \cup S$;
    $k := k + 1$;
  $t := t + 1$;

This approach can be applied to and solved for much larger portfolios, unlike the full MSSP, and is able to capture the relationships between the different programmes, unlike the MSSP and the KDA. In the next section, we will provide a detailed discussion of a simulation study for a smaller example portfolio.

Drug, i	$\text{rev}_i^{\max}$	c_{i1}	c_{i2}	τ_{i1}	τ_{i2}	λ_{i11}	λ_{i21}	λ_{i12}	λ_{i22}	γ_i^L	γ_i^D
1	3100	90	220	4	4	1	2	2	3	19.2	22
2	3250	80	200	3	5	2	2	1	3	19.6	28
3	3300	90	180	3	4	1	2	1	3	20	26

Drug, i	θ_{i1}^{ind}	θ_{i2}^{ind}	α_{i1}	α_{i2}	β_{i1}	β_{i2}	δ_{i1}	δ_{i2}	$\mu_i^{[1,1]}$	$\sigma_i^{[1,1]^2}$
1	1	1	0.1	0.05	0.2	0.1	0.25	0.25	0.25	0.1
2	1	1	0.1	0.05	0.2	0.1	0.25	0.25	0.25	0.1
3	1	1	0.1	0.05	0.2	0.1	0.25	0.25	0.25	0.1

Table 4.1: Parameters used in the MSSP, KDA and aKDA for the example portfolio.

4.5 Results

In this section, we compare the aKDA to the MSSP [39] and the KDA [71] in a simulation study. In order to illustrate the implementation of the methods and provide performance results, we consider an example portfolio containing three programmes. Each programme in the portfolio has two consecutive studies remaining and the parameters used for this example are given in Table 4.1. The planning horizon used is 12 time periods and there are two resource types with maximum availabilities given by $\lambda_1^{\max} = 2$ and $\lambda_2^{\max} = 3$. We specified a prior correlation between θ_2 and θ_3 of 0.6 and specified no correlation, hence no borrowing, between θ_1 and any other θ_i.

Since the success probabilities in the aKDA are calculated in each subproblem, we set the success probabilities in the MSSP and the KDA to be equal to those in the first subproblem, $[1, 1]$, in the aKDA. This gives $\phi_{i1} = 0.6$ and $\phi_{i2} = 0.62$ for all $i \in I$.

4.5.1 Decision trees

The decision trees found using the MSSP, KDA and aKDA for the portfolio with parameters given in Table 4.1 can be found in Figure 4.1. In these decision trees, a green arrow indicates study success and a red arrow indicates study

failure. The boxes in the decision tree indicate the studies that are initiated at the corresponding time point and subproblem and a box containing "NA" indicates that there were no eligible studies to run. The x-axis in each of the decision trees represents time.

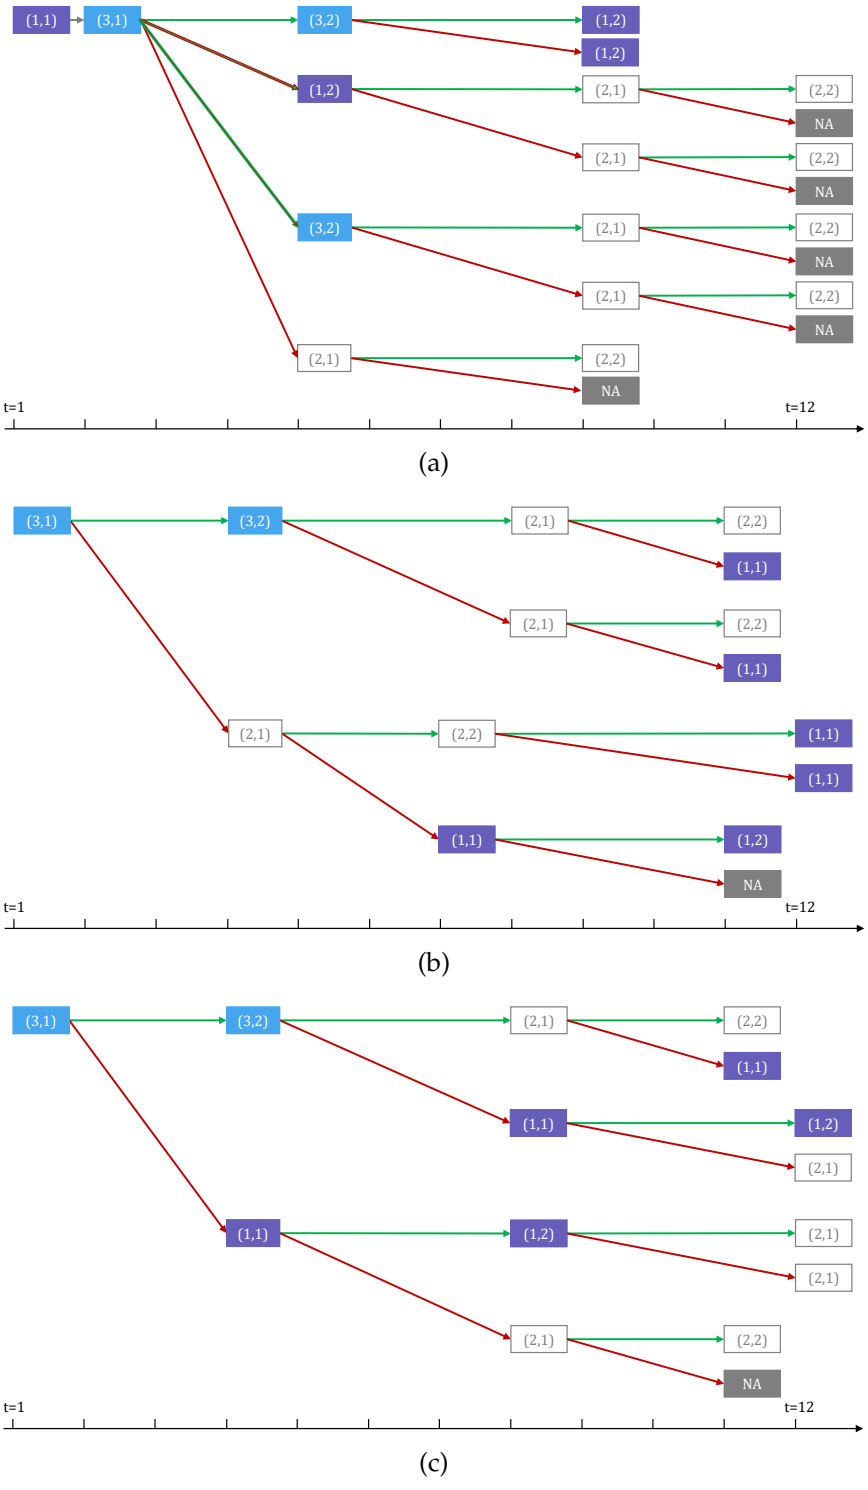

Figure 4.1: Decision trees found using (a) the MSSP, (b) the KDA and (c) the aKDA.

The most noticeable difference between the decision trees is that the MSSP decision tree appears to be denser than the other two trees. This is due to the restrictions that the overscheduling constraint puts on study selection, which was discussed in Section 4.2.3 and is discussed in more detail in [71].

All three approaches select to run study $(3,1)$ in one of the first two time points and $(3,2)$ is selected upon successful completion of $(3,1)$. The MSSP then prioritises the development of Drug 1 over Drug 2, unlike the KDA. This is because the MSSP runs $(3,1)$ and $(1,1)$ simultaneously, making it more profitable to run $(1,2)$ before moving onto $(2,1)$. Note that $(3,1)$ and $(2,1)$ cannot be run simultaneously in any approach since the requirement of resource 1 for $(2,1)$ is equal to the maximum availability of resource 1, $\lambda_{211} = \lambda_1^{\max} = 2$. Also, the KDA is unable to select to run $(3,1)$ and $(1,1)$ simultaneously, as was seen in the MSSP tree, due to the overscheduling constraint. However, the most interesting differences appear later in the decision trees.

We specified a prior correlation of 0.6 between θ_2 and θ_3, both of which were uncorrelated with θ_1. The effect of this can be seen in Figures 4.1(b) and (c). In the KDA decision tree, study $(2,1)$ is selected when $(3,1)$ is unsuccessful. However, in the aKDA tree, study $(1,1)$ is selected when $(3,1)$ is unsuccessful. This is because $\theta_{31}^{\text{ind}} = 1$, therefore the distribution of $\boldsymbol{\theta}$ is updated on completion of this study at $t = 4$. When $(3,1)$ is successful, it is most profitable, and indeed most realistic, to continue development for drug 3. Conversely, when it is unsuccessful, this causes the mean of θ_3 to decrease from the prior mean, $\mu_3^{[1,1]} = 0.25$. This causes the means of all positively correlated θ_i to decrease also. Hence, the mean of θ_2 is decreased, which causes the PoS to decrease, making $(2,1)$ less attractive to run. Thus, the aKDA chooses to run $(1,1)$ over $(2,1)$, which is the choice in the KDA. This is what we would hope to see given that if we believe the outcomes of the studies for drugs 2 and 3 are related, then a failure in one programme would cause us to believe that a failure in the other is more likely. We also see

this upon the unsuccessful completion of $(3, 2)$ in Figure 4.1(c) as compared to Figure 4.1(b).

In terms of the computational time required to produce the decision trees, both the KDA and aKDA took under 10 seconds to run with the KDA taking one second and the aKDA taking seven seconds. The MSSP, however, took 36 seconds to run. Although these times are small, the pattern will remain the same as we increase the number of studies until the MSSP can no longer be solved in reasonable time. This happens quickly due to the non-anticipativity constraints. Here, there were only 27 trial outcome scenarios, therefore 27^2 non-anticipativity constraints, before any reductions on the constraint set are performed. If we increase the portfolio size to $|I| = 4$, each with two studies, then we would have 81 scenarios and 81^2 non-anticipativity constraints. The KDA and aKDA can both be solved for much larger portfolios without becoming too computationally intensive.

4.5.2 Simulation study design

We explored the performance of the three approaches in a simulation study where we simulated study outcomes and found the decisions that would be made under each decision tree given these outcomes. This allowed us to find the NPV of each set of decisions using the calculation given in the objective function of the MSSP [39], which is as follows.

$$NPV = Rv + FRv - Cst$$

$$Cst = \sum_{ijt} cd_t c_{ij} X_{ijt}$$

$$Rv = \sum_{i \in S^I} \sum_{t} \left\{ rev_i^{\max} X_{i|J_i|t} - \gamma_i^D \left(Z_{i|J_i|-1t} + Z_{i|J_i|t} \right) - \gamma_i^L \left(t + \tau_{i|J_i|} \right) X_{i|J_i|t} \right\}$$

$$FRv = \sum_{i \in S^I} \sum_{j} rev_{ij}^{\text{open}} f_{ij} Z_{ij|T|} + \sum_{i \in S^I} \sum_{j \geq |J_i|-1} \sum_{t > |T|-\tau_{ij}} rev_{ijt}^{\text{run}} f_{ij+1} X_{ijt}$$

Here, S^I is the set of successful programmes. Note that, unlike in the full MSSP of the PS approach discussed in Chapter 2, here the variables X_{ijt} and Z_{ijt} are not dependent on the trial outcome scenario, s. This is because, in the MSSP, these equations were used to calculate the NPV in each of the different trial outcome scenarios, $s \in S$. However, here we are using the equations to calculate the NPV under each of the replications of study outcomes for each of the three methods. For simplicity, we have not introduced additional notation to show this, but the variables X_{ijt} and Z_{ijt} are dependent on the method used (MSSP, KDA or aKDA) and the replication. For further discussion of the calculation of the NPV and the parameters involved in this, see Chapter 2.

The main difference that sets the aKDA approach apart from the other approaches is the fact that it is able to capture the relationships between programmes in the portfolio. Therefore, we considered four different correlation structures between the different treatment effects, $\theta_i \, \forall i \in I$, represented by matrices (a) - (d).

$$\text{(a)} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0.6 \\ 0 & 0.6 & 1 \end{pmatrix} \text{(b)} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \text{(c)} \begin{pmatrix} 1 & 0 & 0.6 \\ 0 & 1 & 0 \\ 0.6 & 0 & 1 \end{pmatrix} \text{(d)} \begin{pmatrix} 1 & 0.6 & 0.6 \\ 0.6 & 1 & 0.6 \\ 0.6 & 0.6 & 1 \end{pmatrix}$$

The correlation structure that was used to generate the decision trees was equal to correlation matrix (a), hence this is the case where we expect the aKDA to perform the best. The correlation matrix given in (b) shows no correlation between any of the treatment effects. This is what is assumed by the MSSP and the KDA. The third correlation matrix, (c), represents the case where there is a relationship between two treatment effects, but not the ones that we originally thought. The final correlation matrix, (d), is the case where we were correct about the correlation between θ_2 and θ_3, but other correlations exist that we were not aware of when building the decision trees.

Then each replication in the simulation study goes as follows:

1. Draw a vector of true treatment effects, $\boldsymbol{\theta}^*$, from

$$\boldsymbol{\theta} \sim N_3\left(\begin{pmatrix}0.25\\0.25\\0.25\end{pmatrix}, \begin{pmatrix}0.05 & \rho_{12}^{(x)}0.05 & \rho_{13}^{(x)}0.05\\ \rho_{12}^{(x)}0.05 & 0.05 & \rho_{23}^{(x)}0.05\\ \rho_{13}^{(x)}0.05 & \rho_{23}^{(x)}0.05 & 0.05\end{pmatrix}\right)$$

   where (x) represents the correlation matrix that is being used.

2. Draw the study outcomes, Z_{ij}^*, from

$$Z_{ij} \sim N\left(\theta_i^* V_{ij}, V_{ij}\right).$$

3. Find the decisions made for this set of outcomes under the MSSP, KDA and aKDA.
4. Find the corresponding NPV for this set of outcomes under the MSSP, KDA and aKDA.

We specified the means of $\theta_i \; \forall i \in I$ used in the simulation study to be equal to the reference value that is used in each of the studies, $\delta_{ij} = 0.25$, and the variance to be equal to 0.05 so that the distribution produces realistic values of the treatment effects but is otherwise uninformative.

We can then find the number of times each method leads to the highest NPV and explore the distribution of the NPVs under each approach.

4.5.3 Simulation study results

The results of the simulation study are summarised in Table 4.2 and Figures 4.2 and 4.3.

		Full comparison	Heuristic comparison
(a)	MSSP	55.16 (29.96)	-
	KDA	8.55 (2.56)	17.42 (7.16)
	aKDA	11.67 (0.21)	25.51 (15.77)
	KDA + aKDA	24.61 (19.96)	57.07 (-)
		Full comparison	Heuristic comparison
(b)	MSSP	55.66 (35.14)	-
	KDA	11.81 (3.30)	25.84 (10.94)
	aKDA	8.15 (0.16)	17.37 (10.74)
	KDA + aKDA	24.39 (16.87)	56.79 (-)
		Full comparison	Heuristic comparison
(c)	MSSP	55.98 (30.78)	-
	KDA	17.21 (4.16)	26.06 (16.13)
	aKDA	11.12 (0.19)	17.05 (6.95)
	KDA + aKDA	15.69 (11.09)	56.89 (-)
		Full comparison	Heuristic comparison
(d)	MSSP	55.93 (21.32)	-
	KDA	9.61 (2.87)	18.05 (7.89)
	aKDA	18.87 (0.27)	25.15 (7.58)
	KDA + aKDA	15.60 (10.58)	56.80 (-)

Table 4.2: The percentage of times each method produced the highest NPV in the simulation study with 500000 replications for the correlation matrices (a) - (d). The row "KDA + aKDA" represents the situation where the KDA and the aKDA produced the highest NPV simultaneously. The numbers in brackets give the percentage of times that each method produced the highest NPV, which is also at least 3% higher than the second highest NPV.

Table 4.2 provides the percentage of times each approach led to the highest NPV in the simulation study for 500000 replications. The results are given for the full comparison of the MSSP, KDA and aKDA and also when the two heuristics are compared directly to one another. The reason that we consider the heuristic comparison alongside the full comparison is that the MSSP will not always be a viable option, for example when we are interested in larger portfolios, and so in these cases we are interested in which heuristic performs the best.

The table also shows the percentage of times that each method led to the highest NPV and was simultaneously at least 3% higher than the second highest NPV. One of the things that is apparent from this measure is that, even if we were to specify the correlation incorrectly, there is little to be lost from using the aKDA over the KDA. This is because in over 50% of cases, there is little difference

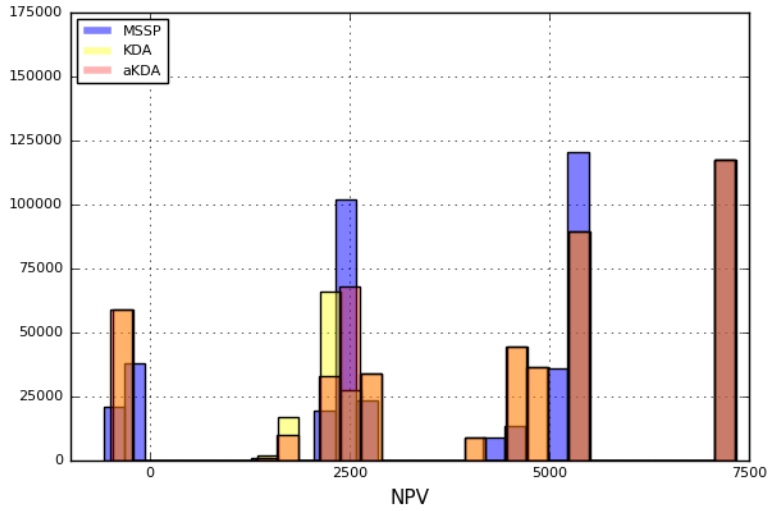

(a)

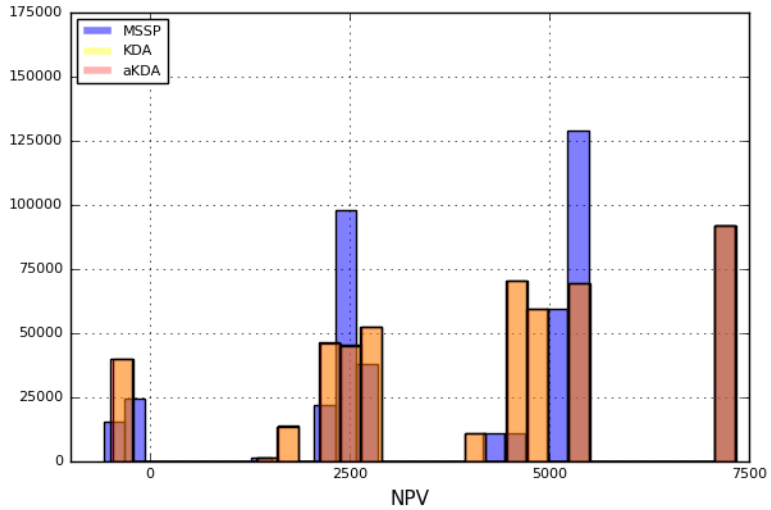

(b)

between the highest and second highest NPVs.

