

University of Oxford



Models and Software for Improving the Profitability of Pharmaceutical Research

Shuo Qu

Exeter College

Thesis submitted to the Mathematical, Physical and Life
Sciences Division for the degree of Doctor of Philosophy of the
University of Oxford.

*Department of Statistics, 1 South Parks Road,
Oxford OX1 3TG*

June 2011

I certify that this is my own work (except where otherwise indicated).

Candidate: Shuo Qu

Signed

Dated

Abstract

Pharmaceutical R&D is time-consuming, extremely costly and involves great uncertainty. Although there is a broad range of literature on statistical issues in clinical trials, there is not much that focuses directly on the modelling of pre-clinical research.

This thesis investigates models and associated software for improving decision-making in this area, building on earlier work by the same research group. We introduce a class of adaptive policies called forwards induction policies for candidate drug selection, and show that these are optimal, with a straightforward solution algorithm, within a restricted setting, and are usually close to optimal more generally. We also introduce an adaptive probabilities model that allows the incorporation of learning from a project's progress into the planning process. Real options analysis in the evaluation of project value is discussed. Specifically, we consider the option value of investing in clinical trials once a candidate drug emerges from pre-clinical research. Simulation algorithms are developed to investigate the probability distributions of the total reward, total cost, profitability index and the required future resource allocations of a pharmaceutical project under a given allocation plan. The ability to simulate outcome distributions means that we can also compare the riskiness of different projects and portfolios of projects.

Acknowledgements

I would like to thank my supervisor, Professor John Gittins, for his firm commitment to my doctoral work, generous donation of time, and for demonstrating how to conduct research with intellectual rigour, clarity of expression and attention to detail. I would also like to acknowledge my debt to him for polishing the ideas, insights and writings in this thesis, through the patient reading and discussion of so many drafts. It is my great honour getting to know Professor John Gittins, and working under his supervision has been such an enjoyable experience. I am and always will be very grateful for his dedication, patience, kindness, understanding and encouragement.

I would like to thank the Department of Statistics and Exeter College, University of Oxford, for providing generous financial support. I feel extremely fortunate to have studied at these two great institutions in the world, where I have met some of the world's greatest minds, grown intellectually and personally, and made life-long friends. The past five years that I spent at these two great institutions makes an amazing chapter in my life. I would like to give my best wishes to the Department of Statistics and to Exeter college, and hope they continue to thrive and serve as the leaders for excellence in teaching and research for years to come.

Finally, my deepest gratitude goes to my mother, for giving me life, unconditional love, unreserved support, and being the greatest mother that I could ever dream of having, for always being there for me through all the ups and downs in my life, for teaching me how to love, to forgive, and to be a decent person, for showing me courage and determination, for her inestimable patience and endless encouragement, and for so much more. There is no way I can ever express my gratitude enough for everything she had given me, without which I would never have come this far today. She lost the battle to cancer on March 2007, but she has always lived in my heart like a guiding star. At this very moment, when I am about to enter the next chapter in my life, I hope she is proud of what I have achieved. *This thesis is dedicated to the loving memory of her.*

Contents

1	Introduction	7
1.1	Pharmaceutical R&D Summary	7
1.2	Drug Discovery and Development	13
1.3	Context of this Work	17
1.4	Contribution of this Thesis	19
1.5	Contents Overview	22
2	Stages, Probabilities, Rewards and Costs	24
2.1	Stages of Drug Discovery	24
2.2	Clinical Trials and Beyond	26
2.3	The Expected Values of Successive Candidate Drugs	27
2.4	Profitability Criteria	29
3	A Forwards Induction Approach to Candidate Drug Selection	31
3.1	FI Policies	32
3.2	FI for the Simplified Problem	33
3.2.1	Infrastructure	33
3.2.2	Solution to the Simplified Problem	38

3.3	FI for the General Problem	41
4	More OPRRA Modelling	45
4.1	The Work-Rate/Progress Relationship	45
4.2	Discounting and Obsolescence	47
4.3	(s, n) Policies	49
4.3.1	Definition of an (s, n) Policy	49
4.3.2	Expected Reward and Expected Cost	50
4.3.3	Optimization Calculations with PI and IRR	56
4.4	Types of Allocation Policy for a Project	58
4.5	Contributions to OPRRA	59
5	The Adaptive Probabilities Model	61
5.1	Estimates for Success Probabilities	62
5.2	Estimates for Hyper-parameters	65
5.3	Adaptive FI Policies	69
6	The Option Valuation Model	73
6.1	Real Options Overview	74
6.2	General Valuation Framework and Numerical Solution	79
6.3	Calculations of Profitability for (s, n) and FI Policies	86
7	Tests on Optimization Procedures	91
7.1	Optimization Procedures	91
7.2	Parameter Values	93
7.3	Numerical Results	94

8	Effects of Model Parameters	101
8.1	Effects of Parameters on PI^*	101
8.1.1	Identifying Important Parameters	101
8.1.2	Factorial Design	107
8.2	Effects of Parameters on s^*	114
9	Conclusions and Desirable Future Work	122
A	Glossary	126
B	Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy	127
B.1	FP Setting	129
B.1.1	Calculation of $PI(CD)$	129
B.1.2	Calculation of $PI(LS)$	129
B.2	AP Setting	131
B.2.1	Calculation of $PI(CD)$	131
B.2.2	Calculation of $PI(LS)$	131
C	Simulation Algorithms	134
C.1	Notation	134
C.2	Assumptions	135
C.2.1	FP Setting	135
C.2.2	AP Setting	136
C.3	Pseudocode	137
C.3.1	(s, n) Simulation	137
C.3.2	FI Simulation in FP Setting	140

C.3.3	FI Simulation in AP Setting	143
D	Bellman Equation	147
D.1	Discrete Time	147
D.2	Continuous Time	149
E	Itô's Formula	150
F	Root Finding Techniques	153
F.1	Newton-Raphson Method	153
F.2	Bisection Method	155
	Bibliography	157

List of Figures

7.1	<i>The distribution of Δ.</i>	96
7.2	<i>Plot of Δ against s^*.</i>	97
7.3	<i>Plots of P_{FI}, P_{FI11} and P_{RFI} against s^*.</i>	98
8.1	<i>On the left: Residual plot of the initial full model, which regresses PI^* against the model parameters. On the right: Profile log-likelihood from a Box-Cox analysis of the initial full model.</i>	103
8.2	<i>The C_p plot.</i>	105
8.3	<i>Residuals vs. fitted for the final model.</i>	105
8.4	<i>Normality check for the final model.</i>	106
8.5	<i>Normal probability plot of effects for the 2_V^{21-12} design.</i>	110
8.6	<i>On the left: Normal probability plot of the residuals. On the right: Plot of residuals versus predicted.</i>	113
8.7	<i>α-W interaction plot.</i>	114
8.8	<i>Plots of 12 explanatory variables against s^*.</i>	115
8.9	<i>Effects of effort allocations.</i>	119

List of Tables

7.1	<i>Input parameter ranges.</i>	95
7.2	<i>Profitability Index values for the selected cases.</i>	99
7.3	<i>Parameter values for the selected cases.</i>	100
8.1	<i>Estimates of regression coefficients.</i>	104
8.2	<i>Low and high values for the selected factors.</i>	108
8.3	<i>Effect estimate summary.</i>	111
8.4	<i>Analysis of Variance.</i>	112
8.5	<i>Output for the cumulative logit model with s^* as the ordered response.</i>	120
8.6	<i>Estimated probabilities describing the effects in the refined model.</i>	121

Chapter 1

Introduction

1.1 Pharmaceutical R&D Summary

This introductory section provides a general discussion on the current problems and challenges faced by the pharmaceutical R&D sector, and possible means of tackling these problems. The material discussed in this section is based on recent reports concerning the current and future state of the pharmaceutical industry by Accenture (2005, 2007), PricewaterhouseCoopers (2008), Bain and Company (2009, 2010), McKinsey (2010) and the references therein. These reports have contributed equally as they all describe the situation in very similar terms.

The world needs the pharmaceutical industry to succeed. According to the World Health Organization, heart disease, stroke, cancer, chronic respiratory disease and diabetes cause more than 60% of all deaths worldwide. With major health problems still unsolved, the pharmaceutical industry needs to prime its innovation pump and achieve better results.

Few industries have a more complex business model than the pharmaceutical sector. The pharmaceutical industry also spends more on research and development, relative to its sales revenue, than almost any other industry. According to an estimate by Bain and Company, the industry's spending on drug R&D stood at £40 billion in 2007, up roughly 56% from 2001. Pharmaceutical R&D is in-

creasingly risky, costly and at times inefficient. Estimates of the average cost of bringing a new drug to market range from £500 million to £1.2 billion, extending over a period of around 10-15 years.

The pharmaceutical R&D process consists of two major phases, discovery and development. Drug discovery (or pre-clinical research) is the process by which drugs are discovered and/or designed. Once a compound has shown its value in discovery, it will begin the process of drug development (or clinical trials). Out of every 5000 potential new development compounds identified during the discovery process, only five are considered safe for testing in human volunteers after pre-clinical evaluations. After three to six years of further clinical testing in patients, on average only one of these compounds ultimately emerges as a marketable and regulatory-approved drug. Then only about 30% of drugs launched recover their risk-adjusted cost of R&D - a rate that is likely to deteriorate further given the increasing demands of payers and access agencies that want more value for their money (Hariharan and Singh, 2010). To cover their costs, pharmaceutical manufacturers traditionally rely on a few blockbuster compounds to make up for the underperformers. But blockbusters are few and far between. In 2009, 19 new molecular entities (NMEs) were approved by the US Food and Drug Administration (FDA), a volume which is much lower than the industry and investors would like it to be.

The pharmaceutical industry has reached a tipping point. With few exceptions, pharmaceutical companies have been unsuccessful in achieving their growth objectives. The industry as a whole has created limited shareholder value since 2000, underperforming the S&P 500's own tepid results. Total shareholder return from the beginning of the decade through year-end 2007 was only 2.7% for the pharmaceutical industry versus 3.6% for the S&P 500 and 5.9% in the relatively low-growth consumer goods sector (Russell and Gjaja, 2008). The pharmaceutical industry has entered a period of significant uncertainty and transition, characterized by rising R&D cost, declining productivity and increasing attrition rates.

Over recent years, the costs of pharmaceutical R&D have soared. According to the 2009 annual report by the European Federation of Pharmaceutical Industries and Associations (EFPIA, 2009), between 2004 and 2008, pharmaceutical R&D

expenditure has risen by about 7% annually. The industry's average R&D expenditure as a percentage of sales stood at 16% in 2008, more than doubled since 1980. The increasingly complex nature of science, the significant cost of ever larger clinical trials and the amount of resources needed to get approval by regulatory authorities are the primary reasons for this exponential increase of R&D costs. The financing of such R&D expenditure requires a sustained and substantial investment.

However, continued growth in R&D spending has appeared to have little effect on the pace at which new drugs are developed. Annual approvals of NMEs by the FDA experienced a pronounced decline since the mid-1990s. In this decline, the total number of NMEs approved each year fell from a high of 53 in 1996 to 19 in 2009. As a result, the average R&D cost per new drug has grown significantly, while prices worldwide are under pressure. R&D productivity, as measured by R&D expenditure per NME license application, declined on average by 21% annually for 10 years since 1998 (Bain and Company, 2009).

More than 80% of compounds that enter clinical trials are destined to fall out of the development pipeline. Attrition rates of this magnitude are the single biggest cause of the industry's productivity challenge. A 2010 McKinsey report (Singh et al., 2010) tracking the attrition rates of more than 3,000 compounds in each phase of development showed that overall clinical success rates fell between 1996 and 2007. Their research shows that success rates in phase *I*¹ have been relatively stable, having stood at around 65% for the past 12 years. After a dip earlier in the analysis period, success rates in phase *III* are recovering, from 50% in 2003 to 60% in 2007, leaving the attrition in phase *II* as the single biggest reason for declining success. The phase *II* success rate has decreased 16% since 1997, leading to a rate of 33% in 2007. The McKinsey analysis also revealed a decline in the success rates in most therapeutic areas. The average decline is 4.5%, but the figure varies by therapeutic area: oncology success rates have dropped by 3.2% and endocrine by 13%. Declining productivity and expensive late-stage failures have led to an urgent need for reassessment of the criteria used to select compounds for progression. Inadequate understanding of the commercial and scientific drivers of these processes is a major obstacle to improved decision making.

¹For a description of the phases of clinical trials see section 1.2.

Drug discovery organizations have not been able to deliver on the expectations they set for themselves. The pressure will continue to increase and intensify as products come off patent and society reaches the limits of willingness and ability to pay for pharmaceutical innovation. Faced with patent expirations, rising expenses, competition from generics and pressure on branded drug prices, big pharma's revenue gap could balloon to almost £60 billion by 2014 (Bain and Company, 2010). If pharmaceutical R&D is to remain at the forefront of medical research, and continue helping patients to live longer, healthier lives, pharmaceutical companies will need to change their fundamental practices and underlying research processes.

For instance, companies need to integrate genomics, proteomics and other technologies to improve target identification and attrition, and to enhance lead optimization and improve clinical trial design so as to speed approval, and shift from broadly targeted drugs to more focused medicines with much higher therapeutic value for the target population. These emerging technologies have the potential to fundamentally change the drug discovery process, by intervening at multiple points to increase the speed and to reduce attrition rates in early discovery, as well as by improving the quality of the compounds. Pharmacogenomics may improve the success rate of clinical trials through use of patient subsets with specific genetic risks and reduced chances of toxicities and side effects. This genetic profiling can reduce clinical trial costs and increase the success rates of drugs after launch. Some drugs that have been shelved may be revived after a re-examination with reference to specific genetic subsets. Genetic variations among individuals may suggest the need for individualized treatments and tailored drugs. Exploiting this opportunity will require companies to introduce genotype-based diagnostics into personalized medicine, which will have broad impact in the industry. And further, using genotype-based elimination of major side effects will actually create even larger blockbuster products than is possible based on today's approach to developing and prescribing medicines.

Creating partnerships and alliances will enable companies to become 'virtual' entities, with streamlined parent companies tied to outside alliances that provide fast access to critical capabilities. They also offer the opportunity to experiment with emerging technologies and facilitate involvement in specific disease or ther-

apeutic areas. Much of the innovation needed to fuel the growth expectations of leading pharmaceutical companies is occurring in the biotechnology² sector. To survive, leading pharmaceutical companies will need to make better use of innovation in this sector. As the competition for value-delivering partnerships with biotechnology firms increases among the pharmaceutical companies, they will need to respond by overcoming some of the existing barriers to effective alliances between the two sectors and by developing new competencies in alliance evaluation and management. Alliances are increasingly popular for the opportunities they offer for collaboration, knowledge sharing and resource sharing and as catalysts for fundamental change. The winners in this intensely competitive environment will be those companies that recognize external innovations as critical to their future success and therefore build the strategy, organization, process and technology elements needed to support and enhance the effectiveness of their alliance activities.

R&D organizations will need to cope with ever-expanding volumes of relevant data crucial to the business. They will need to embed a new, integrated technology system in R&D processes to connect disparate data, information and technology. To work effectively using the new networked approach, R&D organizations need to strengthen their capabilities to support enhanced collaboration and integration, encompassing both internal and external stakeholders, in a manner that accommodates both structured and unstructured interaction. To foster a truly global R&D approach, R&D organizations will need to develop a flexible infrastructure, one that is highly secure, yet also elastic and cost-effective. To utilize a wealth of available internal and external data, R&D organizations will need to enhance their capabilities to assimilate and interpret a wide array of inputs, ranging from sentiment monitoring and social networks to fully digital operational data that is integrated and accessible. The pharmaceutical R&D sectors will also be required to incorporate a broad spectrum of analytics and predictive capabilities, ranging from simulation and modelling tools to visualization techniques that provide the means to absorb massive volumes of data.

One critical competence for increased performance in discovery will be the

²Technological application that uses biological systems, living organisms, or derivatives thereof, to make or modify products or processes for specific use.

need for more rapid and accurate prioritization and decision-making. R&D leaders grapple with decisions about program termination, acceleration, resourcing and prioritization. Project termination decisions are especially difficult and can cost a company hundreds of millions of pounds if made too late. The current high attrition rate in clinical trials suggests that companies have overlooked or ignored key signals, and in some cases made poor decisions about aspects over which they have substantial control. The growing complexity of discovery operations and the increasing importance of managing functional interfaces have led some companies to examine how to best create discovery operating platforms. Companies need to improve their performance and quality of output by aligning the discovery process to overall strategy and integrating new scientific and information technologies with the decision-making and resource management of the organization.

Careful portfolio analysis is the key to making right decisions, and R&D leaders can increase returns by identifying and removing poor performers from the portfolio earlier in development. One of the barriers to implementing portfolio decisions lies in poor linkages to resource planning systems. According to a survey conducted by Accenture (Accenture, 2001), 90% of participating companies have developed resource planning systems for discovery since 1997. However these are rarely directly linked to the portfolio management system. In the majority of cases, the system is used only to track utilization rather than for forecasting or modelling resource usage. The result is that resource allocations often become disconnected from and out of alignment with the portfolio strategy, and supply and demand mismatches of resource arise .

Decision-making itself is a weakness in the industry. Approximately 50% of participants in the 2001 Accenture research considered decision-making within the discovery organization to be ineffective and generally slow. There is a clear need in many companies to define a hierarchy of decision types and identify the bodies best able to take them. One common approach has been to more fully empower project leaders to make decisions, particularly concerning the progression of projects along the discovery process. Success has been slow, caused in part by differing perceptions of who is ultimately responsible for each decision and also through variable interpretations of milestones and inconsistencies in transitioning from one to another. Perhaps most harmfully, there is often a shortage of true

project management skills among discovery scientists and of tools to support their decision-making processes.

As a result, there is a clear and widely recognized need for a flexible, reliable and user-friendly tool which assists the project planning process in discovery. This is what our research and this thesis aims to provide.

1.2 Drug Discovery and Development

New drugs begin in the laboratory with scientists, including chemists and pharmacologists, who identify cellular and genetic factors that play a role in specific diseases. They search for chemical and biological substances that interact with these biological markers and are likely to have drug-like effects. An investigational compound must be tested extensively in the laboratory to ensure it will be safe to administer to humans (Rodda et al., 2001).

Once the disease which is to be treated has been identified a large proportion of discovery projects proceed through the following sequence of operations.

- Before they can develop a medicine, bioscientists must first understand the biology of the disease they are trying to treat. Drugs usually act on either cellular or genetic chemicals in the body, known as targets, which are associated with the relevant disease. Scientists use a variety of techniques to identify and isolate individual targets to learn more about their functions and how they influence disease. To select targets most likely to be useful in the development of new treatments for disease, researchers compare each drug target to others based on their association with a specific disease and their ability to regulate biological and chemical compounds in the body. This is known as target identification and validation. Bioscientists then work out a hypothesis for the way in which a chemical intervention in the body's processes might achieve the desired therapeutic result, and devise tests using animal tissue or live animals in order to screen compounds for relevant activity. We refer to this first stage at which target is identified and validated, as *before screening*.

- Chemists select and/or synthesize compounds which are designed in the hope of finding relevant activity. These compounds are subject to screening tests. The initial screen is usually a high throughput robotic screen on tens of thousands of library compounds. When a compound with a sufficiently high level of activity has been found, other compounds with chemical structures which are similar are synthesized and tested. The original active compound is termed a *lead compound*, and the later compounds with similar structures form a *lead series* (LS). Thus, a lead compound is a lead in the investigation, like a piece of evidence in police work.
- Compounds from a chosen LS are tested in a variety of ways for therapeutic activity and toxicity in the hope of finding one which is sufficiently promising to be nominated as a *candidate drug* (CD), which will proceed to clinical trials. This process is known as *lead optimization*. This is the most resource-intensive stage in drug discovery, requiring considerable input from synthetic chemistry, modelling, disease biology and assay design. The early stages of the lead optimization process are usually focused on increasing antibodies' affinity to their target of interest and achieving desired selectivity. Selectivity requirements vary from target to target and, in particular, between different therapeutic areas. Once a molecule is identified, the next step is to check its ADMET (adsorption, distribution, metabolism, excretion, and toxicity) properties. The real challenge in lead optimization is balancing when to focus on each of the different desirable properties and deciding when to abandon an LS (Hubbard, 2006). Lead optimization typically takes 18-30 months, depending on the complexity of the target biology, the resource deployed and the chemistry of the LS.
- Any subsequent CD selected for development after the first CD is selected from an LS is called a *backup CD*. After finding the first CD from an LS we can look for a CD from a different LS, or look for backup CDs from the same LS, or not look for any more CDs. At most 20% of CDs are eventually marketable, so typically more compounds are screened while a CD is undergoing clinical trials, so that one or more backup compounds may in turn be selected as CDs. These may be from more than one LS, which reduces the risk that they will all fail for the same reason.

In summary, the discovery process can be modelled by five successive stages.

1. Before screening.
2. Looking for an LS.
3. Looking for a backup LS.
4. Optimizing an LS to find a CD.
5. Looking for a backup CD from an LS which has already provided one or more CDs.

To begin the drug development process, pharmaceutical companies need to apply for permission from the appropriate regulatory authorities to begin administration to healthy volunteers or patients. In addition, an institutional or independent review board (IRB) or ethical advisory board must approve the protocol for testing as well as the informed consent documents that volunteers sign prior to participating in a clinical study. An IRB is an independent committee of physicians, community advocates and others that ensures a clinical trial is ethical and the rights of study participants are protected.

Clinical trials also take place in a sequence of stages, usually referred to as phases. Progress to each phase depends on success at all previous phases. In each successive phase, increasing numbers of patients are tested.

Phase *I* clinical studies are designed to verify safety and tolerability of the CD in humans and typically take six to nine months. These are the first studies conducted in humans. A small number of subjects, usually from 20 to 100 healthy volunteers, take the investigational drug for short periods of time.

Phase *II* clinical studies are designed to determine effectiveness and further study the safety of the CD in humans. Depending upon the type of investigational drug and the condition it treats, this phase of development generally takes one to two years. Testing is conducted with up to several hundred patients suffering from the condition the investigational drug is designed to treat. This testing determines safety and effectiveness of the drug in treating the condition and establishes the

minimum and maximum effective doses. Most phase *II* clinical trials are randomly divided into groups, one of which receives the investigational drug, one of which gets a placebo containing no medication, and sometimes there is a third group that receives a current standard treatment to which the new investigational drug will be compared. Most phase *II* studies are double-blinded, meaning that neither the patient nor the research team evaluating the compound know whether a given patient is receiving the investigational drug or the placebo.

Phase *III* studies provide expanded testing of effectiveness and safety of an investigational drug, usually in randomized and blinded clinical trials. In phase *III* studies, safety and efficacy testing is conducted with several hundred to thousands of volunteer patients suffering from the condition the investigational drug treats. The trials are also required to establish the short-term and long-term safety/efficacy balance of the CD. Phase *III* studies on average take two to four years to complete, and approximately two-thirds of CDs under study are eventually submitted for regulatory approval.

At the conclusion of successful pre-clinical and clinical testing, an application is made for marketing approval to the relevant licensing authority. The application must present substantial evidence that the drug will have the effect it is represented to have when people use it under the conditions for which it is recommended. Obtaining approval to market a new drug frequently takes between six months and two years. Marketing approval gives the manufacturer the right to market the drug exclusively under a trademark name for the time remaining of the patent life of the drug. After the patent expires, other companies are permitted to manufacture and launch generic forms of the drug.

Once a drug has been approved and is being marketed, it is studied in a phase *IV* clinical trial to evaluate the side-effects of the new treatment that were not apparent in phase *III* trials, and to find out the long term risks and benefits. Such post-marketing clinical studies are mandatory in some jurisdictions.

Companies spend on average 27% of their R&D budgets on discovery. Clinical trials (phases *I*, *II* and *III*) together account for 54% of R&D budgets, while an additional 14% of R&D investment is spent on phase *IV* trials once the medicine has been approved by the regulatory authorities. In addition, about

5% of the R&D budgets is allocated to the approval process (EFPIA, 2010).

Further information on the drug discovery and development process can be found in Bergman and Gittins (1985), Gittins (1997), and Rodda et al. (2001).

1.3 Context of this Work

The whole R&D process is time-consuming, costly and involves great uncertainty. There is no aspect to which statistical methods cannot contribute. From screening chemicals for activity to forecasting sales, they have made significant contributions to the development of pharmaceutical R&D methodology. There is a broad range of literature on statistical issues in pharmaceutical R&D decision making, which include project evaluation, resource allocation, pharmacoeconomics, portfolio management, and the design and analysis of clinical trials. The journals *Pharmacoeconomics*, *R&D Management*, and *Pharmaceutical Statistics* are good sources for such literature. The earlier literature is reviewed by Bergman and Gittins (1985).

The book edited by Rodda et al. (2001) provides a general guide to statistical methods used in the pharmaceutical industry. The volume comprises contributions from more than 20 statisticians working in the industry. The drug development process is described from the viewpoint of statistical applications. The authors describe studies, data, methods, and results that finally lead to a new drug. The topics cover all aspects of the R&D process from pre-clinical safety assessment to dose finding, safety studies, large clinical trials, and other issues from analysis of health economic data to production and quality assessment.

Walley et al. (2004) examine how pharmacoeconomic analyses can be incorporated sensitively into the drug development process and can offer significant improvements in the discipline and mode of thinking in decision making relating to the use of therapeutic drugs. Miller (2005) discusses the use of five main pharmacoeconomic analytic techniques: clinical trial simulation, option pricing, investment appraisal, threshold analysis and value of information analysis. These can provide useful input into the design of clinical development programmes,

portfolio management and optimal pricing strategy. However all the techniques described are concerned with clinical rather than pre-clinical research.

Senn (1996, 1998) examines various aspects of portfolio management and project prioritization within the pharmaceutical industry. In these papers, probability of success, expected reward, and expected cost are identified as the key factors for evaluating a research project, and the strengths and weaknesses of the profitability (or Pearson) index (defined as the ratio between the expected net present value and the expected cost of the project) for project ranking are also discussed.

The paper by Poh et al. (2001) presents a good comparative study on the commonly used evaluation methods in R&D management for the past two decades. Based on the criteria proposed by the authors, their study shows that the scoring method is the most common method of R&D project evaluation because of its ability to deal with the many dimensions of R&D problems and its simplicity in formulation and use. They also show that Analytic Hierarchy Processing (AHP), when used as an R&D evaluation method, is a close second to the scoring method. AHP is capable of handling multiple objectives for R&D projects and decomposing the problem into a multilevel structure or hierarchy. The decision tree method, as the study reveals, is a popular method because of its accessibility and ability to deal with risk and uncertainty. However, these results are not confined to pharmaceutical R&D, and the methods described are not particularly relevant to pre-clinical research.

The application of real options analysis in the pharmaceutical industry is relatively new but has attracted a growing amount of interest. The theory of real options, introduced by Kester (1984) and Myers (1984), has revolutionized the way academics and practitioners think about investment projects by explicitly incorporating managerial flexibility into the analysis. This flexibility can represent a substantial part of the value of pharmaceutical projects, which are characterized by long time horizons and great uncertainty. Neglecting it can grossly undervalue these investments and induce a miss-allocation of resources. Real options models are based on the assumption that there is an underlying source of uncertainty, such as the price of a commodity or the outcome of a research project. Over time, the

outcome of the underlying uncertainty is revealed and managers can adjust their strategy accordingly. In fact, most pharmaceutical R&D decisions can be seen as options, for example, the decision to enter into a new contract, the decision to terminate a project, or the decision to sell off the patent to another company. The pharmaceutical giant Merck was one of the pioneers in the use of options analysis to value its investment opportunities, where it first used the Black-Scholes option pricing model to evaluate a proposed business relationship with a small biotechnology company. Its use of this technique received considerable attention in business publications (for example, Nichols, 1994; Sender, 1994; Thackray, 1995), and also facilitated the increasing applications of real options in this field. See, for example, Myers and Howe (1997), Childs and Triantis (1999), Loch and Bode-Greuel (2002), Bode-Greuel and Greuel (2005).

It is striking that the statistics literature is concerned almost exclusively with clinical, rather than pre-clinical, research. There is a lack of an established framework, both in the literature and in the pharmaceutical industry, for modelling pre-clinical research. The comment by Miller (2005) for example, that there is currently very little pharmacoeconomic planning in the early stages of drug R&D supports this view. This is where our work fits in. We have modelled the pre-clinical stages of research, in discussion with people working in the industry, in a much more detailed fashion than is to be found elsewhere. Our aim is to evaluate the profitability and facilitate the effective allocation of effort at the different stages of a pre-clinical R&D project.

1.4 Contribution of this Thesis

This thesis is a part of the research on developing models and software for improving the profitability of pharmaceutical R&D carried out by John Gittins and his co-workers in Oxford. Gittins (1996, 1997) proposed a stochastic model to investigate the profitability and resource allocation of a pharmaceutical research project, which is a prototype model for later work. In his model Gittins divides the drug discovery process into four stages and examines the relationships between the profitability index for a project and the number of scientists allocated to each

stage and the number of CDs selected from each LS. The main conclusion from the Gittins model is that larger teams than those which were then typical would lead to a substantial increase in the profitability of a pre-clinical project in many cases. Chen (2004) extends this model by incorporating internal rate of return as a second profitability criterion, and also considers the option value of waiting to start an investment. Yu and Gittins (2007) extend earlier models by considering the possibility of selecting CDs from more than one LS. An important conclusion is that it is often worth incurring the additional cost of optimizing more than one LS in the search for CDs, so as to reduce the risk of a string of CDs which all fail for similar reasons. Charalambous (2009) further develops the models by more detailed modelling of the pre-clinical stages and of clinical trials.

The starting point for this thesis is the model investigated by Charalambous (2009), which is based on the earlier work of Chen (2004) and Yu and Gittins (2007). Chen's model is the basis for the first version of the OPRRA software package, initially written in FORTRAN 90. OPRRA is the acronym for *Optimizing Pharmaceutical Research Resource Allocation*. Later versions of OPRRA have been written in Visual C++ in order to provide a more user-friendly interface. The overall aim of this thesis is to further advance OPRRA, partly through improving the models and solution algorithms and partly through improving the software itself.

All the previous work was based on a class of non-adaptive CD selection policies in a fixed probabilities setting. These policies have been termed (s, n) policies, and are described in detail in section 4.3. They do not adapt flexibly to the record of successes and failures at the various stages of discovery and clinical trial phases as the project progresses. They also only give point estimates of the expected reward, cost and profitability index. In this thesis I have addressed these deficiencies.

One major improvement is the introduction of a class of adaptive CD selection policies called *forwards induction* (FI) policies. The principle of forwards induction is that the next decision is always the one that maximizes the immediate expected reward rate, with no attempt to look further ahead. It thus defines a kind of myopic greedy algorithm. We shall see that FI policies are optimal for a sim-

plified CD selection problem, and are usually close to optimal more generally. FI policies may be implemented in an adaptive probabilities setting, which allows for learning about success rates as the project proceeds. This adaptive setting using Bayes theorem to update success probabilities is described in chapter 5.

I have developed simulation algorithms to investigate the probability distributions of total reward, total cost, profitability index, and required future resource allocations, both for a single pharmaceutical project under a given allocation plan and for a portfolio of projects. These simulation algorithms may be used to investigate both types of CD selection policy under different probabilities settings.

A further improvement is the adoption of option valuation for candidate drugs. This takes into account the uncertainty in the future income to be generated from the new drug by allowing it to fluctuate according to a jump-diffusion process, where obsolescence and catastrophic events that may suddenly reduce the value of the project are modelled as Poisson jumps. It also allows the option of not proceeding to clinical trials when the estimated income turns out to be smaller than anticipated.

These additional features³ have been incorporated into the current version of the OPRRA software. Two standard capital budgeting investment criteria are suggested for evaluating the profitability of research projects and to guide resource allocation decisions in OPRRA. These are the project's *profitability index* (PI) and its *internal rate of return* (IRR). These criteria require a complete accounting of costs, and put pressure on actively reducing them. They also lead to a more efficient allocation of resources and can take into account the sequential nature of the decisions concerned. Their use should lead to disciplined funding decisions at every major stage and ensure that money is spent only on the most attractive projects. The purpose of the OPRRA software is to assist pharmaceutical companies in taking decisions which improve profitability on CD selection and on rates of resource allocation during pre-clinical pharmaceutical research. Pre-clinical research is very different from clinical trials, on which there is a huge literature. Two of the important differences are that it mainly involves chemists and biologists rather than clinicians, and that there are fewer regulatory constraints on how

³Apart from the option valuation model, which will be built into OPRRA in due course.

it is conducted. OPRRA focuses on this relatively unstructured area. It should be used as an aid to inform the periodic meetings at which pre-clinical research projects are selected and prioritized (OPRRA user-guide, Gittins et al., 2011a).

1.5 Contents Overview

In the next chapter we introduce some aspects of the OPRRA model which are prerequisite to what is discussed in chapter 3. These include project stages, probabilities, the expected values of successive candidate drugs, and profitability criteria.

Chapter 3 introduces the forwards induction approach to CD selection problems in pharmaceutical research. We consider two selection problems, one of which is the simplified problem where we restrict our choice of CDs to only one LS. We show in section 3.2 that a forwards induction policy is optimal for this restricted case. The properties of FI policies in the more general problem are discussed in section 3.3. This chapter is based on a paper (Qu and Gittins, 2011) which is to be published in *Advances in Applied Probability*. The paper also includes some of the material from chapter 7.

Chapter 4 gives more detailed modelling of the OPRRA model. The work-rate/progress relationship, discounting and obsolescence are discussed. The various classes of allocation policy are summarized, together with a detailed introduction to the (s, n) class of CD selection policies. Finally, new features that have been added to OPRRA are described at the end of this chapter.

The feedback which comes in the form of the record of successes and failures of a project in the successive stages of discovery, and in the similar records for candidate drugs in each phase of clinical trials, is included in an adaptive probabilities model introduced in chapter 5. This model provides a means for incorporating learning from a research project's progress to date into the planning process. The assumptions of the model are described in section 5.1. We then explain how to estimate the hyper-parameters of the prior distributions in section 5.2. Finally in section 5.3 we discuss how FI policies can be implemented in this

adaptive probabilities setting.

Chapter 6 discusses the option valuation approach in evaluating pharmaceutical R&D investment opportunities. It exploits an analogy with the theory of options in financial markets which permits a much richer dynamic framework than was possible with the traditional theory of investment. A general overview of the theory of real options and its application in R&D decision making is given in section 6.1. The real options model is formulated in section 6.2, where we assume that the value of a new drug is affected by both economic uncertainty and the discontinuous arrivals of rare events from distinct sources. The option value of whether to proceed to clinical trials is calculated. The incorporation of this option value in the calculations of profitability for different CD selection policies is discussed in section 6.3 .

Chapter 7 provides some results comparing the various optimization procedures employed by OPRRA to optimize the profitability index. Conclusions arising from the results are drawn and commented on. Recommendations on the use of OPRRA are then given based on these findings.

OPRRA has been designed to find a resource allocation plan which maximizes the profitability of a pre-clinical research project. To understand how the various parameters affect profitability and the possible interactions between them, a thorough sensitivity analysis on the OPRRA parameters has been carried out in chapter 8. It turns out that interactions between parameters are weak, and outcomes are determined predominantly by a small subset of parameters.

Finally, chapter 9 offers some general conclusions and key findings, and outlines possible future extensions and alternative models.

Chapter 2

Stages, Probabilities, Rewards and Costs

In this chapter we introduce some preliminary modelling of OPRRA which is prerequisite to the forwards induction model we are going to discuss in the next chapter. More detailed modelling follows in chapter 4, which together with this chapter completes the account of the OPRRA model.

2.1 Stages of Drug Discovery

As described in section 1.2 , the process of drug discovery is modelled by five successive stages.

1. Before screening.
2. Looking for an LS.
3. Looking for a backup LS.
4. Optimizing an LS to find a CD.
5. Looking for a backup CD from an LS which has already provided one or more CDs.

Stages 1 and 2 are only carried out once and the project stops if either of them is unsuccessful. The probability that stage i reaches a successful conclusion, given the successful completion of the relevant earlier stages, we write as p_i ($q_i = 1 - p_i$), $i = 1, 3, 4, 5$.

The search for LS may produce more than one LS; we model this by defining $p_{2j} = \text{P}(\text{stage 2 produces } j \text{ LS} \mid \text{stage 1 is successful})$, ($j = 1, 2$), ($q_2 = 1 - p_{21} - p_{22}$). To limit the complexity of our model we restrict the possibility of more than two LS being discovered to the case $j = 2$, and do not explicitly model the possibility that a single application of stage 3 may produce more than one LS. Reasonably realistically we also assume that $p_5 = 1.0$; this means that if we choose to look for backup CDs we can find as many as required.

On successful completion of stage 2 or stage 3, stages 3 and 4 are both possible continuations. We assume that the continuation at this point includes stage 4, and that stage 3 also proceeds, simultaneously, unless either there has been an earlier failure at stage 3 or the decision has been taken not to seek further insurance LS. Stages 3, 4 and 5 can be repeated as often as necessary, except that if stage 3 fails we assume that it is not repeated, for the reason that any repetition would also result in failure. When stages 3 and 4 proceed in parallel we assume that effort is divided between the two stages so that they finish at the same time. This is reasonably realistic as it ensures that if we want to repeat stage 4 with another LS there is an LS available. Otherwise all stages are assumed to occur sequentially and in numerical order.

