

Resource Allocation Decision-Making in Sequential Adaptive Clinical Trials

Alba Claudia Rojas-Cordova

Dissertation proposal submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Industrial and Systems Engineering

Ebru K. Bish, Co-Chair
Niyousha Hosseinichimeh, Co-Chair
Douglas R. Bish
Navid Ghaffarzadegan
Konstantinos P. Triantis

May 3, 2017
Blacksburg, Virginia

Keywords: Clinical trials, Simulation, Stochastic Dynamic Programming, Sequential Decision
Analysis, Bayesian Learning

Copyright 2017, Alba C. Rojas-Cordova

Resource Allocation Decision-Making in Sequential Adaptive Clinical Trials

Alba Claudia Rojas-Cordova

ABSTRACT

Adaptive clinical trials for new drugs or treatment options promise substantial benefits to both the pharmaceutical industry and the patients, but complicate resource allocation decisions. In this dissertation, we focus on *sequential adaptive clinical trials* with binary response, which allow for early termination of drug testing for benefit or futility at interim analysis points. The option to stop the trial early enables the trial sponsor to mitigate investment risks on ineffective drugs, and to shorten the development time line of effective drugs, hence reducing expenditures and expediting patient access to these new therapies. In this setting, decision makers need to determine a *testing schedule*, or the number of patients to recruit at each interim analysis point, and *stopping criteria* that inform their decision to continue or stop the trial, considering performance measures that include drug misclassification risk, time-to-market, and expected profit. In the first manuscript, we model current practices of sequential adaptive trials, so as to quantify the magnitude of drug misclassification risk. Towards this end, we build a simulation model to realistically represent the current decision-making process, including the utilization of the triangular test, a widely implemented sequential methodology. We find that current practices lead to a high risk of incorrectly terminating the development of an effective drug, thus, to unrecoverable expenses for the sponsor, and unfulfilled patient needs. In the second manuscript, we study the sequential resource allocation decision, in terms of a testing schedule and stopping criteria, so as to quantify the impact of interim analyses on the aforementioned performance measures. Towards this end, we build a stochastic dynamic programming model, integrated with a Bayesian learning framework for updating the drug's estimated efficacy. The resource allocation decision is characterized by endogenous uncertainty, and a trade-off between the incentive to establish that the drug is effective early on (exploitation), due to a time-decreasing market revenue, and the benefit from collecting some information on the drug's efficacy prior to committing a large budget (exploration). We derive important structural properties of an optimal resource allocation strategy and perform a numerical study based on realistic data, and show that sequential adaptive trials with interim analyses substantially outperform traditional trials. Finally, the third manuscript integrates the first two models, and studies the benefits of an optimal resource allocation decision over current practices. Our findings indicate that our optimal testing schedules outperform different types of fixed testing schedules under both perfect and imperfect information.

Resource Allocation Decision-Making in Sequential Adaptive Clinical Trials

Alba Claudia Rojas-Cordova

GENERAL AUDIENCE ABSTRACT

Adaptive clinical trials for new drugs or treatment options have the potential to reduce pharmaceutical research and development costs, and to expedite patient access to new therapies. Sequential adaptive clinical trials allow investigators and trial sponsors to terminate drug testing “early,” at interim analysis points, either for benefit or futility reasons. In the first manuscript, we model current practices of sequential adaptive trials, so as to quantify the risk of terminating the development of an effective drug incorrectly. Towards this end, we build a simulation model to realistically represent the current decision-making process. In the second manuscript, we study the financial investment decisions made by the trial sponsor, such as pharmaceutical firms, so as to quantify the impact of interim analyses on a series of performance measures relevant to the firm and the patients. Towards this end, we build a mathematical optimization model that incorporates elements representing the knowledge gained by decision makers on the drug’s efficacy, which is unknown to them at the beginning of the trial. As a result of our analysis, we obtain an optimal strategy to allocate financial resources in a sequential adaptive trial. In the third and final manuscript, we compare the performance of our optimal resource allocation strategy against the performance of the triangular test, a well-known and widely implemented sequential testing methodology, as measured by the aforementioned performance measures.

To my granny Olga and my beloved husband Nikhil

Acknowledgments

Finding the right words to convey my gratitude and appreciation to my co-advisers Drs. Ebru K. Bish and Niyousha Hosseinichimeh is not an easy task. Ebru and Niyousha have been wonderful mentors, collaborators, and friends. My experience as Ebru and Niyousha's advisee has been nothing but enriching and uplifting, and I am deeply grateful to them for their continuous encouragement. Ebru cares deeply about conducting research with honesty and integrity. These two values are present in all her activities, complementing the rigorousness and innovativeness of her research, and making her an exemplary adviser. I look up to her and feel very fortunate to have her as my role model.

Niyousha's dedication, creativity, and generosity were crucial components of my success as a doctoral student, and I cannot thank her enough for the time and effort she invested in advising me. Despite her numerous commitments and activities as an assistant professor, she always provided me with valuable guidance.

I am very thankful to Drs. Douglas R. Bish, Navid Ghaffarzadegan and Konstantinos Triantis for their insightful comments and feedback throughout the different stages of my doctoral research. I enjoyed working with them, and truly appreciate their support.

My deepest gratitude goes to Dr. Ronanth Zavaleta-Mercado, my undergraduate mentor, for encouraging me to pursue graduate studies in the United States. Dr. Zavaleta's thoughtful advice defined my career plans as a researcher, and took me farther than I ever imagined. I am also thankful to Dr. Eileen Van Aken, for her kindness, generosity, and honest advice. Dr. Van Aken's words and support were invaluable to me, and I cannot feel anything but honored to have her as a mentor and friend.

My days as a doctoral student would have not been the same without the weekly calls to my mother Iris, and grandpa Aniceto. I thank them for their continuous support. Most importantly, I wish my grandma Olgui would be able to see me receive my Ph.D. diploma. Nobody believed in me as much as she did, and her love made me strong and confident to face the hardest challenges. Although my life is certainly not the same after her passing, I keep her sweet voice and loving words etched in my memory.

Last but certainly not least, I want to thank my dear husband Nikhil Gupta. In good and bad days, highly productive and insufficiently productive weeks, he was always there for me. He is a blessing and the most precious gift the Lord has put in my life. I would have probably given up

on numerous projects if I had not had the fortune of enjoying his witty jokes, and receiving his encouraging words and unconditional love. I look forward to many more years with you, Nikhil.

Contents

1	Introduction	1
1.1	Background and Motivation	1
1.1.1	The Importance of Pharmaceutical R&D	1
1.1.2	Sequential Adaptive Clinical Trials with Binary Response	2
1.2	Dissertation Outline	3
2	Trial Termination and Drug Misclassification in Sequential Adaptive Clinical Trials	6
2.1	Introduction	6
2.1.1	Context: Phase 3 Clinical Trials	8
2.1.2	Objectives	8
2.2	Background	9
2.2.1	Group Sequential Adaptive Trials and Stopping Rules	9
2.2.2	Critiques Regarding Early Termination	10
2.2.3	False Negatives and the Hot Stove Effect	12
2.3	Methods	12
2.3.1	Model Overview	12
2.3.2	Model Formulation	14
2.3.2.1	Patient Randomization	14
2.3.2.2	Drug Testing Process	14
2.3.2.3	Statistical analysis	15
2.4	Results	17
2.4.1	An Illustrative Case	17
2.4.2	Drug Misclassification Under a Wider Range of Conditions	19
2.4.2.1	False Positives vs. False Negatives with a Constant Interim Analysis Limit	19
2.4.2.2	Misclassification of an Effective Drug with Different Interim Analysis Limits	20
2.4.2.3	False Positives and False Negatives with Different Interim Analysis Limits	22
2.5	Discussion	23
2.5.1	Insights and Implications	23
2.5.2	Limitations and Directions for Future Research	26
2.6	Conclusions	27

3	Optimal Sequential Resource Allocation in Sequential Adaptive Clinical Trials	29
3.1	Introduction and Motivation	29
3.2	Related Literature	33
3.2.1	Group Sequential Adaptive Clinical Trials	33
3.2.2	Dynamic Programming and Information Acquisition	34
3.2.3	Resource Allocation	36
3.3	The Model	37
3.3.1	The Decision Problem	37
3.3.2	Information updating	39
3.3.3	The Optimization Model	40
3.4	Properties of an Optimal Solution	41
3.5	Numerical Analysis and Insights	46
3.5.1	Experimental Design	47
3.5.2	Optimal Resource Allocation in Fixed Duration vs. Sequential Adaptive Trials	49
3.5.3	Optimistic vs. Pessimistic Priors, Risk of Misclassification, and Time-to-market	51
3.5.4	The Impact of Resource Allocation on the Drug Misclassification Risk: A Sensitivity Analysis	53
3.6	Conclusions and Directions for Future Research	55
4	Benefits of Optimal Testing Schedules for Sequential Adaptive Clinical Trials	57
4.1	Introduction and Motivation	57
4.2	Related Literature	61
4.2.1	The Triangular Test and Sequential Trials with Binary Response	61
4.2.2	Dynamic Programming and Information Acquisition	65
4.3	The Optimization Model	66
4.3.1	The Decision Problem	66
4.3.2	Information Updating	67
4.3.3	The Stochastic Dynamic Programming Model	68
4.3.4	Structural Properties of an Optimal Testing Schedule	69
4.3.5	An Optimal Resource Allocation Strategy	69
4.4	Numerical Analysis and Insights	70
4.4.1	Experimental Design	71
4.4.2	Description of Fixed Testing Schedules and Stopping Boundaries	74
4.4.3	Optimal Resource Allocation	75
4.4.4	Performance of Fixed and Optimal Testing Schedules under Perfect Information	76
4.4.5	Fixed and Optimal Testing Schedules under Optimistic and Pessimistic Priors	77
4.5	Discussion and Insights	80
4.6	Conclusions and Directions for Future Research	81

5	Conclusions and Future Research Directions	84
5.1	Summary	84
5.2	Broader Impact	86
5.3	Future Research Directions	88
	Bibliography	90
A	Appendix for Chapter 2	98
A.1	Drug Testing Process	98
A.2	Detailed Description of Simulation Model	99
B	Appendix for Chapter 3	100

List of Figures

1.1	Overview of Chapters 2 - 4	4
2.1	Structure of the patient recruitment decision-making model for a single drug (Phase 3)	13
2.2	False negatives (Hot stove effect) shown in red	18
2.3	Absence of false positives	19
2.4	Drug misclassification rates at the interim analysis point for beneficial and futile drugs	20
2.5	Correct and incorrect classification rates, and expected total number of patients for a beneficial drug under perfect information and with different interim analysis limits	21
2.6	Percentage of drugs with trials terminated early	22
3.1	Testing and drug misclassification regions delimited by thresholds on the drug's estimated probability of efficacy.	44
3.2	The value function versus N_t in an intermediate period t , $0 < t < T$ (increments of 2, 3 and 5 units of N_t in $[10, 35]$, and increments of 10 units of N_t in $[45, 135]$). . .	46
3.3	Optimal first period resource allocation in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)	49
3.4	Optimal resource allocation in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)	50
3.5	Impact of deviations from the optimal resource allocation strategy on misclassification risk in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)	54
4.1	Stopping boundaries for the triangular test under discrete decision-making (Adapted from Whitehead [109])	64
4.2	Triangular test's stopping boundaries and the success criterion used in numerical study	73
A.1	Conceptual representation of the drug testing process for treatment group g , where $g \in \{E(\text{experimental}), C(\text{control})\}$	98

List of Tables

2.1	Model parameters and results for 100 simulations	18
3.1	Drug misclassification risk (%) for fixed duration and sequential adaptive trials . .	51
3.2	Expected time-to-market (in number of periods) for fixed duration and sequential adaptive trials (NA=not applicable, as the drug is misclassified as futile)	51
3.3	Expected profit and costs of pessimism vs. optimism for fixed duration and sequential adaptive trials	51
4.1	Summary of the notation used in a parallel group clinical trial with binary response (Adapted from Whitehead [109])	63
4.2	Fixed testing schedules and stopping boundaries	74
4.3	Optimal first period resource allocation (N_1^*) for the different trial settings	75
4.4	Drug misclassification risk (%) for the different trial settings under perfect information ($p_E = 0.66$)	76
4.5	Expected time-to-market (in number of periods) for the different trial settings under perfect information ($p_E = 0.66$)	76
4.6	Expected profit (in \$ million) for the different trial settings under perfect information ($p_E = 0.66$)	76
4.7	Drug misclassification risk (%) for the different trial settings under optimistic and pessimistic priors (* = fraction of time where the final trial results are inconclusive)	78
4.8	Expected time-to-market (in number of periods) for the different trial settings under optimistic and pessimistic priors (NA=not applicable, as the drug is futile)	78
4.9	Expected profit (in \$ million) for the different trial settings under optimistic and pessimistic priors (* = assuming the resources are reallocated to an effective drug)	78
A.1	Variables and parameters used in the simulation model	99
B.1	Optimal first-period solution (N_1^*) for the two-period adaptive trial (the parameters used are specified in Section 5.1)	107

Chapter 1

Introduction

1.1 Background and Motivation

1.1.1 The Importance of Pharmaceutical R&D

The output of pharmaceutical research and development (R&D) has the potential to impact the quality of human life. The detrimental effects of severe diseases such as cancer, hepatitis C, and AIDS have been significantly mitigated thanks to new medications and therapies. As reported by the Pharmaceutical Research and Manufacturers of America (PhRMA) [80], cancer fatality rates have declined 23% since their peak in the 1990s [66], with approximately 83% of survival gains attributable to new treatments and new medications [97]. Unlike a few years ago, a wide range of treatment options are now available to cure hepatitis C, with cure rates of more than 90% over a course of treatment as short as eight weeks [81]. The battle against HIV/AIDS has also been successful, with a declining fatality rate of 87% [68] and 862,000 premature deaths avoided in the United States (US) alone since the introduction of the highly active antiretroviral treatment (HAART) [56].

Drug R&D also represents a significant component of the US and the worldwide economy. Ding et al. [32] state that pharmaceutical R&D holds a share of 19% of all business spending on R&D worldwide, and that the US alone finances about 36% of the global expenses in pharmaceutical R&D. Despite these high expenditures, the rate of success in drug discovery remains steadily low. The development of new drugs is complex, lengthy, and costly, and is surrounded by high

levels of uncertainty. In addition to the inherent uncertainty in testing outcomes, the highly regulated environment in which the new drug development process takes place amplifies the firms' risk of incurring unrecoverable expenses, as the ultimate decision to commercialize the drug in the US depends entirely on the Food and Drug Administration (FDA). These and other elements set drug development apart from the innovation processes in other technology-intensive industries.

DiMasi et al. [29] have recently estimated the time between the start of clinical testing and submission of a New Drug Application (NDA) to the FDA to be around 80.8 months, or slightly short of 7 years, with clinical trials accounting for the majority of this time. This same report indicates that the estimated out-of-pocket cost of a new drug receiving FDA approval was \$1.39 billion in 2013, representing a 166% increase over the cost of drugs approved in the 1990s [31]. In addition, the capitalized cost estimate of an approved drug, which includes the research and development costs of those drugs that failed to get approval, rose to \$2.56 billion, 145% higher than the capitalized cost of drugs approved in the 1990s. DiMasi et al. [29] affirm that the overall change in the risk profile for new drug development accounted directly for a 47% increase in costs. In the early 2000s, the overall likelihood that a drug that is clinically tested will eventually be approved was estimated to be 21.50%. This probability has dropped down to 11.83% in 2016, which translates into a higher risk of developing drugs that will not receive regulatory approval, and whose development costs may not be recovered by the trial sponsor.

1.1.2 Sequential Adaptive Clinical Trials with Binary Response

This dissertation focuses on Phase 3 group sequential adaptive trials with binary response. In binary response trials, a patient's response to a given therapy can take one of two possible values; these two possible outcomes are generically referred to as *success* and *failure* [109]. For instance, acute bleeding may be controlled within a specific time frame or not, the patient may survive after receiving a certain therapy or not, or blood pressure may fall to a specified level or not.

Clinical trials typically consist of three phases, with Phase 3 being the most resource-intensive phase. The traditional approach to conducting a Phase 3 trial consists of determining an *a priori* fixed sample size, based on a desired statistical significance level and power, as well as the drug's characteristics and the treatment advantage to be observed. Next, data are collected on the full

sample of patients, and the drug is deemed successful or futile.

On the other hand, in a group sequential trial it is common to compare a statistic, summarizing the difference between a control (i.e., a standard therapy or a placebo) and the experimental treatment, against specified critical values at a series of interim analysis points (decision epochs) [99]. These critical values represent a stopping rule for the trial, and are often set based on statistical significance and power, as well as clinical factors [5]. If the statistic's value surpasses the critical value, then the trial stops and the null hypothesis, of no difference between the two therapies, is rejected; if the statistic remains below the critical value, the evidence is insufficient to arrive at a conclusion, and the trial continues until the next interim analysis point [99]. Some of the questions that need to be answered in this new setting include: how many patients to recruit at each interim analysis point, and whether to use conservative or aggressive investment strategies given the uncertainty surrounding the drug's efficacy and the opportunity to modify these strategies based on the interim results.

1.2 Dissertation Outline

The remainder of this dissertation contains three manuscripts. Fig. 1.1 provides an overview of the context and the methodologies used in these, detailed in Chapters 2 - 4. In Chapter 2, we analyze whether the new flexibility offered by sequential adaptive trials, in the form of early termination, translates into a higher risk of drug misclassification (i.e., false positives and false negatives). We build a model that represents a Phase 3 parallel group (i.e., two-armed) sequential adaptive trial with binary response realistically, where the efficacy of a candidate drug is compared against the efficacy of a control. We study the performance of what is known as the “triangular test,” the most widely implemented method to determine whether a trial should stop or continue at an interim analysis point [91, 99, 110], by quantifying the drug misclassification risk. In particular, the triangular test involves an upper and a lower boundary, which inform the decision maker on when to stop the trial early for benefit/effectiveness or futility/ineffectiveness, or to continue testing due to inconclusive outcomes.

Although we analyze both forms of drug misclassification, we mainly focus on the false negative drug classification risk, or the likelihood of incorrectly classifying an effective drug as futile by

Chapter 2	Chapter 3	Chapter 4
Current Practices	Optimal Resource Allocation	Optimal, Adaptive Testing Schedules
	Setting	
Parallel Group Trial	Crossover Trial	Parallel Group Trial
Triangular Test Upper Boundary	Exogenous Upper Boundary	Triangular Test Upper Boundary
Triangular Test Lower Boundary	Optimal Lower Boundary	Optimal Lower Boundary
	Methodologies	
Simulation	Stochastic Dynamic Programming	Stochastic Dynamic Programming
	Bayesian Learning	Bayesian Learning
	Performance Measures	
Drug Misclassification	Drug Misclassification	Drug Misclassification
	Expected Profit	Expected Profit
	Time-to-market	Time-to-market

Figure 1.1: Overview of Chapters 2 - 4

terminating its development early. The false negative drug misclassification risk has been under-examined in the literature, which has mainly focused on the false positive misclassification risk. False negatives are also highly undesirable for both the trial sponsor and the patients, because they may lead to unrecoverable expenses for the sponsor, and unfulfilled needs for the patients. We examine the effects of imperfect information and the interim analysis points that lead to drug misclassification through an extensive Monte Carlo-style sensitivity analysis. Contrary to the literature's focus on false positives, our results suggest that false negatives can be more likely. Based on our analysis, we provide important insights for investigators and other decision makers on the causes of false negatives.

In Chapter 3, we shift our focus to the study of the optimal resource allocation decision in crossover (i.e., single group) Phase 3 group sequential adaptive trials, and quantify the impact of interim analyses on the new drug's misclassification risk and time-to-market, as well as the firm's profit. Towards this end, we build a stochastic dynamic programming model that incorporates Bayesian updates on the drug's efficacy. The resource allocation decision in this setting is characterized by endogenous uncertainty, because the realization of the uncertainty (i.e., the candidate drug's efficacy) depends on the investment decision: if the firm does not allocate any resources to

the trial of a particular drug, the firm will not be able to obtain any information on the drug's true efficacy, similarly, the amount of information that the firm will be able to obtain on the drug's true efficacy depends on the amount of resources allocated to the drug's testing.

The resource allocation decision is also characterized by a trade-off between the incentive to establish that the drug is effective early on (exploitation), due to a time-decreasing market revenue, and the benefit from collecting some information on the drug's efficacy early on the trial (exploration), prior to committing a large budget. In this case, the trial termination decision is made based on an exogenously determined upper boundary and an optimal lower boundary. We analyze the structural properties of an optimal resource allocation strategy and perform a numerical study based on realistic data. We extend our focus beyond drug misclassification risk, and analyze the firm's expected profit and the drug's time-to-market as well. Our study generates key insights on the structure of an optimal resource allocation strategy and on the value of group sequential adaptive trials for new drug development.

In Chapter 4, we compare the performance of the proposed optimal resource allocation strategy with current sequential stopping criteria, which utilize the aforementioned triangular test. In particular, we quantify the firm's expected profit, drug misclassification risk and time-to-market under these two stopping mechanisms. Unlike the triangular test, our optimal resource allocation strategy provides the decision maker with an adaptive patient recruitment and testing schedule that can be adjusted based on financial constraints and newly acquired information on the drug's true efficacy. In this last manuscript, we offer insights on the impact of imperfect information on the candidate drug's success likelihood, and on the value of bias correction and learning. These insights can be useful to various stakeholders, such as policy makers and pharmaceutical firms. Finally, Chapter 5 provides a summary of research findings and offers some future research directions. To improve the presentation, some proofs and algorithmic details are relegated to the Appendix.