In Table 4.2(a), we see the results of the simulation study with true correlation matrix (a). This was the case where the true correlation matrix matches the correlation matrix used to generate the aKDA decision tree. That is, our prior assumptions on the correlation matrix were correct. The MSSP leads to the highest NPV in the majority of cases, which is unsurprising given that the decision tree under this approach, shown in Figure 4.1(a), is more dense than the other decision trees and the fact that this is the original method, not a heuristic.

The heuristic comparison, however, shows the advantage of the aKDA as it leads

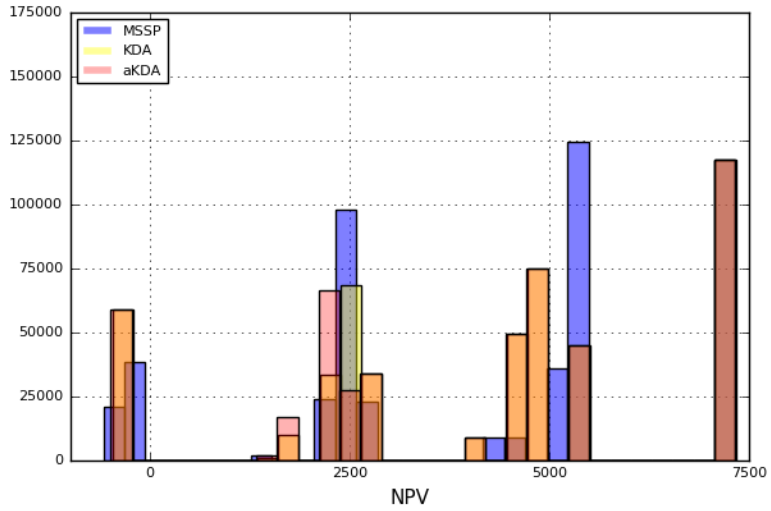

(c)

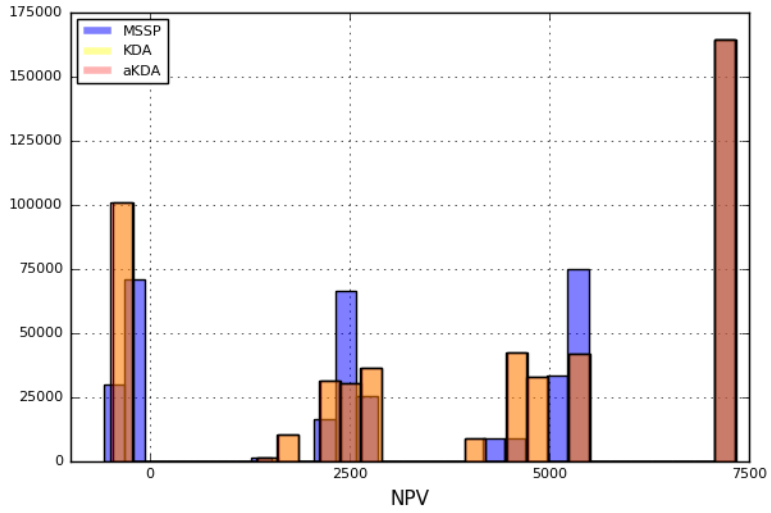

(d)

Figure 4.2: Histograms of the NPV in the simulation study under each approach for the correlation matrices (a) - (d).

to the highest NPV in around 50% more cases than the KDA. We also see the aKDA and the KDA lead to the same NPV in around half of the replications in the heuristic comparison.

It should be noted that the decision trees of the KDA and the aKDA share identical decision paths when studies $(3, 1)$ and $(3, 2)$ are successful, which will lead to any replications with these trial outcomes having the same NPV.

The distribution of the NPVs under the three approaches for this correlation matrix can be seen in Figure 4.2(a). Although there is a lot of overlap between

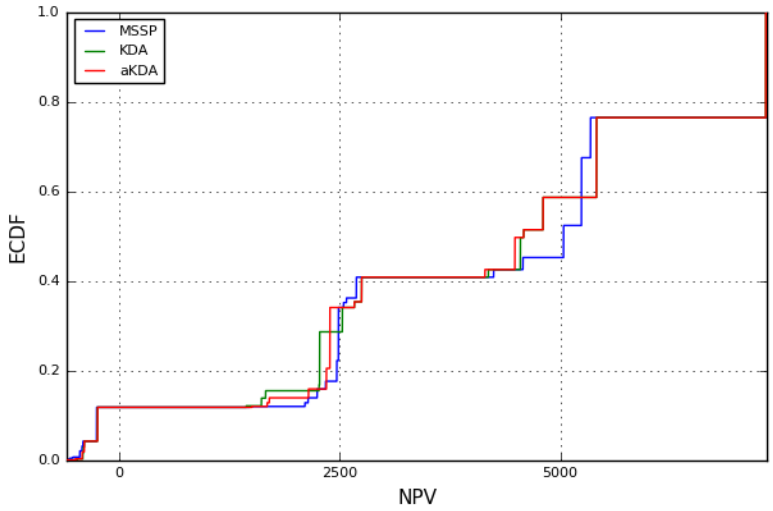

(a)

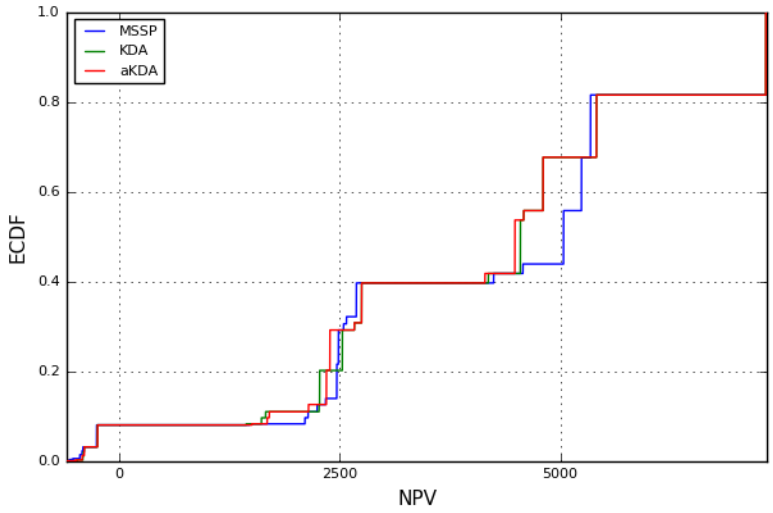

(b)

the KDA and the aKDA, we see that the main contribution to the success of the aKDA is at the NPV of around 2500. This can also be seen in Figure 4.3(a).

In two of the four correlation matrices that were considered, the prior assumption on the correlation between θ_2 and θ_3 was incorrect. In matrix (b), no correlation existed between any $\theta_i \, \forall i \in I$ and in matrix (c) the correlation existed between θ_1 and θ_3. In these two scenarios, the MSSP performs the best, as in scenario (a), but the aKDA performs worse than the KDA. The number of times that the KDA outperforms the aKDA is similar to what we saw for the reverse in scenario (a).

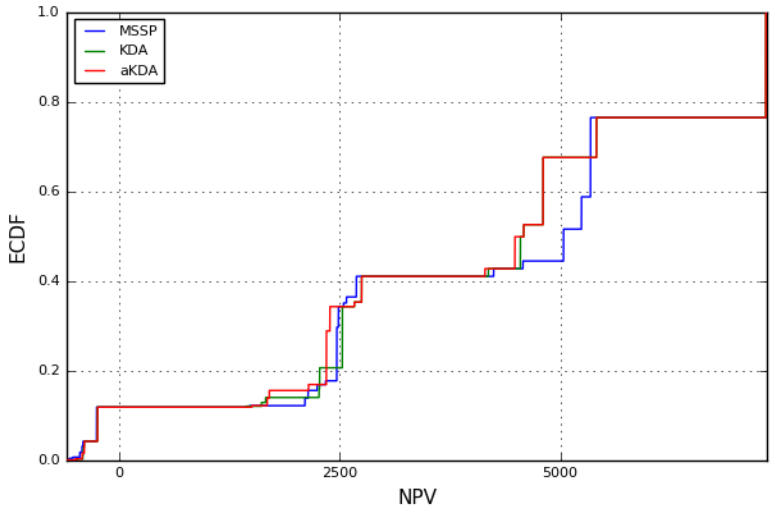

(c)

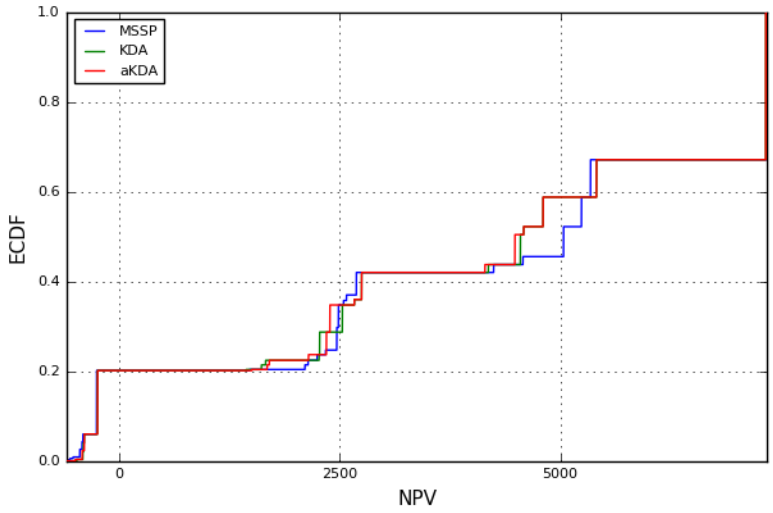

(d)

Figure 4.3: Plots of the empirical cumulative distribution function of the NPV in the simulation study under each approach for the correlation matrices (a) - (d).

However, when we explore the histograms in Figure 4.2 we see that, although the results in Tables 4.2(b) and (c) are similar, the distributions of the NPV are quite different. More specifically, when we are only wrong about the correlation between θ_2 and θ_3, we see less of a difference in the NPVs under the KDA and aKDA in Figure 4.2(b) than we see in Figure 4.2(c), where we were also wrong about the correlation between θ_1 and θ_3.

It is also clear from Figure 4.3 that the MSSP makes most of its gains on the other approaches when the NPV is around 5000, which corresponds to the situation

when almost all studies are successful. This is intuitive, given the decision trees shown in Figure 4.1, as the MSSP is able to run more studies sooner in the planning horizon than the other two approaches due to their overscheduling constraints.

The final correlation matrix that was considered, matrix (d), includes correlation between all of the treatment effects. In this case, our prior assumptions correctly captured one of the relationships but not the other two. Here, we see the aKDA perform the best out of the two heuristics in Table 4.2(d) by a similar margin as was seen in correlation scenario (a). However, the histogram in Figure 4.2(d) shows that the difference in the NPV gained under the aKDA is only marginal compared to correlation scenario (a).

The results of this simulation study show that there is much to be gained from using the aKDA when the prior assumptions on the correlation structure are correct. On the contrary, when the prior assumptions are incorrect, the loss from using the aKDA over the KDA depends on the amount by which we were incorrect. When the relationship specified in the prior assumptions did not exist, there was only a small loss in the NPV. However, when that relationship did not exist and another relationship did, we lose more by using the aKDA. Finally, when we were correct about one of the relationships but failed to capture another, the aKDA performed the best but the differences in the NPV were smaller than when we captured all relationships correctly.

4.5.4 Extended simulation study

In this section, we extend the simulation study described in Section 4.5.2 to explore the results of the KDA and aKDA in more detail. The MSSP was not considered in the extended simulation study due to the computational burden of this approach.

Christian and Cremaschi [71] discuss the sensitivity of the KDA to changes in the input parameters. They note that changes in the revenue, study cost and penalty parameters have little effect on the decision trees of the KDA. We performed a similar analysis and considered the effect of individual changes of +/- 10% in $\text{rev}_i^{\max}$, c_{ij}, γ_i^L and γ_i^D for the example with parameters given in Table 4.1. The only individual parameter that resulted in a different decision tree for the KDA and aKDA from those in Figure 4.1 was $\text{rev}_i^{\max}$. The resulting decision trees that were different to those presented in Figure 4.1 are presented in Figure 4.4.

We then performed the same analysis as presented in Section 4.5.2 for the resulting KDA and aKDA decision trees. The results can be found in Table 4.3 and the histograms and empirical cumulative distribution functions of the results can be found in Appendix C.2.

When we increase $\text{rev}_1^{\max}$ by 10%, the decision trees for the KDA and the aKDA are identical, therefore a comparison is not required in this situation and the results will not be provided. For the other decision trees, we see a similar pattern between the KDA and aKDA trees as in the original example. The aKDA considers the relationship between drugs 2 and 3 and this is reflected in the trees when successes or failures are observed.

Comparison 1 considers the case when we increase $\text{rev}_2^{\max}$ by 10%, Comparison 2 considers the case when we decrease $\text{rev}_2^{\max}$ by 10% and Comparison 3 considers the case when we decrease $\text{rev}_3^{\max}$ by 10%.

In Table 4.3, we see results for Comparison 1 that are similar to those seen in Section 4.5.3. The aKDA outperforms the KDA when the assumption on the correlation between drugs 2 and 3 is correct, scenarios (a) and (d), and less well when there is no correlation between drugs 2 and 3, scenarios (b) and (c).

In Comparisons 2 and 3, the aKDA outperforms the KDA in all four scenarios with the largest margin being for scenarios (a) and (d), as we would expect.

(a)

	Comparison 1	Comparison 2	Comparison 3
KDA	17.33 (8.72)	14.69 (10.05)	15.76 (10.31)
aKDA	25.73 (17.10)	42.11 (18.48)	41.15 (41.15)
KDA + aKDA	56.93	43.20	43.10

(b)

	Comparison 1	Comparison 2	Comparison 3
KDA	25.58 (12.82)	22.92 (15.33)	24.61 (15.62)
aKDA	17.61 (11.85)	34.02 (15.42)	32.46 (32.46)
KDA + aKDA	56.81	43.06	42.93

(c)

	Comparison 1	Comparison 2	Comparison 3
KDA	25.86 (9.61)	23.15 (18.56)	24.50 (10.92)
aKDA	17.20 (8.46)	33.73 (10.04)	32.50 (32.50)
KDA + aKDA	56.94	43.12	42.99

(d)

	Comparison 1	Comparison 2	Comparison 3
KDA	17.56 (9.31)	14.38 (9.39)	15.61 (9.61)
aKDA	25.42 (9.57)	42.50 (9.42)	41.20 (41.20)
KDA + aKDA	27.03	43.12	43.19

Table 4.3: The percentage of times each method produced the highest NPV in the extended simulation study with 500000 replications for the correlation matrices (a) - (d). The row "KDA + aKDA" represents the situation where the KDA and the aKDA produced the highest NPV simultaneously. The numbers in brackets give the percentage of times that each method produced the highest NPV, which is also at least 3% higher than the second highest NPV.

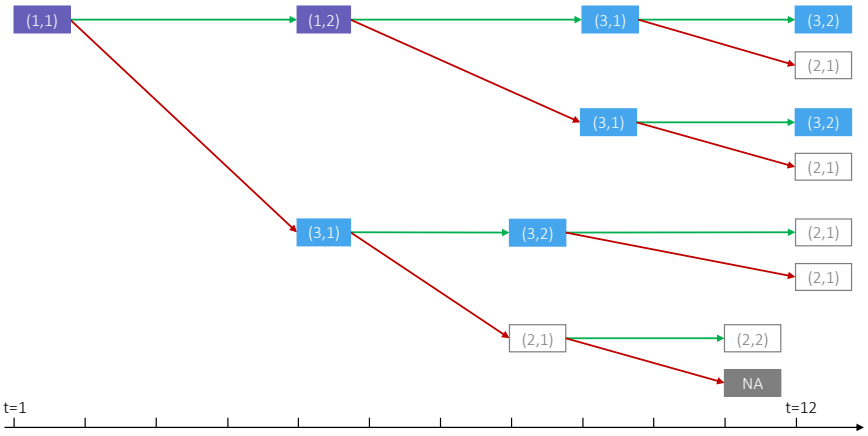
(1,1)
(1,2)
(3,1)
(3,2)
(2,1)
(3,1)
(3,2)
(2,1)
(3,1)
(3,2)
(2,1)
(2,1)
(2,1)
(2,2)
NA
t=1
t=12

(a)

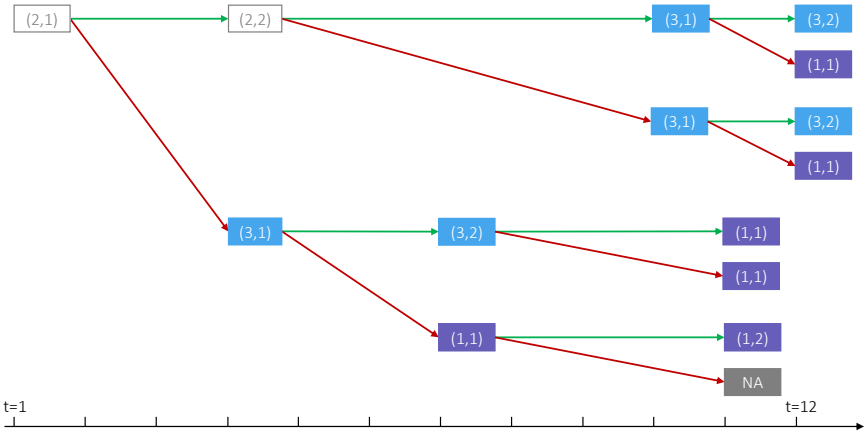
(2,1)
(2,2)
(3,1)
(3,2)
(1,1)
(3,1)
(3,2)
(1,1)
(3,1)
(3,2)
(1,1)
(1,1)
(1,1)
(1,2)
NA
t=1
t=12

(b)