For each stage we have a fixed allocation u_i , measured in terms of the number of senior-scientists, and a fixed duration t_i , ($i = 1, 2, 3, 4, 5$). However, the time required to complete stage 4, given that a first LS has already been successfully optimized, is ρt_4 ($0 < \rho < 1$). An estimate is required for α , the cost per senior scientist per year. An important aspect of the overall resource allocation problem is to choose suitable values for the stage allocations u_i . The role of the effectiveness function governing the relationship between the values of u_i and t_i is discussed in section 4.1.

2.2 Clinical Trials and Beyond

For clinical trials we assume that CDs succeed or fail independently, and that the success probabilities conditional upon successful completion of previous phases are the same for different CDs. Let

$$p_I = P(\text{a CD is successful in phase } I),$$

$$p_{II} = P(\text{a CD is successful in phase } II \mid \text{success in phase } I),$$

$$p_{III} = P(\text{a CD is successful in phase } III \mid \text{success in phases } I \text{ and } II),$$

and $q_J = 1 - p_J$, $J = I, II, III$.

The costs and durations for phase J are c_J and t_J , respectively, $J = I, II, III$. We do not consider variations in either the design of clinical trials or in the allocation of resources to them.

An estimate is also required for W , the expected net present value of the first new drug from the project which completes phase III trials successfully. This is the total expected discounted value of all the cash flows attributable to the new drug if it was available immediately for production.

Since the value of a future sum of money is lower than that of the same sum of money available immediately, an exponential discount rate γ will be used to calculate the expected present value of future expenditure. We assume that sums of money are measured with the effect of inflation removed, thus representing actual purchasing power, and define γ to be the discount rate for money in real terms, corresponding to the weighted average cost of capital (WACC) for the company. Thus the present value of $\mathcal{L}e^{\gamma t}$ which becomes available after t years is $\mathcal{L}1$.

With discount rate γ , the cost in scientist years of employing u scientists for t years is $\int_0^t ue^{-\gamma s} ds = \frac{u}{\gamma}(1 - e^{-\gamma t})$. Given that the cost per scientist per year is α , define $K_i = \alpha \frac{u_i}{\gamma}(1 - e^{-\gamma t_i})$ as the total cost of stage i , $i = 1, 2, 3, 4, 5$.

For income there is additional discounting due to *obsolescence*. Definitions and detailed explanation for different types of obsolescence rates are set out in section 4.2. Denote by $\gamma_1 (> \gamma)$ the discount rate including obsolescence. The expected value of the new drug is discounted at the rate γ_1 whereas the expected expenditure is discounted at the rate γ .

2.3 The Expected Values of Successive Candidate Drugs

The expected values of any second or subsequent candidate drugs from the project depend on the sequence of lead series from which they are selected. To set out this part of our model we first define some notation, as follows.

- U : the therapeutic target,
- S_i : the i^{th} available LS ($i = 1, 2, \dots$),
- D_{ij} : the j^{th} CD from S_i ($i, j = 1, 2, \dots$),
- J_{ij} : the event $\{D_{ij}$ passes successfully through phase J trials $\}$ ($J = I, II, III$).

There are two versions of this part of the model: a *fixed probabilities* (FP) version, and an *adaptive probabilities* (AP) version, for which the success probabilities for the various stages and phases change during the project in response to the record to date of successes and failures. Chapter 5 discusses the AP version, whereas the rest of the thesis is concerned with the FP version.

Model FP

A *good* LS is an LS from which it is possible to get a successful CD. A CD is *successful* if it passes through clinical trials and becomes marketable. Different LS are good or bad independently and with the same probability conditional on the target being achievable. CDs are successful or unsuccessful independently and with the same probability conditional upon them coming from good LS.

Define:

- A : the event $\{\text{the target } U \text{ is in principle achievable}\}$,
- B_i : the event $\{S_i \text{ has no systematic defect which makes it impossible to reach the target } U\}$, or in brief $\{S_i \text{ is good}\}$,
- C_{ij} : the event $\{D_{ij} \text{ passes successfully through phases } I, II \text{ and } III \text{ trials}\}$, note that this means that $C_{ij} = III_{ij}$.

We assume that, for all i, j, m , and n , the following conditions hold.

$$P(B_i|\bar{A}) = P(C_{ij}|\bar{B}_i) = 0.$$

B_i, B_j are independent, given A , with $P(B_i|A) = P(B_j|A)$ ($i \neq j$).

C_{ij}, C_{mn} are independent, given B_i, B_m , with $P(C_{ij}|B_i) = P(C_{mn}|B_m)$ ($(i, j) \neq (m, n)$).

$$P(II_{ij}|\bar{I}_{ij}) = P(III_{ij}|\bar{II}_{ij}) = 0.$$

I_{ij} and II_{ij} are independent of A and B_i conditional on III_{ij} not occurring.

We define $p_a = P(A)$, $p_b = P(B_i|A)$, $p_c = P(C_{ij}|B_i) = P(III_{ij}|B_i)$. Writing p for the probability a CD successfully completes all three phases of clinical trials, it follows that $p = p_a p_b p_c = p_I p_{II} p_{III}$. Input values are required for p_I, p_{II}, p_{III} , p_a and p_b , allowing p_c to be calculated.

Note that in the fixed probabilities model we do not allow feedback from clinical trials, and therefore at the time that a CD is selected we do not know whether it, or any previously selected CD, is from a good or a bad LS.

The additional expected present value of any second or subsequent new drug is smaller than W because of competition from the first new drug, and any others which are earlier. We assume that the additional expected present value of the n^{th} new drug from a project is $\lambda^{n-1}W$ ($0 < \lambda < 1$). The factor λ is called the *sequential discount parameter*.

The following proposition is proved in Charalambous and Gittins (2008).

Proposition 2.1. *The additional expected value of a subsequent CD conditional on the set of previously selected CDs is $pW E[\eta^m|A]$, where $\eta = 1 - (1 - \lambda)p_c$, and $m =$ number of earlier CDs which are either from the same LS, or from another LS which is good.*

For example, if previous CDs have been from LS 1, 2, and 3, from which 2, 3 and 4 CDs, respectively, have previously been selected, and the present CD is from LS 2, then $P(m = 3) = (1 - p_b)^2$, $P(m = 5) = P(m = 7) = p_b(1 - p_b)$, $P(m = 9) = p_b^2$.

Model AP

Uncertainty about the probabilities p_3 and p_4 is expressed by prior probability distributions. These are converted into posterior distributions by Bayes theorem

as stages 3 and 4 are completed, either successfully or unsuccessfully, and hence the posterior expectations of p_3 and p_4 may be calculated. Note that for p_1 , p_{21} , and p_{22} there is no need to include Bayesian learning in this way because the corresponding events occur just once.

In this model the success probabilities for clinical trials vary between lead series. We define $p_{Ii} = P(I_{ij})$, $p_{Ji} = P(J_{ij}|K_{ij})$, ($J = II, III; K = J - 1$). For LS S_i the success probability p_{Ji} of CDs D_{ij} ($i, j = 1, 2 \dots; J = I, II, III$) is drawn from a beta distribution with the parameters $(f_J u_J, f_J(1 - u_J))$. The parameter f_J is known and u_J has a beta prior distribution with the parameters $(r_J m_J, r_J(1 - m_J))$ ($J = I, II, III$). With this hierarchical Bayes structure the posterior expectation for p_{Ji} may be calculated, as described by Consonni and Veronese (1995). In Model AP the events A and B_j are not modelled. This avoids considerable further complexity in the model, and is not too serious as Model AP is itself a way of representing the different characteristics of different LS.

2.4 Profitability Criteria

OPRRA is based on two alternative criteria for profitability, as follows.

- *Profitability Index (PI)* $= E[R]/E[C]$,

where $E[R]$ = expected discounted reward, and $E[C]$ = expected discounted cost. This is one of the standard criteria for the profitability of investment projects (see, for example, Brealey and Myers, 2000). The PI is a reward rate, and may be defined for a CD, or for an LS, as well as for the project as a whole. A portfolio of projects with high PI values leads to a high value for the total expected reward, or net present value (NPV), when the total capital available for investment is limited. Our focus is on the efficient use of resources before clinical trials start. For that reason $E[R]$ is calculated after subtraction of the expected costs of clinical trials, and only the costs incurred before clinical trials start are included in $E[C]$.

- *Internal Rate of Return (IRR)*.

This is the value of the exponential discount rate γ_I for which $E[R] - E[C]$ is equal to zero. In this calculation γ_I is allowed to vary, rather than being set equal to the cost of capital. As its name indicates, IRR is a measure of the average rate at which capital grows within a project or portfolio of projects. Since our goal is to improve resource allocation during discovery rather than during clinical trials, the variable discount rate γ_I is applied only to the pre-clinical stages of a project; during clinical trials discounting is with WACC.

Profitable projects are those for which $PI > 1.0$ and $IRR > WACC$, which are almost equivalent statements.

The profitability index is easier to calculate than the internal rate of return. Consequently solution algorithms are more fully developed for PI than for IRR. The FI policies which we introduce in the next chapter are a class of policies which select CDs according to their PI.

Chapter 3

A Forwards Induction Approach to Candidate Drug Selection

A forwards induction policy is a type of greedy algorithm first proposed by Gittins (1979) (also see Gittins et al., 2011b) for discounted Markov decision processes (MDPs). FI policies are straightforward to implement, and are optimal for a large class of models, especially in stochastic resource allocation. They are discussed in detail by Glazebrook and Gittins (1993) and Glazebrook (1995).

In this chapter we investigate the application of FI policies to the selection of CDs in pharmaceutical research. The measure which we shall use for profitability is the profitability index. The algorithms considered in this chapter seek to maximize the PI for the project. In all cases the expectations are conditional on the previous history of the project. Although our setup is not a discounted MDP, we shall proceed to use PI as the basis for defining FI policies. We shall show that FI policies are optimal for the simplified selection problem in which CDs are to be selected from at most one LS. Although FI policies are not always strictly optimal without the restriction, we shall see in chapter 7 that they still perform well. All algorithms described in this chapter are implemented in OPRRA. Numerical examples of FI policies are given in chapter 7, together with comparisons of the performance of the various plausible policies.

This chapter is based on a paper which will appear in *Advances in Applied*

Probability (Qu and Gittins, 2011).

3.1 FI Policies

We begin by defining some terminology. A *decision point* is a point at which a decision must be taken either to look for a CD to send to clinical trials, in which case we must also decide on the LS from which to look for a CD, or to stop. The *youngest* LS is the LS from which the minimum number of CDs have been selected. $PI(CD)$ is the PI for an additional CD from the youngest already optimized LS, conditional on the set of previously selected CDs. $PI(LS)$ is the PI for starting a new LS, conditional on the set of previously selected CDs. $PI(Proj)$ is the PI for the project as a whole, maximized over all selection policies.

There are two versions of $PI(LS)$, depending on whether stage 3 is to be run alongside stage 4. Choosing between them involves evaluating PI over the CDs chosen from two successive LS. The simpler version, without a concurrent stage 3, is the PI for the LS up to the number k of CDs which gives the maximum PI; k has the property that the k^{th} CD is the last CD from the LS for which the PI is higher than the PI for the whole LS up to that point. A full description of the calculations of $PI(CD)$ and $PI(LS)$ is given in Appendix B.

In an FI policy, the sequence of CDs sent for clinical trials is determined with reference to a *reference PI*, $PI(ref)$. The purpose of the reference PI is that it works as an approximation to the optimal PI for the project, and therefore screens out CDs and LS which are likely to reduce the overall PI. FI policies are more flexible than (s, n) policies, and might be expected to perform better for suitable values of $PI(ref)$.

The FI selection algorithm is as follows.

FI Algorithm

1. Decide a reference PI, $PI(ref)$.
2. At each decision point, compare $PI(CD)$, $PI(LS)$, and $PI(ref)$.

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

3.
 - If $PI(CD)$ is the biggest, take an additional CD from the youngest optimized LS and go back to step 2.
 - If $PI(LS)$ is the biggest, optimize a new LS.
 - (a) If the attempt to find a CD in this LS is successful take the PI-maximizing number of CDs from this LS, then go back to step 2.
 - (b) If the attempt is not successful, go back to step 2.
 - If $PI(ref)$ is the biggest, stop.

3.2 FI for the Simplified Problem

In this simplified problem, we allow CDs to be selected from at most one LS. We also assume that $p_{22} = 0$ and that whenever stage 4 is carried out stage 3 is carried out in parallel. This is for simplicity. The algorithm whose properties are described in chapter 7 ensures that stage 3 is only carried out if there is a possibility that the resulting LS will be used.

The forwards induction algorithm proceeds as before except that after the first LS to be successfully optimized there is no longer the option to optimize a new LS. The main conclusion of this section is that for this modified problem FI is optimal. We start with some further simplified problems and solution algorithms which provide useful building blocks.

Note that in the notation that follows in this section, $|S|$ denotes the cardinality of the set S (i.e, the number of members of S). $S + x$ is a short hand notation for $S \cup \{x\}$, and $S - x$ is a short hand notation for $S \setminus \{x\}$.

3.2.1 Infrastructure

Problem 3.1. For $j = 1, 2, \dots$, let r_j be real numbers and c_j be positive real numbers and let $q_j = \frac{r_j}{c_j}$, where $q_j \geq q_{j+1}$ when $j > 1$. Define $R(k) = \sum_{j=1}^k r_j$, $C(k) = \sum_{i=1}^k c_i$, $Q(k) = \frac{R(k)}{C(k)}$, $Q = \sup_k Q(k)$. The problem is either to show

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

that $Q \leq 0$ or to find k such that $Q(k) = Q$.

The following proposition on ratios of real numbers is the key to solving this and similar problems.

Proposition 3.1. (a) For real numbers x_1, x_2, y_1, y_2 with $y_1 > 0$ and $y_2 > 0$, the following three statements are equivalent: $\frac{x_1}{y_1} < \frac{x_2}{y_2}$, $\frac{x_1}{y_1} < \frac{x_1+x_2}{y_1+y_2}$, and $\frac{x_1+x_2}{y_1+y_2} < \frac{x_2}{y_2}$.
 (b) The same is true with $<$ replaced by $=$ throughout, or by \leq .

Proof. (a) Assume $\frac{x_1}{y_1} < \frac{x_2}{y_2}$, so that $x_1y_2 < x_2y_1$. Adding x_1y_1 to both sides of this inequality we get $x_1y_2 + x_1y_1 < x_2y_1 + x_1y_1$ and hence $\frac{x_1}{y_1} < \frac{x_1+x_2}{y_1+y_2}$. Similarly, adding x_2y_2 to both sides of the inequality instead of x_1y_1 gives us $\frac{x_1+x_2}{y_1+y_2} < \frac{x_2}{y_2}$. Now assume $\frac{x_1}{y_1} < \frac{x_1+x_2}{y_1+y_2}$, thus $x_1y_1 + x_2y_1 < x_1y_1 + x_1y_2$, so that $x_2y_1 < x_1y_2$ and thus $\frac{x_1}{y_1} < \frac{x_2}{y_2}$. A similar proof shows that this also follows from the inequality $\frac{x_1+x_2}{y_1+y_2} < \frac{x_2}{y_2}$, and so the three original inequalities are equivalent.

(b) The proofs are almost identical. \square

For problem 3.1 there is a straightforward solution algorithm, as follows from lemmas 3.1 and 3.2.

Lemma 3.1. For problem 3.1, if $q_{k+1} \left\{ \begin{array}{l} > \\ = \\ < \end{array} \right\} Q(k)$ for some k , then $Q(k + 1) \left\{ \begin{array}{l} > \\ = \\ < \end{array} \right\} Q(k)$.

Proof. If $q_{k+1} > Q(k)$, then $\frac{r_{k+1}}{c_{k+1}} > \frac{R(k)}{C(k)}$. Using proposition 3.1(a), $\frac{r_{k+1}+R(k)}{c_{k+1}+C(k)} > \frac{R(k)}{C(k)}$, hence $Q(k + 1) > Q(k)$. The other two parts of the lemma also follow, in similar fashion, from proposition 3.1. \square

Lemma 3.2. Consider problem 3.1.

- (i) There is at most one k for which $q_k \geq Q(k) > q_{k+1}$. For this k , $Q(k) = Q$.
- (ii) If $Q > 0$ and $q_j < 0$ for some j then there is a k with the above property.
- (iii) If, for some k , $Q(j) \leq 0$, $1 \leq j \leq k$, and $q_{k+1} \leq 0$, then $Q \leq 0$.

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

Proof. If $q_{k+1} \geq Q(k)$, then $Q(k+1) \geq Q(k)$. So either (a) $q_{k+1} \geq Q(k)$, and $Q(k+1) \geq Q(k) \forall k$, or (b) $\exists m$ such that $q_{m+1} < Q(m)$, in which case we define k to be the smallest such m . In case (b) it follows that $q_{n+1} \geq Q(n) \forall n < k$, and $q_{k+1} < Q(k)$. The first of these statements implies that $Q(n+1) \geq Q(n) \forall n < k$. From the second statement it follows from lemma 3.1 that $q_{k+1} < Q(k+1) < Q(k)$. This in turn implies that $q_{k+2} < Q(k+1)$, and hence that $q_{k+2} < Q(k+2) < Q(k+1)$, and the argument extends by induction on n to show that $q_{n+1} < Q(n+1) < Q(n), \forall n > k$.

It follows that in case (a) $\lim_{k \rightarrow \infty} Q(k) = Q$, and in case (b) $Q(k) = Q$. For case (b) it is straightforward to check that the chosen k is the unique value described in the statement of the lemma. In case (a) there is no k for which $q_k \geq Q(k) > q_{k+1}$, and part (i) of the lemma is proved. For part (ii) of the lemma note that under the given conditions it is impossible that $Q(k+1) \geq Q(k) \forall k$, and therefore case (a) does not occur. Part (iii) of the lemma follows from proposition 3.1 and the fact that $q_j \geq q_{j+1}, j > 1$. \square

From lemma 3.2, and with the additional assumption that $q_j < 0$ for some j , it follows that problem 3.1 may be solved by algorithm 3.1.

Algorithm 3.1. Compute in succession $Q(k), k = 1, 2, \dots$, stopping either when $q_k \geq Q(k) > q_{k+1}$, in which case $Q(k) = Q$, or when $Q(j) \leq 0, 1 \leq j \leq k$, and $q_{k+1} \leq 0$, in which case $Q \leq 0$.

The following variants of problem 3.1 are also of interest.

Problem 3.1I. For $x \in$ the finite index set I , which includes 1, let r_x be a real number, c_x be a positive real number, and let $q_x = \frac{r_x}{c_x}$. Let S denote a subset of I , $R(S) = \sum_{x \in S} r_x$, $C(S) = \sum_{x \in S} c_x$, $Q(S) = \frac{R(S)}{C(S)}$ and $Q = \max_{S \in U} Q(S)$, where $U = \{S \subset I, 1 \in S\}$. The problem is to find $S \in U$ such that $Q(S) = Q$.

Problem 3.1W. This is the same as problem 3.1I except that the class of index subsets U is replaced by a sub-class $W \subset U$.

Problem 3.1V. The definition is motivated by the following two lemmas. In these lemmas V denotes the class of index subsets of the form $\{x : q_x > \xi\} + 1$ for some ξ .

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

Lemma 3.3. For any $S \in U \setminus V \exists S' \in U$, such that either $Q(S') > Q(S)$ or $Q(S') = Q(S)$ and $|S'| = |S| + 1$.

Proof. Suppose $S \in U \setminus V$. Then $\exists x \in S - 1$ and $y \in I \setminus S$ such that at least one of the following statements is true.

(i) $q_y > q_x \geq Q(S)$. (ii) $q_y > q_x < Q(S)$. (iii) $q_y = q_x > Q(S)$. (iv) $q_y = q_x < Q(S)$. (v) $q_y = q_x = Q(S)$.

In cases (i) and (iii), let $S' = S + y$. In cases (ii) and (iv), let $S' = S - x$. In all these cases it follows from proposition 3.1 that $Q(S') > Q(S)$. In case (v), let $S' = S + y$ and we have $Q(S') = Q(S)$ and $|S'| = |S| + 1$. \square

Lemma 3.4. For any $S \in U \exists S' \in V$, with $Q(S') \geq Q(S)$.

Proof. We shall suppose the converse of the lemma and show that this leads to a contradiction. Thus suppose $S \in U$ and $\nexists S' \in V$ with $Q(S') \geq Q(S)$, so that in particular $S \notin V$. From lemma 3.3 it follows that $\exists S_1 \in U$ such that either $Q(S_1) > Q(S)$ or $Q(S_1) = Q(S)$ and $|S_1| = |S| + 1$. By our assumption $S_1 \notin V$, so again applying lemma 3.3 it follows that $\exists S_2 \in U$ such that either $Q(S_2) > Q(S_1)$ or $Q(S_2) = Q(S_1)$ and $|S_2| = |S_1| + 1$. Again, by assumption $S_2 \notin V$, and the argument may be repeated to define an infinite sequence with no repeats $S_0 (= S), S_1, S_2, \dots$ such that $S_\tau \in U \setminus V \forall \tau$. However this is impossible as the index set I is finite, as is therefore the class of subsets U . This completes the proof of the lemma. \square

It follows from lemma 3.4 that if W is a class of subsets of I with $V \subset W \subset U$ then $\max_{S \in W} Q(S) = \max_{S \in V} Q(S) = Q$. Problem 3.1V is the same as problem 3.1W with the restriction that $W \supset V$.

Problems 3.1I and 3.1V may be transformed into problems of the form of problem 3.1, as follows. For each triple (r_x, c_x, q_x) change the index x to one of the first $|I|$ integers, leaving the values of the triple unchanged, with the new index values chosen so that r_1, c_1 and q_1 are unchanged and $q_i \geq q_{i+1}, i > 1$. Let q_i be large and negative for $i > |I|$. To complete the solution we carry out algorithm 3.1, and then reverse the index value changes.

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

The next step in this sequence of preliminary problems is to introduce a second subscript, which allows us to model different lead series.

Problem 3.2. For $(i, j = 1, 2, \dots)$, let r_{ij} be real numbers, c_{ij} be positive real numbers, $q_{ij} = \frac{r_{ij}}{c_{ij}}$, where $q_{i1} > q_{i+1,1}$ for $i > 1$, $q_{ij} > q_{ij+1}$ for $(i, j) \neq (1, 1)$, $q_{i1} \rightarrow q < 0$ as $i \rightarrow \infty$, and $q_{ij} \rightarrow q_i < 0$ as $j \rightarrow \infty \forall i$.

Define $R_i(k) = \sum_{j=1}^k r_{ij}$, $C_i(k) = \sum_{j=1}^k c_{ij}$, $R(k_1, k_2, \dots, k_n) = \sum_{i=1}^n R_i(k_i)$,
 $C(k_1, k_2, \dots, k_n) = \sum_{i=1}^n C_i(k_i)$, $Q(k_1, k_2, \dots, k_n) = \frac{R(k_1, k_2, \dots, k_n)}{C(k_1, k_2, \dots, k_n)}$,
 $Q = \sup_A Q(k_1, k_2, \dots, k_n)$, where $A = \{n > 0; k_i > 0, i = 1, 2, \dots, n\}$.

The problem is either to show that $Q \leq 0$ or to find $n^*, k_1^*, k_2^*, \dots, k_{n^*}^*$ such that $Q = Q(k_1^*, k_2^*, \dots, k_{n^*}^*)$.

Lemma 3.5. Problem 3.2 may be regarded as an example of problem 3.1V, with the index pair 11 in place of 1.

Proof. To establish this equivalence we first note that, except that the index set $I = \{ij : 1 \leq i, j\}$ is infinite, problem 3.2 is equivalent to problem 3.1W, with 11 (to be read as ‘one-one’ rather than as ‘eleven’) in place of 1, and W defined by the constraints that if $S \in W$ and $ij \in S$ then $k1 \in S \forall k \leq i$, and $il \in S \forall l \leq j$. To show that problem 3.2 is also an example of problem 3.1V (except for the infinite index set) we need to show that for any ξ the index subset $S_\xi = \{ij : q_{ij} > \xi\} + 11$ belongs to the sub-class W .

For any given ξ this is trivially true if $q_{ij} \leq \xi \forall ij \neq 11$. Suppose that $ij (\neq 11) \in S_\xi$, so that $q_{ij} > \xi$. It is sufficient to show that $q_{xy} \geq q_{ij}$ for every index pair xy which $\in S_\xi$ because $ij \in S_\xi$, and because of the constraints defining W . If $i = 1$ we need to show that $1k \in S_\xi$, $1 \leq k \leq j$. This is true since $11 \in S_\xi$ and $q_{12} \geq q_{13} \geq \dots \geq q_{1j} > \xi$. If $i > 1$ we need to show that $k1 \in S_\xi$, $k \leq i$, and $il \in S_\xi$, $l \leq j$. Again this is true since $11 \in S_\xi$ and $q_{21} \geq q_{31} \geq \dots \geq q_{i1} \geq q_{i2} \geq \dots \geq q_{ij} > \xi$.

To solve problem 3.2 our task is either to show that $Q \leq 0$, which means to show there is no $S \in V$ with $Q(S) > 0$, or to construct the set S of index pairs which maximizes $Q(S)$ for $S \in V$. For these purposes any index pair $ij (\neq 11)$ for which $q_{ij} \leq 0$ may be excluded. This is because if $Q(S) > 0$ for some index set S then $Q(S + ij) < Q(S)$. We also need to note that the restraints on W do

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

not cause the exclusion of ij to necessitate the exclusion of any index pair xy for which $q_{xy} > 0$.

Finally we note that it follows from the inequalities satisfied by q_{ij} , q_i and q that the number of index pairs ij for which $q_{ij} > 0$ is finite. This means that in excluding every index pair ij ($\neq 11$) with $q_{ij} \leq 0$ we reduce the problem to one with a finite index set, and therefore to an example of problem 3.1V, completing the proof of the lemma. \square

Thus we have the following algorithm for solving problem 3.2.

Algorithm 3.2. (i) Reduce the index set to the finite set I consisting of index pairs ij with $q_{ij} > 0$ plus the pair 11.

(ii) Transform the resulting problem 3.1V into the form of problem 3.1 by changing the index set.

(iii) Use algorithm 3.1 to solve the resulting problem 3.1.

(iv) If $Q \leq 0$ for problem 3.1 the same is true for problem 3.2.

(v) If $Q > 0$ for problem 3.1, the same is true for problem 3.2, and we may reverse the index set changes to obtain the maximizing n^* , k_1^* , k_2^* , \dots , k_n^* for problem 3.2.

3.2.2 Solution to the Simplified Problem

Returning to our candidate drug selection problem, note that by restricting the selection of CDs to at most one LS at each decision point we have ensured that there is at most one way in which a project may be continued. After stages 1 and 2, this is by carrying out stages 3 and 4 in parallel if there has not yet been a successful stage 4 or failure at stage 3, and by selecting a new CD from the current LS if there has been a successful stage 4. The two possible continuations are never both available. We recall that: stage 3 is a search for an additional LS and maybe carried out repeatedly until at some point it fails, after which it may not be repeated; stage 4 is the attempt to find a first CD in an LS, and is successful with probability p_4 , after which as many CDs as required may with certainty be found in the same LS.

With the above restriction we can write down formulae for the expected re-

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

ward (after subtracting the cost of clinical trials) and expected cost (before clinical trials) for each successive CD to be selected. Define r_{ij} , c_{ij} and $q_{ij} = \frac{r_{ij}}{c_{ij}}$ to be the expected reward, expected cost and PI for the j^{th} CD selected from LS i . It is a consequence of the fact that there is an unique sequence of events leading to the selection of the j^{th} CD from LS i that the unconditional expected reward and cost are the same as the expected reward and cost conditional on previous history, apart from a common factor. We may therefore, and we shall, define r_{ij} and c_{ij} to be unconditional expectations without changing the value of q_{ij} . The reason for the overlaps of notation with problem 3.2 will soon become clear.

From the description given in chapter 2 it follows that

$$r_{ij} = p_1 p_{21} (p_3 q_4)^{i-1} p_4 [pW e^{-\gamma_1(t_1+t_2+it_4+(j-1)t_5+t_I+t_{II}+t_{III})} \gamma^{j-1} - e^{-\gamma(t_1+t_2+it_4+(j-1)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III})], \quad i, j \geq 1;$$

and that

$$c_{ij} = \begin{cases} K_1 + p_1 e^{-\gamma t_1} K_2 + p_1 p_{21} e^{-\gamma(t_1+t_2)} (K_3 + K_4), & i = j = 1; \\ p_1 p_{21} (p_3 q_4)^{i-1} e^{-\gamma(t_1+t_2+(i-1)t_4)} (K_3 + K_4), & i \geq 2, j = 1; \\ p_1 p_{21} (p_3 q_4)^{i-1} p_4 e^{-\gamma(t_1+t_2+it_4+(j-2)t_5)} K_5, & \forall i, \text{ and } j > 1. \end{cases}$$

With r_{ij} and c_{ij} as above, and using the fact that $\gamma_1 > \gamma$, it is easy to check that all the assumptions of problem 3.2 hold, except that it is not necessarily true that $q_{i1} > q_{i2}$, $i > 1$. We shall use all the notation for problem 3.2 to refer also to our CD selection problem, which we shall call problem 3.3. A deterministic CD selection problem, for example, may be expressed in the form of a sequence $A = \{n > 0; k_i > 0, i = 1, 2, \dots, n\}$.

A deterministic policy is one for which the decision whether to continue or to stop depends deterministically on the number of LS which we have so far tried to optimize, and on the number of CDs so far selected. We could also consider randomized CD selection policies, for which each continue/stop decision is determined randomly. However it is an easy consequence of proposition 3.1 that if the maximum PI for a project is attainable then it is attained by a deterministic policy.

As for problems 3.1 and 3.2 we shall present an algorithm for problem 3.3 which either shows that $Q \leq 0$, so that no CD selection policy produces a positive

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

PI, or finds a CD selection policy A which attains the optimum PI. As for problem 3.2, our task is to maximize $Q(S)$ over the allowed sets S of index pairs. Each included index pair ij now corresponds to the selection of the j^{th} CD from LS i .

As for problem 3.2 we may exclude from S any index pair ij ($j > 1$) with $q_{ij} \leq 0$. By a similar argument we may exclude any index pair $i1$ ($i > 1$) if $q_{i1} \leq 0$ and $q_{i2} \leq 0$. These exclusions mean that we have again reduced the index set to a finite set. However there remains the possibility that $q_{i1} \leq q_{i2} > 0$, unlike problem 3.2, so we have not yet reduced the problem to an example of problem 3.1V.

To proceed further we next note that for any LS i the conditions for problem 3.1 are satisfied with $(r_i, c_i, q_i) = (r_{ij}, c_{ij}, q_{ij})$ and $Q(k) = Q_i(k) = R_i(k)/C_i(k)$, the PI for LS i when k CDs are selected. Also $q_{ij} < 0$ for large j . We may therefore use algorithm 3.1 to either find k_i^{**} such that $Q(k_i^{**}) = \sup_k Q_i(k)$, which we denote by Q_i , or show that $Q_i \leq 0$. Note that it follows from the definitions that $Q_i(k) > Q_{i+1}(k)$, $i > 1$ and $\forall k$, and hence that $Q_i > Q_{i+1}$. Note too that since the number of ij pairs for which $q_{ij} > 0$ is finite it follows that $Q_i > 0$ for at most a finite set of i values.

Since the class of allowed sets S is finite, if $Q > 0 \exists n, k_1, k_2, \dots, k_n$ (all > 0) such that $Q = Q(k_1, k_2, \dots, k_n)$. We define $n^*, k_1^*, k_2^*, \dots, k_n^*$ to be chosen in that order to be the largest values of n, k_1, k_2, \dots, k_n for which this is true. To further simplify problem 3.3 we shall use the following lemma.

Lemma 3.6. $k_i^* \geq k_i^{**}$, $2 \leq i \leq n^*$.

Proof. Note first that $Q_{n^*}(k_{n^*}^*) \geq Q$. If $Q_{n^*}(k_{n^*}^*) < Q$ it follows from the fact that $Q(k_1^*, k_2^*, \dots, k_{n^*}^*) = Q$ that $Q(k_1^*, k_2^*, \dots, k_{n-1}^*) > Q$, using proposition 3.1(a), which contradicts the definition of Q . Thus $Q_{n^*}(k_{n^*}^*) \geq Q$.

Now suppose that $k_i^* < k_i^{**}$ and $1 < i \leq n^*$. We have

$$q_{ik_i^*+1} \geq Q_i(k_i^{**}) \geq Q_{n^*}(k_{n^*}^{**}) \geq Q_{n^*}(k_{n^*}^*) \geq Q.$$

The first inequality follows from proposition 3.1 and the definition of k_i^{**} . The second inequality is equivalent to $Q_i \geq Q_{n^*}$, which is true since Q_i is a decreasing function of i for $i > 1$. The third inequality is by the definition of $k_{n^*}^{**}$. Thus

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

$q_{ik_i^*+1} \geq Q$, so that $Q(k_1^*, k_2^*, \dots, k_i^* + 1, \dots, k_n^*) \geq Q$. This contradicts the assumed maximality of k_i^* and completes the proof. \square

Now note that it follows from proposition 3.1(a) and the definition of Q_i that $Q_i \geq q_{ik_i^{**}+1}$. It thus follows from lemma 3.6 that problem 3.3 may be transformed into an equivalent problem in the form of problem 3.2 as follows:

$r'_{1j} = r_{1j}$, $c'_{1j} = c_{1j}$, $q'_{1j} = q_{1j}$, $\forall j$;
 $r'_{i1} = R_i(1, 2, \dots, k_i^{**})$, $c'_{i1} = C_i(1, 2, \dots, k_i^{**})$, $q'_{i1} = \frac{r'_{i1}}{c'_{i1}}$, and the index pair ij is replaced by $i(j - k_i^{**} + 1)$, $i > 1$ and $j > k_i^{**}$.

We can now write down an algorithm which solves problem 3.3.

Algorithm 3.3. (i) Use algorithm 3.1 to compute k_i^{**} and Q_i , $i = 2, 3, \dots$, stopping when $Q_i < 0$, at which point set $n^{**} = i - 1$.

(ii) For all index pairs ij with $i > n^{**}$ replace q_{ij} by $-M$, where M is large.

(iii) Transform the problem into one of the form of problem 3.2 as just described.

(iv) Use algorithm 3.2 to solve the resulting problem 3.2.

(v) Express the solution in terms of the index set for problem 3.3.

We conclude section 3.2 by noting that the output from algorithm 3.3 defines the decision to be taken at each decision point in the procedure for selecting candidate drugs, and that the policy defined in this way is equivalent to an FI policy with $PI(ref) = Q$, the optimal PI for the project under the restrictions of this section.

3.3 FI for the General Problem

In the general selection problem, CDs are to be selected from an unrestricted number of LS and as a result the expected reward and expected cost for each successive CD to be selected are no longer defined by a deterministic sequence of successively selected CDs as they were in the simplified problem.

Unlike the restricted problem considered in the previous section it is not in general true to say that an FI policy is itself an optimal policy. This is shown by a counter example. Lemma 3.7, on the other hand, provides encouragement for

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

the view that FI policies may be close to optimal, while lemma 3.8 narrows the choices which need to be considered at each step in an FI policy.

Let r_{ij} and c_{ij} now be the expected reward and cost from the j^{th} CD in LS i , conditional on the set of previously selected CDs. Note that this is different from section 3.2, where our notation referred to unconditional expectations. Here r_{ij} and c_{ij} are random variables depending on time, and on the numbers of CDs from previous LS. Let t denote the time at which the j^{th} CD from LS i is selected.

Counter Example. *This shows that if $PI(CD) = PI(LS) > PI(ref) = PI(Proj)$ at a decision point P in an FI policy it does not in general follow that taking an additional CD and trying to optimize a new LS are both decisions leading to the overall maximum PI value for the project, $PI(Proj)$.*

Proof. Let a and b be the expected reward and cost for the additional CD, and let A and B be the expected reward and cost for the new LS. Suppose that $\frac{a}{b} = \frac{A}{B}$. Let $PI(ref) = \frac{a}{b} - \epsilon$, where $\epsilon (> 0)$ is sufficiently small to ensure that after any of the three possible ways in which the project might be continued from point P , $PI(ref) > \max(PI(CD), PI(LS))$, so that the project then terminates under an FI policy. The three possible continuations are

1. an additional CD is chosen from an optimized LS,
2. we try successfully to find a CD from a new LS and then select the number of backup CDs from that LS which maximizes the PI for that LS,
3. we try unsuccessfully to find a CD from a new LS.

It is easy to construct numerical cases satisfying these conditions.