Chapter 2

Trial Termination and Drug Misclassification in Sequential Adaptive Clinical Trials

2.1 Introduction

Drug research and development (R&D) is critical for preserving human health and advancing life quality, but at the same time, involves a high-risk high-reward process. While a major driver of profit and growth in the pharmaceutical industry, drug R&D is an expensive and time-consuming endeavor with major uncertainties [32, 34]. In 2014, the R&D spending of United States-based pharmaceutical companies reached \$53.3 billion [80], with clinical trials representing at least 50% of the total R&D cost [32, 31, 39]. The increasing costs of Phase 2 and 3 trials, along with late-stage failures (i.e., drugs that go through the entire clinical trial process and ultimately fail either for safety issues or due to lack of efficacy), are the main components of R&D expenditures [77, 23, 73], and have led academicians and practitioners to find ways of reducing costs and investment risks during these phases.

Adaptive clinical trials emerged as a path to faster and more efficient drug R&D, promising benefits to both the trial's sponsor and the patients [4, 13, 10, 11]. These trials use different adaptive settings (e.g., stopping early for anticipated benefit or futility, dropping arms or doses, among

others) to enable researchers to efficiently learn about the drug’s efficacy [23, 11]. Further, the impact of adaptive trials has not been limited to the pharmaceutical industry, and has translated into changes in the regulatory process performed by the United States Food and Drug Administration (FDA). The FDA’s Critical Path Initiative suggests modifying clinical trial and program design to reduce attrition, and consequently, increase R&D productivity [106, 107].

Sequential adaptive trials promise major cost savings by allowing early termination of effective or ineffective drugs; however, they can also lead to a higher rate of drug misclassification compared to fixed duration (i.e., traditional) trials, which involve the testing and analysis of a sample of patients of fixed size. The decision to continue or terminate testing after the interim analysis depends on the feedback the decision makers receive from the drug testing process, in terms of the outcomes received from patients that are administered the candidate drug. Drug misclassification (i.e., terminating the development of an effective drug for futility, or terminating the development of an ineffective drug for benefit) is possible, and depends, among others, on the stopping rule, defined prior to the beginning of the trial. While past studies have mainly focused on false positives, limited analysis has been performed on both types of drug misclassification (i.e., false positives and false negatives). In addition, to the best of our knowledge, none of the existing studies utilize simulation models to analyze feedback accuracy and the effects of imperfect prior information within a sequential adaptive trial context.

This paper models the decision makers’ learning process from interim analysis results in a Phase 3 sequential adaptive clinical trial, and analyzes potential drug misclassification in the form of false positives and false negatives. Our model builds upon the experiential learning models proposed by Denrell and colleagues (e.g., Denrell and March [28], Denrell [26], Denrell and Le Mens [27]) in the behavioral decision-making literature. Denrell and March [28] argue that, when one wants to discover the expected value of a risky and novel option, initial bad experiences can inhibit learning and lead to underestimation of the risky option’s expected value; a phenomenon referred to as the “hot stove effect.” We postulate that sequential adaptive clinical trials can exhibit a similar pattern, which magnifies the likelihood of misclassifying effective drugs, and represents an aspect of adaptive clinical trials that is under-examined in the literature.

Our simulation model enables us to study the effects of imperfect information, and to obtain insights on the conditions that increase drug misclassification rate. Simulation models (e.g., Park

et al. [75]) have been used to study healthcare delivery systems and analyze complex decision-making processes, and have been presented as policy management tools. Simulation models (e.g., Larson et al. [58]) have been proposed as effective means to analyze non-linear phenomena emerging from the interdependencies of different elements within a system, and to obtain an intuition and insights able to inform policy making [36]. To the best of our knowledge, our research is the first to analyze drug misclassification resulting from learning and imperfect information, and the occurrence rate of false negatives within the context of sequential adaptive clinical trials.

2.1.1 Context: Phase 3 Clinical Trials

Clinical trials have three phases. In Phase 3, the candidate drug is tested in a large group of patients to establish safety and efficacy, determining tradeoffs between benefits and risks. The candidate drug is tested comparatively against current standard treatments or placebo options to establish efficacy. Phase 3 trials are both the costliest and the longest of all phases [32]. A group sequential design allows for early trial termination due to safety, futility/efficacy, or both, based on interim results [16]. This design presents decision makers with the opportunity to reduce capital commitment—an option that fixed duration trials do not offer. The decision to terminate a trial is very important and time sensitive from a managerial perspective, because it not only affects the current study, but might affect future trials in the same therapeutic area. Further, continuing a trial for too long delays the dissemination of important information and puts participants at unnecessary risk [99]. The decision to terminate or continue a sequential trial is made by the Data and Safety Monitoring Board (DSMB), which will evaluate the trial results in light of various factors, such as the drug’s characteristics, the treatment expected to be observed, and external information.

2.1.2 Objectives

The objectives of this research are the following: (1) to model and simulate the decision to continue or stop a Phase 3 group sequential adaptive trial, (2) to examine the hypothesis that false negatives can be disproportionately penalized in adaptive trials, leading to the hot stove effect, (3) to analyze the effects of imperfect information on drug misclassification rate, and (4) to examine the conditions that increase the drug misclassification rate of an effective drug at an interim anal-

ysis point. We model the statistical analysis (i.e., the quantitative evaluation of the trial’s results using pre-specified thresholds) performed by the DSMB, which translates into a trial termination or continuation decision, through a well-known and widely implemented methodology. Adaptive clinical trials may involve various adaptations besides interim analyses (e.g., changing the drug dosage during Phase 3, eliminating treatment arms, etc.) which we do not intend to analyze in this paper; our focus is on one aspect of group sequential adaptive trials: the decision to stop or continue at an interim analysis point, and its potential effects on drug misclassification.

2.2 Background

2.2.1 Group Sequential Adaptive Trials and Stopping Rules

To establish a new drug’s safety and efficacy, clinical trials are carried out in three phases: in Phases 1 and 2, the drug is tested on small groups of patients, while in Phase 3, statistically significant results need to be generated by testing large groups of patients. Once the safety and efficacy of a candidate drug are proven, a New Drug Application (NDA) can be submitted to regulatory agencies for approval. Upon approval, the drug can be commercialized. Traditionally, Phase 3 clinical trials involve a fixed sample size, which is calculated prior to the start of the trial based on the drug’s characteristics, the treatment advantage to be observed, and a specific statistical significance level and power. Once data on the full patient sample are collected and analyzed, the candidate drug is classified as successful (i.e., effective) or futile (i.e., ineffective) based on a comparison between its performance and the performance of a control therapy. Subsequent approval will be granted if the regulatory agency considers that the evidence provided to support the new drug’s safety and efficacy is sufficient.

Designs that include interim analyses and stopping rules are called “sequential” or “group sequential” [111]. In a sequential test, a statistic or primary measure of efficacy—representing the difference between the experimental treatment and a control is compared against pre-specified critical values (i.e., thresholds) at different points throughout the trial’s course [99]. The critical values make a boundary or stopping rule, whose precise form is determined based on statistical power and significance level, and desired treatment advantage with respect to the primary measure

of efficacy, as well as financial and clinical factors [5]. If the absolute value of the test statistic is higher than the critical value, an appropriate conclusion is drawn (e.g., rejection of the hypothesis of no difference between treatments). If the statistic remains within the test boundary, the evidence is insufficient to come to a conclusion and further analysis is performed at the next interim point [99].

Different forms of stopping rules consider if the experimental therapy is: (1) clearly better than the control (i.e., stopping for benefit), (2) clearly worse than the control, and (3) clearly not going to be proven to be better than the control (i.e., stopping for futility) [111]. Some designs (e.g., O’Brien and Fleming [71], Gordon Lan et al. [37], Goss et al. [38]) fulfill the first and second objectives: stop the trial if the candidate drug is clearly better or clearly worse than the control. Other designs (e.g., the triangular test [20]), satisfy the first and third objectives: stop once it is evident that the experimental therapy is superior to the control, or that the candidate drug will not show to be better than the control. In our model, we use the triangular test as presented in Whitehead [109]. The triangular test is the best known and widely implemented design among all the sequential methodologies based on straight line boundaries [91, 99, 110].

In our model, we consider a single interim analysis point. Researchers suggest that a study with just two interim analyses considerably reduces the average number of patients exposed to an inferior treatment when compared to a fixed duration trial [86]. Further, the literature indicates that although it is possible to take an interim look after every patient, conducting six or more interim analyses is unlikely to generate statistical, practical, or ethical benefits, unless an extremely large difference between the experimental treatment and the control is expected [64, 85]. In fact, the more often one analyzes the accumulating data, the greater the chance of eventually and wrongly detecting a difference between the candidate drug and the control, thus, drawing incorrect conclusions from the trial [99].

2.2.2 Critiques Regarding Early Termination

Stopping rules based on evidence of higher efficacy—to support early termination for benefit—are the most controversial. Some researchers consider that no amount of evidence is excessive to demonstrate efficacy, and that “no stone has to be left unturned” when identifying potential

adverse effects. Thus, trials that stop at the interim analysis for benefit are oftentimes said to have “stopped prematurely” [111]. Zannad et al. [113] state that these trials involve a risk of overlooking significant effects on secondary end points (e.g., mortality), and that it is necessary to carefully consider the potential for knowledge to be lost (or gained) if the trial is stopped early, as well as the direction, magnitude, and strength of benefit.

Various researchers show concerns about the risks associated with stopping a clinical trial for benefit, but do not analyze trials that are stopped for futility to the same extent. Specifically, researchers state that trials can be erroneously terminated for benefit because early stopping may lead to an overestimation of the drug’s effects [5, 113], caused by a phenomenon denominated “random high” [5, 84] or “regression to the truth” [83, 87]. However, clinical trial experts dissent about how stopping guidelines can overcome this bias [53, 33]. Bassler et al. [5] affirm that investigators have an incentive to terminate a trial early for benefit because of honor, publicity, and attention for the trial’s sponsor, making this decision “ethically questionable.”

Stopping for futility represents the major conflict of interest between trial investigators and sponsors, with the former likely motivated to continue collecting valuable data (provided no safety concerns are present), and the latter motivated to stop because of financial reasons [111]. However, stopping for futility has not been criticized as strongly and has not received as much attention as stopping for benefit. If not performed carefully, early trial termination for futility at the interim analyses does not inflate type I error, but may significantly decrease the power, which is equivalent to inflate type II (i.e., false negative) error [42].

Early termination due to futility reasons can lead to positive outcomes, such as preventing future patients from receiving an ineffective treatment, reducing trial duration and costs, and expediting patients’ access to new therapies. However, trials showing modest treatment effects could be mistakenly stopped early in cases where there appears to be no benefit at an interim analysis, thus missing an important clinical effect Jitlal et al. [48]. For instance, the trials ACT I [100] and UKHAN_2 [98], testing cancer drugs, showed unsatisfactory results at the interim analysis but clinically important effects once the full sample size was analyzed.

2.2.3 False Negatives and the Hot Stove Effect

The hot stove effect has been defined by Denrell and March [28] as a phenomenon leading to “a bias against new alternatives that require practice and against alternatives involving risk.” In their work, they present the hot stove effect as a behavioral consequence of experiential learning, which is a form of adaptation based on the concept that organizations and the people in them change their actions based on an evaluation of their experiences [22, 43]. Since the risky alternative is mistakenly believed to be unsuccessful, it is avoided, and belief-correcting information is not obtained.

In contrast, a type II error (false negative) takes place if the researcher fails to reject the null hypothesis—the hypothesis of no treatment difference between the candidate drug and the control—that is actually false in the population. This happens because the results in the sample do not reflect reality in the population, and random error leads to a mistaken conclusion. Type I and type II errors can never be fully avoided; however, the researcher can reduce their likelihood.

The hot stove effect is a behavioral phenomenon that manifests itself in the same way as a false negative, but goes beyond a statistical error caused by randomness. Within a sequential adaptive clinical trial context, the hot stove effect consists in a magnified likelihood of terminating the development of an effective drug at the interim analysis point, on the basis of unsatisfactory results. A type II error or false negative also consists in the erroneous termination of an effective drug’s trial, because the null hypothesis of no difference between treatments is falsely retained. Whereas the false negative rate is fully characterized by the trial’s statistical power, the hot stove effect is a magnification of this rate, caused by the behavioral factors explained earlier.

2.3 Methods

2.3.1 Model Overview

Our unit of analysis is a candidate drug that will undergo Phase 3 testing. Similarly to other models in literature (e.g., Patel et al. [76] and the references therein), we assume that no safety concerns exist, and focus on efficacy. The logic of the model is as follows. We assume the trial has a predetermined *full sample size* based on statistical power and significance. Fig. 1 provides an

overview of the conceptual model: the trial starts, a patient is recruited and assigned to a treatment group during the *patient randomization* process, and receives the control therapy or the candidate drug in the *drug testing process*. The outcome of the *drug testing process* is subject to a *statistical analysis*, which later defines the *interim analysis result*. This result will trigger the *recruitment decision* to continue or terminate of the trial after the interim analysis point, and the subsequent *patient randomization* in case the trial continues and a new patient is recruited.

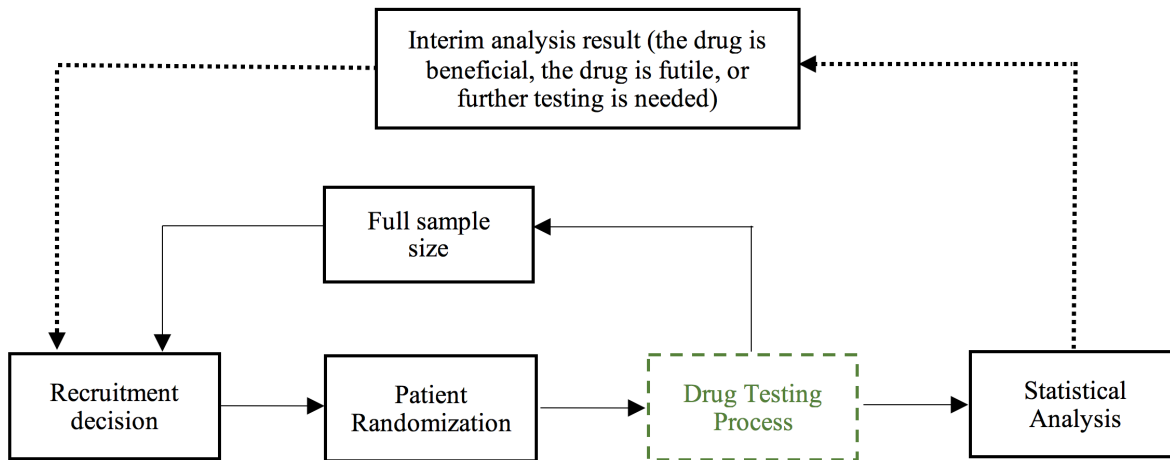


Figure 2.1: Structure of the patient recruitment decision-making model for a single drug (Phase 3)

In the *patient randomization*, we randomly assign each recruited patient to a control group (where the patient is treated with a placebo or standard therapy), or to an experimental group (where the patient is treated with the candidate drug). The *drug testing process* represents a clinical trial with binary response, meaning that every patient will be treated either successfully or unsuccessfully depending on the drug's true probability of efficacy/success, or the control's true probability of efficacy/success, based on the group that the patient was assigned to (see Appendix A for a graphical representation). The performance of the control and the candidate drug are evaluated in a *statistical analysis*, which follows the methodology of the triangular test as presented in Whitehead [109].

The *patient randomization* and subsequent *drug testing process* are triggered by the *recruitment decision*. This decision will be to either stop the trial at an exogenously determined interim analysis point (t), or to continue the trial until the *full sample size* is analyzed. This *full sample size* and the interim analysis point are defined in terms of a specific number of patients. Consistent with

the literature that makes recommendations on the number of interim analyses (see Section 2.2), we use a single interim analysis point. As recommended by the United States Food and Drug Administration [105], we assume the interim analysis point is fixed before the trial starts, and cannot be changed once patients start being recruited for the trial.

The *recruitment decision* depends on the *interim analysis result* (Fig. 1). The *interim analysis result* will be one of the following: (1) the drug is clearly beneficial, (2) the drug is clearly futile, or (3) the testing is inconclusive. Results (1) and (2) will lead to a *recruitment decision* to stop the trial, whereas result (3) leads to a trial continuation decision, as more patients will need to be recruited. The decision to continue or stop the trial cannot be changed at a later point in time. We describe the trial stopping mechanism in the next section.

2.3.2 Model Formulation

We explain the model’s formulation by dividing it into three parts: *patient randomization*, *drug testing process*, and *statistical analysis*.

2.3.2.1 Patient Randomization

Once a patient is recruited, she will be assigned to either the control group, or to the experimental group. We seek to obtain an even assignment of patients, such that the number of patients in the control group and in the experimental group are equal.

2.3.2.2 Drug Testing Process

Once a patient is assigned to a group, she will be tested, and her treatment result deemed successful or unsuccessful shortly after the therapy is administered. The patient’s health status after the treatment will change either due to the effect of the drug, or due to the effect of the control. We model the treatment result observed in a patient in group g (TR_g), where $g \in \{E(\text{experimental}), C(\text{control})\}$, as a binary variable whose value (1: successful outcome, 0: unsuccessful outcome) is a function of a uniformly distributed indicator (J), and the therapy’s true probability of efficacy (P_g), which

is unknown to the decision maker:

$$\begin{aligned} TR_g &= 1 && \text{if } J > (1 - P_g), \\ TR_g &= 0 && \text{otherwise, where } J \sim U[0, 1]. \end{aligned}$$

2.3.2.3 Statistical analysis

In our model, the statistical analysis follows the methodology of the triangular test. This method consists of an upper and a lower boundary. If the upper boundary is exceeded, the drug is deemed beneficial (i.e., better than the control). If the lower boundary is neither matched nor exceeded, the drug is either deemed futile (i.e., not better than the control), or worse than the control. We determine these boundaries, as well as the full sample size, using the guidelines provided in Whitehead [109].

We assume the patient response in this trial is binary, and available soon after the therapy has been administered. The decision maker expects the control and experimental therapies to have true probabilities of efficacy (i.e., success rates) p_C and p_E equal to 0.5 and 0.7, respectively. In addition, we choose a significance level of $\alpha = 0.05$, and power $1 - \beta = 0.90$, which are widely accepted and commonly used in the pharmaceutical industry. We use the log-odds ratio as the parameter to represent treatment difference, and utilize it to compute the reference improvement between treatments θ_R :

$$\theta_R = \log \left\{ \frac{p_E(1 - p_C)}{p_C(1 - p_E)} \right\}. \quad (2.1)$$

We calculate the Z test statistic, measuring treatment difference, and the V statistic, measuring accrued information, as follows:

$$Z = \frac{(N_C S_E - N_E S_C)}{N}, \quad (2.2)$$

$$V = \frac{N_E N_C S F}{N^3}, \quad (2.3)$$

where N_C and N_E are the number of patients treated in the control and experimental groups, respectively, S_C and S_E are the number of patients treated successfully in the control and experimental groups, $N = N_E + N_C$ is the total number of patients treated in both groups, and $S = S_E + S_C$ and $F = N - S$ are the total number of patients treated successfully and unsuccessfully in both groups.

Incorporating our parameters to Eq. 2.1, we obtain a reference improvement level, θ_R , equal to 0.847. A fixed sample trial satisfying the significance level and power requirements chosen earlier requires an amount of information $\theta_R^2 \hat{V} = 10.51$ [109], hence, $\hat{V} = 14.64$. Given that no trial data are available at the planning stage, and the anticipated overall success rate $\bar{p} = \frac{PE + PC}{2}$, we use Eq. 2.3, the value of V we just calculated, and the following assumptions to obtain the *full sample size*: for large samples, $S \approx N\bar{p}$ and $F \approx N(1 - \bar{p})$; since the patient assignment between the control and experimental groups is even, $N_E = N_C = N/2$, thus:

$$\hat{V} = \frac{N\bar{p}(1 - \bar{p})}{4}. \quad (2.4)$$

Substituting the corresponding values of \hat{V} and \bar{p} in Eq. 2.4, the resulting *full sample size* is equal to 244 patients. Since a triangular test will be adopted, we now need to determine the corresponding stopping boundaries.

At the interim analysis point, the Z statistic will be computed using Eq. 2.2, and will be compared against an upper threshold or threshold for benefit (BT), and a lower threshold or threshold for futility (FT). These thresholds are functions of the V statistic, and will be determined based on the chosen significance level and power requirements using the following expressions Whitehead [109]:

$$FT = -5.82 + 0.772V + 0.583\sqrt{V} \quad (2.5)$$

$$BT = 5.82 + 0.257V - 0.583\sqrt{V} \quad (2.6)$$

The outcome of the statistical analysis is represented by the *interim analysis result* (IR_t), which is defined as a binary variable. The value of IR_t is determined as follows:

$$IR_t = 1 \quad \text{if } Z \geq BT \text{ or } Z \leq FT,$$

$$IR_t = 0 \quad \text{otherwise,}$$

where $IR_t = 1$ determines a trial termination decision, and $IR_t = 0$ determines a trial continuation decision. If the Z statistic is larger than or equal to the threshold for benefit (BT), the *recruitment decision* will be to stop the trial for benefit. If Z is lower than or equal to the threshold for futility (FT), the *recruitment decision* will be to stop the trial for futility. If Z falls between the two thresholds, the analysis is considered inconclusive and the *recruitment decision* will be to continue the trial until the *full sample size* is analyzed. Appendix A provides a comprehensive

description of the model formulation.