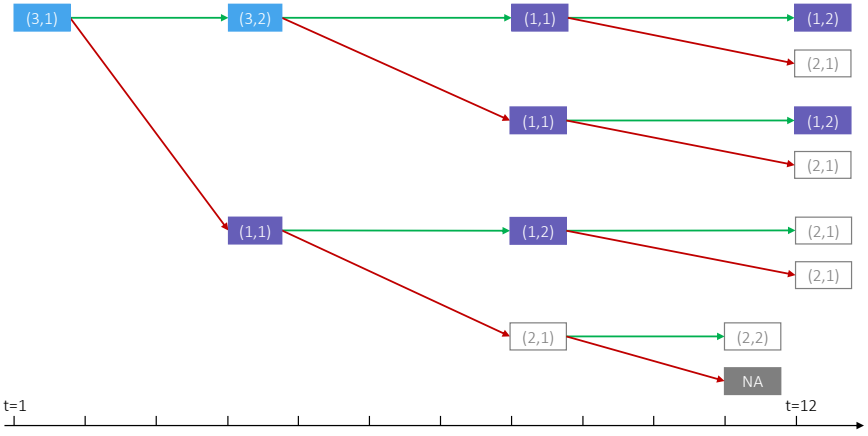
(3,1)
(3,2)
(1,1)
(1,2)
(2,1)
(1,1)
(1,2)
(2,1)
(1,1)
(1,2)
(2,1)
(2,1)
(2,1)
(2,2)
NA
t=1
t=12

(c)

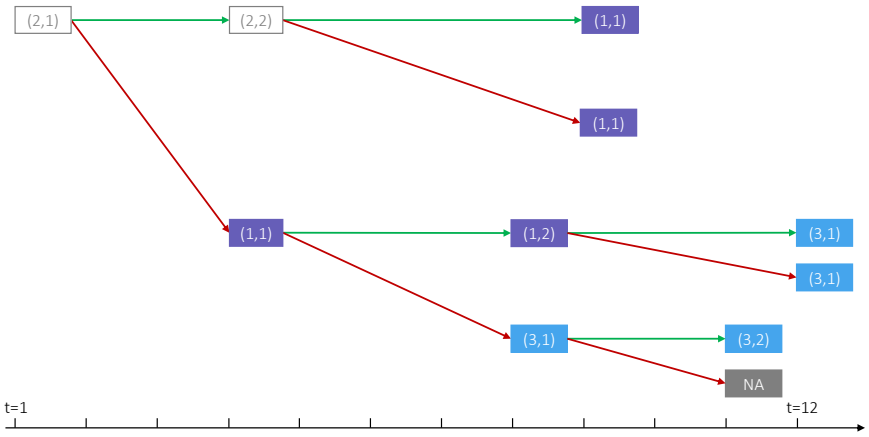

(d)

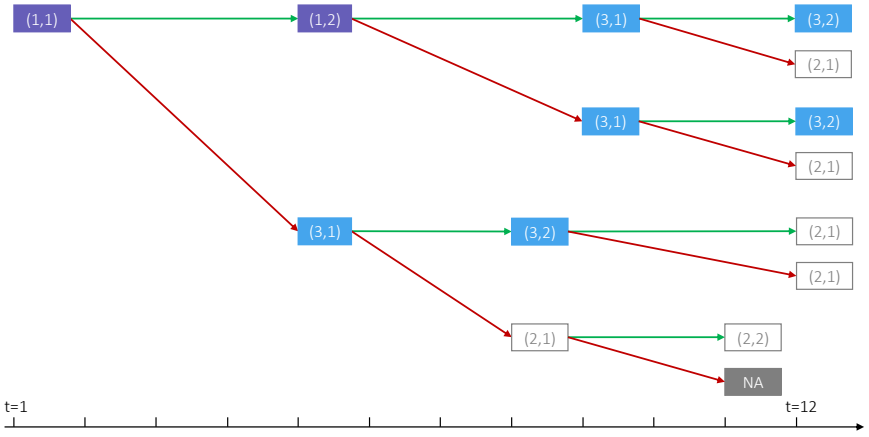

(e)

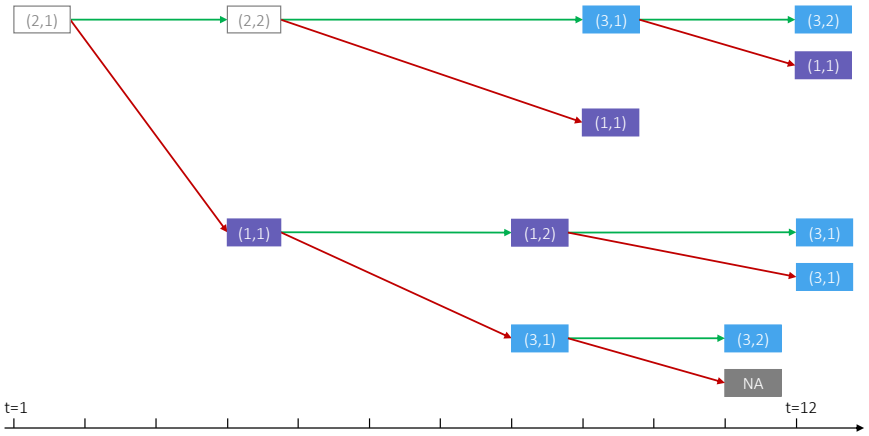

(f)

Figure 4.4: Decision trees found using (a)-(d) the KDA and (e)-(f) the aKDA with (a), (e) $\text{rev}_1^{\max}$ increased by 10%, (b), (f) $\text{rev}_2^{\max}$ increased by 10%, (c) $\text{rev}_2^{\max}$ decreased by 10% and (d), (f) $\text{rev}_3^{\max}$ decreased by 10%.

4.6 Extensions

The presented method is flexible and is able to be modified in order to suit the needs of the decision-maker with minimal effort. In this section, we discuss some of the extensions and modifications that we have considered along with how they can be implemented.

4.6.1 Multiple study outcomes

The original formulations of the KDA and aKDA only consider two study outcomes: success and failure. This means that we are required to find a single value of the study score statistic, Z, that represents the outcome and use this to update our beliefs about the treatment effects, $\boldsymbol{\theta}$. This may be considered as an oversimplification as it does not capture the way that different levels of success, or failure, can influence decisions in the real world. These effects may occur across related programmes and within programmes. For example, if a study showed very promising results, a company might decide to prioritise this programme over other programmes, regardless of original revenue and cost estimates. Similarly, if a study was less successful, a company may choose to abandon development and invest elsewhere, even if it was a "success".

We considered the addition of a third study outcome, "super success", to the aKDA with an aim to capture the effects of different levels of success on the portfolio decisions. This is related to the way that Frewer et al. [72] discuss a "super go" option in the decision to accelerate development in a single drug development programme.

There are several different ways that this could be implemented, based on the requirements of the user, but the framework would remain the same. Previously, we summarised the study outcomes by selecting two values, Z_s and Z_f, which

were then used to update the study success probabilities. In order to find appropriate values for these we considered two truncated normal distributions and the most recent mean, $\mu_i^{[t,k]}$, for the treatment effect, θ_i. These distributions were given in Equation (4.10).

We will instead consider three values, Z_{s^*}, Z_s and Z_f, to correspond to super success, success and failure. Therefore, we must consider three truncated distributions corresponding to these three outcomes, which are given as follows.

$$Z_{s^*} \sim \text{TN}\left(\mu_i^{[t,k]} V_{ij}, V_{ij}; u_{ij}, \infty\right) \qquad Z_s \sim \text{TN}\left(\mu_i^{[t,k]} V_{ij}, V_{ij}; d_{ij}, u_{ij}\right)$$

$$Z_f \sim \text{TN}\left(\mu_i^{[t,k]} V_{ij}, V_{ij}; d_{ij}, \infty\right)$$

The value of u_{ij} corresponds the value of Z_{ij} that would be considered a super success for study (i, j). This can be chosen differently for each study or it could be set to a common multiplier of the critical value, d_{ij}, across all studies. For example, $u_{ij} = 1.5 \times d_{ij}$ for all (i, j). Thus, we have three zones that categorise the study outcomes. Any study with a score statistic less than d_{ij} will be deemed a failure, between d_{ij} and u_{ij} a success, and greater than u_{ij} a super success.

We could also include more study outcomes if it was deemed necessary by the user, or consider different numbers of outcomes for different studies. For example, study success and failure may be considered enough information for a Phase I study, but we might want to consider more study outcomes for the later phases.

We implemented this for our example portfolio detailed in Section 4.5 with super success being defined as a score statistic exceeding $1.25 \times d_{ij}$ and the resulting decision tree is given in Figure 4.5. The generation of this decision tree took eight seconds, compared to seven seconds for the decision tree with two outcomes shown in Figure 4.1(c). This increase, although small, is because there are more

knapsack subproblems to be solved in this approach than the original aKDA. This approach will also be solvable in reasonable time, unlike the MSSP, for much larger portfolios.

In the decision tree in Figure 4.5, the dashed green line represents the outcome of super success. As we would expect, many of the decisions made in the decision tree are the same as was seen in Figure 4.1(c). However, one of the surprising things that we see is that when study $(3, 2)$ is a success, not a super success, the decision tree selects $(1, 1)$ rather than $(2, 1)$. This may seem surprising at first because both the KDA and aKDA decision trees in Figure 4.1 selected $(2, 1)$ on the successful completion of $(3, 2)$. However, due to our definition of super success being the score statistic exceeding $1.25 \times d_{ij}$, the values corresponding to a normal success are actually quite low, with their corresponding estimates of θ_i being less than the minimally important difference, δ_{ij}. Thus, when $(3, 2)$ is only successful, this reduces the mean of θ_2 and so the success probabilities associated with the correlated programme are also reduced.

It should also be noted that the tree does still select to complete all programmes that have been initiated where possible over those that have not, as we would expect in most cases. However, depending on how close the upper bound, u_{ij}, of the success category is to the critical value, d_{ij}, we might not want to continue development after a "successful" study. This would be reflected in the decision tree, which collates all available information and makes a recommendation based on this.

In this example, the definition of the success category corresponded to a marginal success for the sake of illustrative purposes. In real world use, the definition of the categories should be dependent on what the user would find most interesting.

4.6.2 Constraints on sets of studies

Another modification that may be required by a user is related to the nature of a portfolio of combinations. We have discussed the benefit that can be gained from learning across related combination studies and have shown how this can be included in the decision-making process. However, we have not considered some of the potential complications that can arise in a portfolio of combinations compared to a portfolio of single agents.

One of the aspects to consider is the risk that is incurred if many positively correlated programmes are selected to run simultaneously over other uncorrelated programmes. For example, if a Phase II study results in a success that is correlated with many other programmes, the aKDA may select to run the correlated programmes over other uncorrelated programmes. This is beneficial as it is using relevant information in the decision-making procedure and it does still consider the uncertainty in the trial outcomes as the success probabilities are used to calculate the study values in this procedure. However, if the Phase II study showed more promise than is true for the programme, then selecting as many correlated programmes as possible would lead to a potentially high level of loss.

We propose that, to overcome this potential issue in portfolios with large sets of correlated programmes, a constraint set should be imposed that controls the number of correlated programmes that are selected. In order to achieve this, we would introduce another input parameter that details the groups of correlated programmes. We will use the notation g_i to represent the group of programme i. We will use G to denote the set of groups thus $|G|$ gives the total number of groups and $g_i \in G$. We will also introduce a variable that tracks the number of programmes initiated in each group throughout the decision tree. The variable $n_g^{[t,k]}$ gives the number of programmes in group $g \in G$ that have been initiated

prior to subproblem $[t, k]$.

Then, in addition to the constraints given in Equations (4.3) - (4.6), we will include the constraint

$$n_g^{[t,k]} + \sum_{i:g_i=g} \sum_j X_{ij} \leq n_g^{\max} \ \forall g$$

to the aKDA where $n_g^{\max}$ gives the maximum number of programmes of set g that can be initiated within the planning horizon and must be specified by the user in advance.

This could be used in the illustrative example given in Appendix C, which has six different groups of correlated programmes, to ensure that we do not limit development to a small number of groups.

This approach might be considered as too conservative in some settings, for example if all other programmes were abandoned then it would not make sense to limit development on the remaining programmes. However, in reality this situation is unlikely and the addition of this constraint where appropriate would allow a user to modify the aKDA such that it reflects the real world choices that they would be willing to make whilst also considering risk. The aKDA could also be modified further such that this constraint is only included when there are "enough" alternative options.

Another aspect that this constraint set would capture is the benefit of learning about unrelated combinations. If the two most attractive programmes are not correlated and the second one is not able to be selected near the beginning of the planning horizon, for example due to resource constraints, the selection of other programmes correlated to the first might be prioritised over the selection of the second programme and its group of correlated programmes. Therefore, limiting the number of correlated programmes to run in the planning horizon

will also encourage the development of and learning about other uncorrelated programmes.

A final set of constraints on sets of studies that may be included relates to the logistics of combination drug development. In this chapter, we have discussed programmes of a sequential nature for simplicity and to parallel the work presented in [49]. However, this is not always a realistic assumption as programmes may have studies that can be completed simultaneously and programmes may also be reliant upon the successful completion of other programmes or studies. For example, in order to develop a combination therapy, the single agents must also be developed separately.

This can also be included in the aKDA by modifying the definition of what it means to be eligible and ensuring that this information is sent across the subproblems when they are generated.

4.6.3 Decision rules on the PoS

Another aspect of real world decision-making that is not currently captured by the aKDA is that if the PoS of a study dropped below a certain threshold, the company is unlikely to decide to run that study unless there were no other available options. This, as with most of the other modifications discussed, can be implemented using a simple constraint set of the form

$$\phi_{ij} X_{ij} \geq \phi^*_{ij} \ \forall i, j$$

where ϕ^*_{ij} is the threshold that the PoS of study (i, j) must exceed in order to make the decision to run the study. The aKDA could also be modified so that this constraint is disregarded when there are no other options and it may also

lead an empty solution set in some cases if ϕ_{ij}^{*} is too high compared to the actual success probabilities.

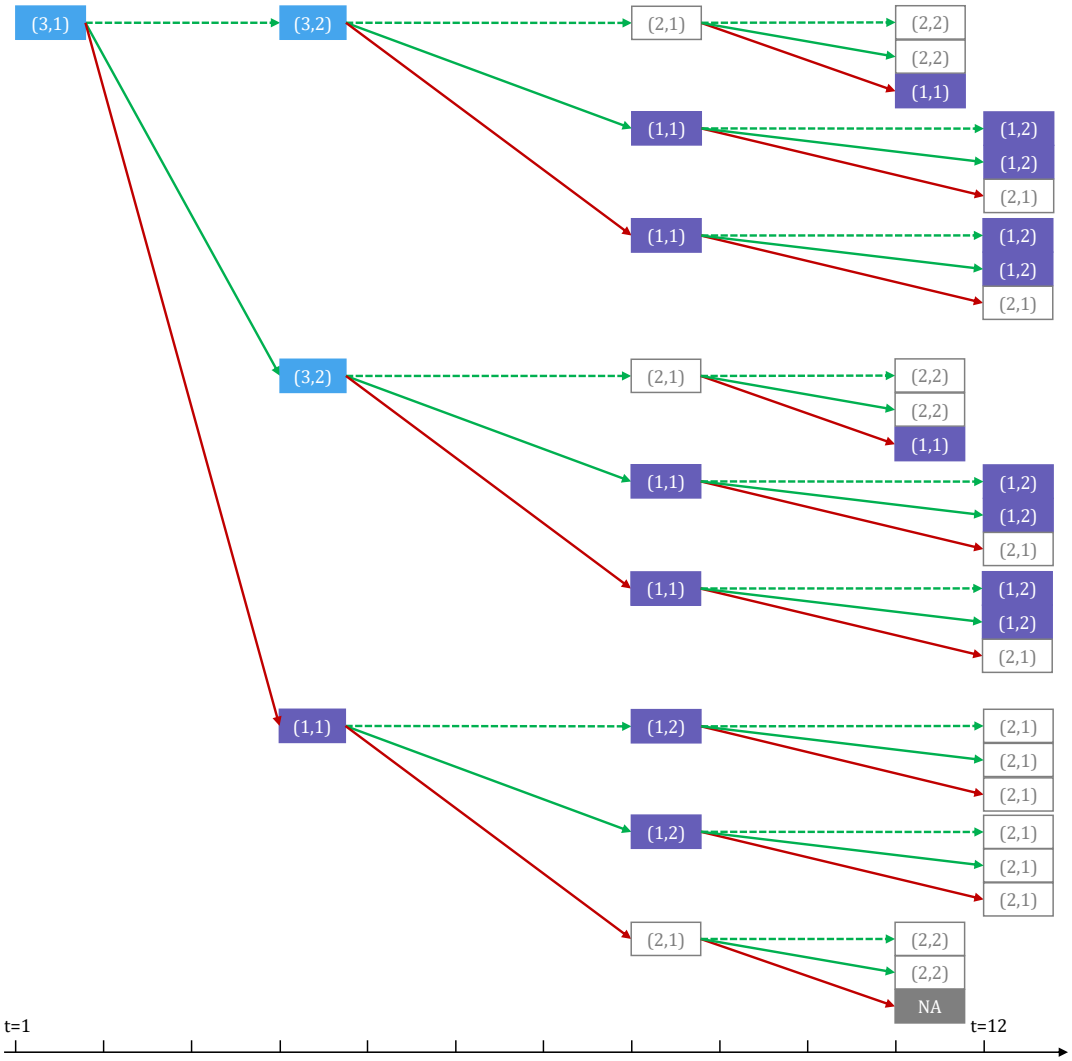

Figure 4.5: Decision tree found using the aKDA with three outcomes considered per study and parameters given in Table 4.1.

4.7 Discussion

In this chapter, we presented a portfolio management method that updates the study success probabilities throughout the procedure whenever a relevant outcome is observed. This allows us to consider the relationships between different combination drug development programmes and the benefit gained from sharing information across related programmes.

In Section 4.5, we used a simple three drug example to show the differences in the decision trees under the full MSSP [39], the KDA [71] and the aKDA. We also presented the results of a simulation study that considered different correlation structures, and highlighted the benefit of the aKDA over the KDA when our assumptions on the correlation structure are correct. We also showed that there is little to be lost by using the aKDA over the KDA when our prior assumptions are incorrect and the aKDA may even perform better than the other approaches in some of these cases, unless there are fundamental discrepancies between our prior correlation structure and the true correlation structure.