Let R and C be the expected reward and expected cost under an FI policy, excluding the contributions which arise after reaching the decision point P , and let θ be the probability that point P is reached. Thus the overall PI from an FI policy is $(R + \theta A)/(C + \theta B)$ if at P we look for a CD in a new LS, and $(R + \theta a)/(C + \theta b)$ if at P we select a CD from an already optimized LS. We have assumed that $\frac{R}{C} < \frac{A}{B} = \frac{a}{b}$, and it therefore follows from proposition 3.1 that if $A > a$ then $(R + \theta A)/(C + \theta B) > (R + \theta a)/(C + \theta b)$, completing the proof. \square

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

Lemma 3.7. *In a sequence of backup CDs from the same LS the PI is decreasing.*

Proof. By definition, $q_{ij+1} = \frac{r_{ij+1}}{c_{ij+1}}$ is the PI for the j^{th} backup CD in LS i . Using proposition 2.1 we have $r_{ij} = e^{-\gamma_1 t} \eta^j v_a - e^{-\gamma t} v_b$, $c_{ij} = e^{-\gamma t} K_5$. Here

$$v_a = pW e^{-\gamma_1(t_5+t_I+t_{II}+t_{III})} E[\eta^N],$$

N = number of previously selected CDs from good LS other than LS i , and

$$v_b = e^{-\gamma t_5} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}).$$

Thus the PI for the j^{th} backup CD in LS i is

$$\begin{aligned} q_{ij+1} &= \frac{r_{ij+1}}{c_{ij+1}} \\ &= \frac{e^{-\gamma_1 t} \eta^j v_a - e^{-\gamma t} v_b}{e^{-\gamma t} K_5} \\ &= \frac{e^{-(\gamma_1 - \gamma)t} \eta^j v_a - v_b}{K_5}, \end{aligned}$$

which is a decreasing function in j , $j \geq 1$. This completes the proof. \square

Lemma 3.7 means that within each optimized LS there is no reason for taking account of the PIs from later CDs when deciding whether or not to select one more CD. To this extent, then, the forwards induction greedy principle is correct.

Our next lemma reduces the computational complexity of implementing an FI policy. Recall that by youngest LS we mean an LS from which the minimum number of CDs have so far been selected.

Lemma 3.8. *The expected value of the next CD from an already optimized LS is maximized by choosing it from the youngest optimized LS.*

Proof. It will be sufficient to show that, for any two optimized LS, we will get a higher expected value if we choose a CD from the younger LS.

Consider selection from two LS, i and j , at time t . Assume that we have so far selected x_i CDs from LS i and x_j CDs from LS j . Let ω_{ij} denote the number

Chapter 3 A Forwards Induction Approach to Candidate Drug Selection

of CDs which have been selected from good LS other than i and j , and let

$$\begin{aligned} m_i &= \omega_{ij} + x_j \mathbf{I}(j) + x_i, \\ m_j &= \omega_{ij} + x_j + x_i \mathbf{I}(i), \end{aligned}$$

where $\mathbf{I}(k)$ is the indicator random variable for the event {LS k is good}, $k = i, j$. Thus, using proposition 2.1, and dropping the condition $|A$ from the notation, which applies to every expectation in this proof,

$$\begin{aligned} r_{ix_{i+1}} &= e^{-\gamma_1 t} E[\eta^{m_i}] v_a - e^{-\gamma t} v_b, \\ r_{jx_{j+1}} &= e^{-\gamma_1 t} E[\eta^{m_j}] v_a - e^{-\gamma t} v_b, \end{aligned}$$

where

$$\begin{aligned} v_a &= pW e^{-\gamma_1(t_5+t_I+t_{II}+t_{III})}, \\ v_b &= e^{-\gamma t_5} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}), \end{aligned}$$

so that

$$\begin{aligned} \frac{r_{ix_{i+1}} - r_{jx_{j+1}}}{e^{-\gamma_1 t} v_a} &= E[\eta^{m_i}] - E[\eta^{m_j}] \\ &= E[\eta^{m_i} - \eta^{m_j}] \\ &= E[\eta^{\omega_{ij} + x_j \mathbf{I}(j) + x_i} - \eta^{\omega_{ij} + x_j + x_i \mathbf{I}(i)}] \\ &= E[\eta^{\omega_{ij}} (\eta^{x_j \mathbf{I}(j) + x_i} - \eta^{x_j + x_i \mathbf{I}(i)})]. \end{aligned}$$

Since ω_{ij} is independent of the events {LS i is good} and {LS j is good}, conditional on the event that the target is achievable, we have

$$E[\eta^{\omega_{ij}} (\eta^{x_j \mathbf{I}(j) + x_i} - \eta^{x_j + x_i \mathbf{I}(i)})] = E[\eta^{\omega_{ij}}] E[\eta^{x_j \mathbf{I}(j) + x_i} - \eta^{x_j + x_i \mathbf{I}(i)}],$$

so that

$$\begin{aligned} E[\eta^{x_j \mathbf{I}(j) + x_i} - \eta^{x_j + x_i \mathbf{I}(i)}] &= p_b \eta^{x_i + x_j} + (1 - p_b) \eta^{x_i} - p_b \eta^{x_i + x_j} - (1 - p_b) \eta^{x_j} \\ &= (1 - p_b) (\eta^{x_i} - \eta^{x_j}). \end{aligned}$$

This is positive if $x_i < x_j$, since $p_b < 1$ and $0 < \eta < 1$. Hence, $r_{ix_{i+1}} - r_{jx_{j+1}} > 0$ if $x_i < x_j$, as required. This completes the proof. \square

Chapter 4

More OPRRA Modelling

This chapter completes the account of the current OPRRA model which we began in chapter 2. Much of this is the same as the version described by Charalambous (2009). The main changes from that version are described at the end of this chapter. The current version of the OPRRA software is written in Visual C++.

4.1 The Work-Rate/Progress Relationship

The model assumes that the time taken to complete each stage in a project is determined by the number of scientists allocated to that stage. The default assumption is that the rate of progress is proportional to the number of scientists allocated. However, there is the possibility that team sizes above some level may be relatively inefficient. This can occur because of difficulties of communication, little space to work, lack of coordination and other reasons which reduce the efficiency of a very large team.

Define:

- $e_i(u_i)$ = effectiveness of a team of u_i scientists for stage i ,
- $\frac{e_i(u_i)}{u_i}$ = relative efficiency of a team of u_i scientists, $i = 1, 2, 3, 4, 5$.

We also make the following assumptions for $i = 1, 2, 3, 4, 5$.

- The amount of effort required to complete each occurrence of stage i is X_i , except that after the first successful occurrence of stage 4 the work required for a repetition of stage 4 is ρX_4 , for some $0 < \rho < 1$. Effort rate u_i is measured in terms of the number of senior scientists allocated. Note that u_i does not have to be an integer.
- The time t_i required to complete stage i is $X_i/e_i(u_i)$, where $e_i(\cdot)$ is a rate of effective working function for stage i ; this function takes account of the fact that the efficiency of a team of scientists depends on its size, so that, for example, there may be a less than 50% reduction in the time to complete a stage if the team size is doubled. Initial input values for u_i and t_i are required, and the link equation means that these also provide values for the model parameter X_i , which is rather less easy to estimate directly.
- No more scientists than cap_i can be allocated to stage i . This cap may simply be designed to avoid extreme solutions. It also needs to be sufficiently low for our modelling assumptions to be realistic.
- A minimum duration $mint_i$ is set for stage i (for example there may be a minimum time needed to complete a series of tests). This is ensured by imposing a second cap $cap2_i = e_i^{-1}(\frac{X_i}{mint_i})$ on the effort rate allocation. If $cap2_i < cap_i$, we will reset $cap_i = cap2_i$.

It seems reasonable to suppose that the efficiency $e(u)/u^1$ decreases for values of u greater than some value u_f of maximal efficiency. It is convenient to scale the function $e(u)$ so that $e(u_f)/u_f = 1$, and efficiency is measured on a scale between 0 and 1. Let

- $dbeff = e(u_{2f})/u_{2f}$, the relative efficiency at $2u_f$,
- $trief = e(u_{3f})/u_{3f}$, the relative efficiency at $3u_f$,
- $quintef = e(u_{5f})/u_{5f}$, the relative efficiency at $5u_f$.

¹Note that for the remainder of this section the subscript i is suppressed.

OPRRA offers two options for calculating the effectiveness of a team of scientists. The first option, proposed by Yu and Gittins (2007), fits $e(u)$ to the values of u_f and $2u_f$, using a function of the form

$$e(u) = \frac{bu^2}{1 + au^r}, \quad (4.1)$$

where $a > 0$ and $1 < r < 2$. This ensures that $e(u)$ is increasing for all u and does not grow faster than linearly with u as u tends to infinity; b is chosen so that

$$\max_u \frac{e(u)}{u} = 1.$$

The other option, proposed by Charalambous (2009), uses cubic interpolation to fit $e(u)$ between the values for u_f , $3u_f$, and $5u_f$. The interpolation formula $f(u)$ is as follows.

$$f(u) = \begin{cases} u & u \leq u_f, \\ \frac{(u-u_f)^2}{3u_f-5u_f} \left[\frac{(u-u_f)e(3u_f)}{(3u_f-u_f)^2} - \frac{(u-3u_f)e(5u_f)}{(5u_f-u_f)^2} \right] \\ + \frac{(u-3u_f)(u-5u_f)}{(u_f-3u_f)(u_f-5u_f)} \left[u - (u-u_f)u_f \left(\frac{1}{u_f-3u_f} + \frac{1}{u_f-5u_f} \right) \right] & u_f < u \leq 5u_f. \end{cases} \quad (4.2)$$

This effectiveness function is for most purposes an improvement on the previous effectiveness function as it fits functions to two points above u_f rather than just one point. It gives a reasonable fit up to $u = 5u_f$. The previous function is unreliable for $u > 2.5u_f$.

4.2 Discounting and Obsolescence

Income from the sales of a drug is likely to be lower if the launch time is delayed, due to the general tendency for better drugs to become available from competitors as time goes by. This effect is described as *obsolescence*. It means that the exponential discount rate for the value of a CD from the research project is higher than the discount rate applied to future costs.

Our obsolescence modelling takes account of the rates of occurrence of three types of events:

- Type A: the registration of a competitive patent,
- Type B: a competitor obtains a licence to market a new drug,
- Type C: changes in the regulatory environment.

First we consider Type A events. Let $\xi_A \Delta t$ denote the probability of any other company registering a competitive patent in a short interval of time Δt , and f_A be the fraction by which the competitor's patent reduces the value of a new drug. It may be shown that the effect of this pattern of occurrence of competitive patents is to multiply the expected value of a new drug which becomes available at time t by the factor $e^{-f_A \xi_A}$, thus increasing the discount rate by $f_A \xi_A$.

The other two classes of events mentioned above are modelled in the same way, with a further contribution to the discount rate for each class. Regulatory changes typically call for additional procedures to be carried out before a drug is given regulatory approval. Thus we have three components $f_A \xi_A$, $f_B \xi_B$ and $f_C \xi_C$ which contribute to ν_1 , the obsolescence component of the discount rate for the value of a new drug.

Pharmaceutical patents protect manufacturers by prohibiting competitors from producing the drug during the period of its patent protection. The length of these periods varies in different countries, though they are usually between 15 and 20 years. However, pharmaceutical companies apply for patent protection before they actually begin producing the drug, thus significant portions of the patent term for a new drug are lost before the drug enters the market. As a result, the effective remaining patent life of a new drug has an important effect on its market value. Denote by Q the factor by which the value of a new drug is reduced by a one year postponement of availability, taking account only of the reduction in unexpired patent life. It follows that $-\log(Q)$ is an additional contribution to the obsolescence rate ν_1 . Hence we have $\nu_1 = f_A \xi_A + f_B \xi_B + f_C \xi_C - \log(Q)$.

When competitors learn of progress in our research project, there may well be an increase both in competitive research and in the similarities between the compounds screened by competitors and those screened in the project. To model this effect of additional self-induced obsolescence, a second discount rate ν_2 is introduced. We estimate ν_2 in a similar way to the estimation of ν_1 .

We define the exponential rate $\gamma_1 = \gamma + \nu_1$ for discounting future rewards without self-induced obsolescence to be the sum of the rate for discounting future costs and the obsolescence rate ν_1 . We denote by $\gamma_2 = \gamma + \nu_2$ the overall discount rate including self-induced obsolescence. The expected value of a new drug is discounted at the rate γ_1 until the completion of the first attempt to find a CD, and from then on at the rate γ_2 . For simplicity of notation we set $\gamma_1 = \gamma_2$ in the expressions throughout the thesis.²

4.3 (s, n) Policies

4.3.1 Definition of an (s, n) Policy

The parameters s and n which define an (s, n) policy have the following meanings:

- s is the number of LS from which we would like to select CDs.
The optimization calculations in OPRRA allow the values $s = 1, 2$ and 3 , so that we consider projects which select CDs from at most three LS. This is to limit the complexity of the required formulae and computations, and is unlikely to be a serious restriction in practice.
- n is the maximum number of LS from which we will attempt to select CDs.
Obviously, $s \leq n$.

The rules for an (s, n) policy are:

1. Stage 3 is carried out in parallel with stage 4 in the case where we find one LS in stage 2, except for the n^{th} time that stage 4 is carried out. When two LS are found in stage 2, stage 3 is carried out in parallel with stage 4 except for the 1^{st} and n^{th} times that stage 4 is carried out.

²Note that OPRRA does not require this assumption. The arguments in chapter 3 are unchanged with $\gamma_1 \neq \gamma_2$. The tests conducted in chapters 7 and 8 do not impose this equality either.

2. If an LS becomes the i^{th} successfully optimized LS, $1 \leq i < s$, selecting k_i CDs from it has precedence over optimizing further LS. After we have obtained k_i CDs from the i^{th} optimized LS, we will never return to this LS for more CDs.
3. If an LS becomes the s^{th} successfully optimized LS before all n attempts have been made, the project will be terminated after selecting k_s CDs from it.
4. The project is terminated after n optimization attempts have been made, even if the number of successfully optimized LS is less than s .

4.3.2 Expected Reward and Expected Cost

In this section we give expressions for the total expected discounted reward (TEDR) and total expected discounted cost (TEDC) for an (s, n) policy. More details and derivations are given by Charalambous (2009).

Let $R(s, n)$ and $C(s, n)$ denote the TEDR and TEDC for policy (s, n) , respectively. Recall that our focus is on the efficient use of resources in pre-clinical research, hence the costs for the clinical trials are included in the calculation for TEDR rather than in TEDC. As we are only optimizing the resource allocation for pre-clinical research these costs are also fixed.

$$(s, n) = (1, n)$$

If the first LS is successfully optimized, the expected value of the first CD is

$$e^{-\gamma_1(t_1+t_2+t_4+t_I+t_{II}+t_{III})}pW - e^{-\gamma(t_1+t_2+t_4)}(c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}). \quad (4.3)$$

The TEDR of the first k_1 CDs is $V_{1a} - V_{1b}$, where

$$\begin{aligned}
 V_{1a} &= e^{-\gamma_1(t_1+t_2+t_4+t_I+t_{II}+t_{III})} pW \sum_{m=1}^{k_1} (\eta e^{-\gamma_1 t_5})^{m-1} \\
 &= e^{-\gamma_1(t_1+t_2+t_4+t_I+t_{II}+t_{III})} pW \left[\frac{1 - (\eta e^{-\gamma_1 t_5})^{k_1}}{1 - \eta e^{-\gamma_1 t_5}} \right], \\
 V_{1b} &= e^{-\gamma(t_1+t_2+t_4)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \sum_{m=1}^{k_1} (e^{-\gamma t_5})^{m-1} \\
 &= e^{-\gamma(t_1+t_2+t_4)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \left[\frac{1 - (e^{-\gamma t_5})^{k_1}}{1 - e^{-\gamma t_5}} \right].
 \end{aligned}$$

We have

$$\begin{aligned}
 R(1, n) &= p_1 p_{21} p_4 V_{1a} \left(\frac{1 - (p_3 y_2)^n}{1 - p_3 y_2} \right) + p_1 p_{22} p_4 V_{1a} \left(1 + y_2 \left(\frac{1 - (p_3 y_2)^{n-1}}{1 - p_3 y_2} \right) \right) \\
 &\quad - p_1 p_{21} p_4 V_{1b} \left(\frac{1 - (p_3 y)^n}{1 - p_3 y} \right) - p_1 p_{22} p_4 V_{1b} \left(1 + y \left(\frac{1 - (p_3 y)^{n-1}}{1 - p_3 y} \right) \right). \\
 &\quad (y = q_4 e^{-\gamma t_4}, y_2 = q_4 e^{-\gamma_1 t_4})
 \end{aligned}$$

The following notation simplifies the expression for $C(1, n)$.

$$C_1 = \alpha \frac{u_1}{\gamma} (1 - e^{-\gamma t_1}) = \text{TEDC for stage 1.}$$

$$C_2 = \alpha p_1 e^{-\gamma t_1} \frac{u_2}{\gamma} (1 - e^{-\gamma t_2}) = \text{TEDC for stage 2.}$$

$$\begin{aligned} C_{31} &= \alpha p_1 p_{21} e^{-\gamma(t_1+t_2)} \frac{u_3}{\gamma} (1 - e^{-\gamma t_4}) \\ &= \text{contribution to TEDC for stage 3 corresponding to one LS found in stage 2.} \end{aligned}$$

$$\begin{aligned} C_{32} &= \alpha p_1 p_{22} e^{-\gamma(t_1+t_2)} \frac{u_3}{\gamma} (1 - e^{-\gamma t_4}) . \\ &= \text{contribution to TEDC for stage 3 corresponding to two LS found in stage 2.} \end{aligned}$$

$$\begin{aligned} C_{41} &= \alpha p_1 p_{21} e^{-\gamma(t_1+t_2)} \frac{u_4}{\gamma} (1 - e^{-\gamma t_4}) . \\ &= \text{contribution to TEDC for stage 4 corresponding to one LS found in stage 2.} \end{aligned}$$

$$\begin{aligned} C_{42} &= \alpha p_1 p_{22} e^{-\gamma(t_1+t_2)} \frac{u_4}{\gamma} (1 - e^{-\gamma t_4}) . \\ &= \text{contribution to TEDC for stage 4 corresponding to two LS found in stage 2.} \end{aligned}$$

$$\begin{aligned} C_{51}(k) &= \alpha p_1 p_{21} p_4 e^{-\gamma(t_1+t_2+t_4)} \frac{u_5}{\gamma} (1 - e^{-k\gamma t_5}) . \\ &= \text{contribution to TEDC for stage 5 corresponding to one LS found in stage 2.} \end{aligned}$$

$$\begin{aligned} C_{52}(k) &= \alpha p_1 p_{22} p_4 e^{-\gamma(t_1+t_2+t_4)} \frac{u_5}{\gamma} (1 - e^{-k\gamma t_5}) . \\ &= \text{contribution to TEDC for stage 5 corresponding to two LS found in stage 2.} \end{aligned}$$

We have

$$\begin{aligned}
C(1, n) &= C_1 + C_2 + C_{31} \left(\sum_{i=1}^{n-1} (p_3 y)^{i-1} \right) + C_{32} \left(\sum_{i=1}^{n-2} (p_3 y)^{i-1} \right) \\
&\quad + [C_{41} + C_{51}(k_1 - 1)] \left(\sum_{i=1}^n (p_3 y)^{i-1} \right) + [C_{42} + C_{52}(k_1 - 1)] \left(1 + y \sum_{i=1}^{n-1} (p_3 y)^{i-1} \right) \\
&= C_1 + C_2 + C_{31} \left(\frac{1 - (p_3 y)^{n-1}}{1 - p_3 y} \right) + I\{n \neq 1\} C_{32} \left(\frac{1 - (p_3 y)^{n-2}}{1 - p_3 y} \right) \\
&\quad + [C_{41} + C_{51}(k_1 - 1)] \left(\frac{1 - (p_3 y)^n}{1 - p_3 y} \right) \\
&\quad + [C_{42} + C_{52}(k_1 - 1)] \left(1 + y \left(\frac{1 - (p_3 y)^{n-1}}{1 - p_3 y} \right) \right),
\end{aligned}$$

where

$$I\{A\} = \begin{cases} 1 & \text{if event A occurs,} \\ 0 & \text{otherwise.} \end{cases}$$

$$(s, n) = (2, n)$$

We now fix $s = 2$ and proceed to give expressions for $R(2, n)$ and $C(2, n)$. As explained in section 2.1, the time required to optimize a subsequent LS, given that a previous LS has been successfully optimized, is ρt_4 ($0 < \rho < 1$). We call the corresponding stage 4ρ to distinguish it from the original stage 4. Since stage 3 may be carried out in parallel with stage 4, we let stage 3ρ be the stage carried out in parallel with stage 4ρ .

The expected value of the first CD from the second successfully optimized LS after we have obtained k_1 CDs from the first successfully optimized LS is

$$\begin{aligned}
&e^{-\gamma(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5+t_I+t_{II}+t_{III})} pW (p_b \eta^{k_1} + 1 - p_b) \\
&- e^{-\gamma(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}). \quad (4.4)
\end{aligned}$$

Let $V_{2a} - V_{2b}$ be the TEDR of the k_2 CDs selected from the second LS, where

$$\begin{aligned}
V_{2a} &= e^{-\gamma_1(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5+t_I+t_{II}+t_{III})} pW (p_b \eta^{k_1} + 1 - p_b) \sum_{m=1}^{k_2} (\eta e^{-\gamma_1 t_5})^{m-1} \\
&= e^{-\gamma_1(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5+t_I+t_{II}+t_{III})} pW (p_b \eta^{k_1} + 1 - p_b) \left[\frac{1 - (\eta e^{-\gamma_1 t_5})^{k_2}}{1 - \eta e^{-\gamma_1 t_5}} \right], \\
V_{2b} &= e^{-\gamma(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \sum_{m=1}^{k_2} (e^{-\gamma t_5})^{m-1} \\
&= e^{-\gamma(t_1+t_2+t_4+\rho t_4+(k_1-1)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \left[\frac{1 - (e^{-\gamma t_5})^{k_2}}{1 - e^{-\gamma t_5}} \right].
\end{aligned}$$

We have

$$R(2, n) = R(1, n) + R_{(2\rho)}(n),$$

where

$$\begin{aligned}
R_{(2\rho)}(n) &= (p_1 p_{21} p_3 p_4^2 + p_1 p_{22} p_4^2) V_{2a} \left(\sum_{j=1}^{n-1} \sum_{i=1}^j (p_3 y_2)^{j-i} (p_3 y_{2\rho})^{i-1} \right) \\
&\quad - (p_1 p_{21} p_3 p_4^2 + p_1 p_{22} p_4^2) V_{2b} \left(\sum_{j=1}^{n-1} \sum_{i=1}^j (p_3 y)^{j-i} (p_3 y_\rho)^{i-1} \right). \\
&\quad (y_\rho = q_4 e^{-\gamma \rho t_4}, y_{2\rho} = q_4 e^{-\gamma_1 \rho t_4})
\end{aligned}$$

To derive expressions for $C(2, n)$ we define $C_{31\rho}$ and $C_{41\rho}$ to be the contributions to the TEDC for stage 3ρ and stage 4ρ respectively, given that we find one LS in stage 2. Similarly we let $C_{32\rho}$ and $C_{42\rho}$ be the contributions to the TEDC for stage 3ρ and stage 4ρ respectively, given that we find two LS in stage 2. Thus we have

$$\begin{aligned}
C_{31\rho} &= \alpha p_1 p_{21} e^{-\gamma(t_1+t_2)} \frac{u_3}{\gamma} (1 - e^{-\gamma \rho t_4}), \\
C_{32\rho} &= \alpha p_1 p_{22} e^{-\gamma(t_1+t_2)} \frac{u_3}{\gamma} (1 - e^{-\gamma \rho t_4}), \\
C_{41\rho} &= \alpha p_1 p_{21} e^{-\gamma(t_1+t_2)} \frac{u_4}{\gamma} (1 - e^{-\gamma \rho t_4}), \\
C_{42\rho} &= \alpha p_1 p_{22} e^{-\gamma(t_1+t_2)} \frac{u_4}{\gamma} (1 - e^{-\gamma \rho t_4}).
\end{aligned}$$

The TEDC for $(s, n) = (2, n)$ is

$$C(2, n) = C(1, n) + C_{(2\rho)}(n),$$

where

$$\begin{aligned} C_{(2\rho)}(n) = & \left[p_4 e^{-\gamma(t_4 + (k_1 - 1)t_5)} (p_3 C_{41\rho} + C_{42\rho}) + p_4 e^{-\gamma((k_1 - 1)t_5 + \rho t_4)} (p_3 C_{51}(k_2 - 1) \right. \\ & \left. + C_{52}(k_2 - 1)) \right] \left(\frac{1}{p_3 y - p_3 y_\rho} \right) \left(\frac{p_3 y - (p_3 y)^n}{1 - p_3 y} - \frac{p_3 y_\rho - (p_3 y_\rho)^n}{1 - p_3 y_\rho} \right) \\ & + p_4 e^{-\gamma(t_4 + (k_1 - 1)t_5)} \left(C_{31\rho} \left(\frac{1}{p_3 y - p_3 y_\rho} \right) \left(\frac{p_3 y - (p_3 y)^{n-1}}{1 - p_3 y} - \frac{p_3 y_\rho - (p_3 y_\rho)^{n-1}}{1 - p_3 y_\rho} \right) \right. \\ & \left. + I\{n \neq 2\} C_{32\rho} \left(1 + (y + y_\rho) \left(\frac{1}{p_3 y - p_3 y_\rho} \right) \times \right. \right. \\ & \left. \left. \left(\frac{p_3 y - (p_3 y)^{n-2}}{1 - p_3 y} - \frac{p_3 y_\rho - (p_3 y_\rho)^{n-2}}{1 - p_3 y_\rho} \right) \right) \right). \end{aligned}$$

$(s, n) = (3, n)$

The expected value of the first CD from the third successfully optimized LS given that we find k_1 and k_2 CDs from the first and the second successfully optimized LS respectively, is

$$\begin{aligned} & e^{-\gamma_1(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5 + t_I + t_{II} + t_{III})} pW \prod_{i=1}^2 (p_b \eta^{k_i} + 1 - p_b) \\ & - e^{-\gamma_1(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I + t_{II})} c_{III}). \end{aligned} \quad (4.5)$$

Let $V_{3a} - V_{3b}$ be the TEDR for k_3 CDs selected from the third LS, where

$$\begin{aligned} V_{3a} &= e^{-\gamma_1(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5 + t_I + t_{II} + t_{III})} pW \prod_{i=1}^2 (p_b \eta^{k_i} + 1 - p_b) \sum_{m=1}^{k_3} (\eta e^{-\gamma_1 t_5})^{m-1} \\ &= e^{-\gamma_1(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5 + t_I + t_{II} + t_{III})} pW \prod_{i=1}^2 (p_b \eta^{k_i} + 1 - p_b) \left[\frac{1 - (\eta e^{-\gamma_1 t_5})^{k_3}}{1 - \eta e^{-\gamma_1 t_5}} \right], \\ V_{3b} &= e^{-\gamma(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I + t_{II})} c_{III}) \sum_{m=1}^{k_3} (e^{-\gamma t_5})^{m-1} \\ &= e^{-\gamma(t_1 + t_2 + t_4 + 2\rho t_4 + (k_1 + k_2 - 2)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I + t_{II})} c_{III}) \left[\frac{1 - (e^{-\gamma t_5})^{k_3}}{1 - e^{-\gamma t_5}} \right]. \end{aligned}$$

We have

$$R(3, n) = R(2, n) + R_{(3\rho)}(n),$$

where

$$\begin{aligned} R_{(3\rho)}(n) &= (p_1 p_{21} p_3^2 p_4^3 + p_1 p_{22} p_3 p_4^3) V_{3a} \left(\sum_{j=1}^{n-2} \sum_{i=1}^j i \cdot (p_3 y_2)^{j-i} (p_3 y_{2\rho})^{i-1} \right) \\ &\quad - (p_1 p_{21} p_3^2 p_4^3 + p_1 p_{22} p_3 p_4^3) V_{3b} \left(\sum_{j=1}^{n-2} \sum_{i=1}^j i \cdot (p_3 y)^{j-i} (p_3 y_\rho)^{i-1} \right). \end{aligned}$$

Also the TEDC for $(s, n) = (3, n)$ is

$$C(3, n) = C(2, n) + C_{(3\rho)}(n),$$

where

$$\begin{aligned} C_{(3\rho)}(n) &= p_3 p_4 e^{-\gamma(\rho t_4 + (k_2 - 1)t_5)} \left[p_4 e^{-\gamma(t_4 + (k_1 - 1)t_5)} (p_3 C_{41\rho} + C_{42\rho}) + p_4 e^{-\gamma((k_1 - 1)t_5 + \rho t_4)} \times \right. \\ &\quad \left. (p_3 C_{51}(k_3 - 1) + C_{52}(k_3 - 1)) \right] \left(\sum_{j=1}^{n-2} \sum_{i=1}^j i \cdot (p_3 y)^{j-i} (p_3 y_\rho)^{i-1} \right) \\ &\quad + p_4^2 e^{-\gamma(t_4 + \rho t_4 + (k_1 + k_2 - 2)t_5)} (p_3 C_{31\rho} + C_{32\rho}) \left(\sum_{j=1}^{n-3} \sum_{i=1}^j i \cdot (p_3 y)^{j-i} (p_3 y_\rho)^{i-1} \right). \end{aligned}$$

4.3.3 Optimization Calculations with PI and IRR

One of the objectives of OPRRA is to find the values of a set of decision variables, for given values of (s, n) , that maximize the project's PI or IRR. The vector of decision variables to be optimized is $\mathbf{u} = (u_1, u_2, u_4, u_5, k_1, \dots, k_s)$, which we call the *allocation variables*. u_3 is not included in the vector of allocation variables as it is linked to u_4 by the requirement that $t_3 = t_4$. We set $t_3 = \frac{X_4}{e_4(u_4)}$ and can obtain u_3 by setting $u_3 = e_3^{-1}\left(\frac{X_3}{t_3}\right)$. The value of t_i may be obtained from setting $t_i = \frac{X_i}{e_i(u_i)}$, $i = 1, 2, 4, 5$. Hence this leaves with us $4 + s$ allocation variables to be optimized.

Define $PI(\mathbf{u}) = \frac{R(\mathbf{u})}{C(\mathbf{u})}$, and γ_I as the IRR for which $G(\gamma_I, \mathbf{u}) = R(\gamma_I, \mathbf{u}) - C(\gamma_I, \mathbf{u}) = 0$. The goal is to find the optimal sets of allocations for PI and IRR, respectively, such that $\frac{\partial PI(\mathbf{u})}{\partial \mathbf{u}} = 0$ and $\frac{\partial \gamma_I}{\partial \mathbf{u}} = 0$.³ We apply the Newton-Raphson algorithm to find these optimal allocations, which we denote in both cases by \mathbf{u}^* , subject to the constraint $u_i \leq cap_i$, ($i = 1, 2, 3, 4, 5$). For this purpose k_1, k_2, \dots, k_s are treated as continuous variables.

For the PI function, the first and second derivatives with respect to u_1 are as follows. Similar expressions give the first and second derivatives with respect to the other allocation variables.

$$\begin{aligned}\frac{\partial PI(\mathbf{u})}{\partial u_1} &= \frac{1}{C(\mathbf{u})^2} \left[C(\mathbf{u}) \frac{\partial R(\mathbf{u})}{\partial u_1} - R(\mathbf{u}) \frac{\partial C(\mathbf{u})}{\partial u_1} \right], \\ \frac{\partial^2 PI(\mathbf{u})}{\partial u_1^2} &= \frac{1}{C(\mathbf{u})^2} \left[C(\mathbf{u}) \frac{\partial^2 R(\mathbf{u})}{\partial u_1^2} - R(\mathbf{u}) \frac{\partial^2 C(\mathbf{u})}{\partial u_1^2} \right] - \frac{2}{C(\mathbf{u})} \frac{PI(\mathbf{u})}{\partial u_1} \frac{\partial C(\mathbf{u})}{\partial u_1}.\end{aligned}$$

For the IRR function, the first and second derivatives to be applied in the Newton-Raphson algorithm with respect to u_1 are as follows, again with similar expressions for other allocation variables.

$$\begin{aligned}\frac{\partial \gamma_I}{\partial u_1} &= - \left[\frac{\frac{\partial G(\gamma_I, \mathbf{u})}{\partial u_1}}{\frac{\partial G(\gamma_I, \mathbf{u})}{\partial \gamma_I}} \right], \\ \frac{\partial^2 \gamma_I}{\partial u_1^2} &= - \left[\frac{\frac{\partial G(\gamma_I, \mathbf{u})}{\partial \gamma_I} \cdot \frac{\partial^2 G(\gamma_I, \mathbf{u})}{\partial u_1^2} - \frac{\partial G(\gamma_I, \mathbf{u})}{\partial u_1} \cdot \frac{\partial^2 G(\gamma_I, \mathbf{u})}{\partial \gamma_I \partial u_1}}{\left(\frac{\partial G(\gamma_I, \mathbf{u})}{\partial \gamma_I} \right)^2} \right].\end{aligned}$$

OPRRA optimizes the allocation variables one at a time, starting with the value set $\mathbf{u}^{(0)} = (u_1^{(0)}, u_2^{(0)}, u_4^{(0)}, u_5^{(0)}, k_1^{(0)}, \dots, k_s^{(0)})$ and working through the variables from u_1 to k_s iteratively. The optimization calculations stop when the PI (or IRR) values resulting from a given number OPTIT (default value = $3(4 + s)$) of successive optimizations with respect to individual variables are all the same. If this convergence condition has not been satisfied before MAXIT (default value = $25(4 + s)$) single variable optimizations have been carried out the procedure is abandoned.

³i.e., $\frac{\partial PI(\mathbf{u})}{\partial u_i} = 0$, $\frac{\partial PI(\mathbf{u})}{\partial k_j} = 0$; $\frac{\partial \gamma_I}{\partial u_i} = 0$, $\frac{\partial \gamma_I}{\partial k_j} = 0$, $i = 1, 2, 4, 5$, $j = 1, 2, 3$.

As a precaution against the possibilities both of a failure to converge and of convergence to a local rather than the global optimum, this procedure is carried out for a total of RANDIT (default value = 10) times, each, apart from the first, from a different randomly-generated starting point. For a more detailed description of the optimization procedure see Yu (2007).

4.4 Types of Allocation Policy for a Project

As we have discussed in chapter 2 and in this chapter, OPRRA offers two profitability criteria, two classes of candidate drug selection policies and two probability settings. The profitability criteria are profitability index and internal rate of return. The CD selection policies are the non-adaptive class of (s, n) policies and the adaptive class of FI policies. Finally the probabilities may be either fixed or adaptive.

The class of (s, n) policies plays a key role in the OPRRA calculations. Under an (s, n) policy OPRRA is able to calculate allocations which maximize either PI or IRR for given values of s and n . In the interests of analytical tractability, it is somewhat restrictive in the sequence of LS and CDs which it allows and is based on fixed probabilities. In all other respects it gives maximum flexibility, and as we shall see in chapters 7 and 8 the restrictions which it imposes often have little influence on profitability.

FI policies on the other hand only use PI as the profitability criterion, but can be implemented with either fixed or adaptive probabilities. OPRRA evaluates FI policies by Monte Carlo simulation. This is the only generally available method, though, as shown in chapter 3, for given stage allocations the PI for a project may be optimized analytically for a restricted FI selection policy for which CDs are selected from at most one LS. Simulation requires much more extensive calculations than evaluation by means of formulae. However, an important advantage of Monte Carlo simulation as compared with calculations based on expected values is that it provides probability distributions of possible outcomes, both as regards profitability and in terms of effort requirements.

An allocation policy is a set of rules concerning the effort allocation at each stage and the sequence of CDs to be selected. To summarize, allocation policy types vary in the following respects:

- the profitability criteria that they use,
- whether these may be evaluated analytically or only by Monte Carlo simulation,
- the rules for sequencing the searches for new LS and CDs,
- whether they are based on fixed probabilities or on probabilities which are updated using Bayes theorem,
- and in the availability of procedures within OPRRA for optimizing the choice of allocation variables.

4.5 Contributions to OPRRA

The following is a list of the improvements to OPRRA which I have made:

- Forward induction policies, a new class of policies for CD selection, both for the simplified problem and the general problem (Chapter 3).
- Calculation by means of Monte Carlo simulation of the probability distributions of rewards, costs, profitability indices and future allocations for individual projects and for a portfolio of projects, for
 - (s, n) policies
 - FI policies both in fixed and adaptive probabilities settingsand Monte Carlo simulation of the obsolescence events (Appendix C).
- Bayesian inference for p_3 , and a numerical procedure to evaluate the hyperparameters f_J , ($J = I, II, III$) (Sections 5.1 and 5.2).

- Additional modelling of contributions to the obsolescence rate, taking account of changes in regulations, as well as the factor by which the value of a new drug is reduced by a one year postponement of availability (Section 4.2).
- A cap on resource allocations and a minimal stage time for each stage (Section 4.1).

Presentations describing the OPRRA software have been given to several pharmaceutical companies, including AstraZeneca, GSK and Pfizer. Discussions on implementation are ongoing. Some of the practical conclusions from this thesis were also presented at the European Operational Research conference in Lisbon in 2010.

Chapter 5

The Adaptive Probabilities Model

This chapter introduces an adaptive probabilities model that allows the incorporation of learning from a research project's progress into the planning process. It takes into account the information accumulated from the discovery phases and from clinical trials and allows constant updates on success probabilities through a hierarchical Bayes structure.