In the next section, we focus our analysis on drug misclassification, and evaluate the model's sensitivity to imperfect information in the form of optimistic and pessimistic prior expectations on the drug's probability of efficacy. For an effective drug, the probability of misclassification can be expressed as $\Pr(Z \leq FT)$. Considering an even assignment between groups, from Eqs. 2.2 and 2.3 we know that the statistics Z and V are both functions of the random variables S_E and S_C :

$$Z = \frac{S_E - S_C}{2},$$

$$V = \frac{(S_E + S_C)(N - S_E - S_C)}{4N},$$

where $S_E \sim \text{Binomial}(N/2, p_E)$, and $S_C \sim \text{Binomial}(N/2, p_C)$. However, deriving $\Pr(Z \leq BT)$ analytically becomes complex due to the structure of the futility threshold (Eq. 2.5), and because the V statistic needed to determine this threshold involves the product of the random variables S_E and S_C . Similar complexity is present in the case of misclassification of an ineffective drug. Because of the complexity of obtaining these probabilities analytically, we utilize a simulation model to examine the effects of imperfect information on drug misclassification rate.

2.4 Results

2.4.1 An Illustrative Case

First, we use the parameter values specified in Table 2.1 to simulate the model for two major scenarios: an effective drug (S1) with a true probability of efficacy larger than the control's, and an ineffective drug (S2) with a true probability of efficacy equal to the control's. These true probabilities are assumed to be unknown to the decision makers, who use their prior expectations to plan the trial and are supposed to learn about the drug's performance by observing the test results. Table 2.1 summarizes the results obtained in these two scenarios.

Table 2.1 presents an illustrative case which we use to evaluate the model's response under different conditions, and quantify the drug misclassification rate for: (1) an effective drug, and (2) a futile drug with an efficacy equivalent to the control's. The interim analysis point is 122 patients, or 50% of the full sample size. The first drug is more effective than the control, thus, either the trial

Table 2.1: Model parameters and results for 100 simulations

	Candidate drug's true probability of efficacy	Full sample size (patients)	Interim analysis point (patients)	Correct classification at interim analysis point	Simulation results (% of drugs classified correctly)
Scenario 1: Effective drug	0.7	244	122	Beneficial (and stop trial)	94%
Scenario 2: Ineffective drug	0.5	244	122	Futile (and stop trial)	100%

Note: In both scenarios, the control's true probability of efficacy is assumed to be equal to 0.5.

should be stopped for benefit at the interim analysis point, or the drug should be evaluated until the full sample size is analyzed. However, we observe the trial stops for futility in 6% of the simulation runs (Fig. 2.2). In these runs, the recruitment stops as soon as the trial reaches the interim analysis point (the vertical dashed line in Fig. 2.2), hence, the ratio of patients treated successfully with the drug does not reflect the drug's true efficacy. Since the drug's true probability of efficacy is larger than the control's, these runs represent false negatives.

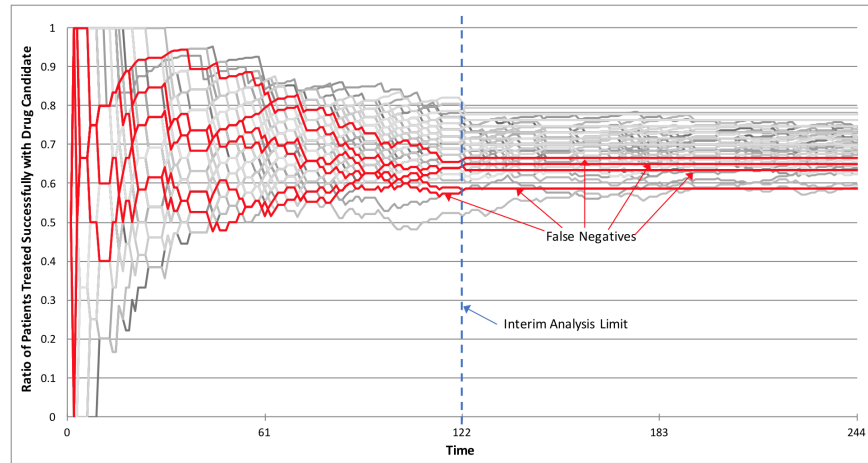


Figure 2.2: False negatives (Hot stove effect) shown in red

The second drug is just as effective as the control, therefore, does not represent any added benefit to the patients and should be classified as futile. Ideally, the trial should be stopped for futility at the interim analysis point in order to allow decision makers to reallocate the resources to another drug, and to stop treating the patients with the candidate drug. Alternatively, though not ideally, the trial could continue until the full sample size is analyzed, and the drug is classified as futile. As shown in Fig. 2.3, the trial does not stop for benefit at the interim analysis point on any of these initial 100 simulation runs, and consequently, the drug is not misclassified. In some cases, the ratio of patients treated successfully with the candidate drug is larger than the drug's

(and the control's) true probability of efficacy, however, the Z statistic does not exceed nor match the threshold for benefit. Hence, the trial continues until the full sample size is collected.

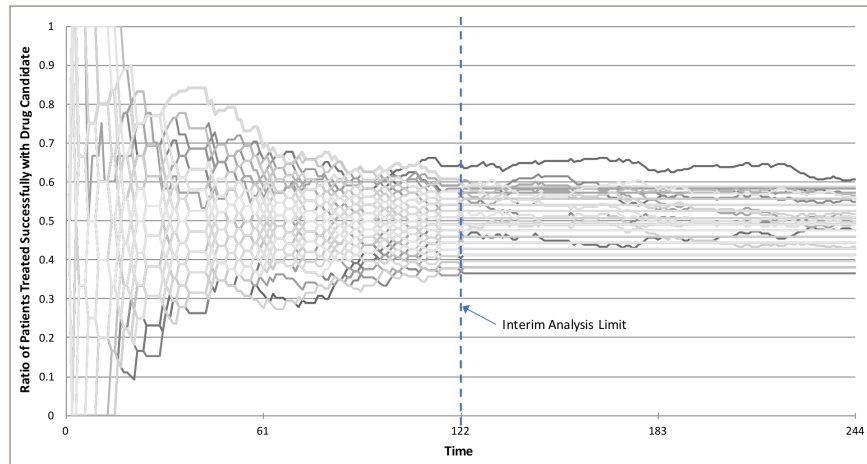


Figure 2.3: Absence of false positives

2.4.2 Drug Misclassification Under a Wider Range of Conditions

We conduct a series of sensitivity analyses on the model's parameters in order to study the effects of imperfect information (i.e., optimistic and pessimistic prior expectations on the drug's efficacy), and the conditions under which the drug misclassification rate increases. We use the parameter values from Table 2.1 and vary the drug's true probability of efficacy/success and the interim analysis point to evaluate misclassification rates. Our analysis shows that false negatives can emerge even when more than 50% of the full sample size has been tested.

2.4.2.1 False Positives vs. False Negatives with a Constant Interim Analysis Limit

We consider a control therapy with a true success probability of 0.50, and analyze the misclassification rate for 1,000 candidate drugs with true probabilities of success ranging between 0.10 and 0.50 (i.e., futile or ineffective), and between 0.51 and 0.90 (i.e., beneficial or effective). In this analysis, the interim analysis point is kept constant and unchanged at 122 patients, or 50% of the full sample size. Our results confirm our postulate that false negatives are indeed disproportionately penalized. Fig. 2.4 shows the percentage of drugs that were classified as beneficial and the percentage of drugs classified as futile at the interim analysis point. The drugs with true success

probabilities larger than 0.5 (to the right of the vertical line in Fig. 2.4) should be classified as beneficial, and drugs with true success probabilities lower than or equal to 0.5, classified as futile.

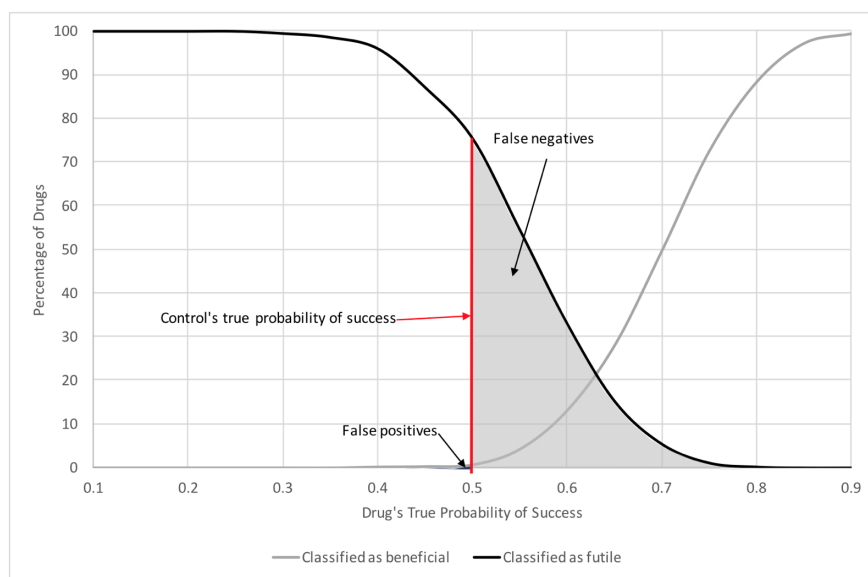


Figure 2.4: Drug misclassification rates at the interim analysis point for beneficial and futile drugs

As shown in Fig. 2.4, the occurrence rate of false negatives (the gray area under the “classified as futile” curve) is significantly larger than the occurrence rate of false positives (the small dark area under the “classified as beneficial” curve). False negatives can arise with effective drugs that match, exceed, or fall below the efficacy expected initially (i.e., 0.7), unlike false positives, which only emerge with drugs with efficacies similar to or slightly lower than the control’s. Although drugs that are much more effective than the control are rarely misclassified, drugs with moderately beneficial effects (i.e., with success probabilities between 0.55 and 0.70) are misclassified as futile in as many as 54% of the simulation runs.

2.4.2.2 Misclassification of an Effective Drug with Different Interim Analysis Limits

The interim analysis limit must be defined before patient recruitment starts. In this section, we vary the interim analysis limit and examine drug misclassification for an effective drug and a control therapy with true success probabilities equal to 0.7 and 0.5, respectively. Fig. 2.5 shows the percentage of drugs that are classified as either beneficial or futile, and the expected total number of patients that will be tested during the trial. We obtain these values by adding up the expected number of patients tested up to the interim analysis, and after the interim analysis. We calculate

the first component by multiplying the interim analysis limit by the percentage of runs where the trial stopped at the interim analysis (either for benefit or for futility), and the second component by multiplying the full sample size by the percentage of runs where the trial continued after the interim analysis.




Figure 2.5: Correct and incorrect classification rates, and expected total number of patients for a beneficial drug under perfect information and with different interim analysis limits

Since we assume the drug is effective, all runs where the drug is classified as futile represent false negatives. Our analysis shows that false negatives can emerge with low, moderate, and high interim analysis limits (Fig. 2.5), and even when the full sample size has been tested. These results do not support the common belief that large samples automatically reduce drug misclassification risk. In fact, we find that the rate of false negatives increases with the interim analysis limit, because the stopping boundaries or thresholds are functions thereof (i.e., functions of the V statistic, measuring accrued information). Moreover, if the chosen interim analysis limit is very low, the sponsor will end up collecting the full sample size. The lowest expected number of tested patients corresponds to an interim analysis limit of approximately 134 patients, or 55% of the full sample size. With this interim analysis limit, the trial stops early for benefit in approximately 55% of the simulation runs.

2.4.2.3 False Positives and False Negatives with Different Interim Analysis Limits

We now expand our analysis and examine drug misclassification with different sample sizes (i.e., interim analysis limits) when testing both effective and ineffective drugs. We simulate trials for drugs with success probabilities ranging from 0.10 to 0.90, and compute the percentage of simulation runs where the trial stops early for benefit, and the percentage of runs where the trial stops early for futility.

Across all interim analysis limits, we find that no more than 2.3% of the ineffective drugs tested are mistakenly classified as beneficial, with a larger proportion of false positives emerging with interim analysis limits equivalent to at least 55% of the sample size (Fig. 2.6a). On the other hand, the proportion of false negatives is significantly larger across the entire range of interim analysis limits studied. As shown in Fig. 2.6b, the percentage of effective drugs that are misclassified as futile—as a result of early trial termination for futility—can be more than 10 times larger than the percentage of ineffective drugs misclassified as beneficial using the same sample size at the interim analysis point.

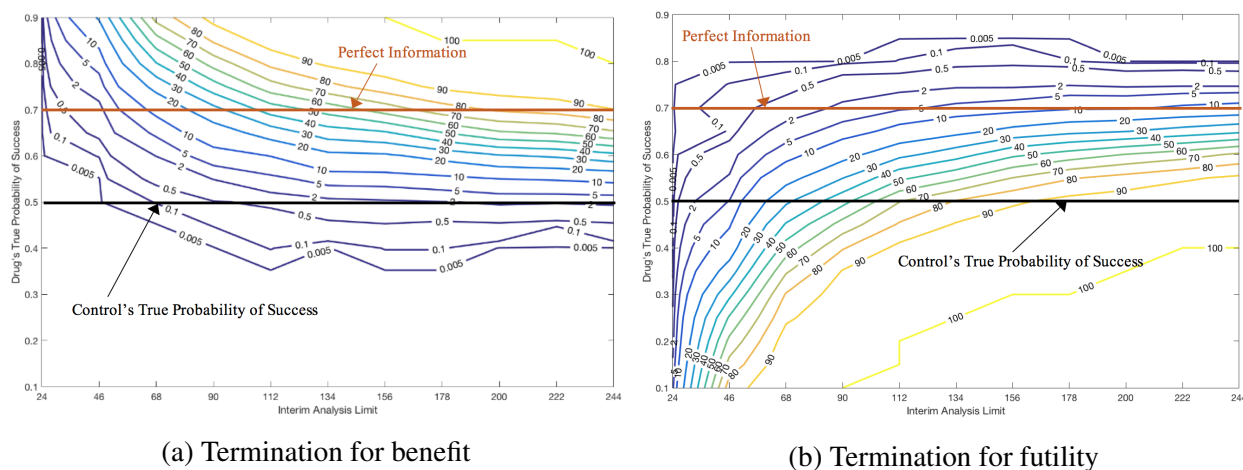


Figure 2.6: Percentage of drugs with trials terminated early

Consistently with the results described earlier, larger interim analysis limits may involve larger drug misclassification rates. Our results also suggest that the overall likelihood of false negatives is significantly larger than the overall likelihood of false positives. These new findings provide us with a better understanding of the phenomenon identified earlier, where we used a fixed interim analysis limit of 122 patients.

In Section 2.4.2.1, we found that false negatives emerged mainly with drugs that were moderately better (i.e., with efficacies 0.05 to 0.20 higher) than the control, whereas false positives emerged only with drugs that were slightly worse (i.e., with efficacies up to 0.02 lower) than the control. Our sensitivity analysis supports and expands these findings, showing that drugs that are not as effective as expected, but able to generate moderately beneficial effects, can be misclassified as futile in as many as 92% of the simulation runs (for a drug with a true success probability of 0.55). Figs. 2.6a and 2.6b also confirm that false negatives are more likely than false positives, regardless of the sample size analyzed at the interim analysis point.

2.5 Discussion

Our research explores the dynamics of early termination of a Phase 3 sequential adaptive clinical trial that uses the stopping boundaries of a triangular test. We use a simulation model to capture the learning process of decision makers, such as the DSMB members, which leads to a trial termination or continuation decision. We analyze the drug misclassification rate, in the form of false positives and false negatives, resulting from this decision. Our model's structure allows us to get an accurate representation of the patient randomization and testing processes, and the early trial termination of both effective and ineffective drugs.

We identify the conditions under which drug misclassification occurs by performing a Monte Carlo-style sensitivity analysis on the sample size tested at the interim analysis point (i.e., the interim analysis limit), and examine the effects of imperfect information by analyzing the misclassification of drugs with various efficacy levels. Our findings suggest the presence of the hot stove effect, as the observed rate of false negatives is higher than the type II error rate targeted with the statistical power chosen when the trial was planned.

2.5.1 Insights and Implications

Our findings suggest that, contrary to most critiques found in literature, group sequential adaptive clinical trials are more susceptible to false negatives than to false positives (Section 2.4.1). Using the triangular test methodology, false negatives can emerge even when testing drugs with high efficacies, which should be able to show significant treatment difference when compared to the

control. The rate of false negatives is more significant when the decision makers' prior expectation on the drug's efficacy is slightly optimistic (i.e., when the candidate drug's efficacy is slightly lower than expected), because the trial is designed to detect a larger treatment difference than the one observed. For instance, an effective drug with a true success probability of 0.65—only 5% lower than expected—was misclassified as ineffective in over 20% of the simulation runs when using interim analysis limits equivalent to at least 50% of the full sample size (Fig. 2.6b).

Although the probability of having a false negative should be directly related to the statistical power chosen ($1 - \beta$), we find that small negative deviations from the expected treatment difference involve much larger false negative rates than the targeted rate β . As shown in Figs. 2.4 and 2.6b, moderately beneficial drugs are overly penalized. This result suggests the presence of the hot stove effect, because the decision makers opt to stop the trial early based on preliminary results that are poorer than expected. In other words, *decision makers are reticent to continue testing the drug if the results at the interim analysis point do not satisfy the treatment difference that they were expecting to observe.*

We also find that, when the sample size analyzed at the interim analysis point is too small, the evidence is likely insufficient to make a conclusion about the drug's efficacy, and the trial stops early for a very small number of drugs (Fig. 2.5). Therefore, the financial savings from adopting a sequential adaptive design are minimal, because most of the time, the full sample size needs to be analyzed in order to classify the drug as successful or futile.

Interestingly, our results suggest that large interim analysis limits (i.e., equivalent to at least 50% of the full sample size) involve higher drug misclassification rates (Figs. 2.5, 2.6a, and 2.6b) because the benefit and futility thresholds increase as the interim analysis limit increases. In the specific case of perfect information (i.e., when the drug's true efficacy matches the decision makers' prior expectation, Fig. 2.5), larger sample sizes lead to higher correct early classification rates, but do not involve any reduction in drug misclassification. The pharmaceutical industry is highly competitive, and revenue from the commercialization of a new drug decreases over time because of patent life reduction and the effects of competitors' actions. Interim analyses on large groups of patients may delay the submission of a drug approval application (e.g. NDA) and, due to the time-decreasing nature of the revenue, may translate into lower profits, while having either a negative or no significant effect on drug misclassification risk.

Under perfect information (Fig. 2.5), the interim analysis of sample sizes larger than or equal to 64% of the full sample size (156 patients) increases the chances of correct classification, but has no effect on misclassification risk reduction. In other words, although the drug may be submitted for approval a little earlier than under a fixed duration setting, the risk of misclassifying it as futile and incurring unrecoverable expenses remains constant or increases once the interim analysis limit exceeds a certain level. If the drug is classified correctly as effective, the additional expenditures of recruiting and analyzing more patients by the interim analysis point may involve a slightly higher revenue. On the other hand, if the drug is misclassified as futile, the entire patient recruitment and analysis expenditures turn into sunk costs, because the drug is not submitted for regulatory approval, and the revenue is never realized. This type of misclassification is particularly important from the patients' perspective, because false negatives prevent them from being treated with novel and effective therapies that are erroneously dismissed during the testing phase.

False positives are also highly detrimental to the patients. However, our results suggest that this type of misclassification is rare compared to false negatives when the trial is planned expecting a beneficial treatment difference between the candidate drug and the control. If the decision makers' prior expectations are highly optimistic, the triangular test will lead to a correct classification of an ineffective drug in the vast majority of cases (Fig. 2.6a). In fact, we observe that a futile drug is early and correctly classified in over 50% of the runs, preventing patients from being treated with an ineffective therapy, and allowing the sponsor to re-allocate its resources. Drugs that are highly ineffective are very likely to be correctly classified as futile early on (i.e. at the interim analysis point), even with small sample sizes (Fig. 2.6b). Drugs with efficacies similar to the control are not terminated as promptly, however, they are classified correctly as futile in at least 70% of the runs with interim analysis limits equivalent to 40% of the full sample size or larger.

False negatives, on the other hand, are much more likely; the highest risk emerges when the decision makers' expectations are slightly optimistic. Our results show that the trial of a drug with true efficacy of 0.6, using an interim analysis limit equal to half of the full sample size, suffers from a false negative rate equal to approximately 0.30 (Fig. 2.6b), thrice the type II error sought (from the chosen power) with the trial's design, which aimed at detecting an efficacy of 0.7. For this and other effective drugs with efficacies lower than the initial expectations, the false negative rate increases with the interim analysis limit. These drugs should be able to show a beneficial treatment

difference when compared to the control, and should be tested beyond the interim analysis point in order to detect their effects. Similar to the ACT I and UKHAN_2 trials, these trials did not show the beneficial effects of the drug at the interim analysis point.

In summary, beneficial drugs with efficacies slightly lower than what the trial is designed to detect (i.e., with efficacies lower than expected) suffer from the highest erroneous trial termination rate and drug misclassification risk. Thus, whenever the characteristics of the drug and the disease it aims to treat permit (e.g., for rare diseases and/or populations with an urgent and important need for new therapies), investigators should consider designing the trial aggressively rather than conservatively, in order to detect a small but beneficial treatment difference. As shown in Figs. 2.6a and 2.6b, the classification of drugs with efficacies higher than expected will likely be correct, hence should not be affected by the adoption of an aggressive trial design. Decision makers should keep these considerations in mind and take into account the consequences of both false positives and false negatives at the time of determining whether or not to act on the stopping rule's verdict, as they may be suffering from a bias against drug candidates whose efficacy shows unsatisfactory at the interim analysis point.

2.5.2 Limitations and Directions for Future Research

In this paper, we model the statistical analysis component of the patient recruitment decision by using an upper and a lower threshold, which represent the stopping boundaries of a triangular test. There exist other group sequential designs with stopping mechanisms different than the one we consider in this paper. In addition to carrying out a quantitative analysis, which we include in our model, DSMB members also incorporate external information and their own experiences into the decision of stopping or continuing a trial. Although our model simplifies the actual recruitment decision-making process, the model's structure and our analysis thereof allow us to provide important insights regarding drug misclassification.