We did also observe in the simulation study that the MSSP outperformed the aKDA in the example presented and this was typically due to the denser decision tree of the MSSP compared to the heuristics. There is potential to improve these heuristics, mainly through the modification of the overscheduling constraint. However, it should be noted that the MSSP becomes unusable for larger portfolios and it is not able to capture the relationships between programmes.

We also presented several extensions that aim to tailor the portfolio management technique to the needs of the user and what is deemed realistic within their decision-making framework. These extensions included the consideration of more than two study outcomes, which was the only option in the existing methods presented by Colvin and Maravelias [39] and Christian and Cremaschi

[49]. This extension allowed for the inclusion of a "super success" that links to the idea of a "super go" in the work presented in [72].

One of the things that should be considered when generating decision trees is the influence of the input parameters, especially if there is uncertainty surrounding the most appropriate point estimates for these. This is true for all three methods, but the aKDA has an additional set of parameters to consider that relate to the distributions of the treatment effects, $\boldsymbol{\theta}$.

In order for the aKDA to perform as intended, the prior means, $\mu_i \, \forall i \in I$, should be specified using similar approaches due to the way that the success probability updating works. For example, they could all be specified using estimates based on historical data or they could all be set to be equal to their associated minimally important difference. The method for updating the success probabilities that was introduced in Chapter 3 considers the distance between the prior mean, μ_i, and the estimate of this based on the observed data, Z_i/V_i. The amount by which the means, $\mu_{i'}$, for correlated treatment effects, $\theta_{i'}$, shift is related to this difference alongside the correlation with the observed treatment effect, $\rho_{ii'}$. Therefore, if one prior mean was selected arbitrarily and this was quite different from the mean of a correlated treatment effect that was selected to be informative, this could cause the updates to be misleading. This is because the prior means being correlated is an underlying assumption of the updating procedure.

For example, let us consider θ_2 and θ_3 from the example outlined in Section 4.5, which have a prior correlation of 0.6. If we set $\mu_3 = \delta_{31} = 0.25$, which represents the belief that there is a benefit whilst also being conservative in the estimate for this, and we arbitrarily selected $\mu_2 = 0$, then the update would go as follows. At the first time point, we still select $(3, 1)$, as in the decision tree shown in Figure 4.1(c). Then we update the distribution of $\boldsymbol{\theta}$ based on the outcome of this study. We only consider two outcomes, failure and success, and generate two subproblems based on these outcomes. Following the method presented

in Section 4.4, we use $E(Z_s) = 28.21$ and $E(Z_f) = 10.81$ as the score statistics representing study success and failure, respectively, in the updating procedure. The updated means based on these values are given by

$$\boldsymbol{\mu}^{[4,1]} = \begin{pmatrix} 0.25 \\ -0.08 \\ 0.12 \end{pmatrix} \quad \text{and} \quad \boldsymbol{\mu}^{[4,2]} = \begin{pmatrix} 0.25 \\ 0.02 \\ 0.28 \end{pmatrix}$$

where $[4,1]$ is the subproblem in which $(3,1)$ was a failure and $[4,2]$ is the subproblem in which $(3,1)$ was a success. Notice that the mean of θ_2 decreases from its already pessimistic prior mean when $(3,1)$ is unsuccessful and when $(3,1)$ is successful, the shift is still only small. Furthermore, since all of the studies in this example have $\theta_{ij}^{\text{ind}} = 1$, the PoS is calculated based on the distribution of θ_i. In this example, the prior success probabilities for the second treatment are given by $\phi_{21} = 0.31$ and $\phi_{22} = 0.32$. After the update, the success probabilities in subproblem $[4,1]$ decrease to $\phi_{21} = 0.19$ and $\phi_{22} = 0.2$, as we would expect given that a positively correlated study was unsuccessful. In scenario $[4,2]$, study $(3,1)$ was successful but the success probabilities for the second programme decrease to $\phi_{21} = 0.3$ and $\phi_{22} = 0.31$. This is because the shift in mean is small, which results in an estimate that is still pessimistic, and there is a reduction in variance, which is considered in the calculation of the PoS.

This example also highlights the fact that the user should consider the success probabilities that the prior means and variances lead to. If these success probabilities are not thought to be realistic then this will encourage a user to think more about the way that they have selected these values and to find more realistic values for them. This will then help to ensure coherence across different programmes.

If the problem of incoherent prior means exists for multiple correlated treatment effects, and multiple updates are performed, then this problem could be exacer-

bated. However, this problem can be easily avoided by ensuring that appropriate prior means are specified and correlations are only included where appropriate, and with appropriate levels. If the user requires additional measures to be put in place to avoid any of the potential pitfalls, however, there are different modifications that can be included to do this.

One approach involves constraining the number of indirect updates that can be performed on the distribution of a treatment effect. This could be tracked simply using a counter for each θ_i that is sent through to each subproblem and in each subproblem we would impose a constraint on the sum of this counter and the decision variables $X_{ij} \forall j \in J_i$.

Alternatively, a more lenient modification would involve stopping the updating of the distribution of θ_i after the final study in programme i, $(i, |J_i|)$, is completed. This could be achieved by including an indicator that is equal to one when a programme is still ongoing and multiplying the correlation, $\rho_{ii'}$, by this indicator each time the distribution of $\boldsymbol{\theta}$ is built in the aKDA. Then, when the programme is completed and the indicator drops to zero, the treatment effect will no longer be correlated with any other treatment effects hence will not be used in any updates. This would make the output of the method more realistic than if we performed unnecessary updates, but this will have no effect on the decision tree of the aKDA as when the programme is completed, we are not going to make any future decisions about it, it would simply affect the vector of means.

Although the motivation for this work lies in combination therapies, the method presented in Section 4.4 could also be applied for portfolios that contain programmes that have treatment effects that are related in some way other than similar combinations. This is because our method for updating the study success probabilities presented in Chapter 3 and used in the aKDA captures the general relationship between the treatment effects of different programmes and does not require specific information relating to combinations.

Chapter 5

Chapter 5

Conclusions

5.1 Overview

In this thesis, we have considered the relationships between similar combination studies and have presented methodology that allows us to capture these relationships and use them to assist portfolio decision-making.

In Chapter 2, we provided a critical discussion of some of the existing methods for decision-making in a pharmaceutical portfolio and compared two methods that draw upon stochastic programming techniques. We concluded that, while both methods capture important aspects of portfolio management and the uncertainty that is inherent to this process, neither method was suitable in its original form to the problem that we were interested in. The ROV approach [36] was able to capture the uncertainty in the value of a programme, but was not able to consider scheduling or the uncertainty in the study outcomes. The PS approach [39], on the other hand, was able to capture the uncertainty in the study outcomes and find the optimal schedules, but this came at a high computational burden. Furthermore, neither of these methods are able to capture the relationships between studies or the way that observing a particular outcome might influence our beliefs regarding the success probability of another study.

We then presented a method that allows the probability of success of a study

to be updated based on the results of a related combination study in Chapter 3. This was achieved by considering both the distribution of the treatment effects of the drugs of interest and the distribution of the score statistic of the observed study. We can then find the conditional distribution of the treatment effects given the observed score statistic and use this distribution to find the updated success probabilities via the calculation of the assurance [23] or other distribution based definitions of the probability of success. We also presented an extension to the method that controls the level of borrowing based on the observed data. This method cannot only be applied in the setting of combinations, but also in other areas where we believe that the study outcomes are related and we would benefit from borrowing across these studies.

The method for updating study success probabilities was combined with a heuristic [71] for the PS approach to update the study success probabilities each time a relevant outcome is observed. This was presented in Chapter 4. This method allows the potential future observations along with the relationships between programmes to guide the decision-making process in a quantifiable way. This method can be used for large portfolios and is quick to solve, easy to interpret and achieves what other, existing pharmaceutical portfolio management techniques cannot in terms of the consideration of relationships between studies.

5.2 Limitations

The main limitations of the methodology presented are related to the specification of the prior distributions. In order to update a study success probability based on the outcome of a related study, a prior correlation between the two studies is required that controls the level of borrowing. This could be seen as a limitation, as the choice for this parameter might not always be clear and there is

no existing methodology for quantifying the true correlation between the studies. Furthermore, if this is specified incorrectly, then the results of the method could be misleading.

Another potential issue relating to the specification of the prior distributions comes from the way that the distributions are updated and the underlying assumptions of this update. The mean of the treatment effect of the unobserved treatment shifts during the update by an amount that is relative to the both the correlation and the difference between the prior mean of the observed study and the estimate from the study. This is related to the assumption that the prior means are also correlated. Thus, if one prior mean appears to be an underestimate of the truth, based on the study data, then the other prior mean will also be assumed to be an underestimate, thus will be shifted upwards during the update. This will usually not be an issue, unless the two prior means are incorrect in opposite directions i.e. one overestimates the truth and the other underestimates the truth. This would be problematic, as the unobserved mean would be shifted further away from the truth during this update.

A final limitation is that our work builds upon a heuristic for the PS approach, which means that the solutions found will not be optimal compared to the solution of the full MSSP. However, the MSSP cannot be solved in reasonable time for portfolios of the size that we require for our problem and so using a heuristic is the only realistic option. There are still potential areas for improvement with regards to the heuristic, however. For example, the overscheduling constraint was noted to be too conservative, thus some improvements could be made here. In addition, the heuristic does not consider penalties in the same way as the full MSSP, which is another aspect that could be modified.

5.3 Further work

We previously discussed the potential issues that can arise when specifying a correlation between studies. Further work in this area could consider developing methodology for the quantification of the correlation between studies based on existing data. This could be used either as a separate tool prior to specifying the input parameters or it could be added to the updating procedure.

The methodology presented in this thesis considers the overall relationship between combination studies, rather than attempting to quantify this based on the separate components. A potential area for further work could include considering the individual components of the combination and modelling the relationships between them in order to estimate the treatment effect of the combination and the correlations between treatment effects of different combinations.

An aspect that could be extended in order to make the portfolio decision-making procedure more realistic is to consider new development programmes that become available as options during the planning horizon, rather than at the beginning of the planning horizon. Given the flexibility of the aKDA, this could be added quite easily if it were required. For example, one could use a dummy study that is forced to be packed into the knapsack at the first node. This dummy study would have a duration that is equal to the time until the time until the new programmes would be available and success probability equal to the probability of it becoming available. It would have zero cost and zero value and would be assigned as a prerequisite to the development programme such that the first study becomes eligible to run on successful completion of the dummy study.

In the simulation study provided in Chapter 4, the full MSSP formulation of the PS approach outperformed the aKDA for the presented example. However, this gain is not meaningful when we consider real world portfolio decision-making problems since the MSSP can only be solved for small portfolios. This

could motivate two areas for further work. The first relates to improving the performance of the KDA and reducing the optimality gap between the solution of the aKDA and the MSSP. Improvements in this area could start with the overscheduling constraint as this constraint leads to sparser decision trees hence suboptimality. The second area of research this could lead to is reducing the complexity and size of the MSSP such that it can be solved for larger portfolios. This would allow us to include the success probability updates into the full MSSP. Some simplifications could include, for example, running programmes without gaps, given that this is often an aspect of the optimal solution. This would lead to some suboptimality, but could reduce the model by quite a large factor. Alternatively, we could include further decision rules that reduce the size of the solution space such as using preference orderings as inputs or other rules that reflect the most likely decisions to be made by the company. While this would still be a heuristic, it would return a solution that is not only close to optimal but also useful to the team given that it reflects their preferences.

Another extension that could be considered for the portfolio management problem is including the use of dynamic programming. Choi et al. [73] discussed the use of dynamic programming for the resource constrained scheduling problem and noted the computational load of this solution method due to the size of the state space for complex problems. They presented an algorithmic framework, dynamic programming in a heuristically confined state space, that aims to tackle this issue. Thus, dynamic programming or frameworks involving dynamic programming could be considered as an alternative solution method to the KDA for the portfolio management problem.

The method for updating the probability of success that was presented in Chapter 3 and used in the portfolio decision-making procedure in Chapter 4 did not consider the updating of the correlation. Furthermore, the robustification that was presented did not use information about the correlation to guide the

updates. Conceptually, it would make sense to use emerging information about the correlation to update these distributions, but this is harder to capture than information on the individual treatments. An area of further work would include the consideration of the correlation and how updating this might inform our decision-making.

Finally, the methods for portfolio decision-making discussed in this thesis only consider the decisions relating to if and when a study should be initiated. These methods do not consider decisions relating to the design of the studies or the effect that the design might have on our decision to run the study. Further work could include developing new designs for combination studies that account for the relationships between combinations and also new methods for decision-making that consider the optimal study designs over the whole portfolio. Patel et al. [37] previously considered the problem of the optimal study designs in a portfolio of Phase III studies and further work could extend this to consider different phases and the relationships between studies when finding the optimal designs.

APPENDIX

Appendix A

Appendix - Chapter 2

A.1 ROV formulation [36]

Additional nomenclature	
M_{ijk_j}	the value of drug i at the beginning of study (i, j) in value scenario k_j
$M_{ijk_j}^{\text{upper}}$	the upper bound of the value of M_{ijk_j}
z_{ijk_j}	the variable introduced to linearise the value constraints
Y_{ijk_j}	the binary continue/abandon decision variable for study (i, j) in value scenario k_j
w_{ijt}	the binary indicator for if study (i, j) begins at time t
$p_{ijk_jk_{j+1}}$	the probability of drug i moving from value scenario k_j to k_{j+1} during study (i, j)
ΔT	the discrete time interval that market movements are considered over

$$
\begin{aligned}
\text{minimise} \quad & \textstyle\sum_{i,j} \sum_{k_j=1}^{N_{ij}} M^{\text{upper}}_{i,j,k_j} \\
\text{subject to} \quad & M^{\text{upper}}_{ijk_j} \geq -c_{ij} + \frac{\sum_{k_{j+1}=1}^{N_{i,j+1}} \phi_{ij} p_{ik_jk_{j+1}} M^{\text{upper}}_{ij+1k_{j+1}}}{(1+r\Delta T)^{\tau_{ij}/\Delta T}} && \forall i, j, k_j \\
& M^{\text{upper}}_{ijk_j} \geq 0 && \forall i, j, k_j \\
\text{maximise} \quad & \textstyle\sum_i M_{i,1,k_1} \\
\text{subject to} \quad & M_{ijk_j} = -c_{ij} Y_{ijk_j} + \frac{\sum_{k_{j+1}=1}^{N_{i,j+1}} \phi_{ij} p_{ik_jk_{j+1}} z_{ik_jk_{j+1}}}{(1+r\Delta T)^{\tau_{ij}/\Delta T}} && \forall i, j, k_j \\
& z_{ik_jk_{j+1}} \leq M^{\text{upper}}_{i,j+1,k_{j+1}} Y_{ijk_j} && \forall i, j, k_j, k_{j+1} \\
& z_{ik_jk_{j+1}} \geq 0 && \forall i, j, k_j, k_{j+1} \\
& M_{i,j+1,k_{j+1}} - M^{\text{upper}}_{i,j+1,k_{j+1}} \left(1 - Y_{ijk_j}\right) \leq z_{ik_jk_{j+1}} && \forall i, j, k_j, k_{j+1} \\
& z_{ik_jk_{j+1}} \leq M_{i,j+1,k_{j+1}} + M^{\text{upper}}_{i,j+1,k_{j+1}} \left(1 - Y_{ijk_j}\right) && \forall i, j, k_j, k_{j+1} \\
& Y_{ijk_j} \leq Y_{i1k_1} && \forall i, j > 1, k_j \\
& Y_{i,j+1,k_{j+1}} \leq \textstyle\sum_{k_j} Y_{ijk_j} && \forall i, j, k_j, k_j \leq k_{j+1} \leq k_j + \frac{\tau_{ij}}{\Delta T} \\
& Y_{i,j,k_j-1} \leq Y_{ijk_j} && \forall i, j, k_j \\
& \textstyle\sum_{i,j} \sum_{k_j}^{N_{ij}} p_{ik_{j-1}k_j} c_{ij} Y_{ijk_j} w_{ijt} \leq B_t && \forall t \\
& M_{ijk_j} \geq 0 && \forall i, j, k_j \\
& Y_{ijk_j} \in \{0, 1\} && \forall i, j, k_j
\end{aligned}
$$

A.2 PS formulation [39]

Additional nomenclature	
X_{ijts}	the binary go/no-go decision variable for study (i, j) at time t in scenario s
Y_{ijts}	the indicator for if study (i, j) has been completed by time t in scenario s
Z_{ijts}	the indicator for if study (i, j) is ready to run at time t in scenario s
Rv_s	the revenue generated in scenario s
FRv_s	the future revenue generated in scenario s if all remaining trials are completed as soon as possible
Cst_s	the costs incurred in scenario s
$p(s)$	the probability of scenario s
$rev_i^{\max}$	the total maximum possible revenue generated by drug i
rev_{ijt}^{run}	the revenue generated on completion of programme i when study (i, j) is ongoing at time $\|T\|$ and started at time t
rev_{ij}^{open}	the revenue generated on completion of the programme for drug i when study (i, j) is ready to run at time $\|T\|$
λ_{ijr}	the resource requirement of study (i, j) of resource type r
$\lambda_r^{\max}$	the level of available resource of type r
f_{ij}	the discounting factor for open revenue
cd_t	the time discounting factor for the time value of money
n_t	the interest rate for a time period
$F^{I,J}(s)$	the set of studies that cannot be conducted in scenario s
$S^I(s)$	the set of successful programmes in scenario s
Ψ	the set containing pairs of scenarios that differ only in the outcome of one study

maximise $\sum_s p(s)\{Rv_s + FRv_s - Cst_s\}$

subject to

(1) $\sum_t X_{ijts} \leq 1$ $\forall i,j,s$

(2) $\sum_i \sum_j \sum_{t-\tau_{ij}<t'\leq t} \lambda_{ijr} X_{ijt's} \leq \lambda_r^{\max}$ $\forall r,t,s$

(3) $\sum_{t'\leq t} X_{ijt's} \leq Y_{ij-1ts}$ $\forall i,j>1,t,s$

(4) $X_{ijts} = 0$ $\forall t,s,(i,j) \in F^{I,J}(s)$

(5) $X_{ijts} = 0$ $\forall i,j,t < \sum_{j'<j} \tau_{ij'},s$

(6) $X_{ijts} - X_{ijts'} \geq -Y_{i^{s,s'},j^{s,s'},t,s}$ $\forall i,j,(s,s') \in \Psi, t>1$

(7) $X_{ijts} - X_{ijts'} \leq Y_{i^{s,s'},j^{s,s'},t,s}$ $\forall i,j,(s,s') \in \Psi, t>1$

(8) $Y_{ijts} = Y_{ijt-1s} + X_{ijt-\tau_{ij}s}$ $\forall i,j,t,s$

(9) $Z_{i11s} = 1 - X_{i11s}$ $\forall i,s$

(10) $Z_{i1ts} = Z_{i1t-1s} - X_{i1ts}$ $\forall i,t>1,s$

(11) $Z_{ijts} = Z_{ijt-1s} + X_{ij-1t-\tau_{ij}s} - X_{ijts}$ $\forall i,j>1,t,s$

(12) $X_{i11s} = X_{i111}$ $\forall i,s$

(13) $Cst_s = \sum_{ijt} cd_t c_{ij} X_{ijts}$ $\forall s$

(14) $Rv_s = \sum_{i\in S^I(s)} \sum_t \{rev_i^{\max} X_{i,\text{PIII},t,s} - \gamma_i^D \times$
$(Z_{i,\text{PII},t,s} + Z_{i,\text{PIII},t,s}) - \gamma_i^L (t + \tau_{i,\text{PIII}}) X_{i,\text{PIII},t,s}\}$ $\forall s$

(15) $FRv_s = \sum_{i\in S^I(s)} \sum_j rev_{ij}^{\text{open}} f_{ij} Z_{ij|T|s} +$
$\sum_{i\in S^I(s)} \sum_{j\in\{\text{PI,PII}\}} \sum_{t>|T|-\tau_{ij}} rev_{ijt}^{\text{run}} f_{ij+1} X_{ijts}$ $\forall s$

(16) $X_{ijts} \in \{0,1\}$ $\forall i,j,t,s$

(17) $Y_{ijts}, Z_{ijts} \in [0,1]$ $\forall i,j,t,s$

where

$f_{ij} = 0.9\left(\frac{rev_i^{\max} - \gamma_i^L |T| - \sum_{j'\geq j} c_{ij'}}{rev_i^{\max} - \gamma_i^L |T|}\right)$

$rev_{ij}^{\text{open}} = rev_i^{\max} - \gamma_i^L \left(|T| + \sum_{j'\geq j} \tau_{ij'}\right)$

$rev_{ijt}^{\text{run}} = rev_i^{\max} - \gamma_i^L \left(t + \sum_{j'\geq j} \tau_{ij'}\right)$

$cd_t = 1 - n_t(t-1)$

We will now provide a brief explanation of the constraints included in the above.