The model was first outlined by Charalambous (2009), who did not, however, incorporate it into OPRRA. We extend this model by allowing feedback from stage 3 as well as implementing FI policies in this setting. The basic forwards induction idea still applies, the only difference being that the calculations of expected rewards and costs are based on the adaptive probabilities model. The procedures discussed in this chapter have been built into the current version of OPRRA.

Section 5.1 sets out the adaptive probabilities model. The posterior distributions of the success probabilities for the various stages are derived given their prior distributions. In section 5.2 we present a method for estimating the hyperparameters of the prior distributions. In section 5.3 we discuss how FI policies can be implemented in this adaptive probabilities setting.

5.1 Estimates for Success Probabilities

The Bayesian approach can be viewed as a formalization of the process of learning from experience, which is a fundamental characteristic of all scientific investigation.

The success probabilities for the different phases of a pharmaceutical project up to now have been assumed to be the fixed point estimates based on previous knowledge. An adaptive probabilities model based on Bayesian learning supplements previous knowledge by allowing the available information to change our opinion about the success probabilities as the project progresses. The use of the adaptive probabilities can therefore enable us to make more informed decisions.

Stages 1 and 2 occur at most once so there is no possibility of Bayesian updating for the probabilities p_1 , p_{21} and p_{22} . Estimates for these probabilities are required, and they stay unchanged throughout the calculations. For all other success probabilities our adaptive probabilities model allows for updating of prior distributions.

The success probabilities p_I , p_{II} and p_{III} are assumed to be different for different LS, and are treated as being randomly sampled from independent probability distributions. Final estimates of these success probabilities are reached by using Bayes theorem, which adjusts the prior distributions according to the record of successes and failures. Beta prior distributions are assumed for all the unknown probabilities, as they form a flexible and mathematically convenient class for quantities constrained to lie between 0 and 1.

Inference for $p_3 = \mathbf{P}(\text{find a backup LS} | \text{success in stages 1 and 2})$.

The probabilities of success ($x = 1$) and failure ($x = 0$) for stage 3 follow a Bernoulli distribution,

$$p(x) = p_3^x(1 - p_3)^{1-x}, \quad x = 0, 1.$$

Stage 3 may be repeated up to the first failure. Let n_3 be the number of stage 3s completed so far, hence the likelihood function may be written as

$$f(\mathbf{x}|p_3) = f(x_1, x_2, x_3, \dots, x_{n_3}|p_3) = p_3^{n_3}.$$

Assume that p_3 has a beta prior distribution,

$$p_3 \sim \text{Beta}(f_3 m_3, f_3(1 - m_3)), \quad 0 < m_3 < 1, f_3 > 0,$$

where $m_3 = E[p_3]$ is the prior expectation for p_3 .

Bayes theorem gives

$$\begin{aligned} P(p_3|\mathbf{x}) &\propto f(x_1, x_2, \dots, x_{n_3}|p_3)P(p_3) \\ &\propto p_3^{n_3} p_3^{f_3 m_3 - 1} (1 - p_3)^{f_3(1 - m_3) - 1} \\ &\propto p_3^{n_3 + f_3 m_3 - 1} (1 - p_3)^{f_3(1 - m_3) - 1}. \end{aligned}$$

It follows that

$$p_3|n_3 \sim \text{Beta}(n_3 + f_3 m_3, f_3(1 - m_3)),$$

thus

$$E[p_3|n_3] = \frac{n_3 + f_3 m_3}{n_3 + f_3}. \quad (5.1)$$

Inference for $p_4 = \mathbf{P}(\text{find a first CD from an LS} \mid \text{success in stages 1 to 3})$.

Let n_4 be the number of LS from which we have tried to find a CD, and y_4 be the number of successfully optimized LS. We have

$$y_4|p_4 \sim \text{Bin}(n_4, p_4),$$

and let p_4 have the prior distribution

$$p_4 \sim \text{Beta}(f_4 m_4, f_4(1 - m_4)), \quad 0 < m_4 < 1, f_4 > 0.$$

Thus $E[p_4]$, the prior expectation for p_4 , is m_4 . Proceeding as for p_3 it follows that

$$p_4|y_4 \sim \text{Beta}(y_4 + f_4 m_4, n_4 - y_4 + f_4(1 - m_4)),$$

so that the posterior expectation for p_4 is

$$E[p_4|y_4] = \frac{y_4 + f_4 m_4}{n_4 + f_4}. \quad (5.2)$$

Inference for $p_{Ji} = \mathbf{P}(\text{success in phase } J \text{ clinical trials for CD from LS } i \mid \text{success in stages 1 to } 5 \text{ and in phase } K, K < J)$. $i = 1, 2, \dots$; $J = I, II, III$.

We assume that for each phase the success probabilities for different LS are different. We relate these probabilities to each other by treating them as a sample drawn from a common beta distribution with an unknown mean, the unknown mean itself having a beta prior distribution. This hierarchical structure provides an ideal framework to reflect the dependence between different LS.

Let n_{Ji} be the number of CDs which have completed phase J clinical trials from LS i , and y_{Ji} be the number of CDs which have successfully completed phase J clinical trials from LS i . Thus

$$y_{Ji} | p_{Ji} \sim \text{Bin}(n_{Ji}, p_{Ji}),$$

$$p_{Ji} | u_J \sim \text{Beta}(f_J u_J, f_J(1 - u_J)),$$

$$u_J \sim \text{Beta}(r_J m_J, r_J(1 - m_J)),$$

where

$$0 < u_J < 1, 0 < m_J < 1, r_J > 0, f_J > 0.$$

Hence $E[p_{Ji} | u_J] = u_J$, $E[u_J] = m_J$, and the prior expectation for p_{Ji} is $E[p_{Ji}] = m_J$.

Assume we have successfully optimized LS $1, 2, \dots, \iota$.

Let $\mathbf{n}_J = (n_{J1}, n_{J2}, \dots, n_{J\iota})$, $\mathbf{y}_J = (y_{J1}, y_{J2}, \dots, y_{J\iota})$, $J = I, II, III$. The derivation of the approximate posterior expectation for p_{Ji} is given by Charalambous (2009). It is based on earlier work by Consonni and Veronese (1995). This gives

$$E[p_{Ji} | \mathbf{y}_J, \mathbf{n}_J] = b_{Ji} a_J m_J + b_{Ji} (1 - a_J) e_J + (1 - b_{Ji}) \hat{p}_{Ji}, \quad (5.3)$$

where

$$b_{Ji} = \frac{f_J}{f_J + n_{Ji}}, c_{Ji} = \frac{f_J + 1}{f_J + n_{Ji}}, a_J = \frac{r_J}{\sum_k c_{Jk} n_{Jk} + r_J}, e_J = \frac{\sum_k c_{Jk} y_{Jk}}{\sum_j c_{Jk} n_{Jk}}, \hat{p}_{Ji} = \frac{y_{Ji}}{n_{Ji}}.$$

5.2 Estimates for Hyper-parameters

Prior experience and judgement should be used to provide probability distributions for the success probabilities p_3 , p_4 and p_J , $J = I, II, III$. These distributions may then be used to calculate plausible values for the hyper-parameters. For the purpose of illustration we give the estimation for f_J ($J = I, II, III$) here. Estimation for f_3 and f_4 proceeds in a similar but less complicated fashion.

Recall that

$$p_{Ji}|u_J \sim \text{Beta}(f_J u_J, f_J(1 - u_J)), \quad (5.4)$$

$$u_J \sim \text{Beta}(r_J m_J, r_J(1 - m_J)). \quad (5.5)$$

Using (5.4) and (5.5) it follows that the mean and variance for p_{Ji} are

$$E[p_{Ji}] = m_J, \text{Var}[p_{Ji}] = m_J(1 - m_J) \frac{r_J + f_J + 1}{(f_J + 1)(r_J + 1)}.$$

Following Consonni and Veronese (1995) we approximate the distribution of p_{Ji} by a beta distribution $\text{Beta}(fm, f(1-m))$ with the same mean and variance. Since $E[p_{Ji}] = m$, and $\text{Var}[p_{Ji}] = \frac{m(1-m)}{f+1}$, this means that

$$m_J = m, \quad (5.6)$$

and

$$m_J(1 - m_J) \frac{r_J + f_J + 1}{(f_J + 1)(r_J + 1)} = \frac{m(1 - m)}{f + 1}. \quad (5.7)$$

Low values for f_J and r_J imply a high variance of p_{Ji} . High values for f_J and r_J mean accurate knowledge of the actual value of p_J . In the absence of any strong reason for a different choice we set $f_J = r_J$. Now solving equation (5.7) gives

$$f_J = f + \sqrt{f^2 + f}. \quad (5.8)$$

We use the method developed by Dorp and Mazzuchi (2000) to calculate the values of f and m . They have shown that a solution exists for the parameters of a beta distribution given any feasible combination of lower quantile and upper quantile values. They describe a numerical procedure which solves for the parameters given these two quantile values. We only give key results here.

Let $X \sim \text{Beta}(fm, f(1 - m))$, i.e.

$$P(X \leq x) = \frac{1}{B(fm, f(1 - m))} \int_0^x u^{fm-1} (1 - u)^{f(1-m)-1} du, \quad (5.9)$$

where $0 < m < 1$ and $f > 0$. The right side of equation (5.9) is the incomplete beta function $B(x|fm, f(1 - m))$ which has no closed form solution.

Definition 5.1. Let $0 < x_q < 1$ and $0 < q < 1$. A random variable X satisfies quantile condition $(x_q, q) \Leftrightarrow P(X \leq x_q) = q$.

Problem 5.1. Solve f and m for $X \sim \text{Beta}(fm, f(1 - m))$ under two quantile conditions (x_{qL}, qL) and (x_{qU}, qU) , where $qL < qU$.

The following theorem ensures that problem 5.1 has a solution for any combination of the two quantile conditions.

Theorem 5.1. There exists a solution (f^*, m^*) for problem 5.1.

The proof involves four steps. In the first step it is proved that, using the notation in (5.9), for a given $f > 0$ and a quantile condition (x_q, q) a unique m° exist such that $X \sim \text{Beta}(fm^\circ, f(1 - m^\circ))$ satisfies that quantile condition. In the second step it is shown that as $f \rightarrow 0$ the parameter $m^\circ \rightarrow (1 - q)$. In the third step it is shown that for $f \rightarrow \infty$ the parameter $m^\circ \rightarrow x_q$. In the fourth step, the theorem is proved. See Dorp and Mazzuchi (2000) for a full account of these steps.

Dorp and Mazzuchi (2000) also describe a numerical procedure based on three different applications of the bisection method¹, *BISECT1*, *BISECT2*, and *BISECT3*, which may be implemented to evaluate the regularized incomplete beta function given by (5.9). It finds the solution to problem 5.1 with a prescribed level of accuracy.

BISECT1(m, f, q) solves for the q^{th} quantile x_q of X . *BISECT2*(x_q, f, q) solves for the parameter m° satisfying the quantile condition (x_q, q) .

BISECT3(x_{qL}, x_{qU}, qL, qU) gives the solution (m^*, f^*) for problem 5.1. A method of determining a starting interval $[a_1, b_1]$ containing f^* is given by STEP1, STEP2, and STEP3 of *BISECT3*.

¹See Appendix F.2 for a description of the bisection method.

BISECT1(m, f, q)

- STEP 1 $j := 1$; Set $[d_1, e_1] = [0, 1]$;
 STEP 2 $x_{q,j} := \frac{d_j + e_j}{2}$; $q_j := B(x_{q,j} | f m, f(1 - m))$;
 STEP 3 If $q_j \leq q$ then $d_{j+1} := x_{q,j}$ and $e_{j+1} := e_j$;
 Else $e_{j+1} := x_{q,j}$; $d_{j+1} := d_j$;
 STEP 4 If $|q_j - q| < \delta$ then stop.
 Else $j := j + 1$; Go to STEP 2;

BISECT2(x_q, f, q)

- STEP 1 $n := 1$; Set $[d_1, e_1] = [0, 1]$;
 STEP 2 $m_{n+1} := \frac{d_n + e_n}{2}$; $q_n := B(x_q | f m_{n+1}, f(1 - m_{n+1}))$;
 STEP 3 If $q_n \leq q$ then $e_{n+1} := m_{n+1}$; $d_{n+1} := d_n$;
 Else $d_{n+1} := m_{n+1}$; $e_{n+1} := e_n$;
 STEP 4 If $|q_n - q| < \delta$ then stop.
 Else $n := n + 1$; Go to STEP 2;

BISECT3(x_{qL}, x_{qU}, qL, qU)

- STEP 1 $k := 1$; $f_{1,k} := 1$;
 STEP 2 $(m^\circ)_{1,k} := BISECT2(x_{qU}, f_{1,k}, qU)$;
 $(x_{qL})_{1,k} := BISECT1((m^\circ)_{1,k}, f_{1,k}, qL)$;
 STEP 3 If $(x_{qL})_{1,k} < x_{qL}$ then $f_{1,k+1} := 2 \cdot f_{1,k}$; Go to STEP 2;
 Else $[a_1, b_1] := [0, f_{1,k}]$;
 STEP 4 $k := 1$; $f_1 := \frac{a_1 + b_1}{2}$;
 STEP 5 $(m^\circ)_k := BISECT2(x_{qU}, f_k, qU)$; $(x_{qL})_k := BISECT1((m^\circ)_k, f_k, qL)$;
 STEP 6 If $(x_{qL})_k < x_{qL}$ then $a_{k+1} := f_k$; $b_{k+1} := b_k$;
 Else $a_{k+1} := a_k$; $b_{k+1} := f_k$;
 STEP 7 If $|(x_{qL})_k - x_{qL}| < \delta$ then stop.
 Else $k := k + 1$; Go to STEP 5;

We now proceed to give numerical examples for the estimations of f_3 , f_4 and f_J ($J = I, II, III$) using the bisection methods just described. The accuracy for δ

bisection methods *BISECT1* and *BISECT2* was set to 10^{-8} . The accuracy in bisection method *BISECT3* was set to 10^{-4} .

For illustration purposes assume that we have the same prior information about p_3 and p_4 , and that both of them are likely to lie between 0.4 and 0.8. Thus let us suppose that

$$\begin{aligned} P(p_3 < 0.4) &= 0.1, P(p_3 < 0.8) = 0.9, \\ P(p_4 < 0.4) &= 0.1, P(p_4 < 0.8) = 0.9, \end{aligned}$$

so that $(x_{qL}, qL) = (0.4, 0.1)$ and $(x_{qU}, qU) = (0.8, 0.9)$.

Recall that

$$p_3 \sim \text{Beta}(f_3 m_3, f_3(1 - m_3)), p_4 \sim \text{Beta}(f_4 m_4, f_4(1 - m_4)).$$

The solutions with the above settings of error tolerances are $m_3 = m_4 = 0.606$ and $f_3 = f_4 = 9.379$.

To complete our illustration we also fix the 10^{th} and 90^{th} quantiles for p_I , p_{II} and p_{III} , bearing in mind that any two distinct pairs of quantile conditions can be used. Suppose that the 10^{th} and 90^{th} quantiles for p_I are 0.4 and 0.7 respectively. Now recall that

$$\begin{aligned} p_{Ii}|u_I &\sim \text{Beta}(f_I u_I, f_I(1 - u_I)), \\ u_I &\sim \text{Beta}(r_I m_I, r_I(1 - m_I)), \end{aligned}$$

and we have approximated the distribution of p_{Ii} by the following beta distribution

$$p_{Ii} \sim \text{Beta}(f m, f(1 - m)),$$

where $m_I = m$ and $f_I = f + \sqrt{f^2 + f}$. Using the bisection method to solve for m and f we obtain $m_I = 0.551$ and $f_I = 35.888$.

Working as above and using the conditions $(x_{qL}, qL) = (0.3, 0.1)$ and $(x_{qU}, qU) = (0.6, 0.9)$ for p_{II} , and $(x_{qL}, qL) = (0.5, 0.1)$ and $(x_{qU}, qU) = (0.8, 0.9)$ for p_{III} , we obtain

$$\begin{aligned} m_{II} &= 0.449, f_{II} = 36.185, \text{ and} \\ m_{III} &= 0.655, f_{III} = 32.544. \end{aligned}$$

5.3 Adaptive FI Policies

The decision rules for FI policies in section 3.1 still apply, but the expected rewards and costs at each stage are now calculated in a Bayesian fashion, where the fixed probability estimates are replaced by the mean values of the prior or posterior probabilities. In addition, $PI(CD)$ is no longer the PI for an additional CD from the youngest optimized LS, but instead from an LS which gives the highest PI for that additional CD. Detailed calculations of relevant expected rewards and costs are set out in Appendix B.2.

At each decision point, a decision has to be made on whether to take a further CD from any of the already optimized LS or to optimize a new LS. Let $\mathbf{p} = (p_{I1}, p_{I2}, \dots, p_{Il}, p_{II1}, p_{II2}, \dots, p_{IIl}, p_{III1}, p_{III2}, \dots, p_{IIIl})$, the vector of clinical trials success probabilities. To avoid undue computational complexity, we assume that when we optimize a new LS, the CDs from that LS have independent and identical chances of being successful conditional on \mathbf{p} . We also ignore the possibility of further feedback during the processing of CDs from a given LS. After the attempt to find a first CD we need to take account of the history of successes and failures, which includes information on the progress of CDs which have started but not yet completed clinical trials.

Let x_{Ji} be the number of CDs which are currently in phase J clinical trials from LS i . We recall that n_{Ji} and y_{Ji} are the numbers of CDs that have completed and successfully completed phase J clinical trials, respectively, from LS i , $J = I, II, III$. Hence we have

$$x_{Ii} = y_{0i} - n_{Ii}, x_{IIi} = y_{Ii} - n_{IIi}, x_{IIIi} = y_{IIi} - n_{IIIi},$$

where y_{0i} is the number of CDs from LS i which have progressed to clinical trials.

Let $Z = z_1 + z_2 + z_3 + z_4$, where

$z_1 =$ number of the $\sum_{k=1}^l x_{Ik}$ CDs in phase I clinical trials that will become marketable drugs;

$z_2 =$ number of the $\sum_{k=1}^l x_{IIk}$ CDs in phase II clinical trials that will become marketable drugs;

$z_3 =$ number of the $\sum_{k=1}^l x_{IIIk}$ CDs in phase III clinical trials that will become marketable drugs;

$z_4 = \sum_{k=1}^{\ell} y_{IIIk}$, the number of successful CDs so far.

We express z_1 , z_2 and z_3 as the sums of Bernoulli random variables, all of which are independent conditional on \mathbf{p} .

$$\begin{aligned} z_1 &= \sum Q_{kj} = \sum_{k=1}^{\ell} \sum_{j=1}^{x_{Ik}} Q_{kj}, \text{ where } Q_{kj} \sim \text{Ber}(p_{Ik}p_{IIk}p_{IIIk}), \\ z_2 &= \sum R_{kj} = \sum_{k=1}^{\ell} \sum_{j=1}^{x_{IIk}} R_{kj}, \text{ where } R_{kj} \sim \text{Ber}(p_{IIk}p_{IIIk}), \\ z_3 &= \sum U_{kj} = \sum_{k=1}^{\ell} \sum_{j=1}^{x_{IIIk}} U_{kj}, \text{ where } U_{kj} \sim \text{Ber}(p_{IIIk}). \end{aligned}$$

The additional expected value of a subsequent CD conditional on the set of previously selected CDs as expressed in proposition 2.1 is replaced by $E[p_{Ii}p_{IIi}p_{IIIi}\lambda^Z]W$, where $E[p_{Ii}p_{IIi}p_{IIIi}\lambda^Z] = E[p_{Ii}p_{IIi}p_{IIIi}\lambda^{z_1+z_2+z_3}]\lambda^{z_4}$. We use the evaluation of $E[p_{Ii}p_{IIi}p_{IIIi}\lambda^{z_1+z_2+z_3}]$ by first order approximation set out in the following lemmas by Charalambous (2009). In the lemmas we use the notation $\mathbf{X} = (X_1, X_2, \dots, X_n)^T$, $E[\mathbf{X}] = (E[X_1], E[X_2], \dots, E[X_n])^T$ and $D[g(\mathbf{X})] = (\frac{\partial g(\mathbf{X})}{\partial X_1}, \frac{\partial g(\mathbf{X})}{\partial X_2}, \dots, \frac{\partial g(\mathbf{X})}{\partial X_n})^T$.

Lemma 5.1. *Suppose $g(\mathbf{X})$ is a scalar linear function of the vector \mathbf{X} . Then*

$$E[g(\mathbf{X})] = g(E[\mathbf{X}])$$

Proof. $g(\mathbf{X})$ is a scalar linear function, so $g(\mathbf{X}) = \mathbf{m}^T \mathbf{X} + b$ where \mathbf{m} is a vector of real numbers and b is a real number. Now

$$E[g(\mathbf{X})] = E[\mathbf{m}^T \mathbf{X} + b] = \mathbf{m}^T E[\mathbf{X}] + b = g(E[\mathbf{X}]),$$

as required. □

Lemma 5.2. *Suppose $g(\mathbf{X})$ is a scalar function which is twice differentiable. Then $E[g(\mathbf{X})] \approx g(E[\mathbf{X}])$ if we use a linear approximation for $g(\mathbf{X})$ in the neighbourhood of $E[\mathbf{X}]$.*

Proof. Taylor's theorem gives

$$\begin{aligned} g(\mathbf{X}) &= g(E[\mathbf{X}]) + D[g(E[\mathbf{X}])]^T(\mathbf{X} - E[\mathbf{X}]) + o(\|\mathbf{X} - E[\mathbf{X}]\|) \\ &\approx D[g(E[\mathbf{X}])]^T \mathbf{X} + g(E[\mathbf{X}]) - D[g(E[\mathbf{X}])]^T E[\mathbf{X}]. \end{aligned}$$

The lemma now follows directly from lemma 5.1. \square

Lemma 5.3.

$$E[p_{I_i} p_{II_i} p_{III_i} \lambda^{z_1+z_2+z_3}] = E[h(\mathbf{p})],$$

where

$$\begin{aligned} h(\mathbf{p}) &= p_{I_i} p_{II_i} p_{III_i} \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{Ik} p_{IIk} p_{IIIk})^{x_{Ik}} \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{IIk} p_{IIIk})^{x_{IIk}} \cdot \\ &\quad \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{IIIk})^{x_{IIIk}}. \end{aligned}$$

Proof.

$$\begin{aligned} E[p_{I_i} p_{II_i} p_{III_i} \lambda^{z_1+z_2+z_3}] &= E[E[p_{I_i} p_{II_i} p_{III_i} \lambda^{z_1+z_2+z_3} | \mathbf{p}]] \\ &= E[p_{I_i} p_{II_i} p_{III_i} E[\lambda^{z_1} | \mathbf{p}] E[\lambda^{z_2} | \mathbf{p}] E[\lambda^{z_3} | \mathbf{p}]] \\ &\quad \text{(using independence of } z_m | \mathbf{p} \text{ and } z_n | \mathbf{p}, m \neq n, m, n = 1, 2, 3) \\ &= E[p_{I_i} p_{II_i} p_{III_i} E[\lambda^{\sum Q_{kj}} | \mathbf{p}] E[\lambda^{\sum R_{kj}} | \mathbf{p}] E[\lambda^{\sum U_{kj}} | \mathbf{p}]] \\ &= E[p_{I_i} p_{II_i} p_{III_i} \prod_{kj} E[\lambda^{Q_{kj}} | \mathbf{p}] E[\lambda^{R_{kj}} | \mathbf{p}] E[\lambda^{U_{kj}} | \mathbf{p}]] \\ &\quad \text{(using independence of } Q_{kj} | \mathbf{p}, R_{kj} | \mathbf{p} \text{ and } U_{kj} | \mathbf{p}) \\ &= E[p_{I_i} p_{II_i} p_{III_i} \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{Ik} p_{IIk} p_{IIIk})^{x_{Ik}} \cdot \\ &\quad \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{IIk} p_{IIIk})^{x_{IIk}} \prod_{k=1}^{\ell} (1 - (1 - \lambda) p_{IIIk})^{x_{IIIk}}] \\ &= E[h(\mathbf{p})], \end{aligned}$$

as required. \square

Lemma 5.4. *If we use a linear approximation to $h(\mathbf{p})$ in the neighbourhood of $E[\mathbf{p}]$ then*

$$E[p_{I_1} p_{I_2} p_{I_3} \lambda^{z_1+z_2+z_3}] \approx h(E[\mathbf{p}]).$$

Proof. The proof is immediate from lemmas 5.2 and 5.3. □

Chapter 6

The Option Valuation Model

The valuation of pharmaceutical R&D projects is a challenging task due to the high degree of uncertainty involved. Estimating the future commercial value of a compound that is still in early development and far from the market is one of the major challenges. The future value of a new drug is driven by a broad range of factors such as its safety and efficacy, competition, expected demand, market share, etc, and it is likely to fluctuate with these value drivers. The high uncertainty about the new drug's value is linked to the desirability of flexibility in operating strategy. This managerial flexibility often has value and it may be substantial.

All of the analysis up to this point assumes a fixed estimate for the value of the first new drug, W , and a commitment to invest in clinical trials regardless of the value of W . In this chapter we relax this assumption by allowing W to follow a stochastic process that reflects the uncertainties resulting from both endogenous and exogenous factors. We assume that our company only invests in clinical trials when the value of W is greater than the expected cost at the time of investment. With W following a stochastic process and an irreversible investment cost, the value of a pharmaceutical R&D investment opportunity can be seen as a *real option*, analogous to a financial option. The ability of option pricing theory to quantify uncertainty and flexibility makes it a very appealing choice for valuing pharmaceutical investment opportunities.

The section that follows gives an overview of real options analysis and its

applications in R&D decision making. In section 6.2 we formulate the model and show how the option value of an investment opportunity may be calculated numerically using the finite difference method. In section 6.3 we discuss how calculations of profitability for both the (s, n) and FI candidate drug selection policies should be modified in order to incorporate option values. The option valuation model will be built into OPRRA in due course.

6.1 Real Options Overview

In an economic environment characterized by rapid change, great uncertainty and competitive interactions, it has become increasingly important that investment evaluation tools and processes can properly account for both uncertainty and a company's ability to react to new information. As new information arrives and uncertainty about the investment's rewards is gradually resolved, management often has the flexibility to alter the initial operating strategy in order to capitalize on favourable future opportunities or to react so as to mitigate losses. Real options has emerged as an approach that addresses this challenge more successfully than traditional capital budgeting techniques because it explicitly accounts for the value of future flexibility (Trigeorgis, 1996). This approach to investment recognizes threats as well as opportunities from uncertain events, and permits a much richer dynamic framework than was possible with traditional investment theory.

It has become clear that traditional capital budgeting tools, such as net present value and other techniques based on discounted-cash-flow (DCF) analysis are not well suited for evaluating risky projects. The traditional DCF approaches essentially involve discounting the expected net cash flows from a project at a risk-adjusted discount rate. Under such approaches, the impact of risk is one-directional: risk is assumed to depress the value of the investment (Trigeorgis, 1999a). As conventionally applied, DCF techniques, which were originally developed to value passive investments in bonds and stocks, were predicated on the implicit assumption of passive management, and allowed no flexibility to adapt and revise later decisions in response to unexpected market developments. Early critics (for example, Beardsley and Mansfield, 1978; Hayes and Abernathy, 1980;

Hayes and Garvin, 1982) recognize that DCF criteria often undervalue investment opportunities, leading to myopic decisions, underinvestment and eventual loss of competitive positions, because they either ignore or do not properly value important strategic considerations. Myers (1976) acknowledges that traditional DCF methods have inherent limitations when it comes to valuing investments with significant operating or strategic options, suggesting that options pricing holds the best promise of valuing such investments. The main weakness in traditional tools is their inability to capture flexibility : a project approval should not be a commitment to go all the way to market introduction, but only to proceed until further information arrives (Trigeorgis, 1996).

Managerial flexibility creates real options (Dixit and Pindyck, 1994). Flexibility is valuable if volatility is present, or in other words, if contingencies arise to which one needs to react. In order to capture project volatility, researchers have proposed an analog of the financial options pricing formula to understand the effect of volatility (Newton and Pearson, 1994). Since the option to invest involves the acquisition of real assets, the theory is termed *real option theory*. A *call* option on an asset gives the right, with no obligation, to acquire the underlying asset by paying a pre-specified price on or before a given maturity date. Similarly, a *put* option gives the right to sell (or exchange) the underlying asset and receive the exercise price. The option is *American* if it can be exercised at any time before the maturity date, or *European* if it can only be exercised on maturity. As with options on financial securities, management's flexibility to adapt its future actions in response to altered future market conditions and competitive reactions expands a capital investment opportunity's value by improving its upside potential, while limiting downside losses relative to the initial expectations of a passive management (Trigeorgis, 1999b). The resulting asymmetry caused by managerial adaptability lies at the heart of the real option's value.

The quantitative origins of real options derive from the seminal work of Black and Scholes (1973) and Merton (1973) in pricing financial options. Cox et al. (1979) consider a simplified valuation of options in discrete-time using a binomial approach. Geske (1979) values a compound option, which in principle may be applied in valuing growth opportunities which become available only if earlier investments are undertaken. The above line of work provides the building blocks

of real options analysis.

The past two decades witnessed a rapidly growing literature on real options. For an overview of real option theory and applications, see, for example, Dixit and Pindyck (1994) and Trigeorgis (1996). The former provides the first detailed exposition of the real options approach to the capital investment decisions of firms, stressing the irreversibility of most investment decisions, and the ongoing uncertainty of the economic environment in which these decisions are made. The latter brings together a wealth of previously scattered knowledge and research on the new flexibility in corporate resource allocation and in the evaluation of investment alternatives brought about by the dynamic paradigm of real options. The real options approach is detailed, coupling a coherent picture of how option theory is used with practical insights into real-world applications.

More recent contributions in this field include Paxson (2003), Schwartz and Trigeorgis (2004) and Guthrie (2009). Despite enormous theoretical progress, the practical usage of real options is nonetheless limited. Real options analysis clearly adds an entirely new dimension to strategic thinking but also presents many practical difficulties. These problems generally fall into three categories: finding a model whose assumptions match those of the project being analyzed, determining the inputs to this model, and being able to mathematically solve the option pricing algorithm (Bowman and Moskowitz, 2001). These practical difficulties may explain, in part, the limited use of real options analysis in strategic planning. One approach to solving this problem of misspecified option valuation models is to create a more advanced, customized option valuation algorithm that better matches the characteristics of the investment proposal. Applied in this manner, real option theory provides a long-term competitive advantage through better decision-making.

The considerable amount of risk and flexibility involved in R&D makes real options valuation an important alternative to other evaluation techniques. Each stage of R&D investment is analogous to a call option involving decisions to invest in further development or commercialize when the outcome is successful and the entire process may be viewed as a series of call options. At any point in the process, there is also an option to abandon the project (analogous to a put option)

so that only the sunken investment is lost. The real options approach can be applied at various levels of sophistication depending on the availability of data and the complexity of the problem.

R&D options are generally European since the decision to invest in further stages depends on the successful completion of previous R&D stages and as a result can not be exercised early. On the other hand, by waiting to introduce the new product, the company may lose first mover or pioneering advantages, especially in markets which are characterized by decreasing product-cycles and growing competition, see for example Urban et al. (1986). Therefore management will exercise R&D options as soon as it is able to do so (Pennings and Lint, 1997).

The ideas underlying the Black and Scholes formula involve setting up a riskless portfolio consisting of a position in the derivative and a position in the stock. For there to be no arbitrage opportunities, the return from the portfolio must be the risk-free interest rate. If the portfolio earned more than this return, arbitrageurs could make a riskless profit by borrowing money to buy the portfolio; if it earned less, they could make a riskless profit by shorting the portfolio and buying risk-free securities. This generalization of option pricing, as refined by Merton (1974, 1977), has become known as *contingent claims analysis*. The basic setting of this approach is an economy with a rich menu of traded assets with different return and risk characteristics. To value a new asset, we try to replicate its return and risk characteristics through a portfolio of existing traded assets. The price of the new asset must then equal the market value of the portfolio. This approach has been applied in the valuation of natural resource investments, such as gold and copper mines and oil deposits, due to the existence of well-developed future markets for those commodities from which essential market information can be extracted. With options on R&D projects, however, the underlying assets are non-traded and finding a twin asset with similar risk characteristics is impractical. In fact, the major risks of R&D projects are typically project specific and cannot be replicated in external markets. In particular, technical risks such as the uncertainty over development and manufacturing costs, and actions of competitors, are not correlated with any asset traded in the financial markets.

An alternative to contingent claims analysis for valuing options on real assets

is the dynamic programming approach. This approach makes no demand on the completeness of financial market. It breaks a whole sequence of decisions into just two components according to Bellman's Principle of Optimality¹: the immediate decision, and a valuation function that encapsulates the consequences of all subsequent decisions. The optimal sequence of decisions is found by working backward. Without the existence of a sufficiently rich set of markets in risky assets, the objective function simply reflects the decision maker's valuations of risk, and is calculated using an empirical discount rate. Although the two approaches to valuing real options make different assumptions about financial markets, and the discount rate that the firm uses to value future cash flows, they are in fact closely related to each other, and lead to identical results in many applications. For the equivalence between contingent claims analysis and dynamic programming, see Dixit and Pindyck (1994). For the reasons described above it is the dynamic programming approach that we will use in our analysis.

A striking problem that arises from the observation that R&D projects are non-traded concerns the estimation of their volatility. There are clearly no historic data from previous projects that can enable us to estimate reliably the volatility of an R&D project. However, it is a well-established fact that the option value is very sensitive to the volatility of the underlying asset. Reasonable estimates of the volatility are therefore required. One possible approach is to use relevant stock volatility to approximate the volatility resulting from an R&D project. Using judgements of senior management to attain reasonable values for volatility is another possible approach (Pennings and Lint, 1997). Large pharmaceutical enterprises such as Merck usually assume an annual project volatility between 40 and 60% (Chen, 2004).

In recent years, Monte Carlo simulation has emerged as a popular approach in dealing with volatility in real options evaluation. It is a powerful technique that allows for considerable flexibility in the specification of uncertainties. Based on assumed probability distributions for each uncertainty, a large number of possible scenarios can be generated for the underlying project value. The real options value is then calculated for each for these scenarios, and the average of these values is discounted back to the present (Triantis, 2001).

¹A brief introduction to Bellman equation is given in Appendix D.

In this chapter we are interested in the European real option under incomplete information in the presence of discontinuous arrival of jumps from multiple sources. We consider the value of an investment opportunity for which the decision is whether or not to proceed to clinical trials in a fixed period of time. We assume that the current value of the new drug and the cost of clinical trials are given with reasonable accuracy.

6.2 General Valuation Framework and Numerical Solution

The total uncertainty underlying the value of the new drug is posited to be the composition of two types of uncertainties: the economic uncertainty which concerns the volatility of the future value of the new drug, due to changes in market size, changes in interest rates, changes in the economic outlook, or other new information that causes marginal changes in the new drug's value, and the technical uncertainty which is associated with the arrival of rare events that have discrete effects on the new drug's value. Examples of these rare events are competitive entry or arrival of substitute products, a new regulatory decision, and expropriation or political risk. We incorporate these uncertainties by allowing W to follow a stochastic process through time. The economic component of the uncertainty is modelled by a standard geometric Brownian motion; the technical component of the uncertainty is modelled by Poisson jump processes reflecting the discrete impacts of the rare events.

Stochastic processes with discrete events in the context of option pricing were first studied by Merton (1976). Later applications in option pricing include Ball and Torous (1985), Amin (1993) and Pennings and Lint (1997). The finance literature has mostly focused on the case of a single source of rare events. Notable exceptions are Jones (1984), who studied hedging of financial options under two classes of jumps, Abraham and Taylor (1993) who considered the pricing of European options when the stock price is subject to both jumps with anticipated arrival and jumps with random arrival, and Martzoukos and Trigeorgis (2002) who valued real options when the underlying asset follows a mixed jump-diffusion process in-

volving various types of rare events.

Here we assume the existence of multiple sources of jumps. It is assumed that the arrivals within each type of jump are independently distributed Poisson processes. Let q_s denote the Poisson process for type s jumps with arrival rate λ_s , $s = 1, 2, \dots, S$. Thus, the probability that no event from source s occurs in the time interval dt is $1 - \lambda_s dt + o(dt)$, and the probability that one event from source s occurs is $\lambda_s dt + o(dt)$, where $o(dt)$ is the asymptotic order symbol meaning that $\psi(h) = o(h)$ if $\lim_{h \rightarrow 0} [\psi(h)/h] = 0$.

If a rare event occurs from any of the sources, the value of the new drug falls by some fixed fraction. Let ϕ_s ($0 \leq \phi_s \leq 1$) be the size of this fraction resulting from the arrival of a rare event from source s . Hence the increment of the Poisson process

$$dq_s = \begin{cases} 0, & \text{with probability } 1 - \lambda_s dt, \\ \phi_s, & \text{with probability } \lambda_s dt. \end{cases} \quad (6.1)$$

The presence of multiple sources of jumps implies that there is an overall Poisson arrival rate $\sum_{i=1}^S \lambda_i$ of rare events, and the probability that a rare event is from source s is $\lambda_s / \sum_{i=1}^S \lambda_i$.