We do not model adverse events, representing an additional source of uncertainty, because a futility stopping rule (such as the one we used in our model) will turn a safety stopping rule unnecessary, as it will be clear that the candidate drug is not superior to the control well before it is proven that it is inferior [111]. However, the DSMB scrutinizes reports of serious adverse events,

milder adverse events, laboratory analysis, etc., to protect patients from continuation of a trial that has become evidently unsafe. Given that our research objectives focus on drug misclassification within the context of a group sequential trial, and on the implications of wrongly terminating the testing of effective drug candidates, adverse events are outside the scope of our work. Including the dynamics of these events on the drug testing process represents an opportunity for future analysis.

Our work can be built upon to analyze drug misclassification in trials with more than one interim analysis point, and trials following other stopping rules, such as the alpha-spending method [50, 57], and the stopping rules proposed by Pocock [85], O’Brien and Fleming [71], Haybittle [41], and Peto et al. [79], among others. Future research could shed light on the likelihood of false positives and false negatives when drugs are tested under these stopping rules, and complement existing work in the medical and biopharmaceutical fields. Another interesting area for future research is the study of the effects of interdependencies between different drugs on the patient recruitment decision. Pharmaceutical firms usually make decisions considering portfolios of multiple drugs, rather than individual projects. We believe feedback loop-based approaches can be used to examine whether these interdependencies reshape the patient recruitment decision, and to identify their effects on drug misclassification risk.

2.6 Conclusions

Group sequential adaptive clinical trials give decision makers the opportunity to stop testing a new drug early, due to benefit or futility. In this paper, we analyze the decision makers’ bias against effective drugs whose testing is terminated for futility, on the basis of interim results suggesting the drug is ineffective (i.e., the hot stove effect). We use a simulation model to capture the dynamics of the triangular test, the best known and widely implemented design among all the sequential methodologies based on straight line boundaries, and analyze the patient recruitment decision, which translates into trial continuation or termination. We first test our model with a drug of moderately beneficial effect, and show that it is incorrectly classified as futile in a significant percentage of simulation runs.

We conduct a series of Monte Carlo-style sensitivity analyses to identify the conditions that affect drug misclassification rate. Our results show that, under the triangular test’s stopping rules,

large sample sizes do not reduce the risk of erroneous drug classification because the benefit and futility thresholds increase with sample size. Interestingly, we find that false negatives are more likely to occur than false positives. This result is in clear contrast with the majority of critiques to early trial termination in the medical literature, which mainly focus on the implications of false positives.

Our sensitivity analysis results show that beneficial drugs with efficacies slightly lower than what the trial is designed to detect suffer from the highest erroneous trial termination and drug misclassification risk. We thus suggest investigators, whenever the characteristics of the drug and the nature of the targeted disease permit, to consider designing group sequential trials aggressively rather than conservatively, in order to detect small but beneficial treatment differences. Further, we suggest decision makers to keep in mind the potential impact of both false positives and false negatives when deciding whether or not to act on the stopping rule's verdict.

Chapter 3

Optimal Sequential Resource Allocation in Sequential Adaptive Clinical Trials

3.1 Introduction and Motivation

The research and development (R&D) process in the pharmaceutical industry is complex and costly, but provides the backbone for profit and growth [32, 34], and also affects the quality of human life due to the population’s constant need for more effective drugs and treatments. Clinical trials represent at least 50% of the total R&D cost [30, 32, 39]. The main components of these high expenditures include the high costs of Phase 2 and 3 trials as well as late-stage failures (i.e., drugs that go through the entire clinical trial process and fail either due to safety issues or due to lack of efficacy, increasing attrition and reducing productivity) [23, 73, 77], which have raised concern among academicians and practitioners.

Adaptive clinical trials promise not only substantial time and cost savings to the industry, but also faster availability of effective treatment options to patients [4, 10, 11, 13]. By utilizing a variety of adaptive settings (e.g., stopping early for anticipated benefit or futility reasons, dropping arms or doses, among others), these trials provide the decision maker with the means to efficiently learn about the drug and thus better manage the clinical trial process, through tools such as Bayesian methodologies [11, 73]. In 2012, the United States President’s Council of Advisors on Science and Technology (PCAST) released a report including recommendations to advance innovation in drug discovery, development, and evaluation. The PCAST Report states that incorpo-

rating new efficiencies into clinical trials is essential for achieving a higher innovation productivity [88]. The shift towards adaptive trials has already impacted the regulatory process performed by the United States Food and Drug Administration (FDA), which has “established a potential new pathway for accelerated approval of drugs.” [45] Specifically, the design of the I-SPY 2 trial, which tested the breast cancer drugs Neratinib and Veliparib (ABT-888) through an adaptive clinical trial, has been used as a basis for accelerated drug approval [35]. The FDA Critical Path Initiative now suggests modifying clinical trial design to reduce attrition and increase R&D productivity [103, 104]. Adaptive clinical trials are fairly new to the pharmaceutical industry, and their implications at different levels of the R&D and policy making processes are not fully understood. Practitioners will have to define investment and resource allocation strategies and evaluate drugs that are tested not only traditionally, but also adaptively [24].

This paper focuses on “group sequential adaptive trials.” A group sequential design allows for early termination of the trial due to safety, futility/efficacy, or both [16]. In these trials, interim analyses (decision epochs) allow decision makers to adjust their prior beliefs on the drug’s probability of success (i.e., the probability that the drug will be deemed effective at the end of the Phase 3 trial and can be submitted for FDA approval). An early termination based on interim results reduces capital commitment, and if the drug is already deemed successful, then it also shortens the time-to-market—options that fixed duration (i.e., “traditional”) trials do not offer. The decision to terminate a trial is very important, because it not only affects the current drug, but might affect future trials planned in the same therapeutic area. On the other hand, continuing a trial too long delays the dissemination of important information and puts patients at unnecessary risk [99].

The aforementioned flexibility, of terminating an adaptive trial earlier than planned, complicates the resource allocation decision (i.e., the number of patient observations to collect each period) throughout the clinical trial. The drug’s probability of success is inherently uncertain, and further, the magnitude of this uncertainty depends on the resource allocation decision itself, which affects the speed of learning: a phenomenon called endogenous uncertainty [18]. This is further compounded by the fact that, in the pharmaceutical industry, any delays in the launch of a new drug or treatment might lead to a reduced patent life (i.e., a shorter market exclusivity period) and/or a smaller market share. Therefore, the resource allocation decision throughout a sequential adaptive clinical trial is characterized by the trade-off between the incentive to establish that the drug is suc-

cessful early on, due to the time-decreasing nature of the revenue, and the benefit from collecting some information on the drug’s efficacy first, prior to committing a large amount of budget to the trial early on.

We consider adaptive clinical trials for which early termination is possible, and study how interim analyses, based on updates to the drug’s probability of success, impact the risk of misclassifying the drug (e.g., in the case of a drug with beneficial effects, the risk of classifying it as futile), the time-to-market, and the firm’s profit. Towards this end, we build a stochastic dynamic programming model that incorporates Bayesian updates on the drug’s probability of success, so that the firm can optimize its resource allocation decision periodically, i.e., the firm can terminate the trial early and re-allocate its resources to another project in case a drug does not perform sufficiently well, or may allocate more resources to speed up the Phase 3 trial, hence expediting the new drug application, if the drug shows promise in the interim analyses. The drug’s probability of success depends on safety and efficacy; in our model, similar to others in literature (e.g., Patel et al. [76] and the references therein), we focus on efficacy and assume that no safety concerns exist. We model the firm’s revenue function for a successful drug as decreasing over time to capture the market competition and reduced patent life as discussed above. The performance of the drug is evaluated based on an exogenously established threshold that indicates the minimum number of successful observations required for the submission of a new drug application to the FDA, hence, for future commercialization.

Our model explicitly considers the aforementioned trade-off between shortening the trial’s duration, through the collection of a large amount of observations early on, so as to launch the drug as early as possible and obtain a higher revenue (i.e., exploitation), and obtaining some information about the drug’s probability of success early on, prior to committing a large amount of resources (i.e., exploration). However, unlike the classical bandit problem, in our setting the exploration phase not only allows the firm to learn about the drug’s efficacy, but also contributes to the achievement of the required number of successful observations, hence to the realization of the revenue (i.e., the exploitation phase). Therefore, our work contributes to the existing literature on sequential exploration problems (i.e., variations to the classical bandit problem); e.g., Bickel and Smith [14], Brown and Smith [15], and to the research stream that focuses on information acquisition under uncertainty within dynamic programming settings (e.g., Décamps et al. [25], Ulu and

Smith [101], Smith and Ulu [94]).

Our formulation leads to a nonlinear optimization problem under a budget constraint. We identify and analyze the structural properties of the firm’s value function. These properties allow us to determine the structure of an optimal resource allocation policy. We also perform a numerical analysis using publicly available data on pharmaceutical R&D revenue and costs. Our analysis generates key insights to decision makers on the following research questions: **1)** How do the optimal resource allocation strategy, the resulting risk of drug misclassification, time-to-market, and the firm’s profit for a sequential adaptive trial compare to those of a fixed duration trial?, **2)** What is the value of increased flexibility, in terms of more interim analysis points, in a sequential adaptive trial?, **3)** What is the effect of undershooting or overshooting the prior on the drug’s efficacy in fixed duration and sequential adaptive trials?, and **4)** What are the parameters that significantly affect the firm’s optimal resource allocation strategy in fixed duration and sequential adaptive trials? We also study how sensitive the drug misclassification risk is to the resource allocation decision for the trial in order to analyze the “hot stove effect” [28], or the bias against risky and novel alternatives that may seem worse than they actually are due to purely stochastic fluctuations.

Our results indicate that the rigidity of fixed duration trials translates into a high drug misclassification risk, long time-to-market, and high costs even when using slightly optimistic or pessimistic priors on the drug’s efficacy. In fact, our findings suggest that adaptive trials with one or two interim analysis points lead to significantly larger profits (with an increase between 8% and 30%) and a significantly smaller probability of drug misclassification (with a reduction of at least 59%) over fixed duration trials, because updates to the drug’s probability of success are likely to correct the decision maker’s initial biases. In the case of a pessimistic prior, an optimal resource allocation in sequential adaptive trials indicates collecting a small number of observations early on, in order to explore the efficacy of the drug and use this information to make decisions in the subsequent periods of the trial; this contrasts with fixed duration trials that may not even be conducted in these cases. Moreover, our analysis underscores another advantage of adaptive trials: deviations from the optimal resource allocation do not significantly increase the drug misclassification risk, whereas in the fixed duration trial, the consequence might be extremely negative.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of

the existing research on adaptive clinical trials, pharmaceutical resource allocation, and sequential decision-making. Section 3 describes our stochastic dynamic programming model, and Section 4 explores important structural properties of an optimal resource allocation strategy. In Section 5, we perform a numerical study based on realistic data on pharmaceutical R&D revenues and costs, and provide insights. Finally, in Section 6 we provide our conclusions and suggest directions for future research. To facilitate the presentation, all proofs are relegated to the Appendix.

3.2 Related Literature

This section is organized in three parts. In the first part, we briefly discuss the main characteristics of drug R&D, as well as of fixed duration and adaptive clinical trials. In the second part, we position our work within the operations research literature that focuses on information acquisition under uncertainty within dynamic programming problems. In the third part, we provide a brief overview of the literature on resource allocation applied to the pharmaceutical industry, with an emphasis on optimization-based approaches.

3.2.1 Group Sequential Adaptive Clinical Trials

Drug development is a linear process comprised of several phases, where a transition to the next phase depends on a favorable outcome from the preceding phase [32]. A New Drug Application cannot be submitted to regulatory agencies and the drug cannot be commercialized unless the drug has been proven to be safe and effective in the treatment of a certain disease or symptoms. Clinical trials typically consist of three phases: Phases 1 and 2 involve testing for safety and efficacy on small groups of patients, while Phase 3 involves large studies with thousands of patients, hence, is significantly more resource-intensive. The traditional approach to conducting Phase 3 clinical trials consists of determining an *a priori* fixed sample size, based on a desired statistical significance level and power, as well as the drug’s characteristics and the treatment advantage to be observed. Next, data are collected on the full sample of patients, and the drug is deemed successful or futile.

On the other hand, in a sequential trial it is common to compare a statistic, summarizing the difference between a control (i.e., a standard therapy or a placebo) and the experimental treatment, against specified critical values at a series of interim analysis points throughout the trial, all of

which must be pre-specified in the trial’s protocol [99]. These critical values represent a stopping rule for the trial, and are often set based on financial and clinical factors [5]. If the statistic’s value surpasses the critical value, then the trial stops and the null hypothesis, of no difference between the two therapies, is rejected; if the statistic remains below the critical value, the evidence is insufficient to arrive at a conclusion, and the trial continues until the next interim analysis point [99]. Various statistical approaches have been proposed for designing group sequential trials (e.g., Whitehead [109], Jennison and Turnbull [47], Peace [78]).

Most of the existing literature focuses on adaptive trial design and implementation issues. In particular, numerous researchers have built upon the foundational works by Berry and colleagues (e.g., Berry [9]) to develop new adaptive clinical trial designs (see Chow and Chang [16] for a review of adaptive design methods and Berry and Fristedt [12] for a review of the bandit problem). For instance, Ahuja and Birge [1] propose a novel adaptive design that uses Bayesian learning to adjust patient allocation to treatments. This design utilizes forward-looking algorithms to fully exploit learning from multiple patients simultaneously, and compared to the best existing implementable adaptive trial design (e.g., Berry [8]), offers improved patient outcomes. Other researchers analyze the exploration vs. exploitation trade-off, mainly within the context of the classical bandit problem (e.g., Macready and Wolpert [60], Madani et al. [61]). In addition, several researchers analyze the impact of interim analyses from clinical, statistical, and ethical standpoints (e.g., Bassler et al. [5], Zannad et al. [113], Mueller et al. [65], Jitlal et al. [48]), but do not evaluate their effect on the resource allocation decision, which is the focus of our paper.

3.2.2 Dynamic Programming and Information Acquisition

Our work is closely related to the literature that focuses on dynamic programming models that incorporate the newly acquired information to address uncertainty in various contexts. This review is not meant to be exhaustive, but rather indicative of the different application areas within this stream of research. For example, McCardle [63] studies the acquisition of a new stationary technology where, in each period, the consumer either rejects/adopts this technology, or waits and collects new information; this paper presents a dynamic programming model in which in each period the firm updates its prior estimate on the new technology’s profitability using a Bayesian

framework. Ulu and Smith [101] build upon McCardle’s work by considering a new technology that is evolving over time, and utilize Bayesian updates on the consumer’s beliefs on the benefits of the new technology. An interesting healthcare application is presented by Kornish and Keeney [54], who derive optimal strategies for the annual influenza vaccine composition decision, where different strains of the virus need to be chosen for vaccine production under uncertainty about the vaccine’s effectiveness.

The sequential nature of decision-making in our application also positions our work within the literature on sequential exploration problems and sequential investment decisions. For example, Cyert et al. [21] analyze sequential investment decisions (i.e., partial investments) and their impact on “downside” risk, and derive an optimal sequential investment policy for a two-period problem. Kolonko and Benzing [52] examine a generalization of the one-armed bandit problem, where within a sequence of N trials, the decision maker needs to choose one of two Bernoulli experiments for each trial under uncertainty on one of the experiment’s probability of success. Bickel and Smith [14], motivated by an oil and gas exploration problem, evaluate sequential exploration strategies with dependent exploration prospects, and utilize a dynamic programming model to determine optimal exploration strategies. Recently, Brown and Smith [15] build on this work and model sequential exploration problems as variations of the classical multi-armed bandit problem, under a similar trade-off as the one present in our setting.

Specifically, in the oil and gas exploration problem, there exists a trade-off between “earning,” or drilling at locations with maximal expected values, and “learning,” or drilling at locations that offer information. Because of probabilistic dependence, the outcomes at drilled sites turn into valuable information about neighboring targets. In our model, the trade-off is between establishing that the drug is successful early on and realizing the revenue, and collecting some information on the drug’s success probability, prior to committing a large amount of budget to the trial early on. Further, in our problem, the “earning” and “learning” efforts overlap, as collecting observations (i.e., testing patients) contributes to both the achievement of the required number of successes (thus, the realization of the revenue) and the information acquisition on the drug’s efficacy.

3.2.3 Resource Allocation

In general, resource allocation decisions have been studied by numerous researchers in a variety of application areas. In the following, we provide some examples of optimization-based approaches applied to the pharmaceutical industry, in the context of R&D project portfolio construction and capacity planning. Colvin and Maravelias [19] and Christian and Cremaschi [17] propose optimization-based approaches for pharmaceutical R&D project pipeline management. Solak et al. [95] develop a multi-stage stochastic program for constructing an R&D project portfolio under endogenous uncertainty; Patel et al. [76] develop an integer programming model as the basis for Bayesian decision analysis to optimize Phase 3 R&D project portfolios.

Other researchers use real options theory in R&D project portfolio optimization (e.g., Rogers et al. [90], Jacob and Kwak [46], Enea and Lo Nigro [34], Kouvelis et al. [55]). In particular, Kouvelis et al. [55] use real options theory to study Phase 3 trials under a stochastic revenue and a stochastic patient enrollment rate, and use a test statistic to make a decision (i.e., stop the trial for futility or efficacy), hence, they do not explicitly model the decision maker's learning process resulting from data updates (clinical observations). In our model, the patient recruitment decision is made optimally, based on a clinical criterion and financial considerations (including a deterministic, time-decreasing revenue function for a successful drug), and is re-optimized in each period utilizing updates on the drug's probability of success, under the assumption that the firm is able to achieve the target recruitment rates. Subramanian et al. [96] incorporate simulation within an optimization framework to construct R&D project portfolios, to dynamically prioritize project activities based on real-time outcomes. In the robust optimization area, Hassanzadeh et al. [40] develop a model to revise and re-optimize the R&D project portfolio.

Capacity planning in the pharmaceutical industry has also received attention in the literature. For example, Oh and Karimi [72] present a mixed-integer linear programming formulation for capacity expansion planning, including decisions to expand existing manufacturing facilities, build new capacity, or outsource. Rajapakse et al. [89] consider similar decisions but develop a framework based on scenario analysis and discrete-event simulation. Some models address R&D portfolio management and manufacturing capacity planning problems simultaneously (e.g., Levis and Papageorgiou [59], Maravelias and Grossmann [62], Papageorgiou et al. [74]). In a recent paper,

Kaminsky and Yuen [49] evaluate the impact of data updates on the capacity investment decision in the pharmaceutical industry, considering a fixed duration clinical trial and a fixed and *a priori* determined patient recruitment rate. The firm needs to decide, under uncertainty on the outcome of the clinical trial, whether or not to invest in production capacity in each period: if sufficient capacity is not built by the end of the trial and if the trial is successful, then the launch of the new drug will be delayed; and if the trial is unsuccessful, then the resources invested on capacity will become a sunk cost. As opposed to their work, in this paper we identify an optimal resource allocation strategy (i.e., number of patient observations to collect in each period of the trial) for a group sequential adaptive trial that allows the decision maker to stop the trial (i.e., stop recruiting patients) at an interim analysis point for benefit or futility reasons, based on a periodically updated estimate of the drug's success probability. Furthermore, we model the resource allocation decision in an adaptive trial as an optimal stopping problem with a variable patient recruitment rate.

To our knowledge, the optimal resource allocation decision for adaptive clinical trials has not been explored, but the need to study it in a rigorous manner has been acknowledged by various researchers. For example, David et al. [24] emphasize the importance of evaluating the impact of different trial designs (such as sequential designs with interim analyses) on financial value, and argue that adaptive trials that involve sample size re-estimation at interim analyses may increase the firm's flexibility to balance risk, cost, and time in clinical trials, which we quantify in this paper.

3.3 The Model

In this section, we discuss the decision problem, and detail the formulation and characteristics of our stochastic dynamic programming model.

3.3.1 The Decision Problem

The firm's decision problem is to determine the resource allocation, i.e., the number of patients to recruit in each period $t \in [1, \dots, T]$, N_t , over the duration of the Phase 3 trial under uncertainty on the drug's probability of efficacy, p , so as to maximize the expected profit from the drug. Let M denote the maximum number of observations (one from each patient) the firm is able to

collect based on a budget of B (i.e., $M = \frac{B}{c}$, where c represents the recruitment and testing cost per patient). The drug's success (i.e., deeming the drug effective in comparison with a placebo or other existing therapy) is assessed through an exogenously determined criterion, represented in terms of the minimum number of patients that need to be treated successfully with the drug during the Phase 3 trial, $\gamma_c(M)$, which is a non-decreasing function of M , and is determined based on statistical significance and power.

If the drug's success criterion is met, then a time-dependent revenue is realized. We model the revenue for a successful drug in period t , g_t , as a strictly decreasing function in t to represent the potential loss in market share and reduction in active patent life due to delays in commercialization. We assume that the revenue for a successful drug in period t , $g_t \gg B, \forall t \in [1, T]$, which is consistent with real-world applications.

A sequential adaptive Phase 3 design provides the firm with the flexibility to terminate the trial earlier than a “traditional” trial of fixed duration (of T periods), due to either proven benefit (i.e., the success criterion is achieved early), or perceived futility (i.e., initial observations from patients indicate a “low” likelihood of success). An “early” success enables the firm to market the drug earlier and realize a higher revenue, while an “early” termination, due to perceived futility, reduces the Phase 3 trial cost. Note that an early termination option is not available for a fixed duration clinical trial, which the firm must carry to completion (i.e., it must collect the number of observations initially determined to achieve the desired statistical power), regardless of what the interim data may indicate, except for the situation in which the interim data indicate safety concerns, an aspect we do not model.