Constraint (1) ensures that each study (i, j) is only run once in each trial outcome scenario, $s \in S$. Constraint (2) is the resource constraint for different resource types, $r \in R$. This constraint ensures that at each time point, $t \in T$, the combined required resources for any ongoing studies or newly initiated studies do not exceed the maximum available resources per time point, $\lambda_r^{\max}$.

Constraint (3) ensures that (i, j) does not start before $(i, j-1)$ is complete. Constraint (4) ensures that studies that had a prerequisite study fail at an earlier time point are not initiated. These studies are included in the set $F^{I,J}(s)$, which may be defined in advance given that the scenarios, $s \in S$, contain information on study outcomes. Constraint (5) ensures that a study is not started earlier in the time frame than by the earliest time that all of its prerequisite studies can be completed. Constraints (4) and (5) help to reduce the number of variables in the model.

Constraints (6) and (7)are the NACs, which were discussed in Chapter 2. The NACs ensure that information regarding trial outcomes is not used in the model until after it has been revealed. For example, if in scenario s, (i, j) is unsuccessful, we cannot use this information until we have run (i, j) and observed this failure. We cannot anticipate future outcomes.

Constraint (8) is used to determine the variable Y_{ijts}, which is used to track when studies are completed and constraints (9) - (11) are used to calculate the variables, Z_{ijts}, that track when a study is eligible to be run. Nonanticipativity at the initial time point is enforced using constraint (12).

The components of the objective function are calculated for each scenario, s, using constraints (13) - (15).

Finally, the bounds of the variables themselves are specified in constraints (16) and (17). Note that Y_{ijts} and Z_{ijts} are able to be defined as continuous variables

due to their calculation being solely dependent on the binary variables, X_{ijts}, ensuring that they will also be binary. For further details, see [39].

Appendix B

Appendix - Chapter 3

B.1 Kalman filter

The Kalman filter is used to consider a dynamic system that we are able to take measurements from and we will consider this system at discrete time intervals. There is often noise within measurements and inputs but we can use filtering to gain information about the state of the system given previous measurements. The method that we will describe here is described in further detail by Anderson and Moore [66].

Let $\mathbf{x}_k$ be the state of the system at time k and let $\mathbf{y}_k$ be a measurement of the system taken at time k that includes noise. We can write

$$\mathbf{x}_{k+1} = \mathbf{F}_k \mathbf{x}_k + \mathbf{G}_k \mathbf{w}_k$$

where $\mathbf{w}_k \sim \text{MVN}(\mathbf{0}, \mathbf{Q}_k)$ describes the input noise. We can write the measurement at time point k as

$$\mathbf{y}_k = \mathbf{H}_k^T \mathbf{x}_k + \mathbf{v}_k$$

where $\mathbf{v}_k \sim \text{MVN}(\mathbf{0}, \mathbf{r}_k)$ describes the output noise.

We are interested in estimating $\mathbf{x}_k$ given the observations $\mathbf{y}_0, ..., \mathbf{y}_k$. In order to simplify the problem and the equations required to perform the filtering,

Anderson [66] modifies the problem to first consider the prediction of $\mathbf{x}_k$ given $\mathbf{y}_0, ..., \mathbf{y}_{k-1}$ before considering the prediction of $\mathbf{x}_k$ given $\mathbf{y}_0, ..., \mathbf{y}_k$.

We will use the subscript $k|j$ to denote time point k based upon everything up until time point j. Then the equations used to make estimates of the state of the system at time k under the Kalman filter are given in [66] as follows.

$$
\begin{aligned}
\hat{\boldsymbol{\mu}}_{0|-1} &= \boldsymbol{\mu}_0 \\
\hat{\boldsymbol{\Sigma}}_{0|-1} &= \boldsymbol{\Sigma}_0 \\
\hat{\boldsymbol{\mu}}_{k+1|k} &= \left(\mathbf{F}_k - \mathbf{K}_k \mathbf{H}_k^T\right) \hat{\boldsymbol{\mu}}_{k|k-1} + \mathbf{K}_k \mathbf{y}_k \\
\mathbf{K}_k &= \mathbf{F}_k \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k \left(\mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k + \mathbf{R}_k\right)^{-1} \\
\hat{\boldsymbol{\Sigma}}_{k+1|k} &= \mathbf{F}_k \left[\hat{\boldsymbol{\Sigma}}_{k|k-1} - \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k \left(\mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k + \mathbf{R}_k\right)^{-1} \mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1}\right] \mathbf{F}_k^T + \mathbf{G}_k \mathbf{Q}_k \mathbf{G}_k^T \\
\hat{\boldsymbol{\mu}}_{k|k} &= \hat{\boldsymbol{\mu}}_{k|k-1} + \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k \left(\mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k + \mathbf{R}_k\right)^{-1} \left(\mathbf{y}_k - \mathbf{H}_k^T \hat{\boldsymbol{\mu}}_{k|k-1}\right) \\
\hat{\boldsymbol{\Sigma}}_{k|k} &= \hat{\boldsymbol{\Sigma}}_{k|k-1} - \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k \left(\mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1} \mathbf{H}_k + \mathbf{R}_k\right)^{-1} \mathbf{H}_k^T \hat{\boldsymbol{\Sigma}}_{k|k-1}
\end{aligned}
$$

Let us now apply this to our problem. In our problem, the "state" of the system is the vector of true treatment effects, $\boldsymbol{\theta}$. This vector does not evolve in time and it does not have any noise, therefore for our system $\mathbf{F}_k = \mathbf{I}$ and $\mathbf{G}_k = \mathbf{0}$.

Suppose that we observe an outcome on combination $A + C$, as before, which we will summarise using the score statistic Z_2 and the Fisher information V_2. Then, in the Kalman filter setting our observation will be Z_2 and we can write this in terms of $\boldsymbol{\theta}$ as

$$Z_2 = (0, V_2)\, \boldsymbol{\theta} + \epsilon$$

where ϵ is our output noise and has distribution $\epsilon \sim N(0, V_2)$. Hence, for our system $\mathbf{H}_k = (0, V_2)^T$, $\mathbf{v}_k = \epsilon$ and $\mathbf{R}_k = V_2$.

Applying the above formulae for the Kalman Filter to our problem, we get the

following.

$$
\begin{aligned}
\hat{\boldsymbol{\mu}}_{0|-1} &= \boldsymbol{\mu}_0 \\
\hat{\boldsymbol{\Sigma}}_{0|-1} &= \boldsymbol{\Sigma}_0 \\
\hat{\boldsymbol{\mu}}_{0|0} &= \hat{\boldsymbol{\mu}}_{0|-1} + \hat{\boldsymbol{\Sigma}}_{0|-1}(0, V_2)^T \left((0, V_2)\hat{\boldsymbol{\Sigma}}_{0|-1}(0, V_2)^T + V_2\right)^{-1} \left(Z_2 - (0, V_2)\hat{\boldsymbol{\mu}}_{0|-1}\right) \\
&= \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}_0(0, V_2)^T \left((0, V_2)\boldsymbol{\Sigma}_0(0, V_2)^T + V_2\right)^{-1} \left(Z_2 - (0, V_2)\boldsymbol{\mu}_0\right) \\
&= \begin{pmatrix} \mu_1 - \frac{\rho_{12}\sigma_1\sigma_2 V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2}Z_2 \\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}Z_2 \end{pmatrix} \\
\hat{\boldsymbol{\Sigma}}_{0|0} &= \hat{\boldsymbol{\Sigma}}_{0|-1} - \hat{\boldsymbol{\Sigma}}_{0|-1}(0, V_2)^T \left((0, V_2)\hat{\boldsymbol{\Sigma}}_{0|-1}(0, V_2)^T + V_2\right)^{-1} (0, V_2)\hat{\boldsymbol{\Sigma}}_{0|-1} \\
&= \boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0(0, V_2)^T \left((0, V_2)\boldsymbol{\Sigma}_0(0, V_2)^T + V_2\right)^{-1} (0, V_2)\boldsymbol{\Sigma}_0 \\
&= \begin{pmatrix} \sigma_1^2 - \frac{V_2\rho_{12}^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2} & \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} \\ \frac{\rho_{12}\sigma_1\sigma_2}{1+V_2\sigma_2^2} & \frac{\sigma_2^2}{1+V_2\sigma_2^2} \end{pmatrix}.
\end{aligned}
$$

This is the same result that we saw when using the GMRF methodology.

B.2 Mixture model

Let R be the discrete random variable with probability mass function

$$p_R(0) = w_0^0 \quad \text{and} \quad p_R(1) = w_1^0$$

and let Σ_0 be the variance matrix of the bivariate normal distribution when the correlation is 0 and Σ_1 be the variance matrix when the correlation is $\rho_{12} > 0$.

The prior distribution for $\boldsymbol{\theta}$ may then be written as the following hierarchical model.

$$R \sim p_R$$

$$\boldsymbol{\theta} \mid R \sim \text{MVN}(\boldsymbol{\mu}, \Sigma_R)$$

Suppose we observe Z_2 where

$$Z_2 \mid \boldsymbol{\theta} \sim \text{MVN}(\boldsymbol{\mu}, \Sigma_R).$$

This information can be represented in a graphical model

$$R \longrightarrow \boldsymbol{\theta} \longrightarrow Z_2$$

Let $f_{\boldsymbol{\theta}|R}(\boldsymbol{\theta} \mid r)$ denote the probability density of $\boldsymbol{\theta}$ given $R = r$ and let $f_{Z_2|\boldsymbol{\theta}}(z_2 \mid \boldsymbol{\theta})$ be the probability density of Z_2 given the value of $\boldsymbol{\theta}$. Then the joint distribution of $(R, \boldsymbol{\theta}, Z_2)^T$ can be written as

$$f(r, \boldsymbol{\theta}, z_2) = p_R(r)\, f_{\boldsymbol{\theta}|R}(\boldsymbol{\theta} \mid r)\, f_{Z_2|\boldsymbol{\theta}}(z_2 \mid \boldsymbol{\theta}).$$

In a standard Bayesian analysis, we would find the posterior distribution for θ_1 given the observed value of Z_2 from

$$f_{\theta_1|Z_2}(\theta_1 \mid z_2) \propto \sum_r \int p_R(r)\, f_{\boldsymbol{\theta}|R}(\theta_1, \theta_2 \mid r)\, f_{Z_2|\boldsymbol{\theta}}(z_2 \mid \boldsymbol{\theta})\, d\theta_2.$$

Alternatively, we may consider the following graphical representation

$$R \longrightarrow \theta_1 \longleftarrow \theta_2 \longrightarrow Z_2$$

and thus write

$$f(r, \boldsymbol{\theta}, z_2) = p_R(r)\, f_{\theta_2}(\theta)\, f_{\theta_1|R,\theta_2}(\theta_1 \mid r, \theta_2)\, f_{Z_2|\theta_2}(z_2 \mid \theta_2).$$

The posterior distribution of R given $Z_2 = z_2$ will be the same as the prior distribution. We will then have

$$f_{\theta_1|Z_2}(\theta_1 \mid z_2) \propto \sum_r p_R(r) \int f_{\theta_2}(\theta_2)\, f_{\theta_1|R,\theta_2}(\theta_1 \mid r, \theta_2)\, f_{Z_2|\theta_1,R,\theta_2}(z_2 \mid \theta_1, r, \theta_2)\, d\theta_2$$

and we can show that

$$\begin{aligned} f_{\theta_1|Z_2}(\theta_1 \mid z_2) \propto\ & P(R=0)\, f_{\theta_1,Z_2|R}(\theta_1, z_2 \mid R=0) / f_{Z_2|R}(z_2 \mid R=0) \\ & + P(R=1)\, f_{\theta_1,Z_2|R}(\theta_1, z_2 \mid R=1) / f_{Z_2|R}(z_2 \mid R=1) \\ =\ & w_0^0 f_{\theta_1,Z_2|R}(\theta_1, z_2 \mid R=0) / f_{Z_2|R}(z_2 \mid R=0) \\ & + w_1^0 f_{\theta_1,Z_2|R}(\theta_1, z_2 \mid R=1) / f_{Z_2|R}(z_2 \mid R=1) \end{aligned}$$

Therefore, under the standard Bayesian updating procedure, the posterior distribution of θ_1 given Z_2 is a mixture of normal distributions that have the same weights as the prior mixture model. Each of the components of the mixture is updated based on the value of R and Z_2.

B.3 Results

	$\mu_1 = 0.5$	$\mu_2 = 0.2$	$\mu_2 = 0.5$	$\mu_2 = 0.8$
Univariate	Mean PoS	0.762	0.762	0.762
	% “Go” (PoS > 0.6)	100	100	100
GMRF	Mean PoS	0.968	0.861	0.621
	% “Go" (PoS > 0.6)	100	100	60.2
Hypothetical	Mean ω_0^1	0.623	0.137	0.625
	Mean PoS	0.837	0.848	0.725
	% “Go" (PoS > 0.6)	100	100	100
Limiting	Mean ω_0^1	0.639	0.153	0.642
	Mean PoS	0.833	0.846	0.727
	% “Go" (PoS > 0.6)	100	100	100

Table B.1: Table showing the results for combination $A + B$ of the simulation study where the true values of θ_1 and θ_2 are given by 0.5 and $\mu_1 = 0.5$ and μ_2 represent the prior means for each combination. Note that the univariate approach does not update the distribution of combination $A + B$ based on the results of combination $A + C$.

	$\mu_1 = 0.8$	$\mu_2 = 0.2$	$\mu_2 = 0.5$	$\mu_2 = 0.8$
Univariate	Mean PoS	0.915	0.915	0.915
	% “Go” (PoS > 0.6)	100	100	100
GMRF	Mean PoS	0.998	0.982	0.906
	% “Go" (PoS > 0.6)	100	100	100
Hypothetical	Mean ω_0^1	0.623	0.137	0.625
	Mean PoS	0.946	0.973	0.919
	% “Go" (PoS > 0.6)	100	100	100
Limiting	Mean ω_0^1	0.639	0.153	0.642
	Mean PoS	0.945	0.972	0.919
	% “Go" (PoS > 0.6)	100	100	100

Table B.2: Table showing the results for combination $A + B$ of the simulation study where the true values of θ_1 and θ_2 are given by 0.5 and $\mu_1 = 0.8$ and μ_2 represent the prior means for each combination. Note that the univariate approach does not update the distribution of combination $A + B$ based on the results of combination $A + C$.

Appendix C

Appendix - Chapter 4

C.1 Motivation

We provide motivation and context for aKDA through the use of an illustrative example that is based on a real world pharmaceutical portfolio, the Roche neuroscience pipeline [70] as of January 2019. From this portfolio, we use the number of programmes and the current stage of each of the programmes. We then specify information, such as the associated clinical trial costs, using estimates from the literature.

The Roche neuroscience pipeline [70] contained 13 programmes, $|I| = 13$, as of January 2019 and, if we take $|J_i|$ to be the number of remaining phases including the current one, then $|J_i|$ ranged from one to three. We used the estimates given by DiMasi et al. [3] as our baseline for the PoS, duration and cost of each of the studies/phases. We also used the sum of the mean clinical period capitalised costs from DiMasi et al. [3] multiplied by 1.25 to give a baseline value for $\text{rev}_i^{\max}$. These baseline values are given in Table C.1 and the full set of input parameters used are given in Table C.2. We split the planning horizon into periods of 10 months, due to the length of the studies considered, and considered a planning horizon of nine time periods, $|T| = 9$, as this is equal to the time taken to complete a full programme in this example. We consider one resource type, $|R| = 1$, which we set to be equal to the study cost divided by the study duration. We specify

a resource availability of 100 units, $\lambda_1^{\max} = 100$, as this will allow Phase I and II studies to be run alongside a Phase III study, but will not allow two Phase III studies to be run simultaneously. These resources could be related to cost, employees or capacity, for example. We will consider the two penalties to be equal to six units of cost per ten month period, as this represents a similar proportion of the maximum revenue as was used by Colvin and Maravelias [39].