The value of the first new drug W is assumed to follow a jump-diffusion process of the form:

$$dW = \beta W dt + \sigma W dz - W \sum_{s=1}^S dq_s. \quad (6.2)$$

Here β is the drift parameter, σ is the volatility of the project, conditional on no arrival of important new information (i.e., no rare event occurs) (Merton, 1976), $dz = \epsilon \sqrt{dt}$ is a standard Brownian motion, where $\epsilon \sim \mathcal{N}(0, 1)$, dq_s ($s = 1, 2, \dots, S$) are the fractional reductions caused by the Poisson processes described in (6.1) with $E[dq_i dq_j] = E[dq_i] E[dq_j]$ for $i \neq j$, and dz and dq_s are also assumed to be independent (so that $E[dz dq_s] = 0, \forall s$).

Equation (6.2) says that W will fluctuate as a geometric Brownian motion, but over each time interval dt there is a small probability that it will drop to some proportion of its original value and it will then continue fluctuating until another

rare event occurs. The presence of the jumps affects the underlying stochastic process in two ways. First, note that the expected percentage rate of change in W is not β , but instead is $E[dW]/Wdt = \beta - \sum_{s=1}^S \phi_s \lambda_s$, thus increases in λ_s reduce the expected rate of capital gain on W by increasing the chances of sudden drops in W . Second, because rare events occur only infrequently, most of the time the variance of dW/W over a short interval of time dt is just that of the Brownian motion part. However, if a rare event occurs, it contributes a very large deviation, so its contribution to the variance cannot be neglected. To examine these two components of the variance, let us set $\beta = 0$ for simplicity. Then

$$\begin{aligned}
 E[dW] &= -W \sum_{s=1}^S \phi_s \lambda_s dt, \\
 E[(dW)^2] &= E[\sigma^2 W^2 (dz)^2 - 2\sigma W^2 \sum_{s=1}^S dq_s dz + 2W^2 \sum_{i \neq j} dq_i dq_j + W^2 \sum_{s=1}^S (dq_s)^2], \\
 &= \sigma^2 W^2 dt - 2\sigma W^2 \sum_{s=1}^S \phi_s \lambda_s dz dt + 2W^2 \sum_{i \neq j} \phi_i \phi_j \lambda_i \lambda_j (dt)^2 + W^2 \sum_{s=1}^S \phi_s^2 \lambda_s dt, \\
 &= \sigma^2 W^2 dt + W^2 \sum_{s=1}^S \phi_s^2 \lambda_s dt,
 \end{aligned}$$

where $(dz)^2 = dt$, and terms in $dz dt$ and $(dt)^2$ are ignored, following the rules for Itô's formula.²

$$\begin{aligned}
 V[dW] &= E[(dW)^2] - \{E[dW]\}^2 \\
 &= \sigma^2 W^2 dt + W^2 \sum_{s=1}^S \phi_s^2 \lambda_s dt - 0 \\
 &= \sigma^2 W^2 dt + W^2 \sum_{s=1}^S \phi_s^2 \lambda_s dt. \tag{6.3}
 \end{aligned}$$

The first term on the right side of equation (6.3), $\sigma^2 W^2 dt$, is the instantaneous variance of dW , which comes from the Brownian motion part of the process, and is conditional on no jump occurring. The second term, $W^2 \sum_{s=1}^S \phi_s^2 \lambda_s dt$, accounts for the possibility of a jump.

²See Appendix E for an introduction to Itô's formula.

Now return to our valuation problem, where the firm at the current time φ has an option to invest in clinical trials at the cost of I in τ years time. Recall that p is the probability that a CD passes through all phases of clinical trials, thus $pW_{\varphi+\tau}$ is the expected value of the new drug at time $\varphi + \tau$. If the expected value of the new drug at time $\varphi + \tau$ is greater than the total expected cost of investment I , the decision will be to proceed to clinical trials, while for values less than I the project will be abandoned. Thus the payoff at time $\varphi + \tau$ is $\max\{pW_{\varphi+\tau} - I, 0\}$. Denote by $F(W_\varphi, t, \tau)$ the expected present value of this future payoff at time t ($t \geq \varphi$),

$$F(W_\varphi, t, \tau) = e^{-\gamma(\varphi+\tau-t)} E[\max\{pW_{\varphi+\tau} - I, 0\}], \quad (6.4)$$

where E denotes the expectation, W_φ is the current value of the new drug, τ is the time to maturity of this investment opportunity, and γ is the discount rate.

Since the investment opportunity, $F(W_\varphi, t, \tau)$, yields no cash flows up to time $\varphi + \tau$ when a CD is ready for submission to clinical trials (or sold to another company), the only return from holding it takes the form of capital appreciation. Hence the Bellman equation for $F(W_\varphi, t, \tau)$, the value of the investment opportunity at time t , is

$$\gamma F dt = E[dF]. \quad (6.5)$$

Equation (6.5) says that over the time interval dt , the total expected return on the investment opportunity, $\gamma F dt$, is equal to its expected rate of capital appreciation.

Applying Itô's formula to the calculation of dF we have

$$dF = \frac{\partial F}{\partial t} dt + \frac{\partial F}{\partial W} dW + \frac{1}{2} \frac{\partial^2 F}{\partial W^2} (dW)^2. \quad (6.6)$$

Here

$$\begin{aligned}
(dW)^2 &= (\beta W dt)^2 + \sigma^2 W^2 (dz)^2 + (W \sum_{s=1}^S dq_s)^2 + 2\beta\sigma W^2 dz dt \\
&\quad - 2\beta W^2 \sum_{s=1}^S dq_s dt - 2\sigma W^2 \sum_{s=1}^S dq_s dz \\
&= 0 + \sigma^2 W^2 dt + 2W^2 \sum_{i \neq j} dq_i dq_j + W^2 \sum_{s=1}^S (dq_s)^2 + 0 - 2\beta W^2 \sum_{s=1}^S dq_s dt \\
&\quad - 2\sigma W^2 \sum_{s=1}^S dq_s dz \\
&= \sigma^2 W^2 dt + 2W^2 \sum_{i \neq j} dq_i dq_j + W^2 \sum_{s=1}^S (dq_s)^2 - 2\beta W^2 \sum_{s=1}^S dq_s dt \\
&\quad - 2\sigma W^2 \sum_{s=1}^S dq_s dz. \tag{6.7}
\end{aligned}$$

Substituting (6.7) and (6.2) into equation (6.6) and taking the expectation we get

$$\begin{aligned}
E[dF] &= \left[\frac{\partial F}{\partial t} + \beta W \frac{\partial F}{\partial W} + \frac{1}{2} \sigma^2 W^2 \frac{\partial^2 F}{\partial W^2} \right. \\
&\quad \left. + \frac{1}{2} W^2 \sum_{s=1}^S \phi_s^2 \lambda_s \frac{\partial^2 F}{\partial W^2} - W \sum_{s=1}^S \phi_s \lambda_s \frac{\partial F}{\partial W} \right] dt. \tag{6.8}
\end{aligned}$$

Now substituting equation (6.8) back into equation (6.5) and dividing through by dt , the value of the investment opportunity $F(W_\varphi, t, \tau)$ must satisfy

$$\frac{1}{2} W^2 (\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s) F_{WW} + W (\beta - \sum_{s=1}^S \phi_s \lambda_s) F_W + F_t - \gamma F = 0, \tag{6.9}$$

subject to the following boundary conditions

$$F(0, t, \tau) = 0, \tag{6.10}$$

$$F(W_\varphi, \varphi + \tau, \tau) = \max\{pW_{\varphi+\tau} - I, 0\}. \tag{6.11}$$

Equation (6.9) is difficult to solve analytically, but it is not difficult to obtain a solution numerically using the finite difference method. The finite difference approach suggested by Schwartz (1977) and Brennan and Schwartz (1978), solves the underlying differential equation by converting it into a difference equation.

There are two alternative ways of implementing the finite difference approach. The first, known as the implicit finite difference method, relates the value of the derivative at time $t + \Delta t$ to three alternative values at time t . The second, known as the explicit finite difference method, relates the value of the derivative at time t to three alternative values at time $t + \Delta t$. We shall use the explicit finite difference method to solve equation (6.9) subject to its associated boundary conditions (6.10) and (6.11). It has the advantage of being computationally simpler and avoids the need to specify further boundary conditions.

To make the explicit finite difference method more computationally efficient and to ensure the convergence of its numerical solution, we begin by making a logarithmic transformation on W . Define

$$Z \equiv \text{Ln}(W), \quad (6.12)$$

$$H(Z, t, \tau) \equiv F(W, t, \tau), \quad (6.13)$$

so that

$$F_W = H_Z e^{-z}, F_{WW} = (H_{ZZ} - H_Z) e^{-2Z}, F_t = H_t. \quad (6.14)$$

Equation (6.9) then becomes

$$\frac{1}{2} \left(\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s \right) H_{ZZ} + \left(\beta - \sum_{s=1}^S \phi_s \lambda_s - \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{2} \right) H_Z + H_t - \gamma H = 0. \quad (6.15)$$

Notice that equation (6.15) unlike (6.9) is a partial differential equation with constant coefficients. This simplifies the numerical analysis, and makes it possible to employ an explicit finite difference approximation to (6.15), whereas the explicit finite difference approximation to (6.9) is in general unstable.³

The finite difference method transforms continuous variables Z and t into discrete variables, and replaces the partial derivatives in equation (6.15) with finite

³For issues relating to the convergence and stability of the explicit difference method, see, for example Geske and Shatri (1985), or Hull and White (1990).

differences. Define Δt and ΔZ to be the discrete increments in the value of t and Z , respectively. A grid consisting of a total of $(N+1)(M+1)$ points is constructed for considering values of H when the time t is equal to

$$\varphi, \varphi + \Delta t, \varphi + 2\Delta t, \dots, \varphi + \tau,$$

and Z is equal to

$$\psi, \psi + \Delta Z, \psi + 2\Delta Z, \dots, \Psi,$$

where $\psi = \text{Ln}(\min W)$ and $\Psi = \text{Ln}(\max W)$. $\min W$ and $\max W$ are the smallest and largest possible values of W considered by the model. In addition, $\text{Ln}(0)$ is undefined so we let $\min W \geq 1$ so that $\psi \geq 0$.

The (i, j) ($i = 0, 1, \dots, N, j = 0, 1, 2, \dots, M$) point on the grid is the point that corresponds to $\varphi + i\Delta t$ and $\psi + j\Delta Z$. We denote by H_{ij} the value of the option H at the (i, j) point. The partial derivatives of H with respect to Z can be approximated as follows,

$$\frac{\partial H}{\partial Z} = \frac{H_{i+1,j+1} - H_{i+1,j-1}}{2\Delta Z}, \quad (6.16)$$

$$\frac{\partial^2 H}{\partial Z^2} = \frac{H_{i+1,j+1} + H_{i+1,j-1} - 2H_{i+1,j}}{\Delta Z^2}, \quad (6.17)$$

and the time derivative is approximated as

$$\frac{\partial H}{\partial t} = \frac{H_{i+1,j} - H_{i,j}}{\Delta t}, \quad (6.18)$$

for $i = 0, 1, \dots, N-1$ and $j = 1, 2, \dots, M-1$.

Substituting equations (6.16), (6.17) and (6.18) into the differential equation (6.15) gives

$$H_{i,j} = \frac{1}{1 + \gamma\Delta t} (aH_{i+1,j-1} + bH_{i+1,j} + cH_{i+1,j+1}), \quad (6.19)$$

where

$$a = \left(-\frac{\beta - \sum_{s=1}^S \phi_s \lambda_s - \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{2}}{2\Delta Z} + \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{2\Delta Z^2} \right) \Delta t,$$

$$b = 1 - \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{\Delta Z^2} \Delta t, \quad \text{and}$$

$$c = \left(\frac{\beta - \sum_{s=1}^S \phi_s \lambda_s - \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{2}}{2\Delta Z} + \frac{\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s}{2\Delta Z^2} \right) \Delta t.$$

For any value of i ($i = 0, 1, \dots, N$), equation (6.19) allows us to solve for H_{ij} in terms of $H_{i+1,j}$. The extreme values $H_{N,j}$ ($j = 0, 1, \dots, M$) are given by the boundary condition (6.11), which says the value of the option F at time $\varphi + \tau$ is $\max\{pW_{\varphi+\tau} - I, 0\}$. This is equivalent to saying that $H(Z_\varphi, \varphi + \tau, \tau) = \max\{pe^{Z_{\varphi+\tau}} - I, 0\}$. Hence,

$$H_{N,j} = \max\{pe^{\psi+j\Delta Z} - I, 0\}, \quad j = 0, 1, \dots, M.$$

The value of the option after transformation, $H(Z_\varphi, \varphi, \tau)$ at current time φ , is by definition H_{0,j^*} , for some particular j^* such that $e^{\psi+j^*\Delta Z} = W_\varphi$. It then remains to use equation (6.19) to calculate H_{0,j^*} starting from $H_{N,j}$ and working backwards.

Notice that the coefficients of (6.19) are independent of j and that $a + b + c = 1$. For the stability and convergence of the explicit solution, a sufficient condition is that the coefficients a , b and c be positive as Δt and $\Delta Z \rightarrow 0$ (Ames, 1977). For this condition to be satisfied it is necessary that Δt and ΔZ be chosen so that

$$(\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s) \frac{\Delta t}{\Delta Z^2} < 1. \quad (6.20)$$

The simplest procedure is to let Δt and ΔZ approach 0 in such a way that $\Delta t/\Delta Z^2$ remains constant and less than $1/(\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s)$. In most cases, a good choice for ΔZ is $\sqrt{3(\sigma^2 + \sum_{s=1}^S \phi_s^2 \lambda_s)\Delta t}$ (Hull, 2008, p. 443).

6.3 Calculations of Profitability for (s, n) and FI Policies

Now we have formulated the option valuation model to value an investment opportunity (i.e., whether to invest in clinical trials in τ years time) that is subject to multiple sources of jumps. To incorporate this option value into the OPRRA model, we set $S = 3$. The three sources of jumps are assumed to be the type A, type B and type C obsolescence events. Recall that f_i denotes the fraction by which the occurrence of a type i obsolescence event reduces the value of the new

drug, and ξ_i is the corresponding arrival rate, $i = A, B, C$. Hence we have

$$\begin{aligned}\phi_1 &= f_A, \phi_2 = f_B, \phi_3 = f_C, \text{ and} \\ \lambda_1 &= \xi_A, \lambda_2 = \xi_B, \lambda_3 = \xi_C.\end{aligned}$$

The value of the investment I should equal the total expected cost of clinical trials, hence $I = c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}$.

With the assumptions set out above, we can now incorporate the option valuation model in the calculations of profitability in the different candidate drug selection policies that we have considered in previous chapters. Note that in all of our previous calculations we evaluate the expected total reward as the net reward after deduction of the cost of clinical trials. This expected net present value can be replaced by $F(W_\varphi, t, \tau)$ as defined in equation (6.4), which itself is an expected maximized net present value. The corresponding expected costs concerning pre-clinical stages stay as they are in previous calculations.

First we consider the calculations for (s, n) policies. The expected net present value of the first CD given that the first LS is successfully optimized in equation (4.3) can be replaced by $F(W_0, 0, t_1 + t_2 + t_4)$, where W is given by equation (6.2) with starting value W_0 , and has types A, B and C obsolescence events as distinct sources of jumps. The option value of the second CD from the same LS is $F(\eta W_0, 0, t_1+t_2+t_4+t_5)$, and for the k^{th} CD the option value is $F(\eta^{k-1} W_0, 0, t_1+t_2+t_4+(k-1)t_5)$.

To generalize the result further we let $T = t_1 + t_2 + t_4$ and use a subscript ‘o’ in the notation that follows to differentiate from corresponding notation used in previous chapters. We can write

$$\begin{aligned}R_o(1, 1) &= p_1(p_{21} + p_{22})p_4[F(W_0, 0, T) + F(\eta W_0, 0, T + t_5) + \dots \\ &\quad + F(\eta^{k_1-1} W_0, 0, T + (k_1 - 1)t_5)],\end{aligned}$$

$$\begin{aligned}
R_o(1, n) &= p_1(p_{21} + p_{22})p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + (m-1)t_5) \right] \\
&\quad + p_1(p_{21}p_3 + p_{22})q_4p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + t_4 + (m-1)t_5) \right] \\
&\quad \dots \\
&\quad + p_1(p_{21}p_3^{n-1} + p_{22}p_3^{n-2})q_4^{n-1}p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + (n-1)t_4 + (m-1)t_5) \right] \\
&= \sum_{j=1}^n p_1(p_{21}p_3^{j-1} + p_{22}p_3^{j-2})q_4^{j-1}p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + (j-1)t_4 + (m-1)t_5) \right]. \\
&\quad \text{(with the convention that } p_3^{-1} = 1 \text{)}
\end{aligned}$$

The value of the first CD from the second successfully optimized LS given that k_1 CDs have been taken from the first successfully optimized LS in equation (4.4) can be written as $F((p_b\eta^{k_1} + 1 - p_b)W_0, 0, T + \rho t_4 + (k_1 - 1)t_5)$, and we have

$$\begin{aligned}
R_o(2, 2) &= p_1(p_{21} + p_{22})p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + (m-1)t_5) \right] \\
&\quad + p_1(p_{21}p_3 + p_{22})q_4p_4 \left[\sum_{m=1}^{k_1} F(\eta^{m-1}W_0, 0, T + t_4 + (m-1)t_5) \right] \\
&\quad + p_1(p_{21}p_3 + p_{22})p_4^2 \left[\sum_{m=1}^{k_2} F((p_b\eta^{k_1} + 1 - p_b)\eta^{m-1}W_0, 0, \right. \\
&\quad \left. T + (k_1 - 1)t_5 + \rho t_4 + (m-1)t_5) \right],
\end{aligned}$$

$$\begin{aligned}
R_o(2, n) &= R_o(1, n) + (p_1p_{21}p_3p_4^2 + p_1p_{22}p_4^2) \left[\sum_{j=1}^{n-1} \sum_{i=1}^j (p_3q_4)^{j-1} \right. \\
&\quad \left(\sum_{m=1}^{k_2} F((p_b\eta^{k_1} + 1 - p_b)\eta^{m-1}W_0, 0, \right. \\
&\quad \left. T + (i-1)t_4 + (j-i+1)\rho t_4 + (k_1 - 2 + m)t_5) \right) \left. \right].
\end{aligned}$$

Finally, for $s = 3$ we have

$$R_o(3, n) = R_o(2, n) + (p_1 p_{21} p_3 p_4^3 + p_1 p_{22} p_4^3) \left[\sum_{j=1}^{n-2} p_3^j p_4^{j-1} \sum_{i=1}^j \times \right. \\ \left. \left(\sum_{m=1}^{k_3} F((p_b \eta^{k_1} + 1 - p_b)(p_b \eta^{k_2} + 1 - p_b) \eta^{m-1} W_0, 0, \right. \right. \\ \left. \left. T + (i-1)t_4 + (j-i+2)\rho t_4 + (k_1 + k_2 + m - 3)t_5) \right) \right].$$

Note that to calculate $R_o(1, n)$, $R_o(2, n)$ and $R_o(3, n)$ we need to apply the explicit finite difference procedure for nk_1 , $nk_1 + n(n-1)k_2/2$ and $nk_1 + n(n-1)k_2/2 + (n-1)(n-2)k_3/2$ times, respectively.

We can calculate the optimal allocations of effort at each successive stage of a pre-clinical R&D project by using the profitability index criterion and the internal rate of return criterion as we did for the initial model described in section 4.3.3. Because the calculation of $F(W_\varphi, t, \tau)$ depends on the discount rate γ , this implies that when optimizing the IRR with respect to the effort allocations we need to calculate $R_o(s, n)$ using the explicit finite difference method at every iteration in the Newton-Raphson procedure. However, for moderate degrees of accuracy (e.g., by setting $\Delta t = 0.001$) the explicit finite difference procedure is reasonably fast. Therefore the above task should not be too computationally challenging.

Next we consider the calculations needed in applying FI policies. In an FI policy, the values of $PI(CD)$, $PI(LS)$ and the reference PI are compared at each decision point. The decision to take an additional CD means that we hold a call option which matures in t_5 years' time. Hence the expected value of the next CD is

$$r_{oCD} = e^{-\gamma t} F(W_t, t, t_5) \\ = e^{-\gamma t} F\left(\prod_{i \in \Theta \setminus S} (p_b \eta^{b_i} + 1 - p_b) \eta^k W_0, t, t_5\right), \quad (6.21)$$

where t is the time at which the decision needs to be taken, and b_i is the number of CDs selected from the i^{th} optimized LS, Θ is the set of all LS from which CDs

have been selected, S is the youngest LS from which we will choose the next CD, and κ is the number of CDs we have taken from the youngest LS.⁴

Similarly, the decision to optimize a new LS means that we will in effect hold a sequence of call options with the first one maturing in t_4 years' time. As we have noted in section 3.1, the value of $PI(LS)$ depends on whether stage 3 is to be carried out alongside stage 4, and determining its value involves evaluating the rewards and costs of two successive LS. The relevant rewards in determining $PI(LS)$ are the expected reward for a single attempt to optimize a new LS, the expected reward for a further attempt to optimize a new LS following a first unsuccessful attempt, and the expected reward for a further attempt to optimize a new LS following a first successful attempt. Denote these expected rewards by r_{o1} , r_{o2} and r_{o3} , respectively.

Hence

$$r_{o1} = e^{-\gamma t} p_4 \sum_{m=1}^{k^*} F\left(\prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) W_0, t, t_4 + (m-1)t_5\right), \quad (6.22)$$

$$r_{o2} = e^{-\gamma t} p_4 \sum_{m=1}^{k^*} F\left(\prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) W_0, t, 2t_4 + (m-1)t_5\right), \quad (6.23)$$

$$r_{o3} = e^{-\gamma t} p_4 \sum_{m=1}^{k^*} F\left(\prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) (p_b \eta^{k^*} + 1 - p_b) W_0, t, t_4 + \rho t_4 + (k^* + m - 2)t_5\right). \quad (6.24)$$

The above calculations are for FI policies in a fixed probabilities setting. Equations (6.21) to (6.24) correspond to equations (B.1), (B.2), (B.3), and (B.4) in Appendix B.1, respectively. The calculations needed in determining $PI(CD)$ and $PI(LS)$ in an adaptive probabilities setting may be carried out in a similar manner.

⁴This notation and the notation that follows is defined in Appendix B.1.

Chapter 7

Tests on Optimization Procedures

In this chapter we report the findings of investigations conducted on OPRRA's profitability index optimization procedures. OPRRA offers three different types of policy that may be used to optimize the PI, namely, the (s, n) policy and the FI policy in its restricted and its general forms. The main aim is to assess and compare the performances of the associated procedures so as to give recommendations on the use of OPRRA.

7.1 Optimization Procedures

The procedures used to evaluate, and in some cases also to optimize, PI are as follows.

- (s, n) expectation algorithm (*SNEA*). This algorithm evaluates PI for given resource rate allocations u_1, u_2, u_3, u_4 and u_5 for the five stages of discovery and given values of s and n . It can optimize the allocations for given s and n using numerical optimization procedures. The values s^* and n^* of s and n which lead to the overall optimal value of PI^* may be determined by repeatedly running the algorithm with different values of s and n .
- FI simulation. This algorithm provides a distribution of outcomes for given

allocations, using FI to select successive CDs. It is used to provide estimates for PI using optimal allocations calculated by *SNEA* for $(s, n) = (s^*, n^*)$ (*FISIM**), and for $(s, n) = (1, 1)$ (*FISIM11*), with the reference PI equal to the optimal PI calculated with *SNEA* for the corresponding (s, n) .

- Restricted FI expectation algorithm (*RFIEA*). This is the algorithm described in section 3.2 for sequential CD selection. In OPRRA there are two extensions to that algorithm. First, the expression for r_{ij} is modified to allow for the possibility that $\gamma_1 \neq \gamma_2$. Second, there is a modification which ensures stage 3 is only carried out if there is a possibility that the resulting LS will be used. It is implemented with allocations optimized using *SNEA* with $(s, n) = (1, 1)$.

As the general FI policy is too complex to be evaluated analytically, we assess the performance of FI policies by simulating the candidate drug selection process.

A simulation of any system or process in which there are inherently random components requires a method of generating numbers that are random. Thus a good random number generator (RNG) is essential for obtaining accurate simulation results. The RNG we use in the simulation study is a mixed generator recommended in the book by Press and Teukolsky (2007), with a period length of 1.8×10^{19} .

A replication is a run of a simulation model that uses a specific stream of random numbers. Different streams of random numbers are used for different replications of simulation in our study to ensure the accuracy of the estimate. Each run of the simulation results in a reward and a cost for the project. Let X_j and Y_j be the reward and cost obtained on the j^{th} run for $j = 1, 2, \dots, n$. Then the X_j 's and Y_j 's are i.i.d. random variables. Thus the sample means $\bar{X} = \frac{\sum_{i=1}^n X_i}{n}$ and $\bar{Y} = \frac{\sum_{i=1}^n Y_i}{n}$ are unbiased estimators of the population means for reward and cost, and the sample variances $S_X^2 = \frac{\sum_{i=1}^n [X_i - \bar{X}]^2}{n-1}$ and $S_Y^2 = \frac{\sum_{i=1}^n [Y_i - \bar{Y}]^2}{n-1}$ are unbiased estimators of the population variances for reward and cost. Denote by $\overline{PI} = \frac{\bar{X}}{\bar{Y}}$ the average *PI* for the project. The corresponding estimated variance for \overline{PI} is derived using the delta method (see, for example Casella and Berger, 2001), which gives

$$\widehat{Var}\left(\frac{\bar{X}}{\bar{Y}}\right) = \frac{1}{n} \left[\frac{\bar{X}^2}{\bar{Y}^4} S_Y^2 + \frac{S_X^2}{\bar{Y}^2} - \frac{2\bar{X}}{\bar{Y}^3} Cov(\bar{X}, \bar{Y}) \right],$$

and therefore the 95% confidence interval for \overline{PI} is

$$\left[\frac{\bar{X}}{\bar{Y}} - 1.96 \sqrt{\widehat{Var}\left(\frac{\bar{X}}{\bar{Y}}\right)}, \frac{\bar{X}}{\bar{Y}} + 1.96 \sqrt{\widehat{Var}\left(\frac{\bar{X}}{\bar{Y}}\right)} \right].$$

We set the number of runs $n = 100,000$ in all of the simulation tests.

7.2 Parameter Values

The OPRRA model requires 33 input parameters. They are: two firm-specific parameters γ and α , six project-specific parameters W , ν_1 , ν_2 , λ , p_a , and p_b , and 25 stage/phase specific parameters p_1 , p_{21} , p_{22} , p_3 , p_4 , t_1 , t_2 , t_3 , t_4 , t_5 , u_1 , u_2 , u_3 , u_4 , u_5 , p_I , p_{II} , p_{III} , t_I , t_{II} , t_{III} , c_I , c_{II} , c_{III} and ρ . In this study p_{22} is set to 0 so that we are able to compare the other procedures with *RPIEA*, in which $p_{22} = 0$.

Throughout the studies in this chapter and in chapter 8 effectiveness function (4.2) is assumed, with a minimal duration of 3 months, a cap of 70 scientists for each stage, a 15% loss of efficiency with $3u_f$ scientists, and a 32% loss of efficiency with $5u_f$ scientists as compared with the maximal efficiency with $u_f = 14$ scientists. To simplify our discussion we also make the broadly realistic assumption that the costs of clinical trials are proportional to their durations, (i.e., $\frac{c_I}{t_I}$, $\frac{c_{II}}{t_{II}}$, and $\frac{c_{III}}{t_{III}}$ are constants), so that the values of t_I , t_{II} and t_{III} are determined by c_I , c_{II} and c_{III} .

500 sets of input parameters are chosen to study the various optimization procedures. Each of the input parameters is randomly and independently drawn from its own specified uniform distribution displayed in Table 7.1. The input parameter ranges are chosen so as to include many of the values which are likely to occur in practice. However, 155 of the simulated projects are excluded from the study due to negative PI^* or optimization failure for SNEA¹. All numbers in

¹The Newton-Raphson procedure is abandoned if the number of iterations exceeds the maximum number of allowed iterations. This accounts for 5 out of the 155 excluded projects

the results are rounded to an accuracy of three decimal places.

7.3 Numerical Results

For the remaining 345 projects, the PI^* ranges from 0.001 to 4.093 and has a mean value of 0.292. Only 20 projects are profitable, with $PI^* > 1$, the vast majority of projects being loss making². 198 projects have $(s^*, n^*) = (1, 1)$, 77 projects have $(s^*, n^*) = (1, n^* > 1)$, 43 projects have $(s^*, n^*) = (2, n^*)$ and 27 projects have $(s^*, n^*) = (3, n^*)$.

For the 57% of projects with $(s^*, n^*) = (1, 1)$, there are just three distinct procedures: *SNEA*, $FISIM11 = FISIM^*$ and *RFIEA*. The optimal solutions from all three solution algorithms take simple forms. They lead to the same decisions on the selection of CDs and give the same value for PI .

When $(s^*, n^*) \neq (1, 1)$, we use *SNEA* to calculate both PI^* and $PI11$, where $PI11$ is the optimal PI for the project when (s, n) is set at $(1, 1)$. Define $\Delta = \frac{PI^*}{PI11}$, the ratio of the overall optimal PI to the optimal PI at $(s, n) = (1, 1)$. Let PI_{FI} , PI_{FI11} and PI_{RFI} be the outputs of $FISIM^*$, $FISIM11$ and *RFIEA*, respectively. We set PI^* as the baseline measure and assess the performance of $FISIM^*$, $FISIM11$ and *RFIEA* by comparing their outputs to the baseline. Define $P_{FI} = \frac{PI_{FI}}{PI^*}$, $P_{FI11} = \frac{PI_{FI11}}{PI^*}$ and $P_{RFI} = \frac{PI_{RFI}}{PI^*}$ to be the performance scores of $FISIM^*$, $FISIM11$ and *RFIEA*. A performance score greater than one means outperforming the baseline, whereas a score less than one means underperforming.

Figure 7.1 plots the distribution of Δ for cases where $(s^*, n^*) \neq (1, 1)$. The mean value of Δ is 1.086, indicating that on average PI^* only outperforms $PI11$ by less than 10%. By plotting Δ against s^* , we see that on average Δ tends to be bigger when s^* is bigger (Figure 7.2). The maximum value of Δ is 1.667, which occurred with $s^* = 3$. (The relationship between the parameter values and s^* is discussed in the next chapter.)

²Note this is broadly in line with the distribution of profitability for commercial pharmaceutical projects, see section 1.1.

Table 7.1: *Input parameter ranges.*

Parameter	Minimum	Maximum
λ	0.5	0.9
u_1	3	15
u_2	6	30
u_3	6	30
u_4	6	30
u_5	6	30
t_1	0.5	4.5
t_2	0.3	3.0
t_3	0.3	3.0
t_4	0.8	7.5
t_5	0.5	4.5
p_1	0.3	0.9
p_{21}	0.3	0.6
p_3	0.4	0.8
p_4	0.4	0.8
p_a	0.84	0.96
p_b	0.4	0.8
ρ	0.7	0.9
ν_1	0.03	0.25
ν_2	0.03	0.25
γ	0.05	0.15
p_I	0.4	0.7
p_{II}	0.3	0.6
p_{III}	0.5	0.8
c_I	0.2 million	0.8 million
c_{II}	0.6 million	2.4 million
c_{III}	6.5 million	26 million
W	500 million	2000 million
α	0.2 million	2 million

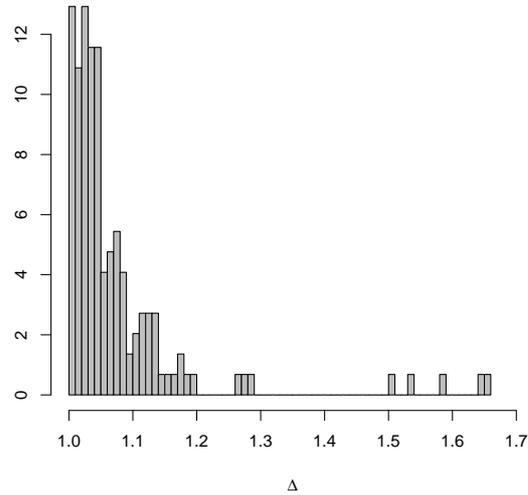


Figure 7.1: The distribution of Δ .

Figure 7.3 shows the plots of P_{FI} , P_{FI11} and P_{RFI} against s^* . The plots reveal that the average performance scores of $FISIM^*$ and $FISIM11$ both appear to be equal to one for all values of s^* , suggesting that $FISIM^*$ and $FISIM11$ on average perform equally as well as $SNEA$. It also suggests that $FISIM11$ is as good as $FISIM^*$ as little difference can be seen between the performances of these two procedures. This is confirmed by a p-value of 0.1 from the two sample paired t-test on the values of PI_{FI} and PI_{FI11} , which does not reject the null hypothesis that the true difference in means is equal to 0. P_{FI} ranges from 0.854 to 1.189. However, cases where $FISIM^*$ either underperforms or outperforms $SNEA$ by 10% are relatively few, (i.e., $P_{FI} < 0.9$ or > 1.1), which only account for about 7% of the cases. Note that PI_{FI} and PI_{FI11} are obtained by Monte Carlo simulation and are thus subject to random variation.

Three linear models are fitted using P_{FI} as the response and Δ as the explanatory variable for $s^* = 1$ (with $n^* > 1$), $s^* = 2$ and $s^* = 3$, respectively. The results of the linear models show that Δ has no effect on P_{FI} when $s^* = 1$ and 2. However, there is strong evidence ($P < 0.0001$) of positive correlation between P_{FI} and Δ when $s^* = 3$. This suggests that the FI simulation method is likely

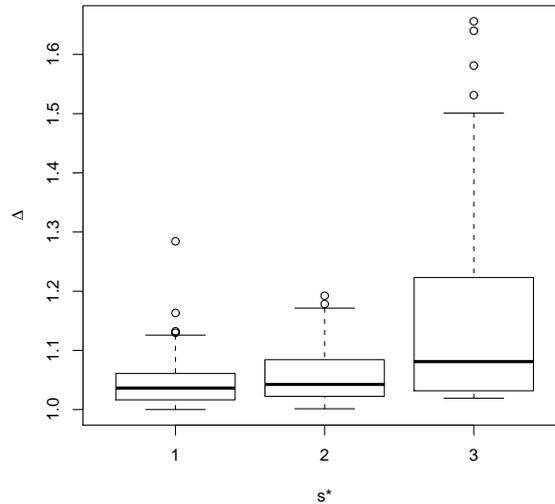


Figure 7.2: Plot of Δ against s^* .

to work better when both Δ and the optimal number of LS to be optimized are big. An effect estimate of 0.161 suggests that a one unit increase in Δ would lead to a 16% increase in performance of $FISIM^*$ relative to $SNEA$ when $s^* = 3$. This is also true for P_{FI11} . We therefore recommend using both $SNEA$ and FI simulation when $s^* = 3$ and Δ is bigger than the average value of 1.1

The performance of $RFIEA$, on the other hand, decreases as s^* increases. The performance score P_{RFI} varies from .998 to 1.006 when $(s^*, n^*) = (1, n^* > 1)$, indicating that $RFIEA$ and $SNEA$ have almost exactly the same performance (to 3 d.p.). This result is a bit surprising as we proved in section 3.2 that $RFIEA$ is optimal when the number of LS to be optimized is restricted to one. Further investigations were conducted on the breakdowns of $RFIEA$ and $SNEA$ results. We have seen that although $RFIEA$ has the advantage of allowing different number of CDs at each attempt to optimize an LS, this advantage is undermined by the heavy discounting of the rewards and costs of later CDs. The effects of later CDs on the overall PI is very small. Hence $RFIEA$ does not outperform $SNEA$ very much in terms of PI. The mean of P_{RFI} is 0.971 and 0.913 when $s^* = 2$ and $s^* = 3$, respectively, though P_{RFI} can get as low as 0.724 when

$s^* = 3$. This means that *RFIEA* is likely to underperform *SNEA* when $s^* > 1$, which is understandable as *RFIEA* is restricted to allow at most one LS.

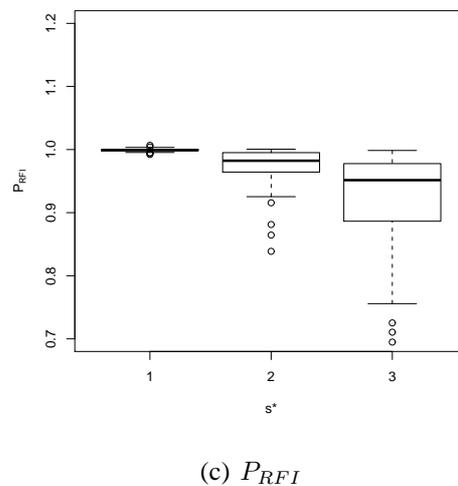
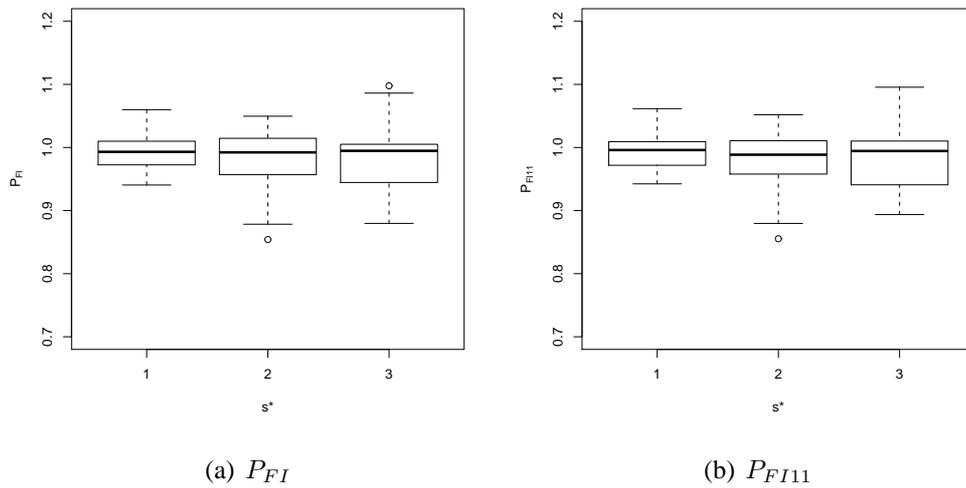


Figure 7.3: Plots of P_{FI} , P_{FI11} and P_{RFI} against s^* .