Consequently, the resource allocation decision in a sequential adaptive clinical trial bears similarities to the classical exploitation vs. exploration problem, i.e., how should the resources be split between the efforts of obtaining information (hence better estimating the unknown probability, p), and performing the experiments, or in our context, recruiting patients so as to achieve the success criterion. The main difference in our setting, however, is that the efforts overlap, meaning that the firm can use the interim data both for information updating purposes as well as for achieving the success criterion. Thus, as discussed previously, our problem revolves around the trade-off between the incentive to reach the minimum number of successful observations as early as possible to realize the maximum revenue, and the benefit resulting from the collection of information

and consequent updates to the decision maker's prior on the drug's probability of success, without committing a large amount of budget early on.

We also model and analyze a fixed duration Phase 3 trial (of T periods), which we use as a benchmark. In the sequential adaptive trial, at the beginning of each period t , the decision maker updates her estimate on p, \hat{p}_t , based on observations collected through the end of period $t - 1$, and determines the number of observations to collect in the current period (i.e., N_t , for $t = 1, 2, \dots, T$), constrained by $R_{t-1} \equiv M - \sum_{j=1}^{t-1} N_j$, i.e., the remaining number of observations that can be collected in periods t, \dots, T . When needed, we use the superscripts F and A to represent a fixed duration and a sequential adaptive trial, respectively.

We assume that the response of each patient to treatment with the candidate drug is independent from the responses of other patients. Let N_t^+ denote the number of successful observations obtained in period t , that is, $N_t^+ \sim \text{Binomial}(N_t, p), \forall t$, where the drug's efficacy probability p is unknown to the decision maker. Therefore, the decision maker models N_t^+ using her estimate, \hat{p}_t , which is updated at the beginning of each period $t, t = 1, \dots, T$, following the stochastic process outlined in Section 3.3.2.

3.3.2 Information updating

We model the decision maker's initial estimate of p at the beginning of the Phase 3 trial (\hat{p}_1) as a Beta prior with parameters λ_0 and β_0 . These parameters may represent preliminary data in the form of results from earlier trial phases, expert opinions, etc., and in case no preliminary data are available/used, one can set $\lambda_0 = \beta_0 = 1$. We define, the stochastic processes $\{\gamma_t, t \geq 0\}$ and $\{\zeta_t, t \geq 0\}$ to respectively correspond to the cumulative number of successful observations, and the cumulative number of unsuccessful observations obtained by the end of period $t, t = 1, \dots, T$:

$$\begin{aligned}\gamma_t &= \sum_{j=1}^t N_j^+ = \gamma_{t-1} + N_t^+, \\ \zeta_t &= \sum_{j=1}^t (N_j - N_j^+) = \zeta_{t-1} + (N_t - N_t^+),\end{aligned}\tag{3.1}$$

where $\gamma_0 = \zeta_0 = 0$. Then, \hat{p}_t , the posterior distribution of p at the beginning of period t , follows a Beta distribution $(\gamma_{t-1}, \zeta_{t-1})$, which the firm uses in the value function in period t . Consequently,

the posterior predictive distribution of the number of successful observations in a given period follows a Beta-binomial distribution (i.e., the likelihood times the prior, resulting in a compound distribution).

3.3.3 The Optimization Model

As explained above, at the beginning of period t , the firm observes γ_{t-1} and R_{t-1} , updates its estimate \hat{p}_t , and determines its resource allocation in period t , N_t . At the end of period t , the realization of the random variable $N_t^+ \sim \text{Binomial}(N_t, p)$ is observed. The event that the drug is successful at the end of period t is given by $\{\gamma_t = \sum_{j=1}^t N_j^+ \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\}$, where $\gamma_c(M)$ is exogenously determined. If, at the end of period t , the drug is deemed successful, then the Phase 3 trial is terminated, and the firm realizes a one-time reward of g_t . Otherwise, the firm continues with the next period of the Phase 3 trial, unless the trial's maximum duration (i.e., period T) or the maximum possible number of observations (M) is reached.

Given this sequence of events and the stopping criteria, we formulate the firm's resource allocation problem in a sequential adaptive clinical trial as a stochastic dynamic programming problem, with the value function in period $t = 1, 2, \dots, T$, given as:

$$V_t(\gamma_{t-1}, R_{t-1} | \gamma_{t-1} < \gamma_c(M)) = \max_{0 \leq N_t \leq R_{t-1}} -cN_t + \begin{cases} g_t, & \text{if } \{\gamma_t \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \\ E[V_{t+1}(\gamma_{t-1} + N_t^+, R_{t-1} - N_t) | \gamma_t < \gamma_c(M)], & \text{if } \{\gamma_t < \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \end{cases}, \quad (3.2)$$

and $V_{T+1}(\gamma_T, R_T | \gamma_T < \gamma_c(M)) = 0, \forall \gamma_T, \forall R_T$,

where the event $\{\gamma_t \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \equiv \{N_t^+ \geq \gamma_c(M) - \gamma_{t-1} | \gamma_{t-1} < \gamma_c(M)\}$. We denote by $N_t^*, t = 1, \dots, T$, the optimal solution to Eq. (2). We assume that the observation collection rate is not constraining, that is, the decision maker can select the number of observations to collect in each period throughout the trial, subject to the remaining budget.

Remark 1. Observe that $\frac{\gamma_c(M)}{M}$, i.e., the ratio of the required number of successful observations to the maximum number of observations that can be collected, acts as an implicit threshold for the trial. Note that we do not use the ratio, $\frac{\sum_{j=1}^t N_j^+}{\sum_{j=1}^t N_j}$, i.e., the ratio of the cumulative number of successful observations to the cumulative number of observations collected by the end of period t , as a success criterion in period t , as this may lead to stopping the trial after only a small number of

observations, hence may lack the required statistical power, and is not consistent with the industry practice nor FDA guidelines [102].

3.4 Properties of an Optimal Solution

The firm utilizes the stochastic process, $\{\gamma_t, t \geq 0\}$, to update its posterior predictive probabilities, which are then used in the value function. Therefore, in what follows, we first study this stochastic process. All proofs can be found in the Appendix.

Lemma 1. $\{\gamma_t, t \geq 0\}$ is a time non-homogeneous Markov process with the following transition probabilities:

$$\Pr(\gamma_1 = k \mid \lambda_0, \beta_0) = \Pr(N_1^+ = k \mid \lambda_0, \beta_0) = \binom{N_1}{k} \frac{\prod_{j=0}^{k-1} (\lambda_0 + j) \prod_{j=0}^{N_1-k-1} (\beta_0 + j)}{\prod_{j=0}^{N_1-1} (\lambda_0 + \beta_0 + j)}, \text{ for } k \in [0, 1, \dots, N_1],$$

and for periods $t = 2, 3, \dots, T$:

$$\begin{aligned} \Pr(\gamma_{t-1+m} = k + \gamma_{t-1} \mid \gamma_{t-1}) &= \Pr\left(\sum_{j=1}^m N_{t-1+j}^+ = k \mid \gamma_{t-1}\right) \\ &= \binom{N_{t,m}^{Cum}}{k} \frac{\prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t,m}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t,m}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)}, \text{ for } k \in [0, 1, \dots, N_{t,m}^{Cum}], m \in [1, \dots, T+1-t], \end{aligned} \quad (3.3)$$

where $N_{t,m}^{Cum} \equiv \sum_{j=1}^m N_{t-1+j}$.

Thus, $\{\gamma_t, t \geq 0\}$ is a Markov process due to the assumption that each clinical experiment (patient) follows an independent and identical Bernoulli distribution with parameter p . Hence, γ_{t-1} is the sufficient statistic for the upcoming transitions. Eq. (3) indicates that the impact of the outcome of any single period on the transition probability decreases over time. In consequence, the stochastic process $\{\gamma_t, t \geq 0\}$ is time non-homogeneous, with a state space that is expanding over time.

Lemma 2. $\Pr(\gamma_t \geq \gamma_c(M) \mid \gamma_{t-q} = \gamma)$ is increasing in each of γ and q , for $\gamma: \gamma \in \mathbb{Z}^+, \gamma < \gamma_c(M)$, $q = 1, \dots, t-1$.

The probability, $\Pr(\gamma_t \geq \gamma_c(M) \mid \gamma_{t-q} = \gamma)$, is the posterior predictive probability of success in period t , that is, it is calculated based on the estimate \hat{p}_t . Then, Lemma 2 is intuitive: the more successes are observed up to time $t-q$, or the larger the value of q is, i.e., the earlier γ successful observations are obtained, the higher the perceived likelihood of success in period t will be.

Lemmas 1 and 2 are similar to those in Kaminsky and Yuen [49], mainly because of the assumption that the clinical experiments are independent and identical. However, Kaminsky and Yuen [49] assume that the firm obtains a single observation in each period, and define the stochastic process as the sum of Bernoulli random variables. In our work, we allow for the possibility of multiple observations in each period, which is endogenously determined, and define the stochastic process as the sum of Beta-binomial random variables. More importantly, we model the adaptive trial setting, where the trial is terminated for benefit or for futility earlier than planned. In contrast, Kaminsky and Yuen [49] consider a fixed duration trial, i.e., without an early termination option.

We next study structural properties of the value function, $V_t(\cdot)$, and of an optimal resource allocation strategy N_t^* , $\forall t$.

Lemma 3. (i) For a given R_{t-1} , if $N_t^*(\gamma_{t-1}) = 0$ for some γ_{t-1} , then $N_t^*(\gamma) = 0$, $\forall \gamma \leq \gamma_{t-1}$.
(ii) For a given γ_{t-1} , if $N_t^*(R_{t-1}) = 0$ for some R_{t-1} , then $N_t^*(R) = 0$, $\forall R \leq R_{t-1}$.

Thus, if it is optimal to not collect any observations in period t with an available budget to cover R_{t-1} observations, then the optimal strategy will remain the same at lower budget levels. Similarly, if it is optimal to not collect any observations in period t when the cumulative number of successes at the end of period $t - 1$ is γ_{t-1} , then the optimal strategy will not change with a lower number of successes realized.

Lemma 4. (i) If for any $0 < t < T$, $N_t^* = 0$, then $N_{t+1}^* = N_{t+2}^* = \dots = N_T^* = 0$.
(ii) If for any $0 < t < T$, $R_{t-1} < \gamma_c(M) - \gamma_{t-1}$, then $N_t^* = N_{t+1}^* = \dots = N_T^* = 0$.

Lemma 4 is a direct consequence of Lemma 3. If the optimal strategy is to not collect any observations in any given period t , then this translates into terminating the trial either for success or futility. This policy is aligned with the pharmaceutical industry's practice: once patient recruitment is suspended, it is rarely reinitiated. Furthermore, if in any period the firm does not have the budget to recruit a sufficient number of patients to satisfy the success criterion, it becomes optimal to terminate the trial for futility.

Next we characterize the structure of an optimal resource allocation solution in a sequential adaptive trial.

Theorem 1. The optimal resource allocation for a sequential adaptive trial in period $t \in [1, \dots, T]$ follows a threshold policy:

if $\gamma_{t-1} \geq \gamma_c(M)$, or $R_{t-1} < \gamma_c(M) - \gamma_{t-1}$, then $N_t^{*A}(\gamma_{t-1}, R_{t-1}) = 0, \forall \gamma_{t-1}, R_{t-1}$;

otherwise:

$$N_t^{*A}(\gamma_{t-1}, R_{t-1}) = \begin{cases} 0, & \text{if } \{\gamma_{t-1} \leq \underline{\gamma}_{t-1}^A\} \equiv \left\{E[\hat{p}_t] \leq \underline{p}_t^A \leq \frac{\gamma_c(M)}{M}\right\}, \\ > 0, & \text{if } \{\gamma_{t-1} > \underline{\gamma}_{t-1}^A\} \equiv \{E[\hat{p}_t] > \underline{p}_t^A\} \end{cases}, \quad (3.4)$$

where $E[\hat{p}_t] = \frac{\gamma_{t-1}}{\gamma_{t-1} + \zeta_{t-1}}$, for $t \in [2, \dots, T]$, and $E[\hat{p}_1] = \frac{\lambda_0}{\lambda_0 + \beta_0}$. Further, the threshold, $\underline{\gamma}_{t-1}^A$, is a function of R_{t-1} , and decreases in each of γ_{t-1} and R_{t-1} .

Theorem 1 indicates that, the larger the cumulative number of successes obtained up to period t , or the larger the remaining budget, the smaller the region is where it is optimal to stop the trial for futility.

Corollary. The optimal resource allocation for a fixed duration trial with T periods follows a threshold policy:

$$N^{*F}(\lambda_0, \beta_0, M) = \begin{cases} 0, & \text{if } \{\lambda_0 \leq \underline{\lambda}_0^F\} \equiv \left\{E[\hat{p}_1] = \frac{\lambda_0}{\lambda_0 + \beta_0} \leq \underline{p}^F \leq \frac{\gamma_c(M)}{M}\right\}, \\ > 0, & \text{if } \{\lambda_0 > \underline{\lambda}_0^F\} \equiv \{E[\hat{p}_1] > \underline{p}^F\} \end{cases},$$

where the threshold $\underline{\lambda}_0^F$ decreases in each of λ_0 and M , and the threshold $\underline{p}^F = \underline{p}_T^A|_{(\gamma_{T-1}=0, R_{T-1}=M)}$.

Next we study how the resource allocation thresholds, $\underline{p}_t^A, t \in [1, \dots, T]$, change over time. Observe that at the beginning of period t , given realizations γ_{t-1} and R_{t-1} , the threshold \underline{p}_t^A is a constant, while the thresholds in the subsequent periods, $\underline{p}_{t+1}^A, \dots, \underline{p}_T^A$, are random variables. We thus present the following result.

Lemma 5. (i) For a sequential adaptive trial, the resource allocation thresholds are such that:

$$\Pr\left(\underline{p}_{t+1} \geq \underline{p}_t(\gamma_{t-1}, R_{t-1})\right) = 1, \text{ for } t \in [1, \dots, T].$$

(ii) For a fixed duration trial, we have that $\underline{p}^F \geq \underline{p}_1^A$.

With the characterization of the decision thresholds, we now compare the drug misclassification risk in the adaptive and fixed duration trials. For this purpose, let $\beta(p) \equiv \Pr(\text{misclassifying a drug having a true probability of efficacy of } p)$.

Theorem 2. Consider an effective drug with a true probability of efficacy of p . Then:

- (i) If $E(\hat{p}_1) \leq \underline{p}_1^A$, then $N_1^{*F} = N_1^{*A} = \dots = N_T^{*A} = 0$, and $\beta^A(p) = \beta^F(p) = 1$.
- (ii) If $\underline{p}_1^A < E(\hat{p}_1) \leq \underline{p}^F$, then $N_1^{*F} = 0$, $N_1^{*A} > 0$, leading to $\beta^A(p) < \beta^F(p) = 1$.
- (iii) If $E(\hat{p}_1) > \underline{p}^F$, then $N_1^{*F} > 0$, $N_1^{*A} > 0$, leading to $\beta^A(p) < 1$ and $\beta^F(p) < 1$.
- (iv) $\exists \bar{p}^F > p$ such that for $E(\hat{p}_1) \geq \bar{p}^F$, $\beta^F(p) > \beta^A(p)$,

where $E(\hat{p}_1) = \frac{\lambda_0}{\lambda_0 + \beta_0}$. To demonstrate Theorem 2, we consider an effective drug, i.e., a drug with an efficacy probability of $p > \frac{\gamma_c(M)}{M}$, as $E[N^+(M)] > \gamma_c(M)$ (the expected number of successful observations in M observations exceeds the success threshold, $\gamma_c(M)$). Fig. 1 depicts the implications of Theorems 1 and 2 for an effective drug, in terms of the testing and no testing regions in a fixed duration trial and in a two-period sequential adaptive trial (with one interim analysis).




Figure 3.1: Testing and drug misclassification regions delimited by thresholds on the drug's estimated probability of efficacy.

Next we discuss the functional form of the value function $V_t(\cdot)$, which is based on the posterior predictive estimate \hat{p}_t . Specifically, we are able to show that the last period's value function, $V_T(\cdot)$, is either first convex and then concave in N_T with a unique maximizer, or $V_T(\cdot) \leq 0, \forall N_T$ (see Fig.2a, and Lemma A in Appendix), which leads to Theorem 3.

Theorem 3. For a sequential adaptive trial, if $N_T^*(\gamma_{T-1}, R_{T-1}) > 0$, then we have:

$$N_T^*(\gamma_{T-1}, R_{T-1}) = \begin{cases} R_{T-1}, & \text{if } \{\underline{\gamma}_{T-1}^A < \gamma_{T-1} < \bar{\gamma}_{T-1}^A\} \equiv \{\underline{p}_T^A < E[\hat{p}_T] < \bar{p}_T^A\} \\ < R_{T-1}, & \text{if } \{\gamma_{T-1} \geq \bar{\gamma}_{T-1}^A\} \equiv \{E[\hat{p}_T] \geq \bar{p}_T^A\} \end{cases}, \quad (3.5)$$

where the lower and upper threshold values, $\underline{\gamma}_{T-1}^A$ and $\bar{\gamma}_{T-1}^A$, are functions of R_{T-1} .

Next we study the value function in any intermediate period t , $0 < t < T$, which is com-

prised of two components: the expected profit in the current period, given by $-cN_t + g_t \cdot \Pr(N_t^+ \geq \gamma_c(M) - \gamma_{t-1} | \gamma_{t-1} < \gamma_c(M))$, and the expected profit to-go, given by $E[V_{t+1}(N_t) | \gamma_{t-1} < \gamma_c(M)] \equiv E[V_{t+1}(\gamma_{t-1} + N_t^+, R_{t-1} - N_t) | \gamma_t < \gamma_c(M)] \cdot \Pr(N_t^+ < \gamma_c(M) - \gamma_{t-1} | \gamma_{t-1} < \gamma_c(M))$. The expected profit function in the current period has the same structure as that of the value function in period T . The properties of the expected profit to-go function are discussed below; see also Fig. 1.

First, observe that if a sufficiently large number of observations is collected (i.e., a large amount of budget is allocated) in the current period, then the trial will likely stop at the end of the current period, because of the small leftover budget for the next period and/or the high likelihood that the cumulative number of successes will reach $\gamma_c(M)$ in the current period rather than in a future period. Thus, the expected profit to-go function decreases as N_t increases over the region where N_t is sufficiently large (see the $N_t \geq 115$ Region in Fig. 2b). However, the expected profit to-go function is not necessarily monotone outside of this region. The second component, $\Pr(N_t^+ < \gamma_c(M) - \gamma_{t-1} | \gamma_{t-1} < \gamma_c(M))$, which corresponds to the likelihood that the trial continues to the next period, is monotone decreasing as N_t increases. However, the first component, $E[V_{t+1}(\gamma_{t-1} + N_t^+, R_{t-1} - N_t) | \gamma_t < \gamma_c(M)]$, is non-monotone in general, as an increase in N_t is likely to translate into a higher N_t^+ , but also into a lower $R_t (= R_{t-1} - N_t)$. In summary, the expected profit to-go function behaves in a non-monotone manner in N_t up to a point after which it starts decreasing (see Fig. 2b and Lemma B in the Appendix).

As a result, the value function in period t , $V_t(\gamma_{t-1}, R_{t-1} | \gamma_{t-1} < \gamma_c(M))$, is not necessarily monotone in N_t . However, our extensive numerical study indicates that the maximizer of the current period expected profit function is often very close (if not equal) to the global maximizer of the $V_t(\cdot)$ function. Indeed, the values of $E[V_{t+1}(N_t) | \gamma_{t-1} < \gamma_c(M)]$ in the interval $\{0 \leq N_t \leq \gamma_c(M) - \gamma_{t-1} - 1\}$, where it behaves in a non-monotonic manner, show little deviation for a wide range of parameter values. In the interval $\{\gamma_c(M) - \gamma_{t-1} \leq N_t\}$, $E[V_{t+1}(N_t) | \gamma_{t-1} < \gamma_c(M)]$ first increases and then decreases, or solely decreases. As a result, the $V_t(\cdot)$ function remains almost flat in $\{0 \leq N_t \leq \gamma_c(M) - \gamma_{t-1} - 1\}$, as the current profit function is shifted up; and in $\{\gamma_c(M) - \gamma_{t-1} \leq N_t\}$, the $V_t(\cdot)$ function increases and then decreases, mainly because the current profit function has this same behavior (Fig. 2c). Therefore, not much can be said on the properties of N_t^* , except the one provided in Theorem 1. Consequently, in our numerical study we determine the optimal resource allocation through an exhaustive search that relies on the results of Theorems 1 and 3, along with the monotonicity

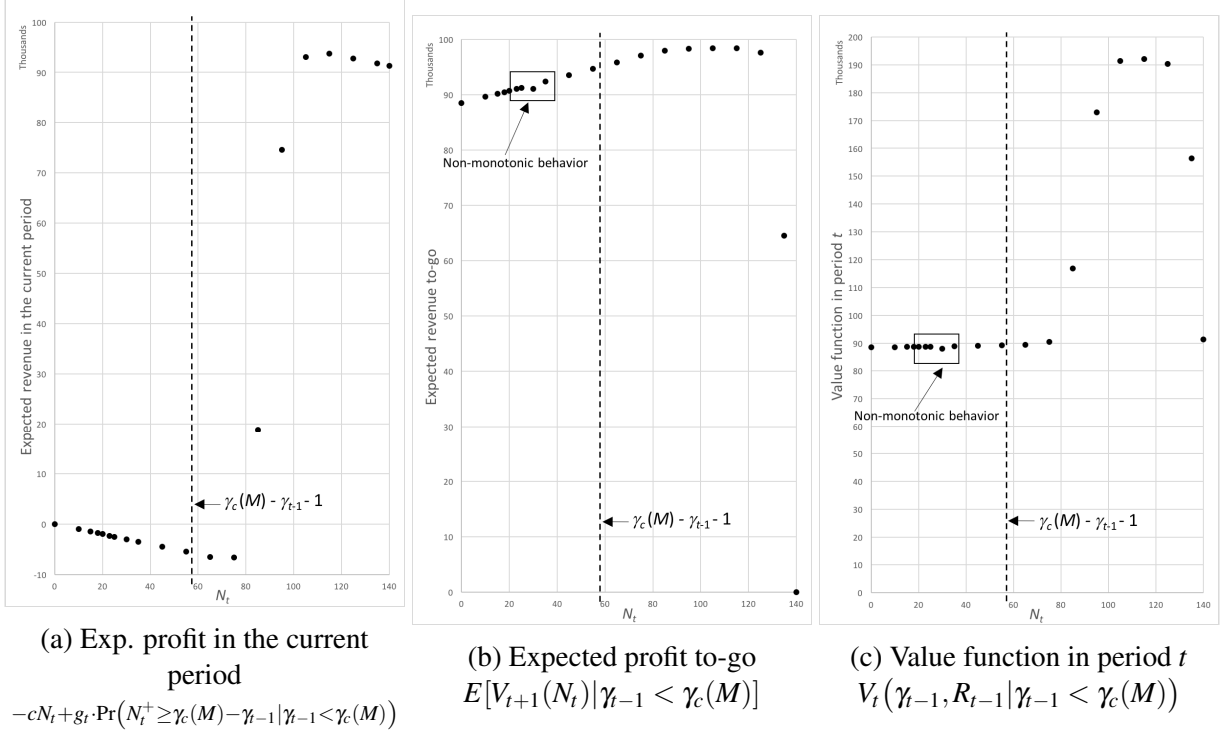


Figure 3.2: The value function versus N_t in an intermediate period t , $0 < t < T$ (increments of 2, 3 and 5 units of N_t in $[10, 35]$, and increments of 10 units of N_t in $[45, 135]$).

properties in Lemmas 3-4 and Lemmas A-B.