	Phase I	Phase II	Phase III
Phase transition probability	0.595	0.355	0.62
Mean out-of-pocket clinical period cost ($m)	25.3	58.6	255.4
Mean clinical period capitalised cost ($m)	49.6	95.3	314.0
Mean time to next phase (months)	19.8	30.3	45.1

Table C.1: Estimates provided by DiMasi et al. [3] The mean out-of-pocket cost represents the actual study cost whereas the capitalised cost takes into account the cost of unsuccessful programmes.

The correlation matrix for the illustrative example is given below.

$$\begin{pmatrix}
1 & 0.5 & 0.4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.5 & 1 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.4 & 0.5 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0.6 & 0.5 & 0.4 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.6 & 1 & 0.6 & 0.4 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.5 & 0.6 & 1 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.4 & 0.4 & 0.5 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.6 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.6 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.4 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.4 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}$$

i	j	$\text{rev}_i^{\max}$	c_{ij}	τ_{ij}	λ_{ij1}	γ_i^L	θ_{ij}^{ind}	PoS	α_{ij}	β_{ij}	δ_{ij}	$\mu_i^{[1,1]}$	$\sigma_i^{2[1,1]}$
1	1	570	25	2	12.5	6	0	0.5			0.5	0.5	0.1
	2		50	3	16.7		1		0.1	0.2			
	3		250	4	62.5		1		0.05	0.1			
2	1	550	20	2	10	6	0	0.5			0.5	0.5	0.1
	2		40	3	15		1		0.1	0.2			
	3		270	4	67.5		1		0.05	0.1			
3	1	610	300	4	75	6	1		0.05	0.1	0.5	0.5	0.05
4	1	560	25	2	12.5	6	0	0.7			0.5	0.5	0.1
	2		60	3	20		1		0.1	0.2			
	3		240	4	60		1		0.05	0.1			
5	1	590	30	2	15	6	0	0.7			0.5	0.5	0.1
	2		50	3	16.7		1		0.1	0.2			
	3		270	4	67.5		1		0.05	0.1			
6	1	600	280	4	70	6	1		0.05	0.1	0.5	0.5	0.05
7	1	610	70	3	23.3	6	1		0.1	0.2	0.5	0.5	0.1
	2		260	4	65		1		0.05	0.1			
8	1	560	50	3	16.7	6	1		0.1	0.2	0.5	0.5	0.1
	2		250	4	62.5		1		0.05	0.1			
9	1	550	240	4	60	6	1		0.05	0.1	0.5	0.5	0.05
10	1	610	300	4	75	6	1		0.05	0.1	0.5	0.5	0.05
11	1	580	50	3	16.7	6	1		0.1	0.2	0.5	0.5	0.1
	2		270	4	67.5		1		0.05	0.1			
12	1	610	300	4	75	6	1		0.05	0.1	0.5	0.5	0.05
13	1	570	70	3	23.3	6	1		0.1	0.2	0.5	0.5	0.1
	2		250	4	62.5		1		0.05	0.1			

Table C.2: Input parameters used for the illustrative example based on the Roche neuroscience pipeline [70] and the baseline values given in Table C.1 based on the work of DiMasi et al. [3]

C.1.1 Discussion

It is clear that we would not be able to use the MSSP formulation of the project scheduling approach to find the optimal schedule or set of programmes for this example. In fact, this method would not be appropriate for many pharmaceutical portfolios due to its inability to be solved for reasonably sized portfolios.

Unlike the full MSSP, the aKDA can be used for this illustrative example with 13 programmes, $|I| = 13$. This makes the heuristic promising in terms of implementation in the real world and also makes it applicable to portfolios of combinations, which are often larger due to their combinatorial nature.

We considered there to be six groups of correlated studies with sizes ranging from one to four and we chose the particular values of μ_i, σ_i^2 and $\rho_{ii'}$ arbitrarily. If we observe the outcome of a study that is contained in a group with size greater than one, then we can apply the methodology from Chapter 3 to update the PoS for the studies in the rest of the group.

The aKDA can be applied to and solved for our illustrative example, unlike the full MSSP, and is able to capture the relationships between the different programmes in the example, unlike the MSSP and the KDA. We specified θ_{ij}^{ind} to be equal to zero for Phase I studies and one for Phase II and III studies. In the next section, we will provide brief details of the results for our illustrative example alongside a detailed discussion of a simulation study for a smaller example portfolio.

C.1.2 Results

In this section, we consider a planning horizon of nine time periods, but only provide the details of the first five in Table C.3 due to the size of the decision tree after this point.

The first column in Table C.3 gives the time point in the planning horizon, which would be represented by the x-axis in a decision tree, and the second column gives the particular subproblem at this time point, which would be represented by the different nodes in a vertical segment of a decision tree. The final two columns are used to record the successes and failures that have been observed in each subproblem, prior to or at the current time point. This study information defines the subproblem and is helpful in understanding the studies that will be eligible in a particular subproblem. The "Selection" column provides the studies that are selected in each subproblem. For example, in the first subproblem at the

Time	Subproblem	Selection	Successes	Failures
1	1	(7,1) (9,1) (11,1)		
4	1	(1,1) (2,1) (8,1)		(7,1) (11,1)
	2	(1,1) (2,1) (8,1)	(11,1)	(7,1)
	3	(2,1) (4,1) (8,1)	(7,1)	(11,1)
	4	(2,1) (4,1) (8,1)	(7,1) (11,1)	
5	1	(4,1) (5,1) (13,1)		(7,1) (9,1) (11,1)
	2	(4,1) (5,1) (13,1)	(9,1)	(7,1) (11,1)
	3	(4,1) (5,1) (13,1)	(11,1)	(7,1) (9,1)
	4	(4,1) (5,1) (13,1)	(9,1) (11,1)	(7,1)
	5	(1,1) (5,1) (13,1)	(7,1)	(9,1) (11,1)
	6	(1,1) (5,1) (13,1)	(7,1) (9,1)	(11,1)
	7	(1,1) (5,1) (13,1)	(7,1) (11,1)	(9,1)
	8	(1,1) (5,1) (13,1)	(7,1) (9,1) (11,1)	

Table C.3: Output of the aKDA for the first five time periods of the illustrative example.

fifth time point, we have observed failures in studies $(7,1)$, $(9,1)$ and $(11,1)$. At this point we then choose to initiate studies $(4,1)$, $(5,1)$ and $(13,1)$.

We see that, when $(7,1)$ is successful, the aKDA chooses to prioritise $(4,1)$ over

$(1,1)$ since $(4,1)$ and $(7,1)$ are correlated, thus $(7,1)$ being successful increases the probability of success of the studies in programme 4. Similar patterns to these are seen later in the decision tree, which we will not show here because the size of the tree increases quickly after the fifth time period due to the number of studies selected thus the large potential number of outcomes and associated subproblems required. The aKDA generates and solves 1709 subproblems to generate the decision tree for this illustrative example and is able to do this in under five minutes.

C.2 Extended simulation study results

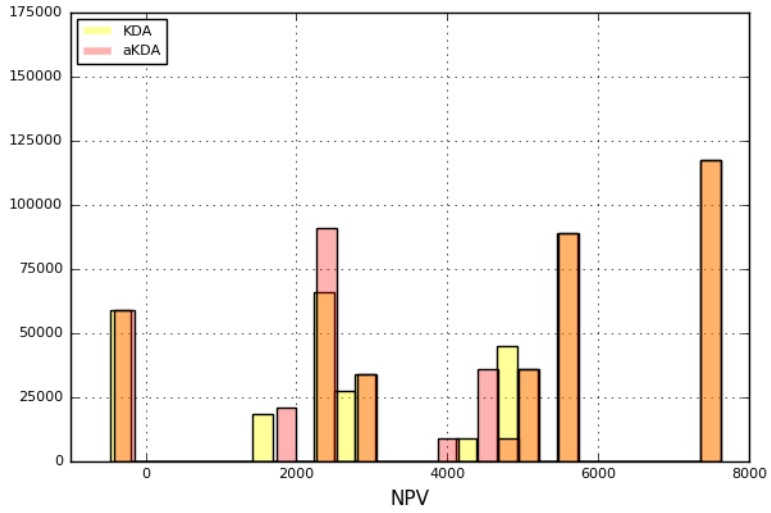
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(a)

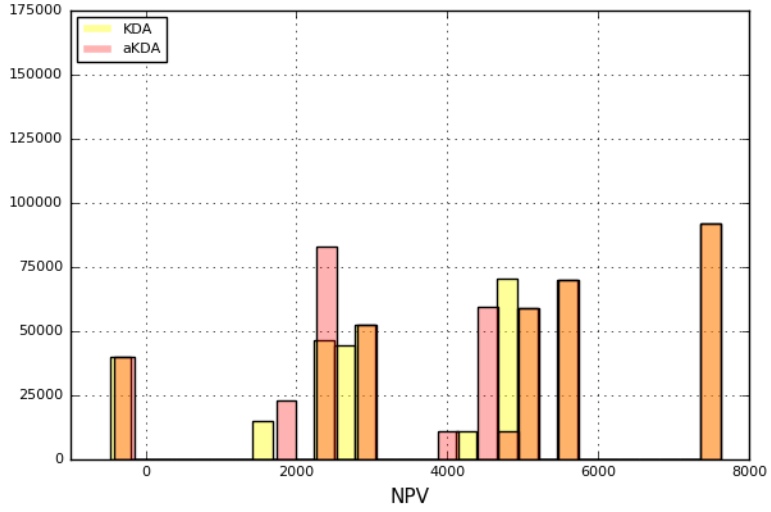
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(b)

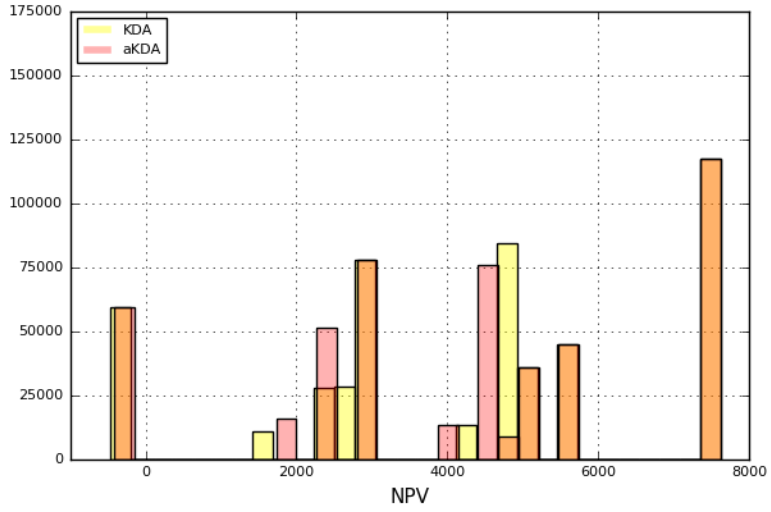

(c)

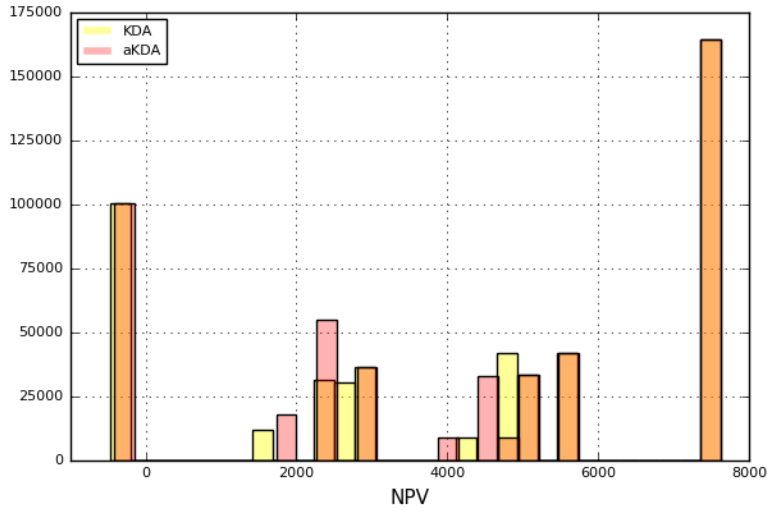

(d)

Figure C.1: Histograms of the NPV in the simulation study under each approach in Comparison 1 for the correlation matrices (a) - (d).

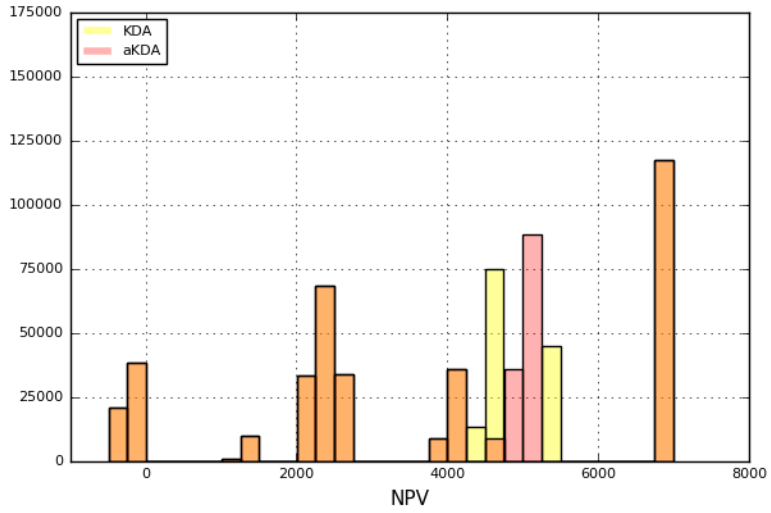
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(a)

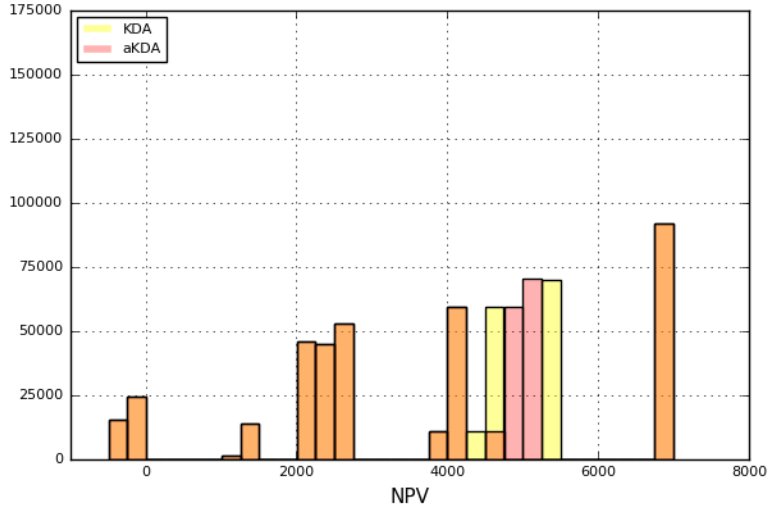
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(b)

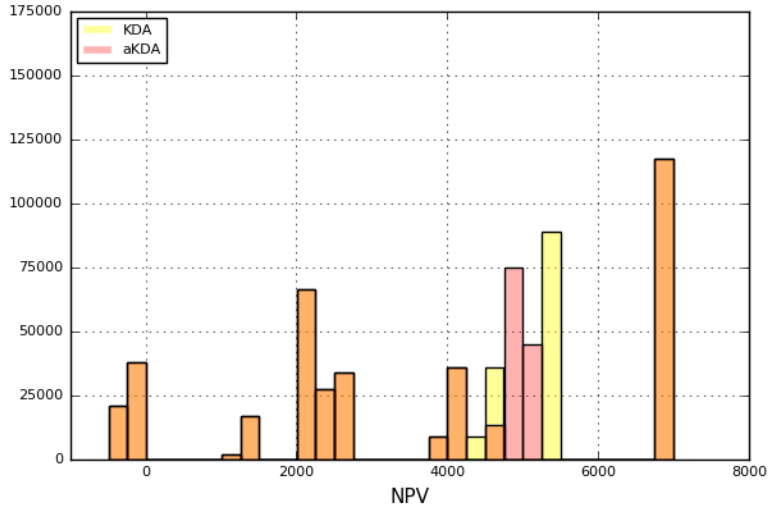

(c)

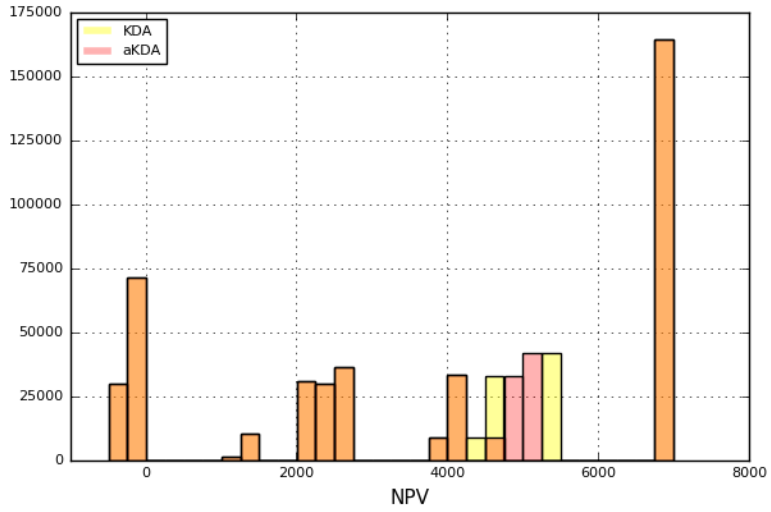

(d)

Figure C.2: Histograms of the NPV in the simulation study under each approach in Comparison 2 for the correlation matrices (a) - (d).