The most striking aspect of the results is the strong performance of *SNEA*. *FISIM** and *FISIM11* on average perform equally as well as *SNEA*, and so does *RFIEA* when $s^* = 1$. There are cases when *FISIM** and *FISIM11* outperform *SNEA*, (i.e., the lower limit of the 95% CI of PI_{FI} or $PI_{FI11} > PI^*$) but these cases are rare (6%) and the level of outperformance is not substantial

(max 19%). This means we can use *SNEA* with confidence. Here we present seven cases selected from the study. Their parameter values are shown in Table 7.3. The results are displayed in Table 7.2.

Based on the results of this study, we give the following recommendation on the use of OPRRA regarding the PI optimization procedures we discussed in this chapter.

- Start with *SNEA*. It is the most comprehensive procedure as it not only calculates the optimal PI but also the optimal allocations. For given values of s and n , *SNEA* gives the optimal allocations and PI, and it is also able to give s^* , n^* and PI^* by repeatedly running the algorithm.
- When $s^* = 1$ interested users can use *RFIEA* to check the optimal number of attempts for stage 4 and the optimal number of CDs to take at each attempt. However the resulting overall optimal PI will not be much different from the PI given by *SNEA*.
- When $s^* = 3$ and $\Delta > 1.1$ we recommend using both *SNEA* and FI simulation, as the latter method is likely to achieve a higher PI in these circumstances.

Table 7.2: Profitability Index values for the selected cases.

Project	s^*	n^*	Δ	SNEA		FISIM and 95% CI			RFIEA	
				(s^*, n^*)	(1, 1)	(s^*, n^*)	(1, 1)		(1, 1)	
1	1	1	1	2.049	*	2.066(2.008 2.123)		*	2.049	
2	1	5	1.126	0.143	0.127	0.143(0.140 0.146)		0.141(0.138 0.143)		0.143
3	1	6	1.135	0.312	0.275	0.311(0.306 0.317)		0.311(0.306 0.316)		0.313
4	2	2	1.032	1.418	1.374	1.448(1.427 1.468)		1.453(1.433 1.474)		1.387
5	2	5	1.046	0.113	0.108	0.115(0.111 0.119)		0.115(0.111 0.119)		0.113
6	3	3	1.146	0.455	0.397	0.458 (0.451 0.465)		0.485(0.478 0.493)		0.402
7	3	4	1.156	0.830	0.718	0.894 (0.876 0.912)		0.888(0.870 0.906)		0.762

Table 7.3: Parameter values for the selected cases.

Project 1	$(s^*, n^*) = (1, 1)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.304, 0.342, 0.788, 0.620, 0.913, 0.652, 0.662, 0.575, 0.624, 1.353, 2.917, 3.215)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.088, 0.057, 0.100, 0.869, 0.796, 0.214m, 1263.526m, 0.541m, 1.750m, 13.930m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (11, 15, 15, 12, 10, 0.594, 0.319, 1.622, 6.933, 3.129)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (21, 19, 12, 50, 53, 0.318, 0.250, 2.084, 0.758)$
Project 2	$(s^*, n^*) = (1, 5)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.565, 0.495, 0.654, 0.464, 0.892, 0.517, 0.528, 0.534, 0.780, 0.618, 2.707, 2.423)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.132, 0.140, 0.228, 0.869, 0.739, 1.155m, 1303.707m, 0.247m, 1.624m, 10.500m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (6, 20, 16, 6, 28, 2.686, 1.913, 0.789, 3.392, 0.818)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (38, 51, 33, 70, 67, 0.484, 0.932, 0.428, 0.462)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (6, 48, 32, 66, 70, 0.503, 0.967, 0.439, 0.453)$
Project 3	$(s^*, n^*) = (1, 6)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.606, 0.438, 0.564, 0.446, 0.871, 0.727, 0.628, 0.506, 0.601, 1.763, 2.935, 1.605)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.073, 0.041, 0.141, 0.611, 0.818, 1.452m, 1959.007m, 0.705m, 1.761m, 6.954m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (3, 26, 17, 17, 9, 2.455, 2.775, 0.674, 1.845, 2.492)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (25, 34, 17, 63, 58, 0.307, 2.240, 0.689, 0.516)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (24, 33, 15, 51, 61, 0.318, 2.291, 0.773, 0.503)$
Project 4	$(s^*, n^*) = (2, 2)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.525, 0.584, 0.459, 0.738, 0.873, 0.634, 0.631, 0.583, 0.747, 0.543, 1.495, 3.231)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.100, 0.086, 0.129, 0.677, 0.748, 0.494m, 1687.803m, 0.217m, 0.897m, 14.001m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (11, 6, 26, 23, 9, 3.610, 0.820, 0.530, 0.830, 4.285)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (32, 19, 37, 65, 49, 1.356, 0.25, 0.403, 0.978)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (32, 19, 33, 55, 48, 1.356, 0.25, 0.439, 0.991)$
Project 5	$(s^*, n^*) = (2, 5)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.360, 0.345, 0.720, 0.682, 0.881, 0.714, 0.580, 0.576, 0.517, 1.095, 3.142, 4.811)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.055, 0.167, 0.186, 0.545, 0.734, 0.520m, 1880.329m, 0.438m, 1.885m, 20.848m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (11, 8, 20, 21, 16, 3.268, 0.380, 1.419, 1.637, 2.245)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (42, 12, 50, 70, 70, 1.007, 0.250, 0.708, 0.753)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (40, 12, 50, 70, 70, 1.041, 0.250, 0.708, 0.753)$
Project 6	$(s^*, n^*) = (3, 3)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.828, 0.589, 0.679, 0.790, 0.894, 0.639, 0.448, 0.529, 0.710, 0.595, 1.698, 2.738)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.133, 0.049, 0.230, 0.650, 0.755, 0.600m, 1717.970m, 0.238m, 1.019m, 11.863m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (3, 25, 13, 9, 11, 4.431, 1.796, 0.768, 2.387, 4.231)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (26, 35, 21, 58, 42, 0.542, 1.370, 0.491, 1.302)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (27, 35, 19, 49, 52, 0.522, 1.373, 0.541, 1.140)$
Project 7	$(s^*, n^*) = (3, 4)$
	$(p_1, p_{21}, p_3, p_4, p_a, p_b, p_I, p_{II}, p_{III}, t_I, t_{II}, t_{III}) = (0.417, 0.546, 0.678, 0.604, 0.955, 0.443, 0.555, 0.464, 0.694, 1.795, 2.680, 4.120)$
	$(\gamma, \nu_1, \nu_2, \lambda, \rho, \alpha, W, c_I, c_{III}, c_{III}) = (0.052, 0.093, 0.130, 0.800, 0.823, 0.453m, 1822.191m, 0.718m, 1.608m, 17.853m)$
	initial inputs $(u_1, u_2, u_3, u_4, u_5, t_1, t_2, t_3, t_4, t_5) = (5, 10, 15, 8, 19, 1.742, 2.859, 0.679, 2.431, 3.061)$
	optimal outputs: $(u_1^*, u_2^*, u_3^*, u_4^*, u_5^*, t_1^*, t_2^*, t_4^*, t_5^*) = (27, 38, 22, 52, 48, 0.342, 0.861, 0.478, 1.483)$
	optimal outputs at $(s, n) = (1, 1)$: $(u_1^{11}, u_2^{11}, u_3^{11}, u_4^{11}, u_5^{11}, t_1^{11}, t_2^{11}, t_4^{11}, t_5^{11}) = (26, 36, 21, 47, 45, 0.348, 0.893, 0.514, 1.541)$

Chapter 8

Effects of Model Parameters

In the previous chapter we came to the conclusion that *SNEA* is the most comprehensive and reliable optimization procedure available and may be recommended with confidence most of the time. In this chapter we further investigate the relationship between the model parameters and the *SNEA* outputs.

The objectives of this study are to identify the sets of parameters that contribute most to the variations of PI^* and s^* under an (s, n) policy, and to estimate the effects of those parameters and their interactions. Knowing which parameters have the most influence on profitability is helpful in project selection and resource allocation.

8.1 Effects of Parameters on PI^*

8.1.1 Identifying Important Parameters

800 randomly generated projects are used in this study. The input parameters of these projects are randomly and independently taken from the intervals shown in Table 7.1. In addition, p_{22} is sampled from the interval $(0.15, 0.3)$. For each project we have (s^*, n^*) and PI^* as the outputs. 224 projects are excluded from the study for the same reason as in chapter 7.

An initial multiple regression model was fitted with the response PI^* and all 30 parameters as predictors. However, the plot of the residuals $\hat{\epsilon}_i$ against the fitted values \hat{y} presents strong evidence of heterogeneous behaviour in the residuals, prompting consideration of a possible change in the structural form of the model (Figure 8.1 left). An analysis of the profile log-likelihood from the Box-Cox transformation suggests a power transform of $\lambda = 0.2$ (Figure 8.1 right).

A full model is refitted with $\sqrt[5]{PI^*}$ as the response, and backward elimination is applied to select the set of important predictors. At each stage we remove the variable with the largest p-value over 0.05. Nine variables $u_1(p = 0.056)$, $t_1(p = 0.078)$, $u_3(p = 0.453)$, $t_3(p = 0.658)$, $p_3(p = 0.131)$, $p_a(p = 0.939)$, $p_b(p = 0.154)$, $\rho(p = 0.105)$, and $p_{22}(p = 0.094)$ appear to be insignificant in explaining the variation in $\sqrt[5]{PI^*}$ and are removed from the model by backward elimination. The resulting model consists of the remaining 21 predictors. Mallows's C_p statistic is used to assess its fit. The C_p statistic is defined as

$$C_p = \frac{RSS_p}{\hat{\sigma}^2} + 2p - n,$$

where $\hat{\sigma}^2$ is the estimate of the variance σ^2 from the full model with all possible explanatory variables, RSS_p is the residual sum of squares (RSS) from a model with p regression coefficients including the intercept, and n is the sample size. If a model lacks important explanatory variables, it will show greater residual variability and thus RSS_p will be large. On the other hand, adding variables in the model would reduce the RSS. The term $2p$ in the formula adds a penalty for having more explanatory variables than necessary. Thus it is a criterion which trades-off fit in terms of RSS against complexity. If a model with p regression coefficients fits, then $E[RSS_p] = (n - p)\sigma^2$ and $E[C_p] \approx p$. A model with a bad fit will have C_p much bigger than p . Models with small C_p statistics are looked at more favourably.

For a model of a given size, the C_p criterion will select the model with the smallest RSS . We plot the C_p statistic for the best model for each size in Figure 8.2. The figure suggests that our resulting model with 21 predictors has the smallest C_p statistic and thus fits well.

Table 8.1 shows the regression output for the model identified by backward

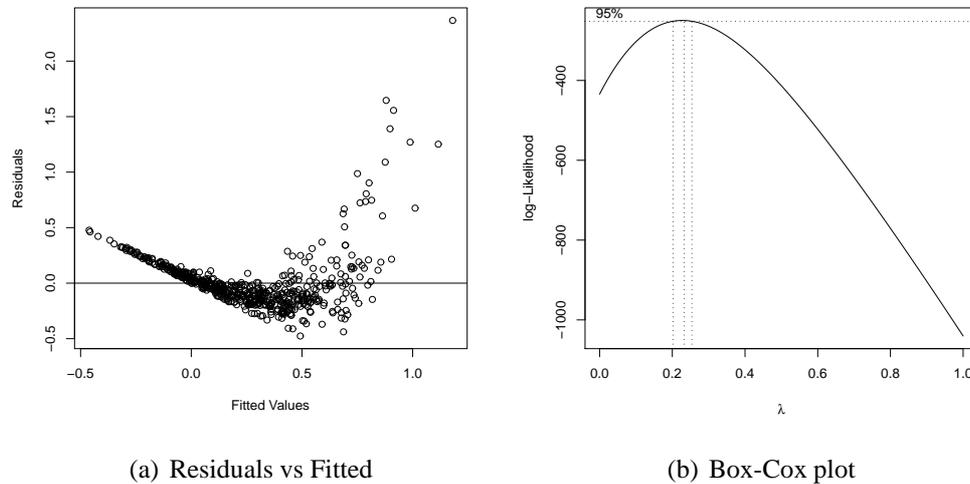


Figure 8.1: *On the left: Residual plot of the initial full model, which regresses PI^* against the model parameters. On the right: Profile log-likelihood from a Box-Cox analysis of the initial full model.*

elimination. The p-values (all < 0.005) show overwhelming significance for all 21 selected variables. We can see from the table that the variables α , u_4 , u_5 , t_4 , t_5 , ν_1 , ν_2 , γ , c_{II} and c_{III} have strong negative effects on profitability, with all p-values being less than $2e-16$, while W has a strong positive effect. The magnitudes of these effects and their interactions will be looked at more thoroughly through a two-level factorial design in the next section.

The diagnostic plot in Figure 8.3 shows that the assumption of constant variance is reasonable. Although Figure 8.4 indicates mild non-normality of errors the consequences can be reasonably ignored because the effects of non-normality are mitigated by large sample size.

Table 8.1: *Estimates of regression coefficients.*

Variable	Coefficient	Standard error	t-statistic	p-value
(Intercept)	1.204e+00	6.535e-02	18.422	< 2e-16
α	-1.422e-01	6.842e-03	-20.783	< 2e-16
λ	9.325e-02	3.122e-02	2.987	0.00294
u_2	-2.371e-03	5.207e-04	-4.552	6.53e-06
u_4	-7.185e-03	5.311e-04	-13.529	< 2e-16
u_5	-4.343e-03	5.185e-04	-8.377	< 2e-16
t_2	-3.483e-02	4.663e-03	-7.470	3.14e-13
t_4	-2.908e-02	1.862e-03	-15.621	< 2e-16
t_5	-2.708e-02	3.159e-03	-8.572	< 2e-16
p_1	1.150e-01	2.104e-02	5.465	7.01e-08
p_{21}	1.467e-01	4.800e-02	3.057	0.00234
p_4	1.840e-01	3.141e-02	5.858	8.02e-09
ν_1	-6.897e-01	5.966e-02	-11.562	< 2e-16
ν_2	-1.562e+00	5.866e-02	-26.623	< 2e-16
γ	-1.460e+00	1.246e-01	-11.713	< 2e-16
W	1.660e-04	8.703e-06	19.074	< 2e-16
c_I	-1.405e-01	2.072e-02	-6.781	3.07e-11
c_{II}	-6.671e-02	6.957e-03	-9.589	< 2e-16
c_{III}	-1.135e-02	6.587e-04	-17.238	< 2e-16
p_I	1.953e-01	4.168e-02	4.686	3.50e-06
p_{II}	2.816e-01	4.185e-02	6.727	4.32e-11
p_{III}	1.693e-01	4.118e-02	4.112	4.52e-05

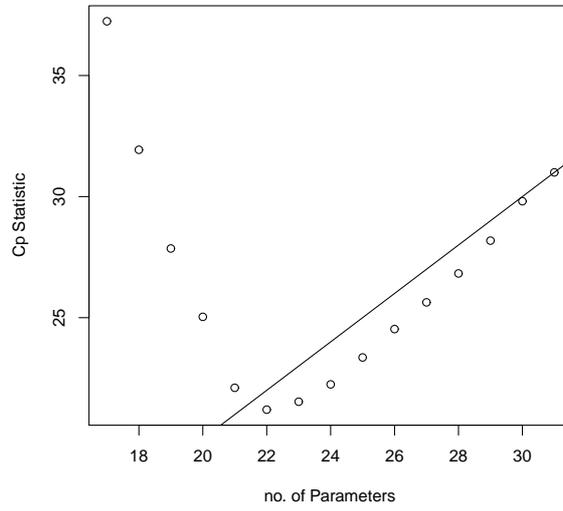


Figure 8.2: *The C_p plot.*

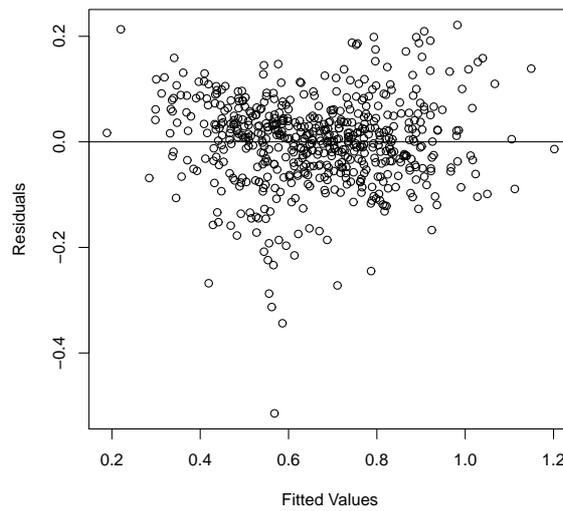


Figure 8.3: *Residuals vs. fitted for the final model.*

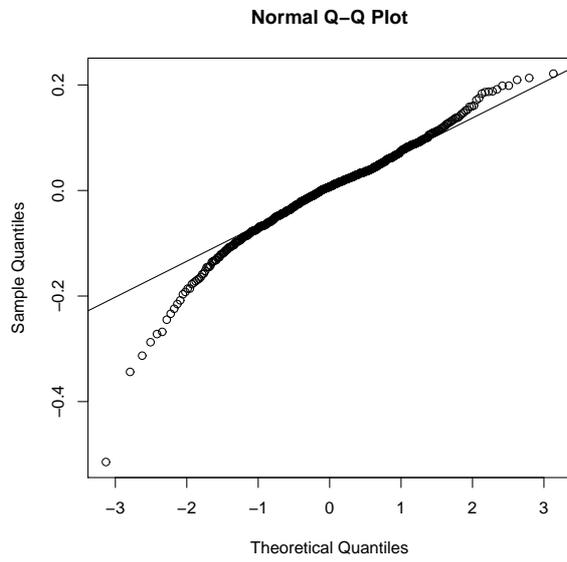


Figure 8.4: *Normality check for the final model.*

8.1.2 Factorial Design

A 2_{V}^{21-12} factorial design calculation was performed on these 21 factors with 512 runs. Resolution V designs are designs in which the main effects and the two-factor interactions do not have other main effects or two-factor interactions as their aliases, thus allowing the unbiased estimation of all the main effects and two-factor interactions, provided that all the three-factor and higher interactions are negligible. The high and low levels for each factor are set at the upper and lower quartiles of the input range. This choice of moderate values is to avoid having negative outputs. We control the remaining variables in the model by setting them at their mean values, giving $(u_1, u_3, t_1, t_3, p_{22}, p_3, p_a, p_b, \rho) = (8, 15, 1.5, 1.5, 0.2, 0.6, 0.9, 0.7, 0.9)$. The 21 factors selected for the design are as follows:

Table 8.2: *Low and high values for the selected factors.*

Factor	Name	Low Level (-1)	High Level (+1)
A	α	0.5 million	1.5 million
B	λ	0.6	0.8
C	u_2	10	20
D	u_4	10	20
E	u_5	10	20
F	t_2	1	2
G	t_4	2.5	5
H	t_5	1.5	3
J	p_1	0.5	0.8
K	p_{21}	0.4	0.6
L	p_4	0.5	0.7
M	ν_1	0.08	0.16
N	ν_2	0.08	0.16
O	γ	0.07	0.12
P	W	500 million	1500 million
Q	c_I	0.3 million	0.6 million
R	c_{II}	1 million	2 million
S	c_{III}	8 million	16 million
T	p_I	0.5	0.7
U	p_{II}	0.4	0.6
V	p_{III}	0.6	0.8

The response PI^* in the 512 runs ranges from 0.0002 to 4.2287, making the ratio of the maximum to minimum responses as high as 17383, which indicates that a transformation is highly likely to be needed. Again an analysis of the profile log-likelihood from the Box-Cox transformation suggests a transform of power $\lambda = 0.2$ would be the most appropriate.

Figure 8.5 presents a normal probability plot of the effect estimates using $\sqrt[5]{PI^*}$ as the response variable. It is a method attributed to Daniel (1959) to select effects in unreplicated factorials, where there are no internal estimates of the errors. The effects that are negligible are normally distributed, with mean zero and constant variance and will tend to fall along a straight line on the plot, whereas significant effects will not lie along the straight line. From examination of the display, the most apparent and important effects are W , α , ν_2 , c_{III} and c_{II} .

Table 8.3 contains the effect estimates and sum of squares for the selected effects. The percent contribution in Table 8.3 is obtained by taking each individual sum of squares and dividing by the total sum of squares and multiplying by 100. Note that the main effects really predominate, accounting for nearly 97 percent of the total variability, whereas the only significant interaction term is α - W , explaining just 0.48 percent of the total variability.

The resulting analysis of variance in Table 8.4 may be used to confirm the magnitude of the selected effects. From Table 8.4 we note that all the effects in the model are highly significant (every p-value < 0.0001).

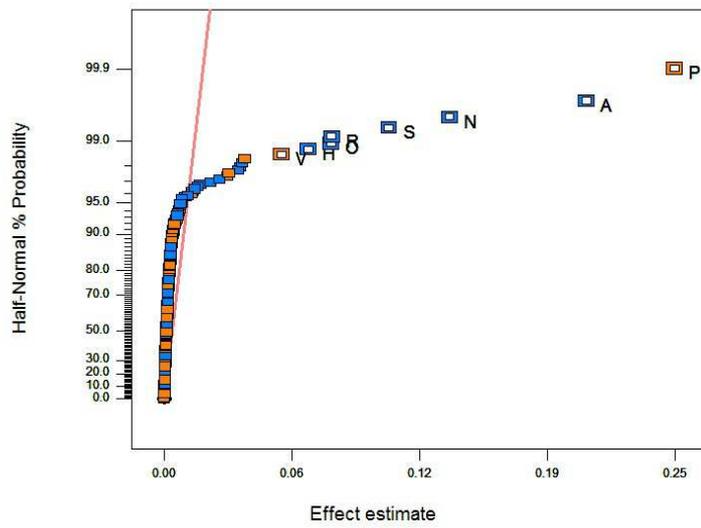


Figure 8.5: Normal probability plot of effects for the 2_{V}^{21-12} design.

Table 8.3: *Effect estimate summary.*

Factor	Effect Estimate	Sum of Squares	% Contribution
P- W	0.25	7.93	31.56
A- α	-0.2	5.25	20.89
N- ν_2	-0.14	2.51	10.00
S- c_{III}	-0.11	1.57	6.25
R- c_{II}	-0.083	0.88	3.50
O- γ	-0.083	0.87	3.46
G- t_4	-0.077	0.75	3.00
D- u_4	-0.075	0.72	2.86
H- t_5	-0.071	0.64	2.57
E- u_5	-0.071	0.63	2.52
U- p_{II}	0.064	0.51	2.05
V- p_{III}	0.057	0.41	1.65
T- p_I	0.054	0.38	1.49
M- ν_1	-0.046	0.27	1.08
L- p_4	0.043	0.23	0.92
Q- c_I	-0.038	0.18	0.72
C- u_2	-0.035	0.15	0.61
F- t_2	-0.033	0.14	0.56
B- λ	0.027	0.093	0.37
K- p_{21}	0.019	0.047	0.19
J- p_1	0.014	0.026	0.10
AP	-0.031	0.12	0.48

Table 8.4: Analysis of Variance.

Source of variation	sum of squares	df	mean square	F_0	p-value
Model	24.31	22	1.105	658.91	<0.0001
P- W	7.93	1	7.93	4728.68	<0.0001
A- α	5.25	1	5.25	3130.59	<0.0001
N- ν_2	2.51	1	2.51	1496.72	<0.0001
S- c_{III}	1.57	1	1.57	936.20	<0.0001
R- c_{II}	0.88	1	0.88	524.75	<0.0001
O- γ	0.87	1	0.87	518.78	<0.0001
G- t_4	0.75	1	0.75	447.23	<0.0001
D- u_4	0.72	1	0.72	429.34	<0.0001
H- t_5	0.64	1	0.64	381.63	<0.0001
E- u_5	0.63	1	0.63	375.67	<0.0001
U- p_{II}	0.51	1	0.51	304.11	<0.0001
V- p_{III}	0.41	1	0.41	244.48	<0.0001
T- p_I	0.38	1	0.38	226.50	<0.0001
M- ν_1	0.27	1	0.27	161.00	<0.0001
L- p_4	0.23	1	0.23	137.15	<0.0001
Q- c_I	0.18	1	0.18	107.33	<0.0001
C- u_2	0.15	1	0.15	89.45	<0.0001
F- t_2	0.14	1	0.14	83.48	<0.0001
B- λ	0.093	1	0.093	55.46	<0.0001
K- p_{21}	0.047	1	0.047	28.03	<0.0001
J- p_1	0.026	1	0.026	15.50	<0.0001
AP	0.12	1	0.12	71.56	<0.0001
Residual	0.82	489	0.0017		
Total	25.13	511			

Figure 8.6 presents a normal probability plot of the residuals and a plot of residuals versus fitted values. The first plot tests the assumption of normality and the latter tests the assumption of constant variance. Both plots are satisfactory, indicating that the usual normality assumptions are satisfied and also lending support to the adequacy of our model and the suitability of the power transform.

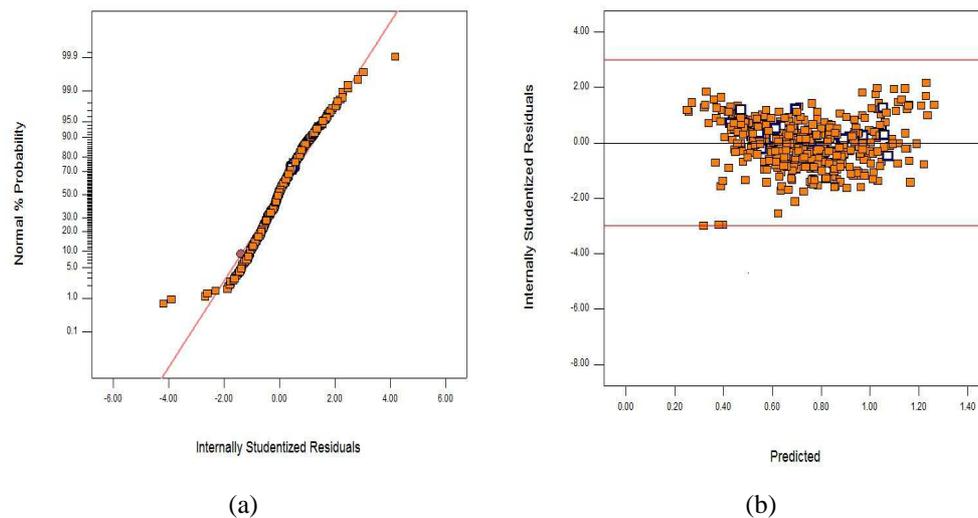


Figure 8.6: *On the left: Normal probability plot of the residuals. On the right: Plot of residuals versus predicted.*

The most important finding of this experiment is the dominance of main effects. Clearly lack of interaction greatly simplifies our interpretations and conclusions. The three most significant main effects in decreasing order are the expected present value W of the first new drug, the annual cost α of employing a senior scientist, and the obsolescence rate ν_2 . Together they account for nearly two-thirds of the total variability in PI^* . The costs of phases *II* and *III* clinical trials and the discount rate γ also prove to be crucial factors that affect the profitability of a project. This confirms our expectation that a discovery project with big potential reward, low research cost, low cost in clinical trials, little competition and small discount rate would be very profitable.

The α - W interaction is plotted in Figure 8.7. The plot shows that profitability is at its highest when the potential reward W is high and cost α is low. It also

reveals that the effect of the cost α is more prominent at higher values of W .

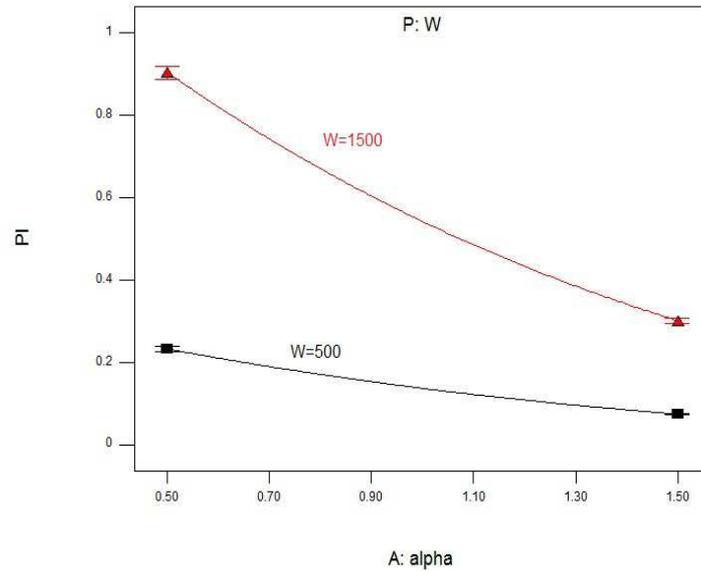
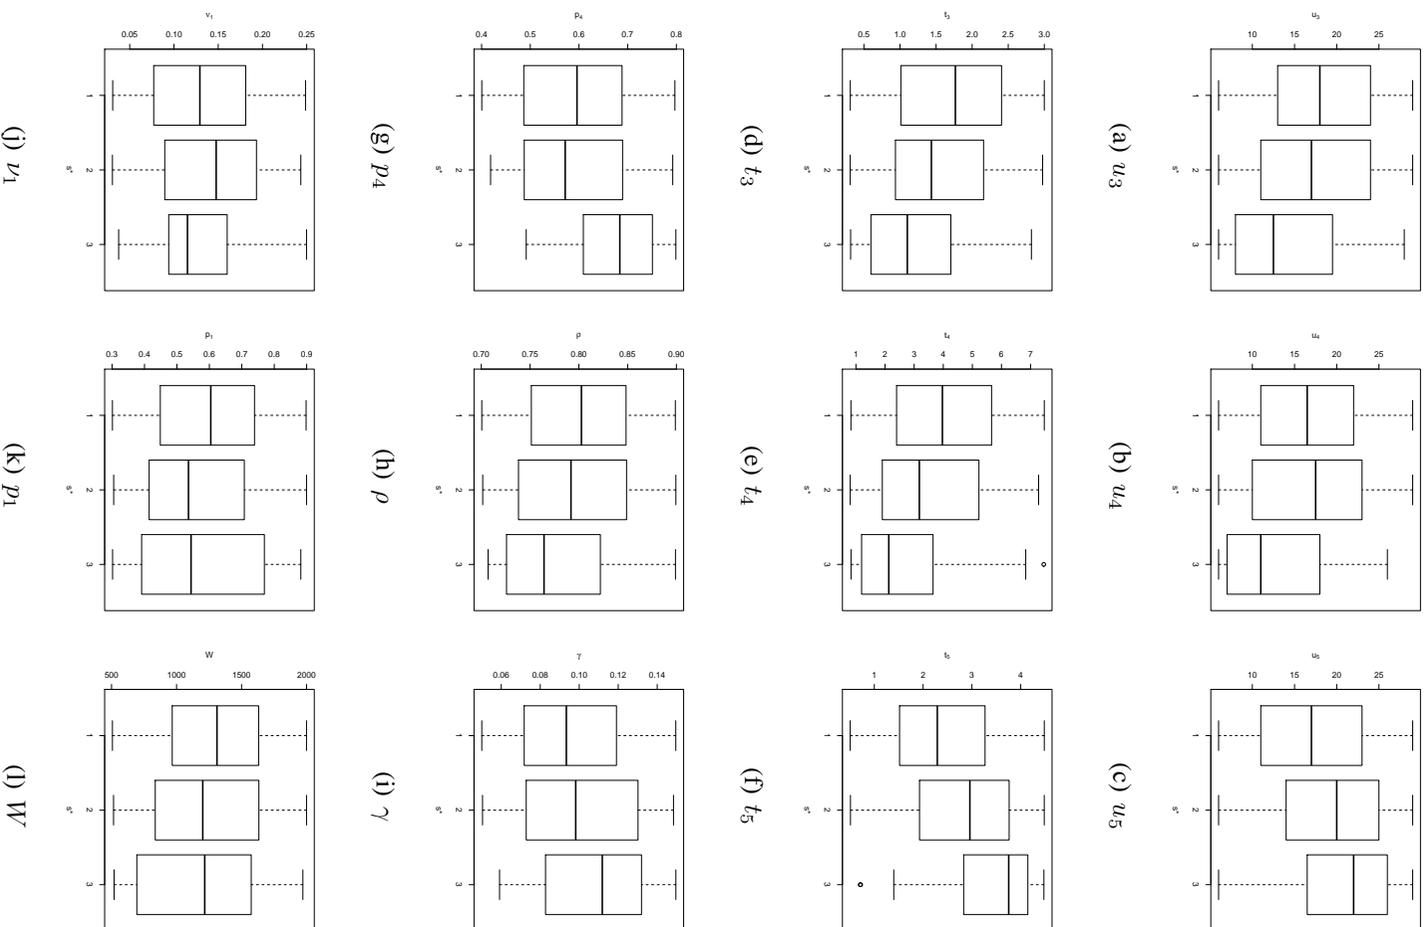


Figure 8.7: α - W interaction plot.

8.2 Effects of Parameters on s^*

The response (s^*, n^*) falls into one of three categories, $(1, n^*)$, $(2, n^*)$ and $(3, n^*)$. We are particularly interested in s^* as its value potentially affects the choice of the optimization procedures we might use. By plotting the other explanatory variables against s^* , we see that there appear to be some obvious relationships between s^* and some of the explanatory variables. There appear to be negative associations between u_3 , u_4 , t_3 , t_4 , ρ and s^* (Figure 8.8 (a), (b), (d), (e), (h)). The smaller these variables are, the bigger s^* tends to be. On the other hand, higher values of u_5 , t_5 and γ appear to be associated with higher values of s^* (Figure 8.8 (c), (f), (j)).

Figure 8.8: Plots of 12 explanatory variables against s^* .

Since the response categories are ordered, we fit a cumulative logits model which accounts for the ordering in the response. Let J denote the number of categories for the response Y . The logits of the cumulative probabilities are defined as

$$\text{logit}[P(Y \leq j)] = \log \frac{P(Y \leq j)}{1 - P(Y \leq j)}, j = 1, \dots, J - 1.$$

Let \mathbf{x} denote the column vector of the explanatory variables. A cumulative logit model incorporating explanatory variables has the form

$$\text{logit}[P(Y \leq j)] = \mu_j + \boldsymbol{\beta}' \mathbf{x}, \quad (8.1)$$

for $j = 1, \dots, J - 1$, for a column vector $\boldsymbol{\beta}$ of parameters that describes the effects of the explanatory variables. The model assumes the same effects $\boldsymbol{\beta}$ for each logit, but different intercepts μ_j . The expression for $P(Y \leq j)$ is in a deliberately simple form. We need to keep in mind that this is actually $P(Y \leq j | \mathbf{x})$, the conditional probability at each fixed value for the explanatory variables. The expression for the cumulative probabilities is

$$P(Y \leq j) = \frac{\exp(\mu_j + \boldsymbol{\beta}' \mathbf{x})}{1 + \exp(\mu_j + \boldsymbol{\beta}' \mathbf{x})}, j = 1, \dots, J - 1.$$

For the cell probabilities themselves,

$$P(Y = j) = \frac{\exp(\mu_j + \boldsymbol{\beta}' \mathbf{x})}{1 + \exp(\mu_j + \boldsymbol{\beta}' \mathbf{x})} - \frac{\exp(\mu_{j-1} + \boldsymbol{\beta}' \mathbf{x})}{1 + \exp(\mu_{j-1} + \boldsymbol{\beta}' \mathbf{x})},$$

with $\mu_0 = -\infty$ and $\mu_J = \infty$.

Maximum-likelihood fitting is used for the cumulative logit model (8.1) and iterative methods are used to solve the likelihood equations and obtain the ML estimates of the model parameters. Based on the model fit, we can conduct statistical inference about the significance of model parameters using the ML estimates. For testing $H_0: \beta_k = 0$, we use the Wald z statistic defined as $z = \frac{\hat{\beta}_k}{SE}$, which (under H_0) has an asymptotic chi-square distribution with one degree of freedom, where SE is the standard error of $\hat{\beta}_k$.