3.5 Numerical Analysis and Insights

In this section, we perform a series of numerical experiments to build upon our earlier findings on the structure of an optimal resource allocation strategy and to answer the following research questions:

1. How do the optimal resource allocation strategy, the resulting risk of drug misclassification, time-to-market, and the firm's profit for a sequential adaptive trial compare to those of a fixed duration trial?
2. What is the value of increased flexibility, in terms of more interim analysis points, in a sequential adaptive trial?
3. What is the effect of undershooting or overshooting the prior on the drug's efficacy in fixed duration and sequential adaptive trials?

4. What are the parameters that significantly affect the firm’s optimal resource allocation strategy in fixed duration and sequential adaptive trials?
5. How sensitive is the drug misclassification risk to deviations from the optimal resource allocation decision for the trial?

For this purpose, we solve the resource allocation problem optimally for a series of problem instances, using parameter values representative of industry data, as detailed in the following sections.

3.5.1 Experimental Design

In order to get insight on the impact of interim analyses, hence updates to the drug’s probability of success, on the optimal resource allocation and the corresponding performance metrics, we determine the optimal strategy in adaptive trial settings with one and two interim analysis points (i.e., two- and three-period problems, respectively), as well as in a fixed duration (single-period) trial. The literature indicates that the gains from interim analyses are the greatest when going from one to two interim analyses, and that having more than five interim analyses is unlikely to generate statistical, practical, or ethical benefits, unless an extremely large difference between treatment with the candidate drug and the control is expected [64, 85]. Furthermore, Pocock [86] states that “a trial with just two interim analyses has considerable advantage over a fixed duration trial in reducing the average number of patients exposed to an inferior treatment.” All three trials have the same total duration (i.e., the length of the fixed duration trial), which is divided into two or three periods of equal length, based on the trial setting. We choose a drug with a *true* probability of efficacy of $p = 0.70$ (i.e., a drug with a moderate beneficial effect), which is unknown to the decision maker. In this setting, misclassifying the drug corresponds to classifying the drug as futile, and is equivalent to a Type II error. We also assume that the observation collection rate is not constraining, that is, the decision maker can select the number of observations to collect in each period throughout the trial.

We use the following set of parameters: an available budget (B) of \$10 million and a per observation (patient) cost (c) of \$50,000 [82], which is representative of the total cost per patient in a Phase 3 oncology trial; the main drivers for the high observation cost are typically the level of

professional resources and the physical infrastructure needed to support the often geographically dispersed Phase 3 trial activities. We consider a time-decreasing revenue with a reduction rate of 5% per period (we also consider reduction rates per period of 1% and 50% in the sensitivity analysis), with a final revenue of $g_3 = \$1$ billion. The revenue figures were chosen based on global net sales per novel active substance (i.e., a molecular or biologic entity or combination product in which at least one element had not been previously approved by the FDA) in the 2005 - 2009 launch cohort [7], and sales expectations for cancer drug Keytruda (Pembrolizumab) [70]. The success criterion in Phase 3 is represented by a minimum number of successful observations of $\gamma_c(M) = 100$, corresponding to $M = \frac{B}{c} = 200$ maximum observations; this threshold exemplifies certain Phase 3 oncology studies [67, 69, 2, 112] after which the drug can be submitted for FDA approval and can be commercialized upon approval.

To evaluate the sensitivity of the optimal solution to the quality of the initial estimate of the drug's efficacy probability, we vary the parameters of $\hat{p}_1, \hat{p}_1 \sim \text{Beta}(\lambda_0, \beta_0)$, such that its expected value, given by $\frac{\lambda_0}{\lambda_0 + \beta_0}$, equals $\{0.1, 0.2, \dots, 0.9\}$, and discuss the $E(\hat{p}_1) = \{0.6, 0.7, 0.8\}$ cases in depth, which respectively cover the perfect information case ($E(\hat{p}_1) = 0.7$), as well as the cases of undershooting and overshooting the prior (i.e., $E(\hat{p}_1) = 0.6$ and $E(\hat{p}_1) = 0.8$, respectively).

We solve the stochastic dynamic programming problem (Section 3.3) to compute the optimal resource allocation for each sequential adaptive trial setting (i.e., two- and three-period problems). In particular, we first compute N_1^* , and then randomly generate a realization of the number of successes N_1^+ , i.e., we generate a realization from a Binomial distribution with parameters N_1^* and the true efficacy probability p . Given the realization of N_1^+ and the remaining budget, $B - cN_1^*$, we then optimize the resource allocation in the remaining period(s), i.e., N_2^* for the one interim analysis setting vs. N_2^* and N_3^* for the two interim analyses setting, depending on the realization of the number of successful observations and the remaining budget (see Remark 2). We use 3,000 replications for each simulation scenario. The optimal solution to the fixed clinical trial setting simply follows from Theorem 1.

The performance measures include: 1) the proportion of time the drug is misclassified, 2) the expected time-to-market, 3) the firm's expected profit, and 4) the expected costs of optimism and pessimism (i.e., $E[\text{Profit}(E(\hat{p}_1) = 0.7)] - E[\text{Profit}(E(\hat{p}_1) = 0.8)]$ and $E[\text{Profit}(E(\hat{p}_1) = 0.7)] - E[\text{Profit}(E(\hat{p}_1) = 0.6)]$, respectively).

3.5.2 Optimal Resource Allocation in Fixed Duration vs. Sequential Adaptive Trials

We first discuss how the firm's optimal resource allocation strategy in sequential adaptive trial variations compares to that in the fixed duration trial in each of the three scenarios: undershooting the prior ($E(\hat{p}_1) = 0.6$), perfect information ($E(\hat{p}_1) = 0.7$), and overshooting the prior ($E(\hat{p}_1) = 0.8$). Fig. 3 depicts the optimal number of observations to be collected in the first period (N_1^*), which, for a fixed trial, also represents the total number of observations to be collected throughout the trial.

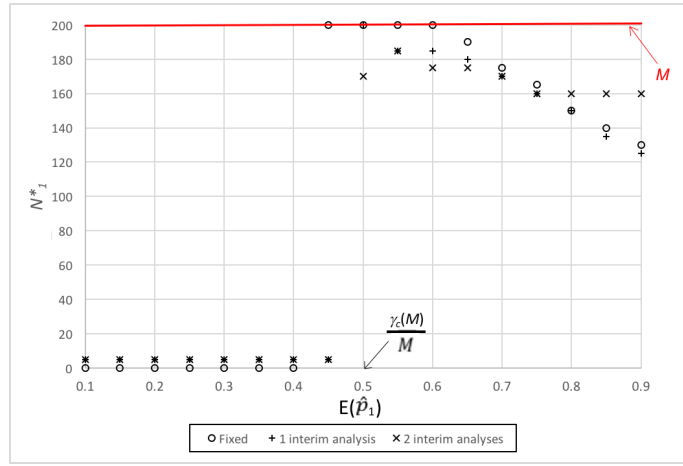


Figure 3.3: Optimal first period resource allocation in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)

Although the revenue is time-decreasing, the optimal strategy in the adaptive trials suggests to initially “test out” or “explore” the drug when the decision maker’s prior is pessimistic (Fig. 3). If the drug’s performance in the first period is promising, then the optimal solution allocates a much larger budget to the trial in the second period, so as to “exploit” the drug’s good performance. Thus, the optimal strategy in the sequential adaptive trials is to act conservatively when the prior on p is low. In the fixed duration trial, however, in the absence of the exploration option, low priors automatically lead to zero budget allocation to the trial, i.e., the drug is immediately classified as futile without conducting a Phase 3 trial.

Fig. 4 depicts the total number of observations for each trial in the three scenarios (i.e., N_1^* for the fixed duration trial, and $N_1^* + E(N_2^*)$ vs $N_1^* + E(N_2^* + N_3^*)$ for the one and two interim analysis variations of the sequential trial). The trade-off between the time-decreasing nature of the revenue




Figure 3.4: Optimal resource allocation in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)

and the flexibility offered by the multi-period nature of adaptive trials, to explore the performance of the drug prior to committing a large amount of budget in the first period, explains why the total number of observations to be collected is the largest for the fixed duration trial in the pessimistic and perfect information scenarios. On the other hand, when using an optimistic prior in the two interim analyses trial, the firm's incentive to realize the revenue early on exceeds the potential benefit of gathering some information about the drug, leading to a larger number of observations than that in the fixed duration trial. As expected, the optimal number of observations decreases as the prior, $E(\hat{p}_1)$, increases in all cases. The total number of observations in the fixed duration trial is the largest among the three scenarios when using perfect information or a pessimistic prior, but the lowest (along with the one interim analysis trial) when using an optimistic prior. As Fig. 4 indicates, the adaptive trials stop at the first interim analysis point in five out of the six scenarios studied, and we analyze this phenomenon in the next section.

Not surprisingly, larger potential revenues produce a “higher tolerance” to low priors. Specifically, there is a wider range of priors that lead to optimal strategies of no investment—or more conservative resource allocation strategies—when the revenue g_3 is lower than \$1 billion and/or the per observation cost c is higher than \$50,000. This happens for both fixed duration and adaptive trials. Thus, in general, the optimal strategy depends on the ratios $\frac{\gamma_c(M)}{M}$ and $\frac{g_t}{c}, \forall t$.

3.5.3 Optimistic vs. Pessimistic Priors, Risk of Misclassification, and Time-to-market

In this section, we analyze the performance of the optimal resource allocation strategy, in terms of the probability of misclassifying the drug, the expected time-to-market, the expected profit, and the costs of optimism versus pessimism in each trial and for each scenario; see Tables 1 - 3.

Table 3.1: Drug misclassification risk (%) for fixed duration and sequential adaptive trials

Setting	$E(\hat{p}_1)$								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Fixed duration	100.00	100.00	100.00	100.00	0.00	0.00	0.00	14.87	95.30
1 interim analysis	41.10	41.10	41.10	41.10	0.00	0.00	0.00	0.00	0.13
2 interim analyses	8.40	8.40	8.40	8.40	0.00	0.00	0.00	0.00	0.00

Table 3.2: Expected time-to-market (in number of periods) for fixed duration and sequential adaptive trials (NA=not applicable, as the drug is misclassified as futile)

Setting	$E(\hat{p}_1)$								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Fixed duration	NA	NA	NA	NA	3.00	3.00	3.00	3.00	3.00
1 interim analysis	3.00	3.00	3.00	3.00	1.50	1.50	1.51	1.74	2.99
2 interim analyses	2.28	2.28	2.28	2.28	1.00	1.00	1.00	1.01	1.01

Table 3.3: Expected profit and costs of pessimism vs. optimism for fixed duration and sequential adaptive trials

Setting	E[Profit] (\$ million)			Cost (\$ million)	
	$E(\hat{p}_1) = 0.6$	$E(\hat{p}_1) = 0.7$	$E(\hat{p}_1) = 0.8$	Pessimism	Optimism
Fixed duration	990.00	991.25	843.83	1.25	147.42
1 interim analysis	1,071.08	1,071.83	1,059.90	0.75	11.93
2 interim analyses	1,099.28	1,099.48	1,099.25	0.19	0.23

Not surprisingly, the fixed duration trial leads to the highest misclassification risk overall: by not allocating any resources in the scenarios with pessimistic priors, the decision maker does not collect any observation and automatically classifies the drug as futile. Interestingly however, when overshooting the prior, the fixed duration trial continues to exhibit a very high misclassification risk, in contrast to the virtually null misclassification risks in the adaptive settings (Table 1), where optimistic priors can often be corrected at the first interim analysis point. An optimistic prior often

translates into an insufficient amount of observations collected in the fixed duration trial, and since the prior is not updated, the required number of successes is not achieved. In contrast, the drug misclassification risk is reduced in both variations of adaptive trials by at least 59%, mainly though the exploration phase when the prior is pessimistic (see Fig. 3).

When using pessimistic priors, the drug misclassification risk in the two interim analyses trial is further reduced by the additional flexibility provided over the one interim analysis trial: in the former, the decision maker collects a small number of observations in the first period, and keeps testing out the drug for a second period in case only a small number of successes are obtained in the first period, whereas the same results in the single interim analysis trial would lead to stopping the trial for futility. When using optimistic priors, in case the required number of successes is not reached in the first period, both adaptive trials give the decision maker the opportunity to collect more observations in the second period and classify the drug as successful, thus showing virtually null misclassification risks.

In terms of the expected time-to-market, the adaptive settings significantly reduce this measure over the fixed duration trial, where the drug is either not tested at all (when using pessimistic priors) or submitted for FDA approval only after the trial's full duration (3 periods). In fact, the single interim analysis trial is the most attractive among the three trial settings, except when the prior is overly pessimistic or overly optimistic (because of the high drug misclassification risk and the long time-to-market, respectively). Recall that we assume that all trials have the same total duration. Therefore, although the two interim analyses trial leads to the shortest time-to-market, the optimal number of observations in this setting needs to be obtained in a third of the total duration—an ambitious goal that may be hard to accomplish.

Finally, as shown in Table 3, the cost of pessimism is only a fraction of the cost of optimism, because overshooting the prior leads the decision maker to collect fewer observations and either misclassify the drug as futile (in the fixed duration trial), or delay product launch (in adaptive trials). Further, the cost of optimism in the fixed duration trial is the largest due to the high misclassification risk caused by overshooting the prior (Table 1). In terms of the expected profit, the fixed duration trial performs the worst, and the single interim analysis trial offers the largest improvement because of its lower misclassification risk and shorter time-to-market. The adaptive trial with two interim analyses is the least sensitive to deviations from perfect information, but shows

only a small improvement in the expected profit over the single interim analysis trial. Importantly, the expected profit increases significantly in both adaptive trials: over the fixed duration trial, the one and two interim analyses trials show a profit increase of 8% and 11%, respectively, in the pessimistic and accurate prior scenarios, and an increase of 26% and 30% when using an optimistic prior.

3.5.4 The Impact of Resource Allocation on the Drug Misclassification Risk: A Sensitivity Analysis

We observe that the drug misclassification risk is high only in the fixed duration trial with pessimistic and optimistic priors, but decreases significantly in the one interim analysis setting. To study the hot stove effect [28], or the bias against risky and novel alternatives that may seem worse than they actually are due to purely stochastic fluctuations, we next study how sensitive the drug misclassification risk is to the resource allocation decision in each trial. In practice, the firm may deviate from the optimal solution either due to incorrect parameter estimation, or due to constraints limiting its resource allocation or the patient recruitment rate, etc. “Positive deviations,” where a larger than optimal number of observations are collected in the first (or only) period of the trial, imply a lower profit but should not negatively impact the risk of drug misclassification. On the other hand, “negative deviations” from the optimal solution may substantially increase the misclassification risk, see Fig. 6, where we vary N_1 in $[10, 200]$ and calculate the resulting misclassification risk.

The misclassification risk in the fixed duration trial is the most sensitive to deviations from the optimal strategy, and interestingly enough, to an optimistic prior. As discussed in Section 5.3, slightly overestimating the prior (i.e., $E(\hat{p}_1)=0.8$) leads to misclassifying the drug in 14.87% of the simulation runs. This explains the high cost of optimism reported in Table 3. Reducing the number of observations by only 3% entails a misclassification risk of 35%, and a reduction in observations of 10% leads to a misclassification risk of 82%.

Both sequential adaptive trials studied are significantly less sensitive to negative deviations from the optimal resource allocation strategy (Fig. 6). Specifically, a deviation of 67% from the optimal strategy translates into a drug misclassification risk of only 3% in the one interim

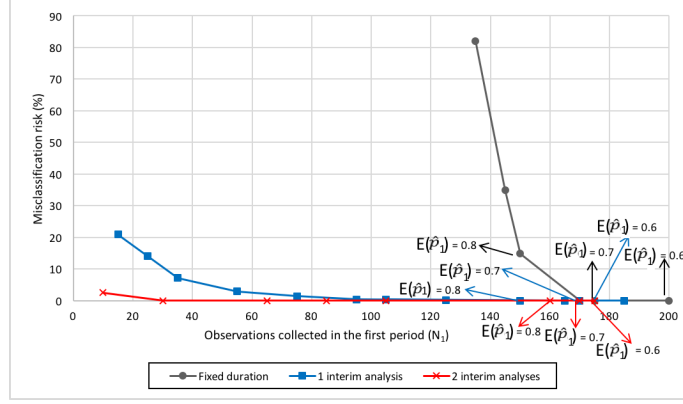


Figure 3.5: Impact of deviations from the optimal resource allocation strategy on misclassification risk in the fixed duration trial and sequential adaptive trial variations (one interim analysis and two interim analyses)

analysis trial for the perfect information scenario, and a misclassification risk of only 0.17% for the optimistic prior scenario. The risk of misclassification is very low in the trial with two interim analysis points, even when significantly deviating from the optimal strategy. If the budget allocated in the first period is insufficient to achieve the number of observations required for drug approval, this can be corrected using an updated prior. By allocating a larger budget in the second period, the required number of observations is likely to be achieved if the drug's true probability of success is sufficiently large. This analysis underscores another power of adaptive trials: deviations from the optimal solution do not significantly increase the drug misclassification risk, whereas in the fixed duration trial, the consequence is extremely negative. Thus, in sequential adaptive trials, the hot stove effect can be mitigated by correcting the decision maker's initial biases, even under non-optimal resource allocation decisions, but this is not the case for fixed duration trials.

We also performed a sensitivity analysis on the revenue, g_t , by varying the per period revenue reduction rate in [1%, 5% (baseline), and 50%] (see Table 4 in the Appendix). Note that the optimal solution in the fixed duration trial does not change with these rates, because the firm realizes the revenue only at the end of the trial (T). In the one interim analysis setting, the optimal solutions show a similar pattern to Fig. 3, where the decision maker first explores the drug's performance when the prior is pessimistic. When the revenue decreases at a rate of 1%, the optimal solutions remain the same as those of the baseline. For a revenue reduction rate of 50%, however, the optimal solutions are slightly higher than those of the baseline when using slightly pessimistic as well as optimistic priors. Given that the revenue reduction rate is very high, the resource allocation

strategy needs to be aggressive—the benefit from realizing the revenue early on exceeds the value of exploring the drug.

3.6 Conclusions and Directions for Future Research

We focus on sequential adaptive clinical trials and study the impact of interim analyses on performance measures relevant to the new drug R&D process, including the risk of drug misclassification, time-to-market, and the firm’s profit. Towards this end, we build a stochastic dynamic programming model that incorporates Bayesian updates on the drug’s probability of efficacy, to optimize the firm’s resource allocation strategy each period. We identify structural properties of an optimal resource allocation strategy and build upon them in a numerical study using publicly available data for drug R&D revenue and costs. Our research suggests that fixed duration trials are significantly more sensitive to pessimistic and optimistic priors on the drug’s probability of efficacy than their adaptive counterparts. In particular, fixed duration trials involve high misclassification risks that may translate into sunk costs. In contrast, adaptive trials, even with one or two interim analysis points, can substantially improve the performance of a Phase 3 trial through shorter time-to-market, low misclassification risks, and higher profits, providing a win-win situation for both the pharmaceutical industry and the patients. Our model also applies to other industries where similar resource allocation problems are present.

Most of the existing literature focuses on adaptive trial design (e.g., allocation of patients to treatment arms or doses, alteration of treatment strategies, etc.). Thus, our work makes a contribution by determining the structure of the decision maker’s optimal resource allocation strategy for adaptive trials under a Bayesian framework. We provide important insights on the value generated by sequential adaptive trials, as well as how they compare to fixed duration trials in terms of the risk of drug misclassification, expected time-to-market, and the firm’s profit. Further, we make a contribution to the existing literature on sequential exploration problems, as well as to the research stream on information acquisition.