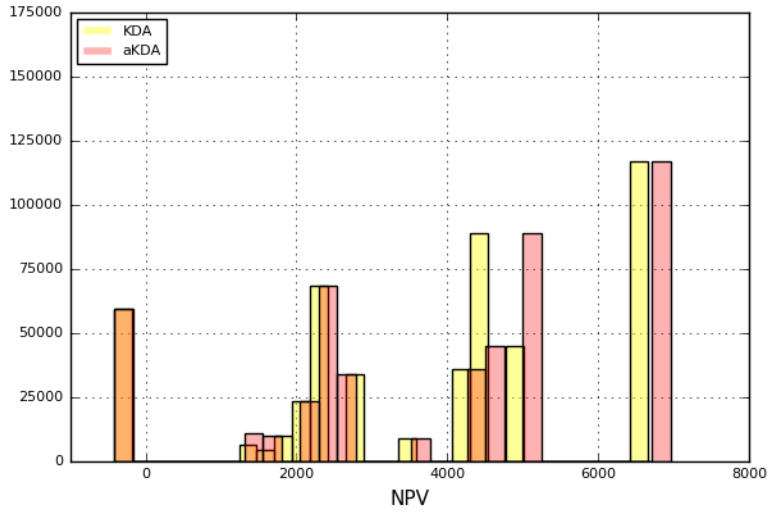
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(a)

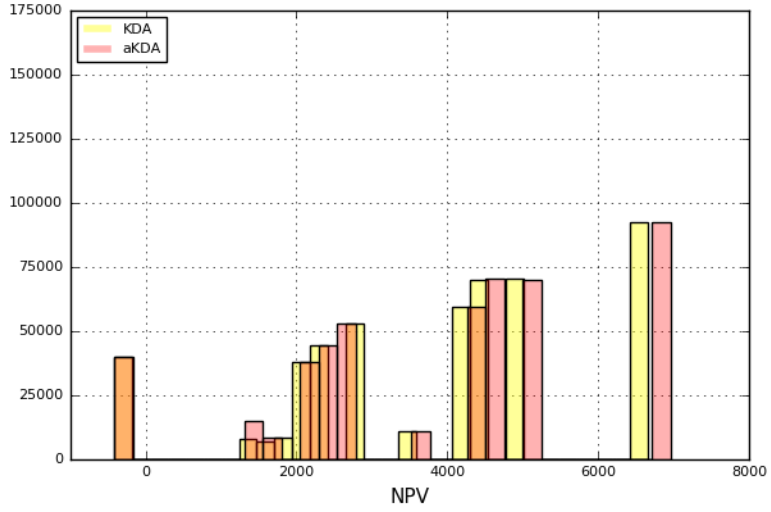
KDA
aKDA
175000
150000
125000
100000
75000
50000
25000
0
0
2000
4000
6000
8000
NPV

(b)

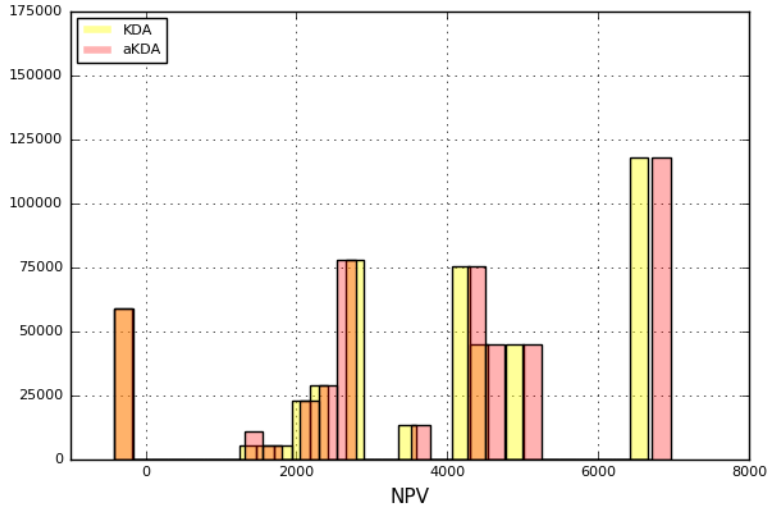

(c)

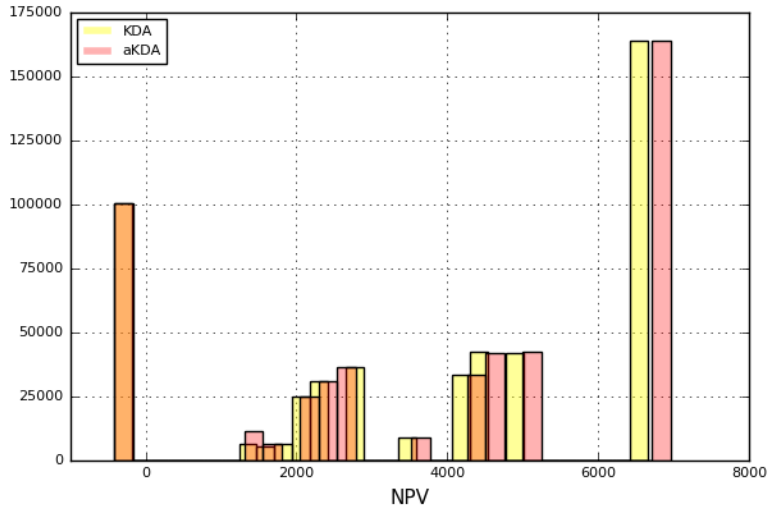

(d)

Figure C.3: Histograms of the NPV in the simulation study under each approach in Comparison 3 for the correlation matrices (a) - (d).

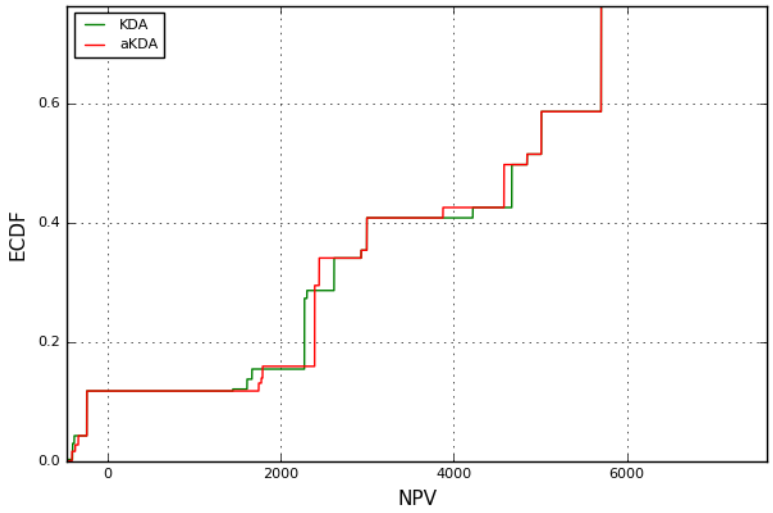
KDA
aKDA
ECDF
0.0
0.2
0.4
0.6
0
2000
4000
6000
NPV

(a)

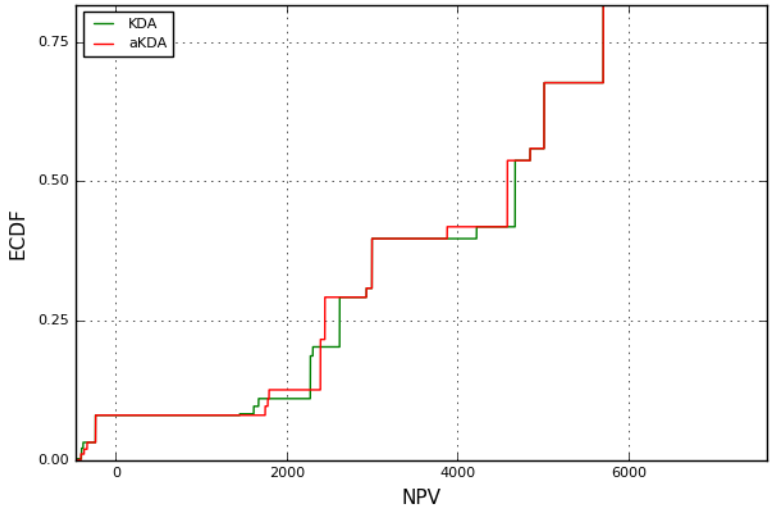
KDA
aKDA
ECDF
0.00
0.25
0.50
0.75
0
2000
4000
6000
NPV

(b)

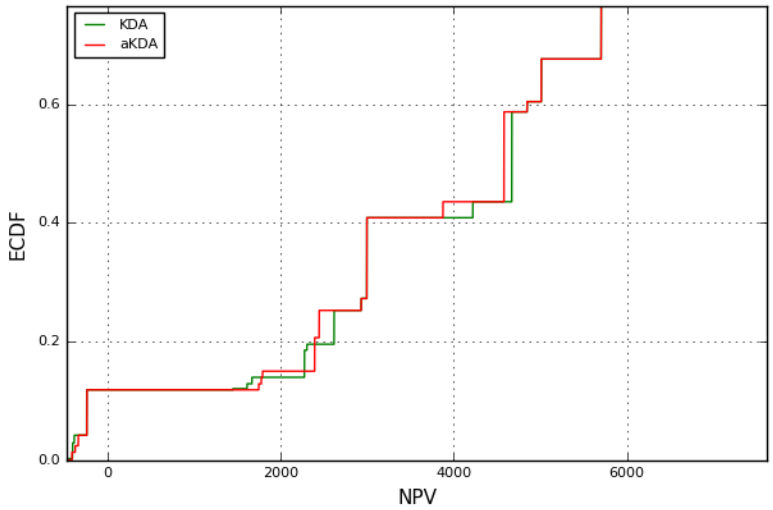

(c)

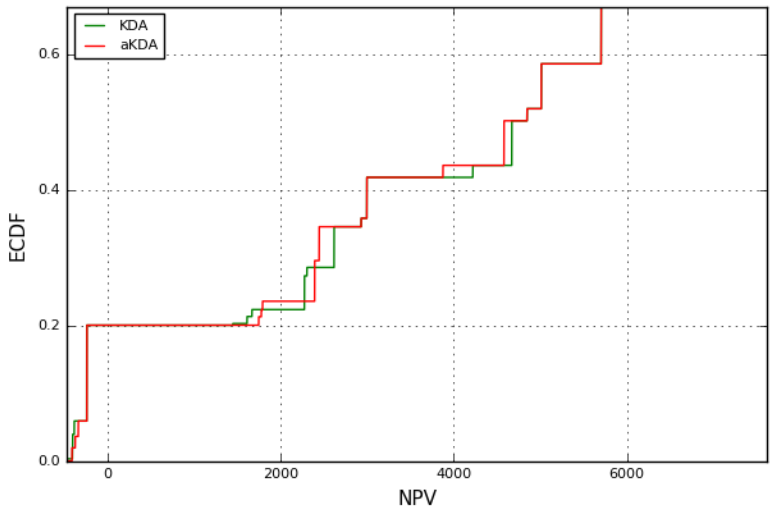

(d)

Figure C.4: Plots of the empirical cumulative distribution function of the NPV in the simulation study under each approach in Comparison 1 for the correlation matrices (a) - (d).

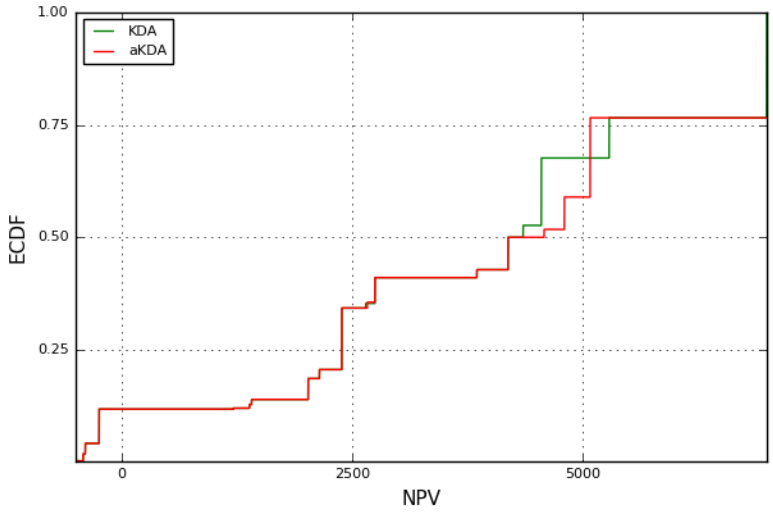
KDA
aKDA
1.00
0.75
0.50
0.25
ECDF
0
2500
5000
NPV

(a)

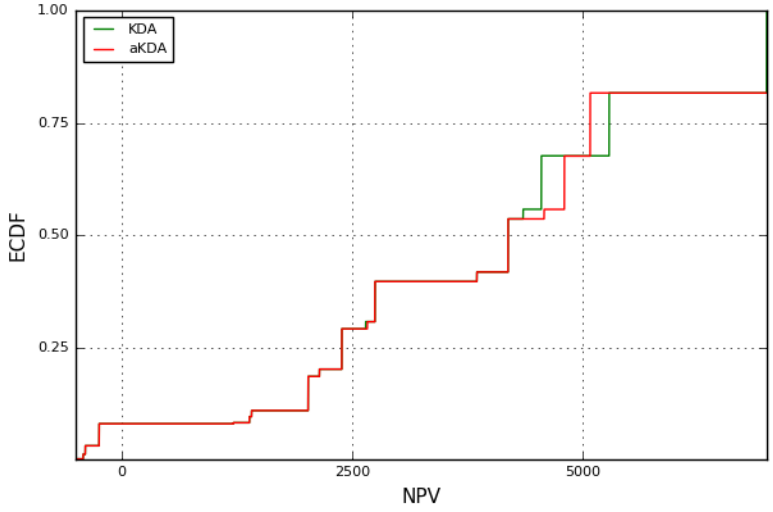
KDA
aKDA
1.00
0.75
0.50
0.25
ECDF
0
2500
5000
NPV

(b)

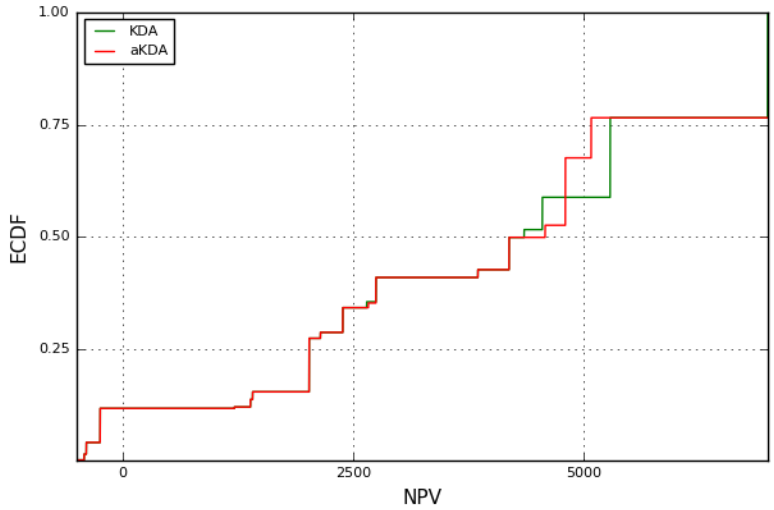

(c)

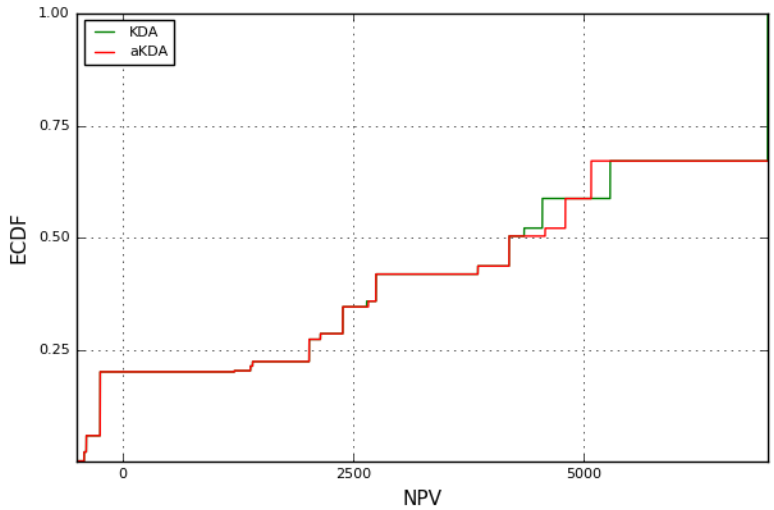

(d)

Figure C.5: Plots of the empirical cumulative distribution function of the NPV in the simulation study under each approach in Comparison 2 for the correlation matrices (a) - (d).

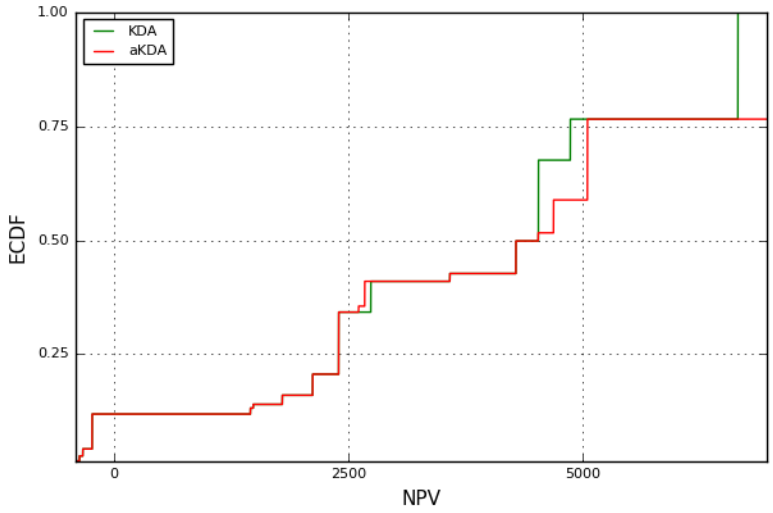
KDA
aKDA
1.00
0.75
0.50
0.25
ECDF
0
2500
5000
NPV

(a)

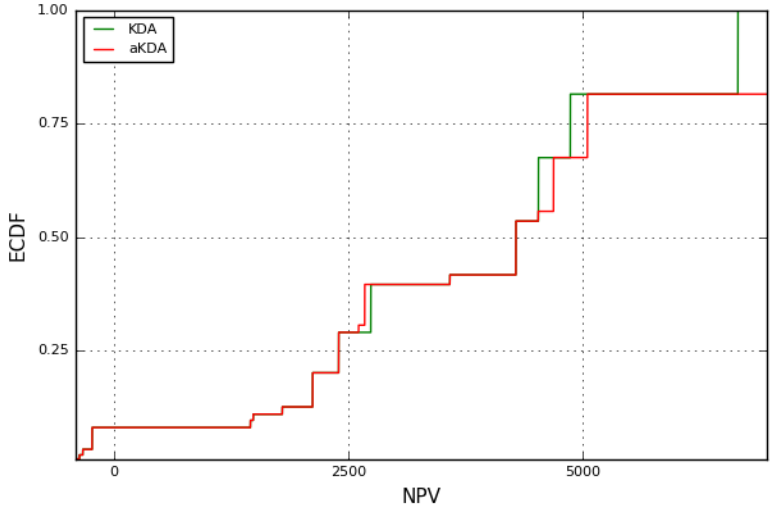
KDA
aKDA
1.00
0.75
0.50
0.25
ECDF
0
2500
5000
NPV