We first fit an initial model of $\text{logit}[P(s^* \leq j)]$, $j = 1, 2$ with all 30 explanatory variables. The maximum-likelihood fit of the initial model suggests that

variables $\alpha, \lambda, u_1, t_1, p_{22}, p_3, p_a, \nu_2, c_I, c_{II}, c_{III}, p_I, p_{II}$ and p_{III} are insignificant ($p > 0.05$) in explaining the variation in s^* . By removing these terms one by one starting from the one with the biggest p-value, and comparing nested models using the likelihood-ratio test, we come to the refined model

$$\begin{aligned} \text{logit}[P(s^* \leq j)] = & \mu_j - 0.0425u_2 + 0.0545u_3 + 0.0522u_4 - 0.0913u_5 - 0.3729t_2 \\ & + 0.8083t_3 + 0.3151t_4 - 0.6853t_5 + 1.7265p_1 - 2.5603p_4 \\ & + 2.211p_b - 4.5960\nu_1 - 12.9241\gamma + 0.0007W, \quad j = 1, 2. \quad (8.2) \end{aligned}$$

Table 8.5 shows the maximum-likelihood fit of the refined model. The effect estimates reflect the tendency for the cumulative probability $P(s^* \leq j)$ to increase as $u_3, u_4, t_3, t_4, p_1, p_b$ and W increase, and to decrease as $u_2, u_5, t_2, t_5, p_4, \nu_1$, and γ increase.

To describe the effect of a variable, we compare $P(s^* = j)$ for a particular j at different values of this variable, such as the minimum and maximum, and control other variables in the model by setting them at their mean values. We illustrate this by describing the effect of u_2 . By substituting the minimum value of u_2 and the mean values of the remaining variables into equation (8.2), we get $P(s^* \leq 1) = P(s^* = 1) = 0.83$, and similarly at the maximum value of u_2 we get $P(s^* = 1) = 0.65$. $P(s^* = 2) = P(s^* \leq 2) - P(s^* = 1)$ stands at 0.15 and 0.3 at the minimum and the maximum values of u_2 respectively. Finally, $P(s^* = 3) = 1 - P(s^* \leq 2)$ increases from 0.02 to 0.04 between the minimum and the maximum values of u_2 . Table 8.6 shows estimated probabilities that help us to interpret the predictor effects.

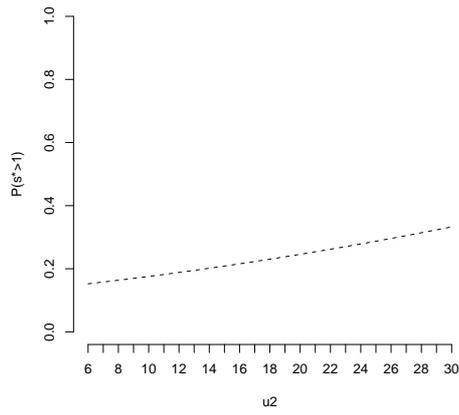
It turns out that the effort required in stages 2, 3, 4 and 5 are important in affecting the value of s^* . In particular, effort allocation in stage 5 has very strong association with the value of s^* . The probability of $s^* = 1$ decreases from 90% to 50% when u_5 increases from the minimum of 6 to the maximum of 30, whereas the probability of $s^* = 2$ increases dramatically from 9% to 40%. Figure 8.9 plots the predicted probabilities of $s^* > 1$ against u_2, u_3, u_4 , and u_5 . We see that the probability of taking more than one LS is predicted to be greater with more effort allocation in stages 2 and 5 and less effort in stages 3 and 4.

Recall that the amount of effort required for each stage is $X_i = e_i(u_i)t_i$

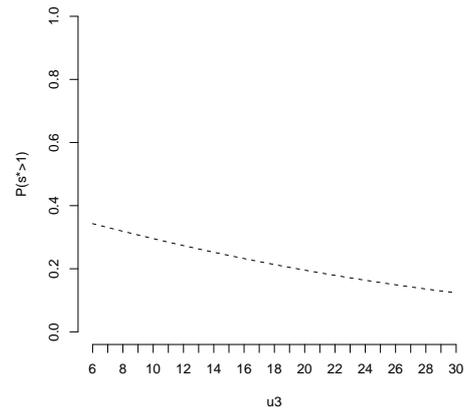
($i = 1, 2, 3, 4, 5$). Hence the effort rate and the stage time have similar effects, as an increase in either u or t leads to an increase in the effort rate X . A big X_2 means a large sunk cost in stage 2 which makes the PI of the first LS relatively small and therefore requires more LS to be optimized to compensate for the sunk cost; while small X_3 and X_4 or big X_5 imply that it takes more effort to produce a backup CD than to optimize a new LS. This means there is a higher incentive to optimize more LS than to take more CDs from a particular LS.

The effects of p_4 and p_b are also straightforward. A high probability of success for stage 4 would motivate more LS being optimized, and this is confirmed by the doubling of $P(s^* = 2)$ from 0.15 to 0.30 when p_4 increases from its minimum to maximum. p_b on the other hand, has almost the opposite effect to p_4 . $P(s^* = 1)$ increases from 0.67 to 0.83 while $P(s^* = 2)$ drops from 0.29 to 0.15, when p_b changes from its minimum to maximum. This is in line with our expectation, as p_b is the probability that an LS is good, and therefore a low value of p_b would result in needing to take more LS.

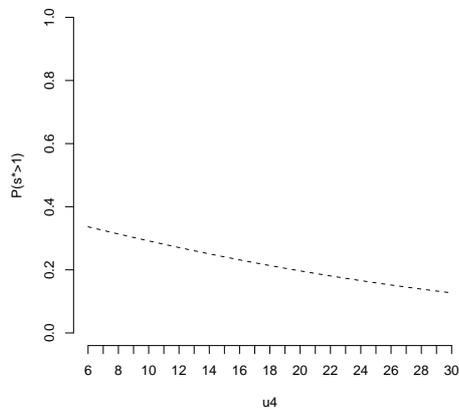
There are no obvious reasons for the dependence of s^* on the values of ν_1 , p_1 and W . We conclude from Table 8.6 that in general the smaller the profitability of a project (i.e., big obsolescence rate ν_1 or small p_1 or W), the more likely it is to be optimal to take more than one LS. This is also confirmed by the plots of ν_1 , p_1 and W against s^* in (j), (k) and (l) in Figure 8.8.



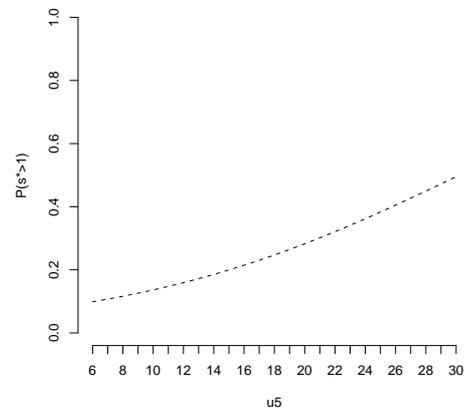
(a)



(b)



(c)



(d)

Figure 8.9: *Effects of effort allocations.*

Table 8.5: Output for the cumulative logit model with s^* as the ordered response.

Variable	Coefficient	Standard Error	Wald z	pr(> z)
μ_1	3.2660	1.3979	4.01	0.0195
μ_2	5.6976	1.4204	2.34	0.0001
u_2	-0.0425	0.0148	-2.88	0.0040
u_3	0.0545	0.0149	3.66	0.0003
u_4	0.0520	0.0153	3.40	0.0007
u_5	-0.0913	0.0155	-5.88	0.0000
t_2	-0.3729	0.1318	-2.83	0.0047
t_3	0.8083	0.1394	5.80	0.0000
t_4	0.3151	0.0570	5.53	0.0000
t_5	-0.6853	0.0979	-7.00	0.0000
p_1	1.7265	0.6123	2.82	0.0048
p_4	-2.5603	0.9030	-2.84	0.0046
p_b	2.2110	0.8946	2.47	0.0135
ν_1	-4.5960	1.7096	-2.69	0.0072
γ	-12.9241	3.4667	-3.73	0.0002
W	0.0007	0.0002	2.84	0.0046
Deviance	703.4			
Log-likelihood	-351.7			

Table 8.6: *Estimated probabilities describing the effects in the refined model.*

		Estimated Probabilities		
		$P(s^* = 1)$	$P(s^* = 2)$	$P(s^* = 3)$
$u_2 =$	min	0.83	0.15	0.02
	max	0.65	0.30	0.04
$u_3 =$	min	0.62	0.33	0.05
	max	0.85	0.13	0.01
$u_4 =$	min	0.64	0.31	0.05
	max	0.86	0.13	0.01
$u_5 =$	min	0.90	0.09	0.01
	max	0.53	0.40	0.07
$p_1 =$	min	0.65	0.30	0.05
	max	0.84	0.14	0.02
$p_4 =$	min	0.84	0.15	0.02
	max	0.65	0.30	0.05
$p_b =$	min	0.67	0.29	0.04
	max	0.83	0.15	0.02
$\nu_1 =$	min	0.83	0.15	0.02
	max	0.65	0.30	0.05
$\gamma =$	min	0.92	0.075	0.0079
	max	0.91	0.076	0.008
$W =$	min	0.65	0.31	0.04
	max	0.84	0.14	0.02

Chapter 9

Conclusions and Desirable Future Work

The discovery and development of new drugs is a long, costly and risky business. It involves many resources and many decisions. In pharmaceutical R&D, as in any business, the set of available options has diverse risk-return characteristics. Obviously, a random selection from this set fails to optimize the return of an R&D portfolio, and so there is a need to address allocative efficiency issues to optimize the use of R&D resources. Allocative efficiency is about maximizing total output by choosing from a range of investment options, it seeks to optimize the distribution of resources across all available projects to enhance the return. Economic analysis explores these allocative efficiency questions by estimating expected rewards and costs and should be integrated into the R&D planning process.

The premise of this thesis is that we are far from optimizing the use of robust economic modelling and analysis in the earlier phases of pharmaceutical R&D and much more can be done. Modelling early in the drug R&D process can provide valuable input into project selection and termination decisions. Drug development is characterized by a highly skewed distribution of returns on investment. It is a small proportion of drugs that in effect cross-subsidizes the bulk of the portfolio. Early modelling and analysis can inform drug developers whether the economic characteristics of projects in their portfolio are more likely to resemble those in the profit-making or loss-making deciles. The amount of money spent on the

entire R&D portfolio can be significantly reduced through the early elimination of unattractive projects and resources saved can be reallocated to more attractive projects. Early modelling and analysis is also likely to provide a solid foundation for communicating product value to external decision-makers further downstream, increasing the likelihood of regulatory approval and commercial success.

This thesis provides detailed modelling of pre-clinical research and contributes to the task of improving the efficiency of resource allocation through better decision making. The models suggested are designed to evaluate profitability and to indicate appropriate allocation levels at the different stages of a pharmaceutical research project. With modelling, despite considerable uncertainty, expected rewards and costs can be estimated. As more data becomes available these models can be updated and results re-calculated.

The software OPRRA described here serves as a management tool to assist in finding allocations which are both feasible and profitable. It enables a planner to explore the consequences of different possible allocation plans in terms of their overall effort requirements, their profitability, and their varying degrees of risk. This is done by fairly sophisticated simulation and sensitivity analysis. It is not a particularly prescriptive tool, nor indeed a comprehensive planning tool. This is partly because it does not model the detailed structure of effort allocation or all the factors which legitimately influence it, and partly because profitability tends not to be very sensitive to moderate changes in allocations. Of course a model can only ever be an imperfect representation of reality, so apparently good allocations derived from OPRRA models can only be suggestive. However, some movement of the resource allocation plan in the direction suggested by OPRRA deserves serious consideration, especially when the apparent scope for increased profitability is high.

The models and policies which we have discussed have demonstrated some interesting and insightful results. The adaptive class of FI policies is able to improve the profitability with respect to the sequence of CDs selected and may be implemented in an adaptive probabilities setting, but does not optimize effort allocations at each stage. The non-adaptive class of (s, n) policies on the other hand, optimizes profitability with respect to the effort allocations at each stage but is

rather restrictive in its choice of the sequence of CDs and LS. However, it turns out that the restriction imposed on s to be less than or equal to 3 in an (s, n) policy has little impact on profitability. The (s, n) policies in most cases perform as well as the more flexible FI policies, making it the most comprehensive and reliable optimization procedure.

Interactions appear to be weak between the parameters that influence profitability of a pharmaceutical project, and profitability is predominately determined by a small number of parameters, such as the value of the first new drug, annual operating cost and obsolescence rates.

Possible Future Work

There are a number of challenges and opportunities that are outside the scope of our current effort but are natural extensions for the future.

Since in this thesis we assumed constant allocation over each stage, one possible extension is to investigate the problem where the number of scientists allocated is allowed to change over time. Some preliminary ideas using calculus of variations to find the optimal allocation can be found in Wilder (2001) and Hoeschler (2002).

The FI policies are policies for selecting a sequence of LS and CDs so as to maximize the PI for a project, for a given set of stage allocations. They are based upon PIs for a new CD from the current LS, for a new LS, and a reference PI. The purpose of the reference PI is that it works as an approximation to the optimal PI for the project, and therefore screens out CDs and LS which are likely to reduce the overall PI. We would like to be able to carry out similar calculations for IRR, but this is less straightforward computationally.

The real options model presented here is simplified in a number of ways. First, the cost of investment, I , is assumed to be fixed and to be made upfront. In pharmaceutical R&D, the cost of the investment can be as uncertain as the future payoff, particularly for lengthy projects. There is uncertainty that relates to the physical difficulty of completing the project. The actual costs and completion time become apparent as the project proceeds. Costs may be greater or less than

anticipated if impediments arise or if the work progresses faster than planned, but the total cost of investment is only known for certain when the project is complete. Other uncertainties arise when the price of labour, material, etc. needed to build a project fluctuate unpredictably, or when government regulations change the cost of investment. Investment is also likely to be made continuously until the project is complete, and there is usually a maximum rate at which the firm can invest. Hence a useful extension would be to model both the payoff and the cost as stochastic processes, incorporating a maximum rate of investment. Relevant work concerning uncertain and continuous investment costs can be found in Majd and Pindyck (1987) and in Pindyck (1993).

Second, only the *option to invest* is considered. Another possible extension is to evaluate the *option to abandon early* and the *growth options*, both of which are valuable in pharmaceutical R&D decision making. The option to abandon is present at every stage of a pharmaceutical R&D project. It is often exercised if the results from the latest development phase do not warrant continuation. The valuation of the option to abandon is analogous to a put option. The impact of this option is to reduce the overall risk. Growth options refer to the unanticipated opportunities that arise from investments. It is important to include growth options in the calculation on profitability of the early stage investments, which enables the more profitable follow-on investments. Growth options have considerable importance in the context of early-stage pharmaceutical R&D.

As we have already noted, no model is going to be completely adequate. However, decision-making with the help of a model is likely to be better than without it, whether it leads to faster termination of uneconomic projects or the allocation of more resources to attractive projects.

Appendix A

Glossary

- AP = Adaptive Probabilities.
- CD = Candidate Drug.
- FI = Forwards Induction.
- FP = Fixed Probabilities.
- IRR = Internal Rate of Return.
- LS = Lead Series.
- OPRRA = Optimizing Pharmaceutical Research Resource Allocation.
- PI = Profitability Index.
- WACC = Weighted Average Cost of Capital.

Appendix B

Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy

In this appendix we show how to calculate the profitability index values for an additional CD and for a new LS in an FI policy as discussed in chapters 3 (FP setting) and 5 (AP setting).

Let r_{CD} and c_{CD} be the expected reward and expected cost for the next CD, respectively. We have

$$PI(CD) = \frac{r_{CD}}{c_{CD}}.$$

Denote by c_0 the cost of the project to date, made up of the costs of the stages which have been completed, and let $r_0 = c_0 PI(ref)$, so that r_0 is an estimate for the reward which would on average accrue with an expenditure of c_0 .

The PIs which are relevant in determining $PI(LS)$ are listed below. There are two versions of each of the relevant PIs, one of which excludes the previous history of the project, and one for which account is taken of the previous history.

- PI (a single attempt to optimize an LS with no parallel stage 3).
Call this $PI_{LS1} = \frac{r_1}{c_1}$. The version which allows for previous history is $PI1 = \frac{r_0+r_1}{c_0+c_1}$.

Appendix B Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy

- PI (a further attempt to optimize an LS with no parallel stage 3, following a first unsuccessful attempt to optimize an LS together with a successful stage 3).

Call this $PILS2 = \frac{r_2}{c_2}$. Extending our notation in an obvious way, the version which allows for previous history is $PI2 = \frac{r_0+r_2}{c_0+c_2}$.

- PI (a further attempt to optimize an LS with no parallel stage 3, following a first successful attempt to optimize an LS together with a successful stage 3).

Call this $PILS3 = \frac{r_3}{c_3}$. $PI3 = \frac{r_0+r_3}{c_0+c_3}$.

- PI (an attempt to optimize an LS together with stage 3, followed by one more attempt to optimize an LS if stage 3 succeeds and the first attempt fails).

Call this $PILS4 = \frac{r_1+q_4p_3r_2}{c_1+q_4p_3c_2}$. $PI4 = \frac{r_0+r_1+q_4p_3r_2}{c_0+c_1+q_4p_3c_2}$.

- PI (an attempt to optimize an LS together with stage 3, followed by one more attempt to optimize an LS if stage 3 succeeds and the first attempt succeeds).

Call this $PILS5 = \frac{r_1+p_4p_3r_3}{c_1+p_4p_3c_3}$. $PI5 = \frac{r_0+r_1+p_4p_3r_3}{c_0+c_1+p_4p_3c_3}$.

- PI (an attempt to optimize an LS together with stage 3, followed by one more attempt to optimize an LS if stage 3 succeeds, and whether or not the first attempt succeeds).

Call this $PILS6 = \frac{r_1+q_4p_3r_2+p_4p_3r_3}{c_1+q_4p_3c_2+p_4p_3c_3}$. $PI6 = \frac{r_0+r_1+q_4p_3r_2+p_4p_3r_3}{c_0+c_1+q_4p_3c_2+p_4p_3c_3}$.

The rewards r_1 , r_2 and r_3 are rewards for the PI -maximizing number of CDs following successful optimization of the corresponding LS. The corresponding costs are c_1 , c'_1 (which alone includes the cost of a parallel stage 3), c_2 and c_3 .

Condition A:

$$PI(ref) < PILS6 \text{ and } PI4 - PI(ref) > PI1 - PI(ref) - 0.1|PI1 - PI(ref)|.$$

Condition B:

$$PI(ref) < PILS6 \text{ and } PI5 - PI(ref) > PI1 - PI(ref) - 0.1|PI1 - PI(ref)|.$$

If either of conditions A and B holds then $PI(LS) = PILS6$, otherwise $PI(LS) = PILS1$.

Appendix B **Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy**

The first part of condition A checks that on average FI is not going to leave the next LS unused. The second part of condition A checks that the immediate relative improvement in PI above $PI(ref)$ if we do stage 3 with stage 4 is at least 90% of the gain if we were to do stage 4 without stage 3. The choice of 90% is arbitrary. Other percentages less than 100% will sometimes do better. The rationale for condition B is similar, now referring to the situation when the first attempt to optimize an LS succeeds.

B.1 FP Setting

B.1.1 Calculation of $PI(CD)$

Define:

- S : the youngest LS from which we will choose the next CD.
- κ : number of CDs selected from LS S so far.
- Θ : the set of all LS from which CDs have been selected.
- b_i : the number of CDs that have been selected from LS i .

Thus

$$\begin{aligned} r_{CD} &= e^{-\gamma_1(t+t_5+t_I+t_{II}+t_{III})} pW \prod_{i \in \Theta \setminus S} (p_b \eta^{b_i} + 1 - p_b) \eta^\kappa \\ &\quad - e^{-\gamma(t+t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}), \quad (\text{B.1}) \\ c_{CD} &= e^{-\gamma(t+t_5)} K_5. \end{aligned}$$

B.1.2 Calculation of $PI(LS)$

The expressions for $r_1, r_2, r_3, c_1, c'_1, c_2,$ and c_3 are as follows.

Appendix B **Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy**

$$r_1 = p_4 \left[e^{-\gamma_1(t+t_4+t_I+t_{II}+t_{III})} pW \prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) \sum_{m=1}^{k^*} (\eta e^{-\gamma_1 t_5})^{m-1} - e^{-\gamma(t+t_4)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right], \quad (\text{B.2})$$

$$c_1 = e^{-\gamma t} K_4 + p_4 e^{-\gamma(t+t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5,$$

$$c'_1 = e^{-\gamma t} (K_3 + K_4) + p_4 e^{-\gamma(t+t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5.$$

$$r_2 = p_4 \left[e^{-\gamma_1(t+2t_4+t_I+t_{II}+t_{III})} pW \prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) \sum_{m=1}^{k^*} (\eta e^{-\gamma_1 t_5})^{m-1} - e^{-\gamma(t+2t_4)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right], \quad (\text{B.3})$$

$$c_2 = e^{-\gamma(t+t_4)} K_4 + p_4 e^{-\gamma(t+2t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5.$$

$$r_3 = p_4 \left[e^{-\gamma_1(t+t_4+\rho t_4+(k^*-1)t_5+t_I+t_{II}+t_{III})} pW \prod_{i \in \Theta} (p_b \eta^{b_i} + 1 - p_b) (p_b \eta^{k^*} + 1 - p_b) \times \sum_{m=1}^{k^*} (\eta e^{-\gamma_1 t_5})^{m-1} - e^{-\gamma(t+t_4+\rho t_4+(k^*-1)t_5)} (c_I + p_I e^{-\gamma t_I} c_{II} + p_I p_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right], \quad (\text{B.4})$$

$$c_3 = e^{-\gamma(t+t_4+(k^*-1)t_5)} K_4 + p_4 e^{-\gamma(t+t_4+\rho t_4+(k^*-1)t_5)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5,$$

where and k^* is the PI-maximizing number of CDs to be selected from the new LS, and k^* is the PI-maximizing number of CDs for the preceding successfully optimized LS.

B.2 AP Setting

B.2.1 Calculation of $PI(CD)$

$$\begin{aligned}
 r_{CD} &= \max_i r_{iCD} = r_{i^*CD}, \text{ where} & (B.5) \\
 r_{iCD} &= e^{-\gamma_1(t+t_5+t_I+t_{II}+t_{III})} E[p_{Ii}p_{IIi}p_{IIIi}\lambda^{z_1+z_2+z_3}]\lambda^{z_4}W - e^{-\gamma(t+t_5)}(c_I \\
 &\quad + E[p_{Ii}]e^{-\gamma t_I}c_{II} + E[p_{Ii}]E[p_{IIi}]e^{-\gamma(t_I+t_{II})}c_{III}), \\
 c_{CD} &= e^{-\gamma(t+t_5)}K_5.
 \end{aligned}$$

The evaluation of $E[p_{Ii}p_{IIi}p_{IIIi}\lambda^{z_1+z_2+z_3}]$ is given in section 5.3.

B.2.2 Calculation of $PI(LS)$

The expressions for $r_1, r_2, r_3, c_1, c'_1, c_2$ and c_3 in the adaptive probabilities setting are as follows. Note that because we are trying to optimize a new LS, there is no update on the success probabilities in clinical trials for that LS yet, hence their

Appendix B Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy

prior expectations m_I , m_{II} and m_{III} are used in the calculations.

$$\begin{aligned}
 r_1 &= E[p_4] \left(e^{-\gamma_1(t+t_4+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W \sum_{m=1}^{k^*} E[\lambda^{\sum_{j=0}^{m-1} Q_j}] (e^{-\gamma_1 t_5})^{m-1} \right. \\
 &\quad \left. - e^{-\gamma(t+t_4)} (c_I + m_I e^{-\gamma t_I} c_{II} + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right) \\
 &\quad (Q_j \sim \mathbf{Ber}(m_I m_{II} m_{III}), j = 1, 2, \dots; Q_0 = 0.) \\
 &= E[p_4] \left(e^{-\gamma_1(t+t_4+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W \times \right. \\
 &\quad \sum_{m=1}^{k^*} \left((1 - (1 - \lambda) m_I m_{II} m_{III}) e^{-\gamma_1 t_5} \right)^{m-1} - e^{-\gamma(t+t_4)} (c_I + m_I e^{-\gamma t_I} c_{II} \\
 &\quad \left. + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right), \tag{B.6}
 \end{aligned}$$

$$c_1 = e^{-\gamma t} K_4 + E[p_4] e^{-\gamma(t+t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5,$$

$$c'_1 = e^{-\gamma t} (K_3 + K_4) + E[p_4] e^{-\gamma(t+t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5.$$

$$\begin{aligned}
 r_2 &= E[p_4] \left(e^{-\gamma_1(t+2t_4+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W \sum_{m=1}^{k^*} E[\lambda^{\sum_{j=0}^{m-1} Q_j}] (e^{-\gamma_1 t_5})^{m-1} \right. \\
 &\quad \left. - e^{-\gamma(t+2t_4)} (c_I + m_I e^{-\gamma t_I} c_{II} + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right) \\
 &= E[p_4] \left(e^{-\gamma_1(t+2t_4+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W \times \right. \\
 &\quad \sum_{m=1}^{k^*} \left((1 - (1 - \lambda) m_I m_{II} m_{III}) e^{-\gamma_1 t_5} \right)^{m-1} - e^{-\gamma(t+2t_4)} (c_I + m_I e^{-\gamma t_I} c_{II} \\
 &\quad \left. + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right), \tag{B.7}
 \end{aligned}$$

$$c_2 = e^{-\gamma(t+t_4)} K_4 + E[p_4] e^{-\gamma(t+2t_4)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5.$$

Appendix B **Calculations of $PI(CD)$ and $PI(LS)$ for an FI Policy**

$$\begin{aligned}
r_3 &= E[p_4] \left(e^{-\gamma_1(t+t_4+\rho t_4+(k^*-1)t_5+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W E[\lambda^{\sum_{j=0}^{k^*} Q_j}] \times \right. \\
&\quad \left. \sum_{m=1}^{k^*} E[\lambda^{\sum_{j=0}^{m-1} Q_j}] (e^{-\gamma_1 t_5})^{m-1} - e^{-\gamma(t+t_4+\rho t_4+(k^*-1)t_5)} \times \right. \\
&\quad \left. (c_I + m_I e^{-\gamma t_I} c_{II} + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right) \\
&= E[p_4] \left(e^{-\gamma_1(t+t_4+\rho t_4+(k^*-1)t_5+t_I+t_{II}+t_{III})} E[m_I m_{II} m_{III} \lambda^{z_1+z_2+z_3}] \lambda^{z_4} W \times \right. \\
&\quad \left. (1 - (1 - \lambda) m_I m_{II} m_{III})^{k^*} \sum_{m=1}^{k^*} ((1 - (1 - \lambda) m_I m_{II} m_{III}) e^{-\gamma_1 t_5})^{m-1} \right. \\
&\quad \left. - e^{-\gamma(t+t_4+\rho t_4+(k^*-1)t_5)} (c_I + m_I e^{-\gamma t_I} c_{II} + m_I m_{II} e^{-\gamma(t_I+t_{II})} c_{III}) \right), \tag{B.8}
\end{aligned}$$

$$c_3 = e^{-\gamma(t+t_4+(k^*-1)t_5)} K_4 + E[p_4] e^{-\gamma(t+t_4+\rho t_4+(k^*-1)t_5)} \sum_{m=1}^{k^*} (e^{-\gamma t_5})^{m-1} K_5.$$

Appendix C

Simulation Algorithms

This appendix gives the simulation algorithms used to investigate the probability distributions of rewards, costs and profitability indices for the (s, n) and FI policies discussed in section 4.5.

C.1 Notation

- f : counter for the number of successfully optimized LS.
- g : counter for the number of optimization attempts.
- k_i : the required number of CDs to be taken from the i^{th} optimized LS in an (s, n) policy.
- b_i : counter for the number of CDs found in the i^{th} optimized LS.
- a_i : counter for the number of successful CDs from the i^{th} optimized LS.
- t : time taken so far.
- R : total reward so far.
- K : total cost so far.

- X : indicator for whether a target is attainable. $X = 1$ means the target is attainable; $X = 0$ otherwise.
- Y : indicator for whether an LS is good. $Y = 1$ means the LS is good; $Y = 0$ otherwise.
- Z : indicator for whether a CD is successful. $Z = 1$ means the CD is successful; $Z = 0$ otherwise.
- L : indicator for whether this is the last LS to be optimized. $LS = 1$ means it is the last LS to be optimized; $L = 0$ otherwise.
- d : a vector of times at which future decisions are made.
- v : counter for the number of decisions made.

C.2 Assumptions

C.2.1 FP Setting

For each CD going to clinical trials we proceed as follows.

1. If this is the first CD of the whole project ever going to clinical trials, we simulate whether the target is attainable by randomly generating the indicator X , $P(X = 1) = p_a$. An unattainable target indicates that none of the CDs from the project would be successful in passing phase *III* clinical trials and become a marketable drug.
2. If this is the first CD from some LS and the target is attainable, we simulate whether this LS is good by generating the indicator Y , $P(Y = 1) = p_b$. A bad LS indicates that none of the CDs from that LS will pass phase *III* clinical trials and become marketable.
3. If a CD is from a good LS and the target is attainable, we simulate whether this CD will be successful by generating the indicator Z , $P(Z = 1) = p_c$. Thus with probability p_c a CD will pass all phases of clinical trials

and become marketable given that it is from a good LS and the target is attainable. In this case we add the net reward $e^{-\gamma_1(t+t_I+t_{II}+t_{III})}\lambda\sum_{i=1}^f a_i W - e^{-\gamma t}(c_I + e^{-\gamma t_I}c_{II} + e^{-\gamma(t_I+t_{II})}c_{III})$ to R .

4. For a CD which is not successful, we determine the phase at which failure occurs by assuming that it occurs at phase I , II , or III with probabilities q_I/Q , $p_I q_{II}/Q$ and $p_I p_{II} q_{III}/Q$, respectively, where $Q = 1 - p_a p_b p_c$. These outcomes lead to reductions in the total reward R of $e^{-\gamma t}c_I$, $e^{-\gamma t}(c_I + e^{-\gamma t_I}c_{II})$ and $e^{-\gamma t}(c_I + e^{-\gamma t_I}c_{II} + e^{-\gamma(t_I+t_{II})}c_{III})$, respectively. This is assuming that the phase at which failure occurs is independent of whether it is due to an unattainable target, an attainable target but a bad LS, or an attainable target and a good LS but a defective CD, and also independent of the phase at failure for other CDs.

C.2.2 AP Setting

For each CD going to clinical trials we proceed as follows.

1. The success probabilities p_{Ii} , p_{IIi} and p_{IIIi} for clinical trials are different for different LS. Their posterior expectations $E[p_{Ii}]$, $E[p_{IIi}]$ and $E[p_{IIIi}]$ are updated using Bayes theorem whenever new information becomes available.
2. A CD from LS i has probability $E[p_{Ii}]$ of passing phase I trial, and whether it succeeds or not there is a reduction of $e^{-\gamma t}c_I$ from the total reward R . If it proceeds to phase II trial there is another reduction of $e^{-\gamma(t+t_I)}c_{II}$ from the total reward and it will pass phase II trial with probability $E[p_{IIi}]$. If again it passes through phase II trial and proceed to phase III , there will be a further reduction of $e^{-\gamma(t+t_I+t_{II})}c_{III}$ from the total reward R , and it will pass phase III trial with probability $E[p_{IIIi}]$. If the CD passes phase III trial a net reward of $e^{-\gamma_1(t+t_I+t_{II}+t_{III})}\lambda\sum_{i=1}^f a_i W$ is added to R .
3. The successes and failures of each CD at each phase are recorded. This information will affect not only the success probabilities for CDs from the current LS, but also the success probabilities for CDs from other LS as well.

C.3 Pseudocode

The pseudocode described in this section is a compact and high-level description of the C++ simulation algorithms used in OPRRA intended for human reading. It omits details that are not essential for human understanding of the algorithm, such as variable declarations, system-specific code and subroutines. Some necessary explanations are given below the code, which are in italics and start with #.

C.3.1 (s, n) Simulation

Step 1. Set $f, g, t, b_f, a_f, R, K, \mathbf{d}$ and v to 0.

This is to initialize all parameter values to 0.

Step 2. Increase K by K_1 , t by t_1 .

With probability:

p_1 , go to step 3;

q_1 , stop and output R and K .

Stage 1 is certain to happen, so there are always cost spent and time taken, whether we succeed or not. If stage 1 succeeds, we will go to the next steps. If stage 1 fails, the simulation stops. Whenever a realization is complete, we stop and output the values of the total reward R and total cost K .

Step 3. Increase K by K_2 , t by t_2 , and with probability :

p_{21} , go to step 4;

p_{22} , go to step 7;

q_2 , Stop and output R and K .

Step 4. Check $f = 0$?

If $f = 0$, go to step 5;

Else go to step 6.

This is to check whether an LS has been optimized previously. If there has been previous success in optimizing an LS, the time taken for stages 3 and 4 will be pt_4 .

Step 5. Increase g by 1, and check $g = n$?

If $g = n$, increase K by $e^{-\gamma t} K_4$, t by t_4 , go to step 8;

Else increase K by $e^{-\gamma t}(K_3 + K_4)$, t by t_4 , go to step 10.

When $g = n$, it means it is the last attempt to optimize an LS, so stage 4 will be carried

out alone. Otherwise, stages 3 and 4 will be carried out in parallel.

Step 6. Increase g by 1, and check $g = n$?

If $g = n$, increase K by $e^{-\gamma t} K_{4\rho}$, t by ρt_4 , go to step 8;

Else increase K by $e^{-\gamma t}(K_{3\rho} + K_{4\rho})$, t by ρt_4 , go to step 10.

Step 7. Increase g by 1, K by $e^{-\gamma t} K_4$, t by t_4 , and check $g = n$?

If $g = n$, go to step 8;

Else go to step 9.

We only come to this step if two LS are found in stage 2, so that stage 3 is not carried out in parallel at the first attempt to optimize an LS.

Step 8. With probability:

p_4 , increase f by 1, b_f by 1, set $L = 1$, go to step 15;

q_4 , stop and output R and K .

This is when it comes to the last attempt to optimize an LS. The simulation ends if the last attempt fails.

Step 9. With probability :

p_4 , increase f by 1, b_f by 1, go to step 11;

q_4 , go to step 5.

We only come to this step from step 7, and we only go to step 7 when two LS were found in stage 2. If we succeed in optimizing a new LS, we send the first CD to clinical trials and select more backup CDs from this LS. After finding the required number of backup CDs, we will return to optimize another LS. If we fail, we go back to step 5 directly to optimize another LS.

Step 10. With probability:

$p_3 p_4$, increase f by 1, b_f by 1, go to step 11;

$p_3 q_4$, go back to step 4;

$q_3 p_4$, increase f by 1, b_f by 1, set $L = 1$, go to step 15;

$q_3 q_4$, stop and output R and K .

When stages 3 and 4 are carried out in parallel there could be 4 possible outcomes :

Outcome 1, both stages 3 and 4 are successful. The first CD will be sent to clinical trials and more backup CDs will be searched for from that LS. Then a new LS will be attempted when enough backup CDs from the current LS are found.

Outcome 2, stage 3 is successful but stage 4 fails. Attempt to optimize another LS.

Outcome 3, stage 3 fails but stage 4 succeeds. The required number of total CDs will be taken from that LS but no more attempt to optimize any further LS.

Outcome 4, both stages 3 and 4 fail. Simulation terminates.

Step 11. Check $f = s$?

If $f = s$, set $L = 1$, go to step 15;

Else go to step 15 directly.

Step 12. Check $L = 1$?

If $L = 1$, go to step 13;

Else go to step 14.

Step 13. Check $b_f = k_f$?

If $b_f = k_f$, stop and output R and K .

Else increase b_f by 1, K by $e^{-\gamma t} K_5$, t by t_5 , and go to step 21.

Step 14. Check $b_f = k_f$?

If $b_f = k_f$, go to step 6;

Else increase b_f by 1, K by $e^{-\gamma t} K_5$, t by t_5 , and go to step 21.

Step 15. Check $f = 1$?

If $f = 1$, go to step 16;

Else go to step 17.

Step 16. With probability:

p_a , set $X = 1$, go to step 18;

$1 - p_a$, Set $X = 0$, go to step 20.

Step 17. Check $X = 0$?

If $X = 0$, go to step 20;

Else go to step 18.

If the target is unattainable, none of the CDs from the project will be successful.

Step 18. With probability:

p_b , set $Y = 1$, go to step 19.

$1 - p_b$, set $Y = 0$, go to step 20.

Step 19. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

With probability:

p_c , increase R by $e^{-\gamma(t+t_I+t_{II}+t_{III})} \lambda \sum_{i=1}^f a_i W - e^{-\gamma t}(c_I + e^{-\gamma t_I} c_{II} + e^{-\gamma(t_I+t_{II})} c_{III})$,

a_f by 1, go to step 12;

$1 - p_c$, go to step 20.