An important limitation of our model is the assumption of a time-varying observation collection (i.e., patient recruitment) rate, which is not constrained, meaning the firm is able to recruit the required number of patients in each period. Studying this problem under a limiting observa-

tion collection rate is an interesting direction for future research. Further, the patient recruitment rate may be stochastic, especially for trials involving rare diseases. Our model can be further expanded to consider a stochastic revenue as well as an observation measurement error (i.e., patient outcomes may not be directly observable) that may affect the drug's misclassification risk. Finally, the analysis of portfolios of drugs tested adaptively—including a series of candidate drugs at different developmental stages—represents another important future research direction, as most pharmaceutical firms will perform their resource allocation at a portfolio, rather than a drug, level.

Chapter 4

Benefits of Optimal Testing Schedules for Sequential Adaptive Clinical Trials

4.1 Introduction and Motivation

Drug research and development (R&D) is a lengthy, complex, and costly process, surrounded by multiple sources of uncertainty affecting its outcomes. Since the drug's eligibility for commercialization is determined by the US FDA, or its European counterpart, the European Medicines Agency (EMA), pharmaceutical firms face significant risks in developing new drugs. Consequently, decision makers need to carry out thorough analyses in order to ensure that the firm's limited resources are allocated to the most promising drug candidates. Most importantly, the outcomes of drug R&D impact the quality and duration of human life. New medications and therapies have mitigated the effects of serious diseases such as cancer, hepatitis C, and AIDS, and have the potential to cure newly discovered diseases.

Adaptive clinical trials promise to expedite patients' access to new treatment options, and to significantly reduce the high costs of drug R&D [4, 10, 11, 13]. DiMasi et al. [29] estimated the out-of-pocket cost of a new drug receiving FDA approval to be around \$1.39 billion in 2013, which represents a 166% increase over the cost of drugs approved in the 1990s [31]. Further, the capitalized cost for the pharmaceutical industry, which includes the costs of failures, or drugs that did not succeed in getting regulatory approval, is also very high. The capitalized 2016 estimated cost reached \$2.56 billion, representing a 145% increase over the capitalized cost of drugs approved

in the 1990s. Clinical trials also account for the majority of the time elapsed between the start of clinical testing and the submission of a New Drug Application (NDA) to the FDA, which was estimated to be around 80.8 months, or slightly short of seven years.

In light of the promise of adaptive clinical trials, the US FDA has modified its drug evaluation procedures and put expedited approval mechanisms in place [45]. Through the Critical Path Initiative, the FDA recommends changing clinical trial design to increase R&D productivity and increase the number of viable drugs under development throughout the R&D process [103, 104]. Since adaptive clinical trials involve novel designs that are not fully understood [105], their implications at the resource allocation and policy making levels need to be examined. In particular, the need for novel investment strategies that leverage the flexibility of adaptive designs has been highlighted by David et al. [24], who emphasize the potential financial effects of these trials.

In this paper, we focus on sequential adaptive clinical trials, and specifically, on parallel group settings, where one group (arm) of patients receives an experimental drug, and the other group receives a control therapy (i.e., a standard treatment or a placebo). In these trials, the efficacy of the experimental drug is compared against the control's at different interim analysis points, giving decision makers the opportunity to terminate the trial early if the clinical results indicate that the experimental drug is more effective (i.e., stopping for benefit), or less effective than or equivalent to the control (i.e., stopping for futility). Unlike fixed sample size trials, early stopping reduces capital commitment, and can shorten the drug's development time line (and consequently, the time-to-market). Early stopping decisions impact not only the drug whose trial is terminated, but also other treatments in the same therapeutic area that may be developed in the future. Further, continuing the development of an ineffective drug can put trial participants at unnecessary risk, and in the case of an effective drug, can delay patient access to the drug as well as the dissemination of important results [99].

Sequential adaptive clinical trials require decision makers to determine, first, the number of patients that need to be recruited and tested (i.e., the number of responses that need to be collected) in each period, and second, the stopping criteria that will trigger a trial termination decision, either for benefit or futility reasons. The first decision requires the firm to determine a testing schedule, which needs to be used in conjunction with success and failure criteria that inform the second decision. Our objective is to compare the performance of the optimal testing schedule and optimal stop-

ping criteria with the termination practices used in sequential trials, by utilizing various measures relevant to the industry, i.e., expected profit, expected time-to-market, and drug misclassification risk. Towards this end, we model current practices via fixed testing schedules and stopping criteria based purely on statistical significance and power (e.g., the triangular test). In particular, we study fixed testing schedules, representing aggressive, neutral, and conservative attitudes towards risk, and pair them with the upper and lower stopping boundaries of the triangular test—the most widely implemented sequential method based on straight line boundaries [91, 99, 110].

To derive optimal, adaptive testing schedules, we solve the stochastic dynamic programming problem developed in Rojas-Cordova and Bish [92] (see Chapter 3). Specifically, the optimal testing schedules are derived considering the exploration versus exploitation trade-off. In particular, the exploration phase corresponds to obtaining information about the experimental drug’s unknown success probability, prior to allocating a large amount of resources to the trial early on. On the other hand, the exploitation phase corresponds to the testing phase, so as to identify a successful drug as soon as possible, in order to obtain a larger revenue, which is decreasing in time due to active patent life loss, potential market share reduction, etc. These optimal testing schedules already provide the decision maker with an implicit futility stopping criterion, as the optimal decision will be to not recruit any patients in a given period if the likelihood that the drug will be successful is lower than some optimally-determined threshold. On the other hand, to ensure statistical significance and power, the optimization model’s success criterion corresponds to the triangular test’s upper stopping boundary, which represents the minimum number of successes needed in the experimental group to submit the drug for FDA approval and future commercialization.

This research makes a contribution by quantifying the benefits of optimal testing schedules for parallel group sequential adaptive clinical trials. The triangular test and other sequential methodologies provide stopping criteria that respond to statistical significance and power requirements, but do not prescribe a testing schedule (i.e., the proportion of the full patient sample size that needs to be tested in each period), nor take financial considerations (e.g., trial budget) into account, and do not allow for learning. In contrast, optimal testing schedules can be modified during the trial’s course, based on financial constraints as well as newly acquired information on the experimental drug’s efficacy.

Our numerical study generates key insights to different stakeholders on the following research

questions: **1)** How do fixed testing schedules compare to optimal, adaptive testing schedules in parallel group sequential adaptive trials, in terms of drug misclassification risk, time-to-market, and the firm's profit?, **2)** What are the benefits of learning (i.e., incorporating newly acquired information on the experimental drug's efficacy), adaptive schedules, and utilizing an optimal stopping criteria?, and **3)** How do optimistic and pessimistic priors on the drug's efficacy alter the performance metrics in parallel group fixed sample size and parallel group sequential adaptive trials?.

As stated above, the optimal resource allocation strategy provides both a testing schedule and a futility stopping criterion, while incorporating financial considerations and allowing for learning. Learning happens through a Bayesian updating mechanism that utilizes the new information to adjust the decision maker's estimate on the drug's success probability. We show how the optimal resource allocation strategy can be combined with statistical stopping methodologies such as the triangular test, in order to fulfill statistical significance and power requirements when classifying the experimental drug as successful. We show that the proposed augmented methodology offers substantial benefits to the trial's sponsor, the trial participants, and the patients in need of the new therapy. With this work, we make a contribution to the research stream on sequential exploration and sequential investment problems, as well as to healthcare policy literature on parallel group clinical trials. Our insights can also be helpful outside the pharmaceutical R&D context, where other problems of sequential decision-making under uncertainty are present.

The remainder of this paper is organized as follows. Section 2 describes the methodology of the triangular test, and provides an overview of the related literature on dynamic programming and information acquisition. Section 3 details our information updating mechanism, discusses our stochastic dynamic programming model and the resulting optimal resource allocation strategy, and describes the application of our optimization model to a parallel group sequential clinical trial with binary response. In Sections 4 and 5, respectively, we perform a numerical study based on realistic data, and provide insights based on our results. Finally, we provide our conclusions and suggest future research directions in Section 6.

4.2 Related Literature

This section is divided into two parts. In the first part, we describe the methodology of the triangular test as applied to sequential clinical trials with binary response. In the second part, we provide a brief overview of the literature on dynamic programming methodologies focusing on information acquisition under uncertainty.

4.2.1 The Triangular Test and Sequential Trials with Binary Response

In the following, we provide an overview of the main characteristics of parallel group sequential trials with binary response, and discuss the triangular test's methodology as described in Whitehead [109]. We first discuss the general terminology used throughout the paper.

In sequential clinical trials, the efficacy of a control (i.e., a standard therapy or a placebo) is compared with the efficacy of the candidate drug, based on an appropriately selected measure of the treatment difference between these two therapies, θ . In a parallel group study, each recruited patient is randomly assigned to either the control or the experimental group, and is administered only one of these two therapies. A binary response applies to trials in which the patient response is binary: success or failure. For example, acute bleeding may be controlled within a specific time frame or not, the patient may survive after receiving a certain therapy or not, or blood pressure may fall to a specified level or not. In a sequential clinical trial, the trial continuation or termination decision is made periodically based on the accumulated results thus far from the trial. Specifically, because of practical constraints, the trial's results may not be evaluated on a continuous basis, but rather at a few specific decision epochs during the study; such trials are commonly referred to as group sequential trials.

The vast majority of sequential procedures belong to one of two main categories: the boundaries approach and the repeated significance testing approach. In the repeated significance testing approach, the decision maker performs a series of conventional statistical analyses with significance levels that are adjusted based on the patient sample size analyzed at each interim analysis point. These methodologies are based solely on the p-value (i.e., the probability of finding the observed, or more extreme, results when the null hypothesis of a study question is true), and do not consider the expected value of the treatment difference between the experimental and control

therapies. The triangular test belongs to the boundaries approach, as its stopping boundaries incorporate the expected treatment difference, and maintain both statistical significance and power levels. Unlike most repeated significance testing methodologies, which only contemplate stopping for benefit, the triangular test includes an upper boundary (to stop for benefit) and a lower boundary (to stop for futility). Furthermore, the triangular test has been compared with other sequential stopping methodologies, and has been shown to offer the largest reduction in sample size [93], and to provide an early determination of the treatment effect, whatever it might be [44].

The boundaries approach is commonly used to design sequential clinical trials, i.e., to make a decision regarding the null hypothesis, $H_0 : \theta = 0$, with a desired statistical significance level α , and a desired statistical power $1 - \beta$. This approach relies on two metrics: the test statistic Z , which measures the cumulative advantage of the experimental treatment over the control (i.e., the efficient score for θ under the null hypothesis), and the test statistic V , which indicates the amount of information about θ contained in Z (i.e., Fischer's information). Variables Z and V are plotted against each other until certain stopping boundaries are crossed. Specifically, if the Z statistic's value crosses one of the boundaries, then the trial stops and the null hypothesis of no difference between the two therapies is rejected; if the Z statistic remains within the area delimited by the boundaries, then the results are deemed insufficient to arrive at a conclusion, the trial continues to the next decision epoch [99], and the same decision process is repeated.

In what follows, we describe the triangular test in the context of a parallel group study with binary response. Consistent with the statistical literature (e.g., Whitehead [109], Wald [108]), we utilize the following probability model. We assume that patient responses are binary and are independent from each other, and that treatment on each patient in the experimental and control groups will be successful with probabilities p_E and p_C , respectively. Since p_E is unknown, the decision maker has an estimate thereof, \hat{p}_E . In contrast, p_C is assumed to be known with certainty, because the control corresponds to either a placebo or standard therapy, whose treatment effects have already been studied. Let N denote the total number of patients recruited, with N_E patients in the experimental group, and N_C patients in the control group, i.e., $N = N_E + N_C$. Let S_k and F_k respectively denote the number of successful and unsuccessful responses in N_k patients, for $k \in \{E, C\}$; see Table 4.1. for a summary of the notation.

The desired treatment difference between the control and the experimental therapies is mea-

Table 4.1: Summary of the notation used in a parallel group clinical trial with binary response
(Adapted from Whitehead [109])

Therapy	Control	Experimental	Total
Number of responses	N_C	N_E	N
Number of successes	S_C	S_E	S
Number of failures	F_C	F_E	F

sured by the log-odds ratio $\theta_R = \ln \left\{ \frac{\hat{p}_E(1 - p_C)}{p_C(1 - \hat{p}_E)} \right\}$. The corresponding Z and V statistics are then derived as follows:

$$Z = \frac{(N_C S_E - N_E S_C)}{N}, \quad (4.1)$$

$$V = \frac{N_E N_C S F}{N^3}. \quad (4.2)$$

At the trial planning stage, it is commonly assumed that N is large, thus $S \approx N\bar{p}$ and $F \approx N(1 - \bar{p})$, where $\bar{p} = \frac{\hat{p}_E + p_C}{2}$. For an equal split of the number of patients among the experimental and control groups (i.e., $N_E = N_C = N/2$), the full sample size, N , is obtained as follows:

$$N = \frac{4\hat{V}}{\bar{p}(1 - \bar{p})}, \quad (4.3)$$

where \hat{V} corresponds to the ratio of the value of the information required by a fixed sample test satisfying the desired significance, α , and power, β , to θ_R^2 . The value of the information required by fixed sample tests for various values of α and $1 - \beta$ are tabulated, see e.g., [109].

Given the full sample size, N , the statistics Z and V are utilized to derive the triangular test's upper and lower boundaries, which correspond to the lines defined by $Z = a + bV$ and $Z = -a + 3bV$, respectively, where the parameters a and b are given by:

Case where $\beta = \alpha$:

$$a = \frac{2}{\theta_R} \ln \left(\frac{1}{2\alpha} \right), \quad (4.4)$$

$$b = \frac{1}{4} \theta_R. \quad (4.5)$$

Case where $\beta \neq \alpha$:

θ_R in Eqs. (4.4) and (4.5) needs to be replaced with a corrected value, θ'_R , given below [6]:

$$\theta'_R = \theta_R \left[\frac{2\Phi^{-1}(1 - \alpha)}{\Phi^{-1}(1 - \alpha) + \Phi^{-1}(1 - \beta)} \right], \quad (4.6)$$

where $\Phi(x)$ denotes the standard normal cumulative distribution function.

When decision-making (monitoring) is performed on a continuous basis (i.e., after each patient observation), the stopping mechanism is fully characterized by the linear upper and lower boundaries, as previously explained. In practice, however, decision-making occurs in discrete times, and this requires an adjustment to the triangular test, through modification of the upper and lower stopping boundaries, resulting in “inner boundaries,” often referred to as “Christmas tree boundaries” because of their shape; see Fig. 4.1 for linear boundaries and the corresponding inner boundaries.

In the discrete monitoring case, the stopping rule is as follows: at the i^{th} interim analysis, the statistics Z_i and V_i are calculated using Eqs. (4.1) and (4.2) based on data observed thus far, as well as an upper stopping limit, u_i , and a lower stopping limit, l_i , where $u_i > l_i$:

$$u_i = a + bV_i - 0.583\sqrt{(V_i - V_{i-1})}, \quad (4.7)$$

$$l_i = -a + 3bV_i + 0.583\sqrt{(V_i - V_{i-1})}, \quad (4.8)$$

where the term $0.583\sqrt{(V_i - V_{i-1})}$ is a correction factor for the discrete monitoring case, with $V_0 = 0$.




Figure 4.1: Stopping boundaries for the triangular test under discrete decision-making (Adapted from Whitehead [109])

If $Z_i \geq u_i$, then the trial stops and the null hypothesis, of no difference between therapies, is rejected, with the conclusion that the experimental drug is more effective than the control. If $Z_i \leq l_i$, then the trial stops with two possible conclusions, depending on the value of V_i : 1) that the experimental drug is less effective than the control, if the i^{th} interim analysis takes place after a small

patient sample has been observed (shown with the dashed part of the lower stopping boundary in Fig. 4.1), or 2) that the therapies are equivalent, thus, with the acceptance of the null hypothesis, if the i^{th} interim analysis is performed after a larger patient sample has been observed (shown with the solid part of the lower stopping boundary in Fig. 4.1).

4.2.2 Dynamic Programming and Information Acquisition

This research fits within the stream of dynamic programming models that incorporate newly acquired information to tackle uncertainty. For instance, McCardle [63] presents a dynamic programming model in which in each period the firm updates its estimate on a new technology's profitability using a Bayesian framework, and makes a decision on whether or not to adopt the new technology. Within the healthcare application arena, an example is given by Kornish and Keeney [54], who derive optimal strategies for the selection of different strains of the influenza virus for vaccine production, under uncertainty on the vaccine's efficacy.

Our work is also positioned within the sequential exploration and sequential investment decision-making literature. For example, within a sequence of a number of trials, Kolonko and Benzing [52] examine a generalization of the one-armed bandit problem to study the selection of one of two Bernoulli experiments for each trial, under uncertainty on one of the experiment's success probability. Within the context of an oil and gas exploration problem, Bickel and Smith [14] build a dynamic programming model to derive optimal sequential exploration strategies considering dependent exploration prospects. Brown and Smith [15] extend this work and model sequential exploration problems as variations of the classical multi-armed bandit problem, under an exploration versus exploitation trade-off somewhat similar to the one that characterizes our problem.

In most sequential investment problems, there exists a trade-off between exploration activities, where information about the value of the investment prospects is gathered, and exploitation activities, which offer an actual payoff or reward. In our model, there exists a trade-off between deeming the candidate drug as effective/successful early on, thus realizing the revenue by commercializing it, and gathering information about the drug's probability of success, before committing a large budget to the trial. What sets our problem apart from traditional sequential investment problems is that there exists an overlap between the exploration and exploitation activities, because collecting

patient responses contributes to both the realization of the revenue, by contributing to the achievement of a specific number of successes, and the acquisition of information on the drug's efficacy.

4.3 The Optimization Model

In this section, we describe the decision problem, the information updating component of our optimization model, and our stochastic dynamic programming model, as applied to a parallel group sequential adaptive clinical trial, and discuss the characteristics of an optimal testing schedule.

4.3.1 The Decision Problem

In this paper, we apply and expand the stochastic dynamic programming model developed in Rojas-Cordova and Bish [92] to model a parallel group sequential trial with binary response, which consists of an experimental group and a control group. In this model, the firm's decision problem is to determine the optimal number of patients to recruit in each period (i.e., the optimal testing schedule), over the duration of a Phase 3 trial under uncertainty on the drug's success probability, p , so as to maximize the expected profit from the drug. We let N_t denote the number of patients to recruit in each period $t \in [1, \dots, T]$, which is split equally between the control group and the experimental group, respectively denoted by subscripts C and E , i.e., $N_{E_t} = N_{C_t} = N_{S_t} = \frac{N_t}{2}$.

The testing schedule is constrained by M , or the maximum number of patients the firm is able to recruit and test based on a budget of B (i.e., $M = \frac{B}{c}$, where c denotes the recruitment and testing cost per patient). The drug's success, based on the comparison between its efficacy and a control's, is assessed through an exogenously determined criterion, $\gamma_c(M)$, which is a non-decreasing function of M . If this criterion is met, then a time-dependent revenue is realized. The revenue for a successful drug in period t , g_t , is modeled as a strictly decreasing function in t to represent the potential reduction in active patent life and market share caused by commercialization delays. Consistent with the reality of drug R&D, we assume that the revenue for a successful drug in period t far exceeds the budget available for the corresponding phase of the trial, i.e., $g_t \gg B, \forall t \in [1, \dots, T]$.

Following the design of a parallel group sequential adaptive trial, we consider an experimental group and a control group. We assume that patient responses in both groups are binary and are inde-

pendent from each other. Each patient response in group k , where $k \in \{E \text{ (experimental)}, C \text{ (control)}\}$, is represented by a Bernoulli distribution with probability p_k . Given that the control consists of a standard therapy or a placebo, whose efficacy is well-understood, we assume that the control's success probability, p_C , is known to the decision maker. However, the experimental drug's success probability, p_E , is unknown, and needs to be estimated. Depending on the testing schedule, the number of successful observations (successes) in each group k in period t follows a binomial distribution with parameters N_{k_t} and $p_k, k \in \{E, C\}$. We let N_t^+ denote the difference between the number of successful observations in the experimental group and the number of successful observations in the control group, i.e., $N_t^+ = (S_{E_t} - S_{C_t})^+$, where S_{E_t} and S_{C_t} correspond to the number of successful observations in period t in the experimental and control groups, respectively.

In this setting, the probability p is given by $(p_E - p_C)^+$, and the decision maker updates her estimate on p, \hat{p}_t , at the beginning of each period t , based on the patient responses observed through the end of period $t - 1$. The decision maker will then utilize \hat{p}_t to determine the number of patients to recruit and observe in the current period (i.e., N_t , for $t = 1, 2, \dots, T$), constrained by $R_{t-1} \equiv M - \sum_{j=1}^{t-1} N_j$, or the remaining number of patients that can be recruited and observed in periods t, \dots, T . At the end of period t , the realization of the random variable N_t^+ is observed. If, at the end of period t , the drug is deemed successful, then the trial is terminated, and the firm realizes a one-time payoff of g_t . Otherwise, the firm continues with the next period of the trial, unless the trial's maximum duration (i.e., period T) or the maximum possible number of patients that can be recruited (M) is reached.

4.3.2 Information Updating

As previously stated, random variables S_{E_t} and S_{C_t} each follow a binomial distribution with parameters $\left(\frac{N_t}{2}, p_E\right)$, and $\left(\frac{N_t}{2}, p_C\right)$, respectively. Then, we approximate N_t^+ as Binomial $\left(\frac{N_t}{2}, p\right)$, $\forall t$, where the difference between the drug's and the control's success probabilities, $p = (p_E - p_C)^+$, is unknown to the decision maker. Therefore, the decision maker models p_t , the success probability for N_t^+ , using her estimate, $\hat{p}_t = \hat{p}_E - p_C$, which is updated at the beginning of each period $t, t = 1, \dots, T$, in accordance with the following stochastic process. We model the decision maker's initial estimate of p at the beginning of the trial (\hat{p}_1) as a Beta prior with parameters λ_0 and β_0 .