(b)

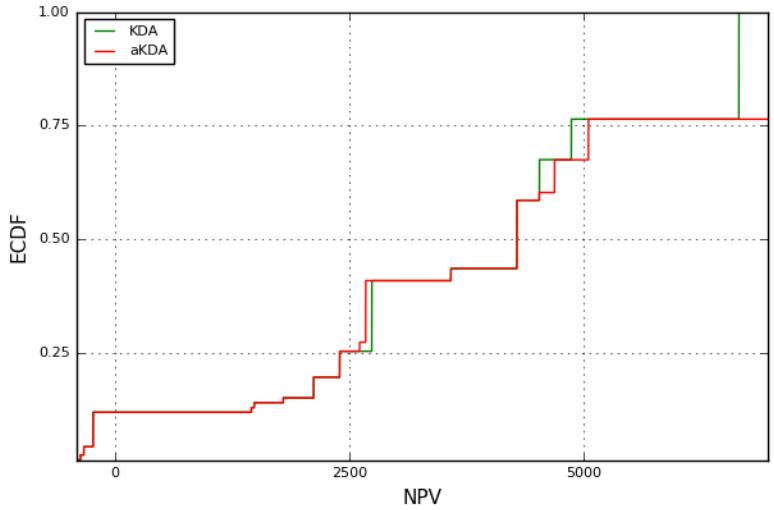

(c)

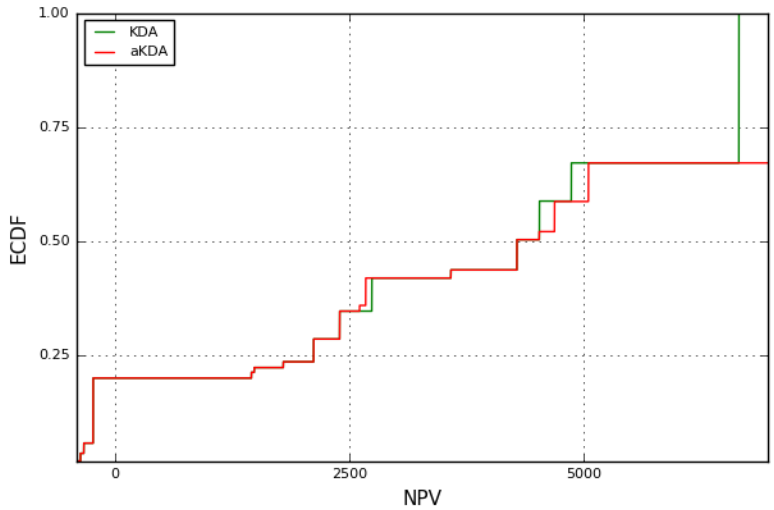

(d)

Figure C.6: Plots of the empirical cumulative distribution function of the NPV in the simulation study under each approach in Comparison 3 for the correlation matrices (a) - (d).

Bibliography

[1] 'Rationalizing combination therapies'. In: *Nature Medicine (Editorial)* 23.10 (2017).

[2] E. Whitmore. *Development of FDA-Regulated Medical Products - Prescription Drugs, Biologics, and Medical Devices*. American Society for Quality (ASQ), 2004.

[3] J. A. DiMasi, H. G. Grabowski, and R. W. Hansen. 'Innovation in the pharmaceutical industry: New estimates of R&D costs'. In: *Journal of Health Economics* 47 (2016), pp. 20–33.

[4] T. R. Fleming. 'Surrogate Endpoints And FDA's Accelerated Approval Process'. In: *Health Affairs* 24.1 (2005), pp. 67–78.

[5] A. Grignolo and S. Pretorius. 'Phase III trial failures: costly, but preventable'. In: *Applied Clinical Trials* 26.8/9 (2016), pp. 36–43.

[6] D. Thomas, J. Burns, J. Audette, A. Carroll, C. Dow-Hygelund, and M. Hay. *Clinical Development Success Rates 2006-2015*. Tech. rep. 2016.

[7] J. Arrowsmith and P. Miller. 'Trial Watch: Phase II and Phase III attrition rates 2011-–2012'. In: *Nature Reviews Drug Discovery* 12.8 (2013), pp. 569–569.

[8] M. S. Glickman and C. L. Sawyers. 'Converting cancer therapies into cures: Lessons from infectious diseases'. In: *Cell* 148.6 (2012), pp. 1089–1098.

[9] T.-C. Chou. 'Drug Combination Studies and Their Synergy Quantification Using the Chou-Talalay Method'. In: *Cancer Research* 70.2 (2010), pp. 440–446.

[10] J. B. Fitzgerald, B. Schoeberl, U. B. Nielsen, and P. K. Sorger. 'Systems biology and combination therapy in the quest for clinical efficacy'. In: *Nature Chemical Biology* 2.9 (2006), pp. 458–466.

[11] S. H. Podolsky and J. A. Greene. 'Combination Drugs — Hype, Harm, and Hope'. In: *New England Journal of Medicine* 365.6 (2011), pp. 488–491.

[12] S. Loewe. 'The problem of synergism and antagonism of combined drugs'. In: *Arzneimittelforschung* 3 (1953), pp. 285–290.

[13] C. I. Bliss. 'The calculation of microbial assays'. In: *Bacteriology Reviews* 20.4 (1956), pp. 243–258.

[14] J. Foucquier and M. Guedj. 'Analysis of drug combinations: current methodological landscape'. In: *Pharmacology Research and Perspectives* 3.3 (2015).

[15] J. Woodcock, J. P. Griffin, and R. E. Behrman. 'Development of Novel Combination Therapies'. In: *New England Journal of Medicine* 364.11 (2011), pp. 985–987.

[16] R. D. Lopes and R. A. Harrington. *Understanding Clinical Research*. McGraw-Hill Education, 2013.

[17] J. O'Quigley, M. Pepe, and L. Fisher. 'Continual Reassessment Method: A Practical Design for Phase 1 Clinical Trials in Cancer'. In: *Biometrics* 46.1 (1990), pp. 33–48.

[18] P. F. Thall and S.-J. Lee. 'Practical model-based dose-finding in phase I clinical trials: Methods based on toxicity'. In: *International Journal of Gynecological Cancer* 13.3 (2003), pp. 251–261.

[19] J. Kim and M. J. Schell. 'Modified Simon's minimax and optimal two-stage designs for single-arm phase II cancer clinical trials'. In: *Oncotarget* 10.42 (2019), pp. 4255–4261.

[20] J. M. S. Wason and T. Jaki. 'Optimal design of multi-arm multi-stage trials'. In: *Statistics in Medicine* 31.30 (2012), pp. 4269–4279.

[21] J. Whitehead. *The design and analysis of sequential clinical trials*. Rev. 2nd ed. Chichester: J. Wiley & Sons, 1997.

[22] C. Chuang-Stein. 'Sample size and the probability of a successful trial'. In: *Pharmaceutical Statistics* 5.4 (2006), pp. 305–309.

[23] A. O'Hagan, J. W. Stevens, and M. J. Campbell. 'Assurance in clinical trial design'. In: *Pharmaceutical Statistics* 4.3 (2005), pp. 187–201.

[24] Z. Su. 'Assessing the success probability of a Phase III clinical trial based on Phase II data'. In: *Contemporary Clinical Trials* 31.6 (2010), pp. 620–623.

[25] F. Liu. 'An extension of Bayesian expected power and its application in decision making'. In: *Journal of Biopharmaceutical Statistics* 20.5 (2010), pp. 941–953.

[26] J. G. Ibrahim, M.-H. Chen, M. Lakshminarayanan, G. F. Liu, and J. F. Heyse. 'Bayesian probability of success for clinical trials using historical data'. In: *Statistics in Medicine* 34.2 (2015).

[27] M. Wang, G. F. Liu, and J. Schindler. 'Evaluation of program success for programs with multiple trials in binary outcomes'. In: *Pharmaceutical Statistics* 14.3 (2015), pp. 172–179.

[28] K. Rufibach, P. Jordan, and M. Abt. 'Sequentially updating the likelihood of success of a Phase 3 pivotal time-to-event trial based on interim analyses or external information'. In: *Journal of Biopharmaceutical Statistics* 26.2 (2016), pp. 191–201.

[29] Y. Wang, H. Fu, P. Kulkarni, and C. Kaiser. 'Evaluating and utilizing probability of study success in clinical development'. In: *Clinical Trials* 10.3 (2013), pp. 407–413.

[30] N. Stallard, J. Whitehead, and S. Cleall. 'Decision-making in a phase II clinical trial: A new approach combining Bayesian and frequentist concepts'. In: *Pharmaceutical Statistics* 4.2 (2005), pp. 119–128.

[31] A. Dmitrienko and M.-D. Wang. 'Bayesian predictive approach to interim monitoring in clinical trials'. In: *Statistics in Medicine* 25.13 (2006), pp. 2178–2195.

[32] C. Chuang-Stein, S. Kirby, J. French, K. Kowalski, S. Marshall, M. K. Smith, P. Bycott, and M. Beltangady. 'A Quantitative Approach for Making Go/No-Go Decisions in Drug Development'. In: *Biostatistics* 45 (2011), pp. 187–202.

[33] T. Sabin, J. Matcham, S. Bray, A. Copas, and M. K. B. Parmar. 'A Quantitative Process for Enhancing End of Phase 2 Decisions'. In: *Statistics in Biopharmaceutical Research* 6.1 (2014), pp. 67–77.

[34] A. O'Hagan and J. E. Oakley. *SHELF: the Sheffield Elicitation Framework (version 3.0)*. (`http://tonyohagan.co.uk/shelf`). School of Mathematics and Statistics, University of Sheffield, UK, 2016.

[35] K. Rufibach, H. U. Burger, and M. Abt. 'Bayesian predictive power: choice of prior and some recommendations for its use as probability of success in drug development'. In: *Pharmaceutical Statistics* 15.5 (2016), pp. 438–446.

[36] M. J. Rogers, A. Gupta, and C. D. Maranas. 'Real Options Based Analysis of Optimal Pharmaceutical Research and Development Portfolios'. In: *Industrial & Engineering Chemistry Research* 41.25 (2002), pp. 6607–6620.

[37] N. R. Patel, S. Ankolekar, Z. Antonijevic, and N. Rajicic. 'A mathematical model for maximizing the value of phase 3 drug development portfolios incorporating budget constraints and risk'. In: *Statistics in Medicine* 32.10 (2013), pp. 1763–1777.

[38] M. J. Rogers, C. D. Maranas, and M. Ding. 'Valuation and design of pharmaceutical R&D licensing deals'. In: *AIChE Journal* 51.1 (2005), pp. 198–209.

[39] M. Colvin and C. T. Maravelias. 'A stochastic programming approach for clinical trial planning in new drug development'. In: *Computers & Chemical Engineering* 32.11 (2008), pp. 2626–2642.

[40] M. Colvin. 'Optimization Problems Under Endogenous Uncertainty Observation: Application to Pharmaceutical Research and Development Planning'. PhD thesis. University of Wisconsin-Madison, 2010.

[41] G. E. Blau, J. F. Pekny, V. A. Varma, and P. R. Bunch. 'Managing a portfolio of interdependent new product candidates in the pharmaceutical industry'. In: *Journal of Product Innovation Management* 21.4 (2004), pp. 227–245.

[42] W. L. Winston. *Operations research : applications and algorithms*. 4th ed. Belmont, California: Thomson/Brooks/Cole, 2004.

[43] M. Lubin and I. Dunning. 'Computing in Operations Research Using Julia'. In: *INFORMS Journal on Computing* 27.2 (2015), pp. 238–248.

[44] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. 'Julia: A Fresh Approach to Numerical Computing'. In: *SIAM Review* 59.1 (2017), pp. 65–98.

[45] J. Puchinger, G. R. Raidl, and U. Pferschy. 'The Multidimensional Knapsack Problem: Structure and Algorithms.' In: *INFORMS Journal on Computing* 22.2 (2010), pp. 250–265.

[46] A. Caprara, D. Pisinger, and P. Toth. 'Exact Solution of the Quadratic Knapsack Problem'. In: *INFORMS Journal on Computing* 11.2 (1999), pp. 125–137.

[47] V. Boskamp, A. Knoops, F. Frasincar, and A. Gabor. 'Maximizing revenue with allocation of multiple advertisements on a Web banner'. In: *Computers & Operations Research* 38.10 (2011), pp. 1412–1424.

[48] K. Kim, B.-I. Kim, and H. Cho. 'Multiple-choice knapsack-based heuristic algorithm for the two-stage two-dimensional cutting stock problem in the

paper industry'. In: *International Journal of Production Research* 52.19 (2014), pp. 5675–5689.

[49] B. Christian and S. Cremaschi. 'Heuristic solution approaches to the pharmaceutical R&D pipeline management problem'. In: *Computers & Chemical Engineering* 74 (2015), pp. 34–47.

[50] M. Colvin and C. T. Maravelias. 'R&D pipeline management: Task interdependencies and risk management'. In: *European Journal of Operational Research* 215.3 (2011), pp. 616–628.

[51] C. W. Schmidt and I. E. Grossmann. 'Optimization Models for the Scheduling of Testing Tasks in New Product Development'. In: *Industrial & Engineering Chemistry Research* 35.10 (1996), pp. 3498–3510.

[52] G. E. Rotstein, L. G. Papageorgiou, N. Shah, D. C. Murphy, and R. Mustafa. 'A Product Portfolio Approach in the Pharmaceutical Industry'. In: *Computers & Chemical Engineering* 23 (1999), S883–S886.

[53] V. A. Varma, J. F. Pekny, G. E. Blau, and G. V. Reklaitis. 'A framework for addressing stochastic and combinatorial aspects of scheduling and resource allocation in pharmaceutical R&D pipelines'. In: *Computers & Chemical Engineering* 32.4-5 (2008), pp. 1000–1015.

[54] A. Sundaramoorthy, J. M. B. Evans, and P. I. Barton. 'Capacity Planning under Clinical Trials Uncertainty in Continuous Pharmaceutical Manufacturing, 1: Mathematical Framework'. In: *Industrial & Engineering Chemistry Research* 51.42 (2012), pp. 13692–13702.

[55] P. Kodukula and C. Papudesu. *Project valuation using real options : a practitioner's guide*. eng. Fort Lauderdale, Florida: J. Ross Publishing, 2006.

[56] G. Lo Nigro, A. Morreale, and G. Enea. 'Open innovation: A real option to restore value to the biopharmaceutical R&D'. In: *International Journal of Production Economics* 149 (2014), pp. 183–193.

[57] M. Colvin and C. T. Maravelias. 'Scheduling of testing tasks and resource planning in new product development using stochastic programming'. In: *Computers & Chemical Engineering* 33.5 (2009), pp. 964–976.

[58] M. Colvin and C. T. Maravelias. 'Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming'. In: *European Journal of Operational Research* 203.1 (2010), pp. 205–215.

[59] A. O'Hagan and J. W. Stevens. 'Bayesian assessment of sample size for clinical trials of cost-effectiveness'. In: *Medical Decision Making* 21.3 (2001), pp. 219–230.

[60] J. Whitehead. 'Designing Phase II Studies in the Context of a Programme of Clinical Research'. In: *Biometrics* 41.2 (1985), pp. 373–383.

[61] B. R. Saville and S. M. Berry. 'Efficiencies of platform clinical trials: A vision of the future'. In: *Clinical Trials* 13.3 (2015), pp. 358–366.

[62] J. Baselga, J. Cortés, S.-B. Kim, S.-A. Im, R. Hegg, Y.-H. Im, L. Roman, J. L. Pedrini, T. Pienkowski, A. Knott, E. Clark, M. C. Benyunes, G. Ross, and S. M. Swain. 'Pertuzumab plus Trastuzumab plus Docetaxel for Metastatic Breast Cancer'. In: *New England Journal of Medicine* 366.2 (2012), pp. 109–119.

[63] E. A. Perez, C. Barrios, W. Eiermann, M. Toi, Y. H. Im, P. Conte, M. Martin, T. Pienkowski, X. Pivot, H. A. Burris III, J. A. Petersen, S. Stanze, A. Strasak, M. Patre, and P. Ellis. 'Trastuzumab emtansine with or without pertuzumab versus trastuzumab plus taxane for human epidermal growth factor receptor 2-positive, advanced breast cancer: Primary results from the phase III MARIANNE study'. In: *Journal of Clinical Oncology* 35.2 (2017), pp. 141–148.

[64] S. M. Swain, J. Baselga, S.-B. Kim, J. Ro, V. Semiglazov, M. Campone, E. Ciruelos, J.-M. Ferrero, A. Schneeweiss, S. Heeson, E. Clark, G. Ross,

M. C. Benyunes, and J. Cortés. 'Pertuzumab, Trastuzumab, and Docetaxel in HER2-Positive Metastatic Breast Cancer'. In: *New England Journal of Medicine* 372.8 (2015), pp. 724–734.

[65] H. Rue and L. Held. *Gaussian Markov Random Fields: Theory and Applications*. Boca Raton, Florida: Chapman and Hall/CRC, 2005.

[66] B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Mineola, New York: Dover Publications, 1979.

[67] M. Goldstein and D. Wooff. *Bayes Linear Statistics: Theory and Methods*. John Wiley & Sons, Ltd, 2007.

[68] I. Wadsworth, L. V. Hampson, and T. Jaki. 'Extrapolation of efficacy and other data to support the development of new medicines for children: A systematic review of methods'. In: *Statistical Methods in Medical Research* 27.2 (2018), pp. 398–413.

[69] T. Sabin, J. Matcham, A. Copas, and M. K. B. Parmar. 'Assessing End of Phase 2 Decision Criteria'. In: *Statistics in Biopharmaceutical Research* 7.3 (2015), pp. 230–237.

[70] *Roche Neuroscience Pipeline*. `https://www.roche.com/research_and_development/who_we_are_how_we_work/pipeline.htm`. Accessed: 2019-01-17.

[71] B. Christian and S. Cremaschi. 'Variants to a knapsack decomposition heuristic for solving R&D pipeline management problems'. In: *Computers & Chemical Engineering* 96 (2017), pp. 18–32.

[72] P. Frewer, P. Mitchell, C. Watkins, and J. Matcham. 'Decision-making in early clinical drug development'. In: *Pharmaceutical Statistics* 15.3 (2016), pp. 255–263.

[73] J. Choi, M. J. Realff, and J. H. Lee. 'Dynamic programming in a heuristically confined state space: a stochastic resource-constrained project scheduling application'. In: *Computers & Chemical Engineering* 28 (2008), pp. 1039–1058.