Types A, B, and C obsolescence events each follow a Poisson process with arrival rates

ξ_A , ξ_B and ξ_C , respectively. We generate the number of obsolescence events N_A , N_B and N_C for each type happened during the time interval Δt . Therefore the value of W is reduced to $W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Step 20. With probability:

p_I/Q , increase R by $e^{-\gamma t}c_I$, go to step 12;

$p_{II}q_{II}/Q$, increase R by $e^{-\gamma t}(c_I + e^{-\gamma t_I}c_{II})$, go to step 12;

$p_{III}q_{III}/Q$, increase R by $e^{-\gamma t}(c_I + e^{-\gamma t_I}c_{II} + e^{-\gamma(t_I+t_{II})}c_{III})$, go to step 12.

Step 21. Check $X = 0$ or $Y = 0$?

If $X = 0$ or $Y = 0$, go to step 20;

Else go to step 19.

When the target is unattainable or the LS is bad, none of the selected CDs will be successful.

C.3.2 FI Simulation in FP Setting

Step 1. Set f , t , b_f , a_f , R , K , \mathbf{d} and v to 0.

Step 2. Increase K by K_1 , t by t_1 .

With probability:

p_1 , go to step 3,

q_1 , stop and output R and K .

Step 3. Increase K by $e^{-\gamma t}K_2$, t by t_2 .

With probability:

p_{21} , go to step 5,

p_{22} , go to step 4,

q_2 , stop and output R and K .

Step 4. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$ and its corresponding k^* .

If $PI(LS) > PI(ref)$,

Increase K by $e^{-\gamma t}K_4$, t by t_4 , and with probability:

p_4 , increase f by 1, b_f by 1, set $L = 0$, go to step 9,

q_4 , go to step 5.

If $PI(LS) < PI(ref)$, stop and output R and K .

k^* is the PI -maximizing number of CDs to be selected from the next LS if it is successfully optimized, as described on page 130. When we optimize a new LS we need to compare its PI with the reference PI , and the attempt to optimize the LS is only made if the PI for the new LS is greater than the reference PI .

Step 5. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$ and its corresponding k^* .

If $PI(LS) > PI(ref)$,

If condition A or B holds, go to step 8; Else go to step 7.

If $PI(LS) < PI(ref)$, stop and output R and K .

Whenever we calculate $PI(LS)$ or $PI(CD)$ we need to calculate W first. For detailed calculations of $PI(LS)$ and $PI(CD)$, and explanations for conditions A and B refer to Appendix B.

Step 6. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$, its corresponding k^* and $PI(CD)$.

If $PI(LS) > \max(PI(CD), PI(ref))$,

if either condition A or B holds, go to step 8; Else go to step 7.

If $PI(CD) > \max(PI(LS), PI(ref))$, go to step 17.

If $PI(ref) > \max(PI(CD), PI(LS))$, stop and output R and K .

This is the point at which we need to decide whether to take an additional CD , to optimize a new LS , or to stop. If the decision is to take a new LS , we need to decide whether to do stage 3 alongside stage 4 depending on whether condition A or B holds. If the decision is to take an additional CD , we will take that CD and will return to this step afterwards.

Step 7. If $f = 0$, increase K by $e^{-\gamma t} K_4$, t by t_4 .

Else increase K by $e^{-\gamma t} K_{4\rho}$, t by ρt_4 .

With probability:

p_4 , increase f by 1, b_f by 1, set $L = 1$, go to step 9,

q_4 , stop and output R and K .

Step 8. If $f = 0$, increase K by $e^{-\gamma t}(K_3 + K_4)$, t by t_4 .

Else increase K by $e^{-\gamma t}(K_{3\rho} + K_{4\rho})$, t by ρt_4 .

With probability:

p_3p_4 , increase f by 1, b_f by 1, set $L = 0$, go to step 9,

p_3q_4 , go to step 5,

q_3p_4 , increase f by 1, b_f by 1, set $L = 1$, go to step 9,

q_3q_4 , stop and output R and K .

Step 9. Check $f = 1$?

If $f = 1$, go to step 10.

Else go to step 11.

Step 10. With probability:

p_a , set $X = 1$, go to step 12,

$1 - p_a$, Set $X = 0$, go to step 14.

Step 11. Check $X = 0$?

If $X = 0$, go to step 14.

Else go to step 12.

Step 12. With probability:

p_b , set $Y = 1$, go to step 13,

$1 - p_b$, set $Y = 0$, go to step 14.

Step 13. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

With probability:

p_c , increase R by $e^{-\gamma(t+t_I+t_{II}+t_{III})} \lambda^{\sum_{i=1}^f a_i} W - e^{-\gamma t}(c_I + e^{-\gamma t_I} c_{II} + e^{-\gamma(t_I+t_{II})} c_{III})$,

a_f by 1, go to step 15,

$1 - p_c$, go to step 14.

Step 14. With probability:

q_I/Q , increase R by $e^{-\gamma t} c_I$,

$p_I q_{II}/Q$, increase R by $e^{-\gamma t}(c_I + e^{-\gamma t_I} c_{II})$,

$p_I p_{II} q_{III}/Q$, increase R by $e^{-\gamma t}(c_I + e^{-\gamma t_I} c_{II} + e^{-\gamma(t_I+t_{II})} c_{III})$.

Go to step 15.

Step 15.

If $L = 0$,

if $b_f < k^*$, increase K by $e^{-\gamma t} K_5$, t by t_5 , b_f by 1, go to step 16.

Else go to step 6.

If $L = 1$,

If $b_f < k^*$, increase K by $e^{-\gamma t} K_5$, t by t_5 , b_f by 1, go to step 16.

Else increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$. Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(CD)$, if $PI(CD) > PI(ref)$, go to step 17; Else stop and output R and K .

If the current LS is the last LS to be optimized, we will decide whether to take an additional CD or to stop after taking the optimal number k^ CDs from this LS. If the PI for an additional CD is still greater than the reference PI, we take that CD, otherwise we stop.*

Step 16. If $X = 0$ or $Y = 0$, go to step 14.

Else go to step 13.

Step 17. Increase K by $e^{-\gamma t} K_5$, t by t_5 , b_i by 1, go to step 16.

C.3.3 FI Simulation in AP Setting

Step 1. Set $f, t, b_f, a_f, R, K, \mathbf{d}$ and v to 0.

Step 2. Increase K by K_1 , t by t_1 .

With probability:

p_1 , go to step 3,

q_1 , stop and output R and K .

Step 3. Increase K by $e^{-\gamma t} K_2$, t by t_2 .

With probability:

p_{21} , go to step 5,

p_{22} , go to step 4,

q_2 , stop and output R and K .

Step 4. Let $p_4 = m_4$.

Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$ and its corresponding k^* .

If $PI(LS) > PI(ref)$,

increase K by $e^{-\gamma t} K_4$, t by t_4 , and with probability:

p_4 , increase f by 1, b_f by 1, set $L = 0$, go to step 9,

q_4 , go to step 5.

If $PI(LS) < PI(ref)$, stop and output R and K .

This is the first attempt to optimize an LS , so we use the prior expectation m_4 as the success probability for stage 4 since there is no Bayesian update yet .

Step 5. If $n_3 = 0$, let $p_3 = m_3$. If $n_4 = 0$, let $p_4 = m_4$;

Else calculate $E[p_3]$ and $E[p_4]$ and let $p_3 = E[p_3]$, $p_4 = E[p_4]$.

Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$ and its corresponding k^* .

If $PI(LS) > PI(ref)$,

if condition A or B holds, go to step 8; Else go to step 7.

If $PI(LS) < PI(ref)$, stop and output R and K .

If there is no update on the success probabilities for stages 3 and 4 yet, their prior expectations m_3 and m_4 are used. Otherwise their posterior expectations $E[p_3]$ and $E[p_4]$ are used. .

Step 6. Calculate $E[p_3]$ and $E[p_4]$ and let $p_3 = E[p_3]$, $p_4 = E[p_4]$. Calculate $E[p_{Ii}]$, $E[p_{IIi}]$, $E[p_{IIIi}]$ for $i = 1, 2, \dots, f$. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Calculate $PI(LS)$, its corresponding k^* , $PI(CD)$ and let $c = i^*$.

If $PI(LS) > \max(PI(CD), PI(ref))$,

if either condition A or B holds, go to step 8; Else go to step 7.

If $PI(CD) > \max(PI(LS), PI(ref))$, let $p_I = E[p_{Ic}]$, $p_{II} = E[p_{IIc}]$ and $p_{III} = E[p_{IIIc}]$, go to step 13.

If $PI(ref) > \max(PI(CD), PI(LS))$, stop and output R and K .

For the estimations of $E[p_3]$, $E[p_4]$, $E[p_{Ii}]$, $E[p_{IIi}]$, and $E[p_{IIIi}]$ refer to section 5.1. For the calculations of $PI(LS)$ and $PI(CD)$ refer to appendix B.2. $PI(CD)$ is the PI for an additional CD from LS i^* , $1 \leq i^* \leq f$, which gives the highest PI .

Step 7. If $f = 0$, increase K by $e^{-\gamma t} K_4$, t by t_4 .

Else increase K by $e^{-\gamma t} K_{4p}$, t by ρt_4 .

Increase n_4 by 1, with probability:

p_4 , increase f by 1, b_f by 1, set $c = f$, $L = 1$,
 calculate $E[p_{Ic}]$, $E[p_{IIc}]$, $E[p_{IIIc}]$ and let $p_I = E[p_{Ic}]$, $p_{II} = E[p_{IIc}]$ and $p_{III} = E[p_{IIIc}]$, go to step 9,
 q_4 , stop and output R and K .

Step 8. If $f = 0$, increase K by $e^{-\gamma t}(K_3 + K_4)$, t by t_4 .

Else increase K by $e^{-\gamma t}(K_{3\rho} + K_{4\rho})$, t by ρt_4 .

Increase n_3 by 1, and n_4 by 1, with probability:

p_3p_4 , increase f by 1, b_f by 1, set $c = f$, $L = 0$, calculate $E[p_{Ic}]$, $E[p_{IIc}]$, $E[p_{IIIc}]$
 and let $p_I = E[p_{Ic}]$, $p_{II} = E[p_{IIc}]$ and $p_{III} = E[p_{IIIc}]$, go to step 9,

p_3q_4 , go to step 5,

q_3p_4 , increase f by 1, b_f by 1, set $c = f$, $L = 1$, calculate $E[p_{Ic}]$, $E[p_{IIc}]$, $E[p_{IIIc}]$
 and let $p_I = E[p_{Ic}]$, $p_{II} = E[p_{IIc}]$ and $p_{III} = E[p_{IIIc}]$, go to step 9,

q_3q_4 , stop and output R and K .

Step 9. Increase R by $-e^{-\gamma t}c_I$, with probability:

p_I , go to step 10;

q_I , got to step 12.

Step 10. Increase R by $-e^{-\gamma(t+t_I)}c_{II}$, with probability:

p_{II} , go to step 11;

q_{II} , got to step 12.

Step 11. Increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$.

Increase R by $-e^{-\gamma(t+t_I+t_{II})}c_{III}$, with probability:

p_{III} , increase R by $e^{-\gamma(t+t_I+t_{II}+t_{III})}\lambda^{\sum_{i=1}^f a_i}W$, a_f by 1, go to step 12;

q_{III} , got to step 12.

Step 12.

If $L = 0$,

if $b_f < k^*$, increase K by $e^{-\gamma t}K_5$, t by t_5 , b_f by 1, go to step 9.

Else go to step 6.

If $L = 1$,

If $b_f < k^*$, increase K by $e^{-\gamma t}K_5$, t by t_5 , b_f by 1, go to step 9.

Else increase v by 1. Let $\mathbf{d}[v] = t$, and $\Delta t = \mathbf{d}[v] - \mathbf{d}[v - 1]$.

Generate N_A from $\text{Pois}(\xi_A \Delta t)$, N_B from $\text{Pois}(\xi_B \Delta t)$, and N_C from $\text{Pois}(\xi_C \Delta t)$.

Let $W = W(1 - f_A)^{N_A}(1 - f_B)^{N_B}(1 - f_C)^{N_C}$. Calculate $PI(CD)$,

if $PI(CD) > PI(ref)$, go to step 13; Else stop and output R and K .

Step 13. Increase K by $e^{-\gamma t} K_5$, t by t_5 , b_c by 1, go to step 9.

Appendix D

Bellman Equation

This appendix discusses the equality condition used to calculate the option value in chapter 6. It is based on Dixit and Pindyck (1994, chapter 4).

A Bellman equation is a necessary condition for optimality associated with dynamic programming. It writes the value of a decision problem at a certain point in time as the sum of the payoff from the initial choice and the optimal value of the remaining decision problem that results from the initial choice. This breaks a dynamic optimization problem into simpler subproblems.

D.1 Discrete Time

Let the state at time t be x_t . At time t a choice of actions is available to the firm. We denote the action chosen by u_t . The set of available actions depends on x_t . The state and control at time t in turn affect the firm's immediate profit flow, which we denote by $\pi_t(x_t, u_t)$. The discount factor between any two periods is $1/(1 + \gamma)$, where γ is the discount rate.

Denote by $F_t(x_t)$ the expected net present value of all the firm's cash flows—when the firm makes all decisions optimally from this point onwards. When the firm chooses the control variable u_t , it gets an immediate profit flow $\pi_t(x_t, u_t)$. At the next period $t + 1$, the state will be x_{t+1} , and optimal decisions thereafter will yield

$F_{t+1}(x_{t+1})$. This is random from the perspective of period t , so we must take its expected value $E_t[F_{t+1}(x_{t+1})]$. Discounting back to period t , the sum of the immediate payoff and the continuation value is

$$\pi_t(x_t, u_t) + \frac{1}{1 + \gamma} E_t[F_{t+1}(x_{t+1})].$$

The Bellman's principle of optimality states: *An optimal policy has the property that, whatever the initial action, the remaining choices constitute an optimal policy with respect to the subproblem starting at the state that results from the initial actions.* Here the optimality of the remaining choices u_{t+1}, u_{t+2}, \dots , is subsumed into the continuation value $E_t[F_{t+1}(x_{t+1})]$, so only the immediate control u_t remains to be chosen optimally. Thus

$$F_t(x_t) = \max_{u_t} \left\{ \pi_t(x_t, u_t) + \frac{1}{1 + \gamma} E_t[F_{t+1}(x_{t+1})] \right\}. \quad (\text{D.1})$$

Equation (D.1) is called the *Bellman equation*, or the *fundamental equation of optimality*. To reiterate, the first term on the right-side is the immediate profit, the second term constitutes the continuation value, and the optimum action this period is the one that maximizes the sum of these two components.

In the setting of an infinite horizon, the value function $F_t(x_t)$ is common to all periods, although it will be evaluated at different points x_t . The Bellman equation for any t becomes

$$F(x_t) = \max_{u_t} \left\{ \pi_t(x_t, u_t) + \frac{1}{1 + \gamma} E_t[F(x_{t+1})] \right\}.$$

Since x_t and x_{t+1} could be any of the possible states, we write them in general form as x and x' . Then for all x , we have

$$F(x) = \max_u \left\{ \pi(x, u) + \frac{1}{1 + \gamma} E[F(x') | x, u] \right\}, \quad (\text{D.2})$$

where we have now denoted the expectation as conditioned on the knowledge of the current period's x and u . This is the Bellman equation for the recursive dynamic programming problem.

D.2 Continuous Time

Suppose now each time period is of length Δt . Ultimately we are interested in the limit where Δt goes to zero and time is continuous. We write $\pi(x, u, t)$ for the rate of the profit flow, so that the actual profit over the time period of length Δt is $\pi(x, u, t)\Delta t$. Similarly, let γ be the discount rate per unit time, so the total discounting over an interval of length Δt is by the factor $1/(1 + \gamma\Delta t)$.

The Bellman equation (D.2) now becomes

$$F(x, t) = \max_u \{ \pi(x, u, t)\Delta t + \frac{1}{1 + \gamma\Delta t} E[F(x', t + \Delta t)|x, u] \}.$$

Multiply by $(1 + \gamma\Delta t)$ and rearrange to write

$$\begin{aligned} \gamma\Delta t F(x, t) &= \max_u \{ \pi(x, u, t)\Delta t(1 + \gamma\Delta t) + E[F(x', t + \Delta t) - F(x, t)] \} \\ &= \max_u \{ \pi(x, u, t)\Delta t(1 + \gamma\Delta t) + E[\Delta F] \}. \end{aligned}$$

Dividing by Δt and letting it go to zero, we get

$$\gamma F(x, t) = \max_u \{ \pi(x, u, t) + \frac{1}{dt} E[dF] \}, \quad (\text{D.3})$$

where $(1/dt)E[dF]$ is the limit of $E[\Delta F]/\Delta t$. This is the Bellman equation for continuous time.

On the right left side of equation (D.3) we have the normal return per unit time that a decision maker, using γ as the discount rate, would require for holding the asset. On the right side, the first term is the immediate payout or dividend from the asset, while the second term is its expected rate of capital gain. Thus the right side is the expected total return per unit time from holding the asset. The maximization with respect to u means that the current operation of the asset is being managed optimally, bearing in mind not only the immediate payout but also the consequences for future values.

Appendix E

Itô's Formula

This appendix is referred in the calculation of the stochastic differential dF in equation (6.6) in chapter 6. It is based on Øksendal (2002, chapter 4).

Theorem (The 1-dimensional Itô formula).

Let X_t be an Itô process given by

$$dX_t = udt + vdz_t.$$

Let $f(x, t)$ be twice continuously differentiable on $[0, \infty) \times \mathbf{R}$.

The Itô formula gives

$$df(X_t, t) = \frac{\partial f}{\partial t}(X_t, t)dt + \frac{\partial f}{\partial x}(X_t, t)dX_t + \frac{1}{2} \frac{\partial^2 f}{\partial x^2}(X_t, t)(dX_t)^2, \quad (\text{E.1})$$

where $(dX_t)^2 = (dX_t) \cdot (dX_t)$ is computed according to the rules

$$dt \cdot dt = dt \cdot dz_t = dz_t \cdot dt = 0, \quad dz_t \cdot dz_t = dt. \quad (\text{E.2})$$

Sketch of proof. Using Taylor's theorem we get

$$\begin{aligned} f(X_t, t) &= f(X_0, 0) + \sum_j \Delta f(X_j, t_j) = f(X_0, 0) + \sum_j \frac{\partial f}{\partial t} \Delta t_j + \sum_j \frac{\partial f}{\partial x} \Delta X_j \\ &\quad + \frac{1}{2} \sum_j \frac{\partial^2 f}{\partial t^2} (\Delta t_j)^2 + \sum_j \frac{\partial^2 f}{\partial t \partial x} (\Delta t_j)(\Delta X_j) + \frac{1}{2} \sum_j \frac{\partial^2 f}{\partial x^2} (\Delta X_j)^2 + \sum_j R_j, \end{aligned}$$

where $\frac{\partial f}{\partial t}$, $\frac{\partial f}{\partial x}$ etc. are evaluated at the points (X_{t_j}, t_j) , $\Delta t_j = t_{j+1} - t_j$, $\Delta X_j = X_{t_{j+1}} - X_{t_j}$, $\Delta f(X_j, t_j) = f(X_{t_{j+1}}, t_{j+1}) - f(X_{t_j}, t_j)$ and $R_j = o(|\Delta t_j|^2 + |\Delta X_j|^2)$ for all j .

If $\Delta t_j \rightarrow 0$, then

$$\begin{aligned} \sum_j \frac{\partial f}{\partial t} \Delta t_j &= \sum_j \frac{\partial f}{\partial t}(X_j, t_j) \Delta t_j \rightarrow \int_0^t \frac{\partial f}{\partial t}(X_s, s) ds, \\ \sum_j \frac{\partial f}{\partial x} \Delta X_j &= \sum_j \frac{\partial f}{\partial x}(X_j, t_j) \Delta X_j \rightarrow \int_0^t \frac{\partial f}{\partial x}(X_s, s) dX_s. \end{aligned}$$

Moreover, since u and v are elementary we get

$$\begin{aligned} \sum_j \frac{\partial^2 f}{\partial x^2} (\Delta X_j)^2 &= \sum_j \frac{\partial^2 f}{\partial x^2} u^2 (\Delta t_j)^2 + 2 \sum_j \frac{\partial^2 f}{\partial x^2} uv (\Delta t_j) (\Delta z_j) \\ &+ \sum_j \frac{\partial^2 f}{\partial x^2} v^2 \cdot (\Delta z_j)^2. \end{aligned} \quad (\text{E.3})$$

The first two terms here in equation (E.3) tend to 0 as $\Delta t_j \rightarrow 0$. We claim the last term tends to

$$\int_0^t \frac{\partial^2 f}{\partial x^2} v^2 ds \quad \text{as } \Delta t_j \rightarrow 0.$$

To prove this, put $a(t) = \frac{\partial^2 f}{\partial x^2}(X_t, t)v^2$ and $a_j = a(t_j)$ and consider

$$E \left[\left(\sum_j a_j (\Delta z_j)^2 - \sum_j a_j \Delta t_j \right)^2 \right] = \sum_{ij} E[a_i a_j ((\Delta z_i)^2 - \Delta t_i) ((\Delta z_j)^2 - \Delta t_j)]. \quad (\text{E.4})$$

If $i \neq j$ then $((\Delta z_i)^2 - \Delta t_i)$ and $(\Delta z_j)^2 - \Delta t_j$ are independent so the corresponding terms in equation (E.4) vanish in this case. So we are left with

$$\begin{aligned} \sum_j E[a_j^2 ((\Delta z_j)^2 - \Delta t_j)^2] &= \sum_j E[a_j^2] \cdot E[(\Delta z_j)^4 - 2(\Delta z_j)^2 \Delta t_j + (\Delta t_j)^2] \\ &= \sum_j E[a_j^2] \cdot (3(\Delta t_j)^2 - 2(\Delta t_j)^2 + (\Delta t_j)^2) = 2 \sum_j E[a_j^2] \cdot (\Delta t_j)^2 \\ &\rightarrow 0 \quad \text{as } \Delta t_j \rightarrow 0. \end{aligned}$$

In other words, we have established that

$$\sum_j a_j (\Delta z_j)^2 \rightarrow \int_0^t a(s) ds \quad \text{as } \Delta t_j \rightarrow 0,$$

and this is often expressed by the striking formula

$$(dz_t)^2 = dt. \tag{E.5}$$

The argument above also proves that $\sum R_j \rightarrow 0$ as $\Delta t_j \rightarrow 0$. That completes the proof of the Itô formula. \square

Appendix F

Root Finding Techniques

This appendix describes the two standard root finding techniques mentioned in this thesis. The Newton-Raphson method is used to find the optimal PI and IRR with respect to allocation variables (section 4.3.3). The bisection method is used to calculate the parameter values of a beta distribution given two quantile conditions (section 5.2).

F.1 Newton-Raphson Method

Perhaps the most celebrated of all one-dimensional root-finding routines is the Newton-Raphson method. This method requires the evaluation of both the function $f(x)$ and the derivative $f'(x)$, at arbitrary points x . The Newton-Raphson formula consists geometrically of extending the tangent line at a current point x_i until it crosses zero, then setting the next guess x_{i+1} to the abscissa of that zero crossing. Algebraically, the method derives from the Taylor series expansion of a function $f(x)$ at some point x_0 ,

$$f(x) \approx f(x_0) + (x - x_0)f'(x_0) + \frac{(x - x_0)^2}{2}f''(x_0) + o(|x - x_0|^2).$$

Setting the quadratic and higher terms to zero and solving the linear approximation of $f(x) = 0$ for x gives

$$x_1 = x_0 - \frac{f(x_0)}{f'(x_0)}.$$

Subsequent iterations are defined in a similar manner as

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}, \quad (\text{F.1})$$

the process is repeated until a sufficiently accurate value is reached.

Next we consider the algorithm's rate of convergence. Within a small distance ϵ of x , the function and its derivative are approximately

$$\begin{aligned} f(x + \epsilon) &= f(x) + f'(x)\epsilon + \frac{f''(x)}{2}\epsilon^2 + \dots, \\ f'(x + \epsilon) &= f'(x) + f''(x)\epsilon + \dots \end{aligned} \quad (\text{F.2})$$

By the Newton-Raphson formula,

$$x_{i+1} = x_i - \frac{f(x_i)}{f'(x_i)}, \quad (\text{F.3})$$

so that

$$\epsilon_{i+1} = \epsilon_i - \frac{f(x_i)}{f'(x_i)}.$$

When a trial solution x_i differs from the true root by ϵ_i , we can use (F.2) to express $f(x_i)$, $f'(x_i)$ in (F.3) in terms of ϵ_i and derivatives at the root itself. The result is a recurrence relation for the deviations of the trial solutions

$$\epsilon_{i+1} = -\epsilon_i^2 \frac{f''(x_i)}{2f'(x_i)}. \quad (\text{F.4})$$

Equation (F.4) says that Newton-Raphson converges quadratically. Near a root, the number of significant digits approximately doubles with each step. This very strong convergence property makes Newton-Raphson the method of choice for any function whose derivative can be evaluated efficiently, and whose derivative is continuous and nonzero in the neighborhood of a root.

The Newton-Raphson method can also be used to find the minimum or maximum of a function $f(x)$. The derivative is zero at the minimum or maximum, so the minima or maxima can be found by applying Newton-Raphson method to the derivative $f'(x)$. The iteration becomes:

$$x_{n+1} = x_n - \frac{f'(x_n)}{f''(x_n)}. \quad (\text{F.5})$$

F.2 Bisection Method

The bisection method is a root-finding method which repeatedly bisects an interval then selects a subinterval in which a root must lie for further processing. It is very simple and robust, but it is relatively slow.

Two initial points a and b which bracket the root are required such that

$$f(a)f(b) < 0.$$

By the intermediate value theorem the continuous function f must have at least one root in the interval $[a, b]$. The method now divides the interval in two by computing the midpoint $c = (a + b)/2$. Evaluate $f(c)$ and examine its sign. Unless c is itself a root, there are now two possibilities. If $f(a)f(c) > 0$ replace a by c ; otherwise if $f(b)f(c) > 0$ replace b by c . After each iteration the interval containing the root decreases by a factor of two. We continue until we have a bracket sufficiently small. If after n iterations the root is known to be within an interval of size ϵ_n , then after the next iteration it will be bracketed within an interval of size

$$\epsilon_{n+1} = \epsilon_n/2.$$

Thus, we know in advance the number of iterations required to achieve a given tolerance in the solution,

$$n = \log_2 \frac{\epsilon_0}{\epsilon},$$

where $\epsilon_0 = b - a$ is the size of the initial interval and ϵ is the desired ending tolerance.

If the interval $[a, b]$ happens to contain more than one root, bisection will find one of them.

Bibliography

- Abraham, A. and Taylor, W. M. (1993). An event option pricing model with scheduled and unscheduled announcement effects. *Working Paper*.
- Accenture (2001). High performance drug discovery. Technical report, Accenture, www.accenture.com.
- Accenture (2005). Rethinking innovation in pharmaceutical R&D. Technical report, Accenture, www.accenture.com.
- Accenture (2007). The pursuit of high performance through research and development. Technical report, Accenture, www.accenture.com.
- Ames, W. F. (1977). *Numerical Methods for Partial Differential Equations*. Academic Press, New York.
- Amin, K. I. (1993). Jump diffusion option valuation in discrete time. *Journal of Finance*, 48:1833–1863.
- Bain and Company (2009). Bringing pharma R&D back to health. Technical report, Bain and Company, www.bain.com.
- Bain and Company (2010). Changing pharma’s innovation DNA. Technical report, Bain and Company, www.bain.com.
- Ball, C. A. and Torous, W. N. (1985). On jumps in common stock prices and their impact on call option pricing. *Journal of Finance*, 40:155–173.
- Beardsley, G. and Mansfield, E. (1978). A note on the accuracy of industrial forecasts. *Journal of Business*, 51:127–135.

- Bergman, S. W. and Gittins, J. C. (1985). *Statistical Methods for Pharmaceutical Research Planning*. Marcel Dekker, New York.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economics*, 81:637–659.
- Bode-Greuel, K. M. and Greuel, J. M. (2005). Determining the value of drug development candidates and technology platforms. *Journal of Commercial Biotechnology*, 11:155–170.
- Bowman, E. H. and Moskowitz, G. T. (2001). Real options analysis and strategic decision making. *Organization Science*, 12:772–777.
- Brealey, R. A. and Myers, S. C. (2000). *Principles of Corporate Finance*. McGraw Hill, New York, 6th edition.
- Brennan, M. J. and Schwartz, E. S. (1978). Finite difference method and jump processes arising in the pricing of contingent claims. *Journal of Financial and Quantitative Analysis*, 13:461–474.
- Casella, G. and Berger, R. (2001). *Statistical Inference*. Duxbury, 2nd edition.
- Charalambous, C. (2009). Models and software for improving the profitability of pharmaceutical research. Dphil thesis, University of Oxford.
- Charalambous, C. and Gittins, J. C. (2008). Optimal selection policies for a sequence of candidate drugs. *Advances in Applied Probability*, 40:359–376.
- Chen, B. P. K. (2004). Prioritization of research projects in the pharmaceutical industry. Dphil thesis, University of Oxford.
- Childs, P. D. and Triantis, A. J. (1999). Dynamic R&D investment policies. *Management Science*, 45:1359–1377.
- Consonni, G. and Veronese, P. (1995). A Bayesian method for combining results from several binomial experiments. *Journal of the American Statistical Association*, 90:935–944.
- Cox, J., Ross, S., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7:229–263.

- Daniel, C. (1959). Use of half-normal plots in interpreting factorial two-level experiments. *Technometrics*, 1:311–341.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press, New Jersey.
- Dorp, J. R. V. and Mazzuchi, T. A. (2000). Solving for the parameters of a beta distribution under two quantile constraints. *Journal of Statistical Computation and Simulation*, 67:189–201.
- EFPIA (2009). The pharmaceutical industry in figures 2009. Technical report, European Federation of Pharmaceutical Industries and Associations, www.efpia.org.
- EFPIA (2010). The pharmaceutical industry in figures 2010. Technical report, European Federation of Pharmaceutical Industries and Associations, www.efpia.org.
- Geske, R. (1979). The valuation of compound options. *Journal of Financial Economics*, 7:63–81.
- Geske, R. and Shatri, K. (1985). Valuation of approximation: A comparison of alternative approaches. *Journal of Financial and Quantitative Analysis*, 20:45–72.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society*, 41:148–177.
- Gittins, J. C. (1996). Quantitative methods in the planning of pharmaceutical research. *Drug Information Journal*, 30:479–487.
- Gittins, J. C. (1997). Why crash pharmaceutical research? *R&D Management*, 27:79–85.
- Gittins, J. C., Charalambous, C., Chen, B., Oreskovich, A. M., Qu, S., and Yu, J. Y. (2011a). *OPRRA User-guide*. Department of Statistics, University of Oxford.

- Gittins, J. C., Glazebrook, K. D., and Weber, R. (2011b). *Multi-armed bandit allocation indices*. Wiley, 2nd edition.
- Glazebrook, K. D. (1995). Stochastic scheduling and forwards induction. *Discrete Applied Mathematics*, 57:145–165.
- Glazebrook, K. D. and Gittins, J. C. (1993). The performance of forwards induction policies. *Stochastic Processes and their Applications*, 46:301–326.
- Guthrie, G. (2009). *Real options in theory and practice*. Oxford University Press, Oxford.
- Hariharan, R. and Singh, N. (2010). What gives drugs the x factors? Technical report, McKinsey, www.mckinseyquarterly.com.
- Hayes, R. H. and Abernathy, W. J. (1980). Managing our way to economic decline. *Harvard business review*, 58:67–77.
- Hayes, R. H. and Garvin, D. A. (1982). Managing as if tomorrow mattered. *Harvard Business Review*, 60:70–79.
- Hoeschler, M. (2002). Optimization of pharmaceutical research. Master's thesis, University of Oxford.
- Hubbard, R. (2006). *Structure-Based Drug Discovery: An Overview*. RSC Publishing, Cambridge.
- Hull, J. C. (2008). *Options, Futures, and other Derivatives*. Pearson Prentice Hall, New Jersey, 7th edition.
- Hull, J. C. and White, A. (1990). Valuing derivative securities using the explicit finite difference method. *Journal of Financial and Quantitative Analysis*, 25(1):87–100.
- Jones, E. P. (1984). Option arbitrage and strategy with large price changes. *Journal of Financial Economics*, 13:91–113.
- Kester, W. C. (1984). Today's options for tomorrow's growth. *Harvard Business Review*, 62:153–160.

- Loch, C. H. and Bode-Greuel, K. M. (2002). Evaluating growth options as sources of value for pharmaceutical research projects. *R&D Management*, 31:231–248.
- Majd, S. and Pindyck, R. S. (1987). Time to build, option value, and investment decisions. *Journal of Financial Economics*, 18:7–27.
- Martzoukos, S. H. and Trigeorgis, L. (2002). Real options with multiple sources of rare events. *European Journal of Operational Research*, 136:696–706.
- McKinsey (2010). Invention reinvented: Mckinsey’s perspectives on pharmaceutical R&D. Technical report, McKinsey, www.mckinseyquarterly.com.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4:141–183.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:124–144.
- Merton, R. C. (1977). On the pricing of contingent claims and the Modigliani-Miller theorem. *Journal of Financial Economics*, 5:147–175.
- Miller, P. (2005). Role of pharmacoeconomic analysis in R&D decision making. *Pharmacoeconomics*, 23:1–12.
- Myers, S. C. (1976). *Modern developments in financial management*. Praeger.
- Myers, S. C. (1984). Finance theory and financial strategy. *Interfaces*, 14:126–137.
- Myers, S. C. and Howe, C. D. (1997). A life-cycle financial model of pharmaceutical R&D. *Working Paper, Program on the Pharmaceutical Industry*.
- Newton, D. P. and Pearson, A. W. (1994). Application of option pricing theory to R&D. *R&D Management*, 24:83–89.
- Nichols, N. A. (1994). Scientific management at Merck: An interview with CFO Judy Lewent. *Harvard Business Review*, 72:88–99.

- Øksendal, B. K. (2002). *Stochastic Differential Equations: An Introduction with Applications*. Springer, 5th edition.
- Paxson, D. A. (2003). *Real R&D options*. Butterworth-Heinemann, Boston, MA.
- Pennings, E. and Lint, O. (1997). The option value of advanced R&D. *European Journal of Operational Research*, 103:83–94.
- Pindyck, R. S. (1993). Investments of uncertain cost. *Journal of Financial Economics*, 34:53–76.
- Poh, K. L., Ang, B. W., and Bai, F. (2001). A comparative analysis of R&D project evaluation methods. *R&D Management*, 31:63–75.
- Press, W. H. and Teukolsky, S. (2007). *Numerical Recipes*. Cambridge University Press, Cambridge, 3rd edition.
- PricewaterhouseCoopers (2008). Pharma 2020: Virtual R&D-Which path will you take? Technical report, PricewaterhouseCoopers, www.pwc.co.uk.
- Qu, S. and Gittins, J. C. (2011). A forwards induction approach to candidate drug selection. *Advances in Applied Probability*, 43(3).
- Rodda, B., Miller, S., and Krause, A. (2001). *Applied Statistics in Pharmaceutical Industry*. Springer-Verlag, New York.
- Russell, M. and Gjaja, M. (2008). The future of pharma: Adjusting the pharma R&D model. Technical report, Boston Consulting Group, www.bcg.com.
- Schwartz, E. S. (1977). The valuation of warrants: Implementing a new approach. *Journal of Financial Economics*, 4:79–93.
- Schwartz, E. S. and Trigeorgis, L. (2004). *Real Options and Investment Under Uncertainty: Classical Readings and Recent Contributions*. MIT Press.
- Sender, G. L. (1994). Option analysis at Merck. *Harvard Business Review*, 72:92.
- Senn, S. (1996). Some statistical issues in project prioritization in the pharmaceutical industry. *Statistics in Medicine*, 15:2689–2702.

- Senn, S. (1998). Further statistical issues in project prioritization in the pharmaceutical industry. *Drug Information Journal*, 32:253–259.
- Singh, N., Zimmel, R., and Denari, A. (2010). The anatomy of attrition. Technical report, McKinsey, www.mckinseyquarterly.com.
- Thackray, J. (1995). A Merck case study. *Planning Review*, 23:47.
- Triantis, A. (2001). Real options: State of the practice. *Journal of applied corporate finance*, 14:8–20.
- Trigeorgis, L. (1996). *Real options: Managerial Flexibility and Strategy in Resource Allocation*. MIT Press.
- Trigeorgis, L. (1999a). *Real Options and Business Strategy: Applications to Decision Making*. Risk Books, London.
- Trigeorgis, L. (1999b). *Real Options and Investment under Uncertainty*. The MIT Press, Cambridge.
- Urban, G. L., Carter, T., Gaskins, S., and Mucha, Z. (1986). Market share rewards to pioneering brands: An empirical analysis and strategic implications. *Management Science*, 32:645–659.
- Walley, T., Haycox, A., and Boland, A. (2004). *Pharmacoeconomics*. Churchill Livingstone, London.
- Wilder, S. (2001). Investigating the optimal allocation of resources for pharmaceutical research projects. Master's thesis, University of Oxford.
- Yu, J.-Y. (2007). Models and software for improving the profitability of pharmaceutical research. Dphil thesis, University of Oxford.
- Yu, J.-Y. and Gittins, J. C. (2007). Models and software for improving the profitability of pharmaceutical research. *European Journal of Operational Research*, 189:459–475.