These parameters may represent preliminary data in the form of expert opinions, results from earlier trial phases, etc., and in case no preliminary data are available or used, we can set $\lambda_0 = \beta_0 = 1$. We define the stochastic processes $\{\gamma_t, t \geq 0\}$ and $\{\zeta_t, t \geq 0\}$ as follows:

$$\begin{aligned}\gamma_t &= \sum_{j=1}^t X_j = \gamma_{t-1} + X_t, \\ \zeta_t &= \sum_{j=1}^t \left(\frac{N_j}{2} - X_j \right) = \zeta_{t-1} + \left(\frac{N_t}{2} - X_t \right),\end{aligned}\tag{4.9}$$

where $X_t = (S_{E_t} - S_{C_t})^+$, and $\gamma_0 = \zeta_0 = 0$. Then, \hat{p}_t , the posterior distribution of p at the beginning of period t , follows a Beta distribution $(\gamma_{t-1}, \zeta_{t-1})$, which is used by the firm in the value function in period t . Thus, the posterior predictive distribution of the difference between the number of successes in the experimental group and the number of successes in the control group in any given period follows a Beta-binomial distribution (corresponding to a binomial likelihood and a Beta prior). The event that the drug is successful at the end of period t is given by $\{\gamma_t = \sum_{j=1}^t X_j \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\}$, where $\gamma_c(M)$ is an exogenously determined parameter.

4.3.3 The Stochastic Dynamic Programming Model

Given the sequence of events and the stopping criteria explained in Section 4.3.1, we formulate the firm's testing schedule (i.e., resource allocation decision) in a sequential adaptive clinical trial as a stochastic dynamic programming problem, with the following value function in period $t = 1, 2, \dots, T$:

$$V_t(\gamma_{t-1}, R_{t-1} | \gamma_{t-1} < \gamma_c(M)) = \max_{0 \leq N_t \leq R_{t-1}} -cN_t + \begin{cases} g_t, & \text{if } \{\gamma_t \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \\ E[V_{t+1}(\gamma_{t-1} + N_t^+, R_{t-1} - N_t) | \gamma_t < \gamma_c(M)], & \text{if } \{\gamma_t < \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \end{cases},\tag{4.10}$$

and $V_{T+1}(\gamma_T, R_T | \gamma_T < \gamma_c(M)) = 0, \forall \gamma_T, \forall R_T$,

where the event $\{\gamma_t \geq \gamma_c(M) | \gamma_{t-1} < \gamma_c(M)\} \equiv \{N_t^+ \geq \gamma_c(M) - \gamma_{t-1} | \gamma_{t-1} < \gamma_c(M)\}$. We denote by $N_t^*, t = 1, \dots, T$, the optimal solution to Eq. (4.10). We assume that the patient recruitment and observation rate is not constraining, or equivalently, that the decision maker can choose the number of patients to recruit and observe in each period throughout the trial, subject to the budget available at that point.

4.3.4 Structural Properties of an Optimal Testing Schedule

In order to study the optimal testing schedule, we first provide a brief overview of the structural properties of an optimal solution, and refer the reader to [92] for details.

As explained above, at the beginning of period t , the firm observes γ_{t-1} and R_{t-1} , updates its estimate \hat{p}_t , and determines its resource allocation in period t , $N_t = 2N_{S_t}$. At the end of period t , the realization of the random variable $N_t^+ \sim \text{Binomial}\left(\frac{N_t}{2}, p\right)$ is observed. Therefore, the firm relies on the stochastic process $\{\gamma_t, t \geq 0\}$ to update the posterior predictive probabilities that will be used in the value function. $\{\gamma_t, t \geq 0\}$ is a time non-homogeneous Markov process with Beta-binomial transition probabilities that change over time. As previously discussed, it follows that the decision maker's estimate \hat{p}_t is modeled with a Beta distribution with parameters γ_{t-1} and ζ_{t-1} . Thus, the decision maker's posterior predictive probability in period t is a Beta-binomial distribution with parameters $\frac{N_t}{2}$, γ_{t-1} , and ζ_{t-1} (i.e., a compound distribution given the Binomial likelihood and the Beta prior).

Lemmas 1 and 2 of Chapter 2 continue to hold in the setting studied in the current chapter, that is, $\Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma)$ is increasing in each of γ and q , for $\gamma: \gamma \in \mathbb{Z}^+, \gamma < \gamma_c(M), q = 1, \dots, t-1$. In other words, the more successes are observed up to time $t-q$ in the experimental group, or the larger the value of q (i.e., the earlier γ unique successes are obtained), the higher the perceived likelihood of success in period t will be. Similarly, Lemmas 3 and 4 of Chapter 2 also continue to hold, consequently, if for any $0 < t < T$, $N_t^* = 0$, then $N_{t+1}^* = N_{t+2}^* = \dots = N_T^* = 0$, and if for any $0 < t < T$, $R_{t-1} < \gamma_c(M) - \gamma_{t-1}$, then $N_t^* = N_{t+1}^* = \dots = N_T^* = 0$. If the optimal strategy is to not recruit and test any patients in any given period t , then patient recruitment and testing will not be reinitiated in future periods, and the trial will be terminated permanently for benefit or futility reasons. Likewise, if in any period the available budget is insufficient to recruit and test the number of patients needed to achieve the required success criterion, then the optimal strategy is to terminate the trial for futility.

4.3.5 An Optimal Resource Allocation Strategy

In the following, we use the superscripts A and F to denote adaptive and fixed duration trials. The optimal resource allocation for a sequential adaptive trial in period $t \in [1, \dots, T]$ follows a threshold

policy:

if $\gamma_{t-1} \geq \gamma_c(M)$, or $R_{t-1} < \gamma_c(M) - \gamma_{t-1}$, then $N_t^{*A}(\gamma_{t-1}, R_{t-1}) = 0, \forall \gamma_{t-1}, R_{t-1}$;

otherwise:

$$N_t^{*A}(\gamma_{t-1}, R_{t-1}) = \begin{cases} 0, & \text{if } \{\gamma_{t-1} \leq \underline{\gamma}_{t-1}^A\} \equiv \{E[\hat{p}_t] \leq \underline{p}_t^A\} \\ > 0, & \text{if } \{\gamma_{t-1} > \underline{\gamma}_{t-1}^A\} \equiv \{E[\hat{p}_t] > \underline{p}_t^A\} \end{cases}, \quad (4.11)$$

where $E[\hat{p}_t] = \frac{\gamma_{t-1}}{\gamma_{t-1} + \zeta_{t-1}}$, for $t \in [2, \dots, T]$, and $E[\hat{p}_1] = \frac{\lambda_0}{\lambda_0 + \beta_0}$. Further, the threshold, $\underline{\gamma}_{t-1}^A$, is a function of R_{t-1} , and decreases in each of γ_{t-1} and R_{t-1} .

The optimal resource allocation for a fixed duration trial with T periods also follows a threshold policy:

$$N^{*F}(\lambda_0, \beta_0, M) = \begin{cases} 0, & \text{if } \{\lambda_0 \leq \underline{\lambda}_0^F\} \equiv \left\{E[\hat{p}_1] = \frac{\lambda_0}{\lambda_0 + \beta_0} \leq \underline{p}^F\right\} \\ > 0, & \text{if } \{\lambda_0 > \underline{\lambda}_0^F\} \equiv \{E[\hat{p}_1] > \underline{p}^F\} \end{cases},$$

where the threshold $\underline{\lambda}_0^F$ decreases in each of λ_0 and M , and the threshold $\underline{p}^F = \underline{p}_T^A|_{(\gamma_{T-1}=0, R_{T-1}=M)}$. A description of how the resource allocation thresholds, $\underline{p}_t^A, t \in [1, \dots, T]$, change over time can be found in [92].

4.4 Numerical Analysis and Insights

We carry out a series of numerical experiments to answer the following research questions:

1. How do fixed testing schedules compare to optimal, adaptive testing schedules in parallel group sequential adaptive trials, in terms of drug misclassification risk, time-to-market, and the firm's profit?
2. What are the benefits of learning (i.e., incorporating newly acquired information on the experimental drug's efficacy), adaptive schedules, and utilizing an optimal stopping criteria?
3. How do optimistic and pessimistic priors on the drug's efficacy alter the performance metrics in parallel group fixed sample size and parallel group sequential adaptive trials?

Towards this end, we solve the resource allocation problem optimally for different problem instances, using parameter values representative of industry data and industry practice, as described

in the following sections.

4.4.1 Experimental Design

To get insight on the effect of updates to the drug's probability of success on the optimal testing schedule and on the performance metrics selected, we derive the optimal strategy in sequential adaptive trial settings with one and two interim analyses (i.e., two- and three-period problems, respectively), as well as in a fixed sample size (single-period) trial. According to the relevant literature, a clinical trial with two interim analyses shows significant advantage over a fixed duration trial in lowering the average number of patients treated with an inferior therapy [86]. Further, the gains from interim analyses are the greatest when going from one to two interim analyses, and a trial with more than five interim analyses is unlikely to generate practical, statistical, or ethical benefits, unless an extremely large difference between the efficacy of the candidate drug and the efficacy of the control is expected [64, 85]. The performance measures we consider include: 1) the proportion of time the drug is misclassified, 2) the expected time-to-market, and 3) the firm's expected profit.

We assume that all trials have the same maximum duration (i.e., the length of the fixed sample size trial), which is divided into two or three periods of equal length, based on the setting. The decision maker knows that the control has a success probability of $p_C = 0.50$, and expects the experimental drug to have a success probability of \hat{p}_E . In what follows, we describe the calculations for the base case scenario, of $\hat{p}_E = 0.66$ (corresponding to the case where the decision maker considers an observed difference of at least $\hat{p}_E - p_C = 0.16$ to be satisfactory). Similar calculations hold in scenarios with different values of \hat{p}_E . We utilize a statistical significance level $\alpha = 0.05$, and a statistical power level $1 - \beta = 0.90$, which are commonly used in practice. We assume that the patient recruitment rate is not limiting, that every recruited patient will generate either a successful or an unsuccessful response, and that this response is available soon after the therapy has been administered.

To evaluate the effects of optimistic and pessimistic priors on the drug's efficacy, we either: 1) vary the decision maker's expectation on the success probability of the experimental drug $\hat{p}_E \in \{0.56, 0.66, 0.76\}$, where these scenarios correspond to a pessimistic prior, perfect information,

and an optimistic prior, respectively, or 2) vary the experimental drug's true success probability $p_E \in \{0.56, 0.66, 0.76\}$, and keep the control's known success rate (p_C) and the decision maker's expectation on the success probability of the experimental drug (\hat{p}_E) unchanged at 0.50, and 0.66, respectively. In 2), the first scenario ($p_E - p_C = 0.06$) corresponds to an optimistic prior, while the second ($p_E - p_C = 0.16$) and third ($p_E - p_C = 0.26$) scenarios represent the perfect information and pessimistic prior cases. We use setting 1) only in Section 4.4.3, and setting 2) in all the remaining sections.

We use the following set of parameters: a per patient cost (c) of \$50,000 [82], which is representative of the total per patient cost in a Phase 3 oncology trial, and an available budget (B) of \$10 million. We consider a time-decreasing revenue with a per period reduction rate of 5%, with a final revenue of $g_3 = \$1$ billion. These figures are selected based on global net sales per novel active substance (i.e., a biologic or molecular entity or combination product in which one or more elements had not previously received FDA approval) in the 2005 - 2009 launch cohort [7], and sales forecasts for drug Keytruda (Pembrolizumab) [70].

To determine the full sample size, N , which represents the number of patients recruited and observed in a fixed sample size trial as well as the *maximum* sample size available in a sequential adaptive setting, we calculate the desired treatment difference, θ_R , measured by the log-odds ratio, and the value of the information required, $\theta_R^2 \hat{V}$, (see Section 4.2.1), leading to $\theta_R = 0.663$, $\theta_R^2 \hat{V} = 10.51$ [109], and $\hat{V} = 23.889$, hence, $N = 392.259$ from Eq. (4.3). This sample size, which fulfills statistical significance and power requirements, is slightly smaller than the maximum number of patients that the firm can afford to recruit and analyze, $M = \frac{B}{c} = 400$. Therefore, we consider a full sample size of $M = 400$ throughout this numerical study.

The success criterion, $\gamma_c(M)$, representing the minimum number of successes $(S_E - S_C)^+$, corresponds to the triangular test's upper stopping boundary, $Z = a + bV$. Following the triangular test's methodology described in Section 4.2.1, and using $\theta_R = 0.663$, $\alpha = 0.05$, and $\beta = 0.10$, the corrected treatment difference is given by $\theta'_R = 0.746$, with the parameters $a = 6.176$ and $b = 0.186$, i.e., the upper stopping boundary is then given by $Z = 6.176 + 0.186V$. Next, using Eq. (4.1) and the upper stopping boundary derived, we determine the required number of successes in the experimental group, S_E , that would be needed to classify a fixed sample size trial, with a total of $M = 400$ patients, as successful, with $N_S = 200$ patients in each of the experimental and control

groups, and $E[S_C] = N_S p_C = 100$ successes in the control group. Based on the resulting number of experimental group successes $S_E = 121.564 \approx 122$, we then determine the minimum number of successes required, $\gamma_c(M) = S_E - E[S_C] = 122 - 100 = 22$. Since the triangular test's upper boundary is an increasing function of V , hence, an increasing function of the observed number of patients (see Figs. 4.1 and 4.2), we utilize Eq. (4.2), $N_S = N_E = N_C = 200$, $S = S_E + E[S_C] = 222$, and the full sample size of $M = 400$ to determine $V = 24.697$ and the corresponding benefit threshold $u_1 = 6.176 + 0.186V = 10.770$ for the first period; this ensures that the success criterion used in the stochastic dynamic programming model, $\gamma_c(M) = 22$, and the triangular test's upper stopping boundary, are equivalent. The triangular test's benefit thresholds for the subsequent periods are determined using Eq. (4.7).




Figure 4.2: Triangular test's stopping boundaries and the success criterion used in numerical study

We solve the stochastic dynamic programming problem and obtain the optimal resource allocation for a fixed sample size trial and each sequential adaptive trial setting. We first compute N_1^* (i.e., the number of patients to be recruited and observed in the first period), and generate a random realization of the number of successes in the experimental group, N_1^+ , from a Binomial distribution with parameters N_1^* and $p = p_E - p_C$. Given the random realization of N_1^+ and the leftover budget, $B - cN_1^*$, we then optimize the number of patients recruited and tested in the remaining period(s), that is, N_2^* for the single interim analysis setting, and N_2^* and N_3^* for the two interim analyses setting, depending on the number of successes realized and the leftover budget. In the fixed sample size trial, we determine whether or not the drug is successful based on the realization of N_1^+ . We

generate 3,000 replications for each simulation scenario.

4.4.2 Description of Fixed Testing Schedules and Stopping Boundaries

For each trial setting, we compare the performance of the optimal, adaptive testing schedules, corresponding to the optimal resource allocations, paired with the success criterion of $\gamma_c(M) = 22$, against the performance of a series of pre-defined fixed testing schedules, paired with the triangular test's upper and lower stopping boundaries, and the corresponding benefit and futility thresholds; see Table 4.2. The fixed testing schedules are described in terms of the percentage of the full sample size that is recruited and analyzed in each period. For example, the aggressive fixed testing schedule in the single interim analysis trial consists of recruiting and analyzing 70% of the full sample size ($N_1 = 280$ patients) in the first period, and the remaining 30% ($N_2 = 120$ patients) in the second period. To determine the values of the triangular test's statistics Z and V (see Section 4.2.1), we compute S_{E_t} and S_{C_t} utilizing the randomly generated number of unique successes in the experimental group in period t , N_t^+ , and the following expressions: $E[S_{C_t}] = p_c N_t$, and $S_{E_t} = E[S_{C_t}] + N_t^+$.

Table 4.2: Fixed testing schedules and stopping boundaries

Setting	Testing schedule (Percent allocation to each period (%))			Threshold in period t	
	Aggressive	Neutral	Conservative	Benefit	Futility
Fixed sample size	100	100	100	$u_1 = 10.770$	$l_1 = -6.176 + 0.559V$
1 interim analysis	70/30	50/50	30/70	$u_1 = 10.770$ $u_2 = 6.176 + 0.186V_2 - 0.583\sqrt{(V_2 - V_1)}$	$l_1 = -6.176 + 0.559V_1 + 0.583\sqrt{V_1}$ $l_2 = -6.176 + 0.559V_2 + 0.583\sqrt{(V_2 - V_1)}$
2 interim analyses	50/25/25	30/40/30	25/25/50	$u_1 = 10.770$ $u_2 = 6.176 + 0.186V_2 - 0.583\sqrt{(V_2 - V_1)}$ $u_3 = 6.176 + 0.186V_3 - 0.583\sqrt{(V_3 - V_2)}$	$l_1 = -6.176 + 0.559V_1 + 0.583\sqrt{V_1}$ $l_2 = -6.176 + 0.559V_2 + 0.583\sqrt{(V_2 - V_1)}$ $l_3 = -6.176 + 0.559V_3 + 0.583\sqrt{(V_3 - V_2)}$

In the pessimistic prior and perfect information scenarios, misclassification corresponds to of deeming the experimental drug futile, or ineffective, whereas misclassification in the optimistic prior scenario consists of deeming the drug successful, or effective. Notice that in the optimistic prior scenario, the expected time-to-market and expected profit measures are not relevant, as the drug's true efficacy does not meet the minimum required success probability of 0.66. Therefore, in this scenario, we use the opportunity cost for this trial, i.e., we assume that the resources that are to be allocated to the experimental drug, which is not effective, can be reallocated to a different, effective drug, immediately after the first experimental drug is correctly classified as ineffective,

and calculate the firm’s expected profit by multiplying the corresponding reward g_t by the proportion of time the first drug is correctly classified as futile. This value represents the payoff realized by the firm from the commercialization of the second drug, resulting from resource reallocation in light of the first drug’s trial results.

4.4.3 Optimal Resource Allocation

We first compute the optimal resource allocation strategy in a fixed sample size trial and the two sequential adaptive trial variations in three different scenarios: perfect information ($\hat{p}_E = p_E = 0.66$), optimistic prior ($\hat{p}_E = 0.76$), and pessimistic prior ($\hat{p}_E = 0.56$). Table 4.3 reports the optimal number of patients to be recruited and analyzed in the first period (N_1^*), which also represents the overall number of patients observed in the fixed sample size trial.

Table 4.3: Optimal first period resource allocation (N_1^*) for the different trial settings

Setting	\hat{p}_E		
	0.56	0.66	0.76
Fixed sample size	0	400	304
1 interim analysis	10	320	250
2 interim analyses	10	290	230

Consistent with the findings of Rojas-Cordova and Bish [92], the optimal resource allocation strategy in sequential adaptive trials suggests to “explore” or “test out” the drug when the decision maker’s prior is pessimistic. In the fixed sample size trial, a pessimistic prior leads to no investment and an automatic drug misclassification, as the drug’s true success probability satisfies the minimum efficacy level required. Across all trial settings, the budget allocations for both pessimistic and optimistic priors are lower than those in the perfect information scenario; in adaptive trials, however, the decision maker has the opportunity to correct her prior belief and re-optimize her investment decision in the subsequent periods.

4.4.4 Performance of Fixed and Optimal Testing Schedules under Perfect Information

In this section, we analyze the performance of an optimal, adaptive testing schedule, measured by the experimental drug misclassification risk, the expected time-to-market, and the firm's expected profit, and compare it against the performance of aggressive, neutral, and conservative fixed testing schedules, which are paired with the triangular test's stopping boundaries (see Table 4.2). The decision maker has perfect information on the experimental drug's efficacy, thus, her estimate on the drug's success probability is given by $\hat{p}_E = p_E = 0.66$. Tables 4.4 - 4.6 report the selected performance measures of each testing schedule.

Table 4.4: Drug misclassification risk (%) for the different trial settings under perfect information ($p_E = 0.66$)

Setting	Testing schedule			
	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	1.37	–	1.37	–
1 interim analysis	1.80	32.20	52.83	4.90
2 interim analyses	1.67	83.53	62.57	30.13

Table 4.5: Expected time-to-market (in number of periods) for the different trial settings under perfect information ($p_E = 0.66$)

Setting	Testing schedule			
	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	3.00	–	3.00	–
1 interim analysis	1.75	1.73	2.77	3.00
2 interim analyses	1.36	1.75	2.20	2.97

Table 4.6: Expected profit (in \$ million) for the different trial settings under perfect information ($p_E = 0.66$)

Setting	Testing schedule			
	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	966.33	–	966.33	–
1 interim analysis	1,031.67	707.54	458.20	931.00
2 interim analyses	1,054.94	156.72	371.83	679.90

Since the fixed sample size trial involves a single period (i.e., no interim analyses), and the optimal resource allocation is equal to the fixed sample size of 400 patients (see Table 4.3), the

optimal and fixed testing schedules perform equivalently in each of the aspects studied. In terms of drug misclassification and expected time-to-market, the optimal testing schedule performs better than all of the fixed testing schedules in both of the sequential trials studied (Tables 4.4 and 4.5). The optimal testing schedule shows the highest improvement in drug misclassification risk, with respect to the fixed testing schedules, in the sequential setting with two interim analyses.

The percentages shown in Table 4.4 include the fraction of the time that the trial ends with inconclusive results under the triangular test's stopping boundaries. Recall that these stopping boundaries delimit three regions: stopping for benefit, stopping for futility, and inconclusive testing. Therefore, there exists the possibility that the trial's result remains inconclusive after the firm exhausts the financial resources available for the experimental drug's testing, and the trial reaches its maximum duration. In contrast, the optimal testing schedule guarantees a conclusive classification of the experimental drug (i.e., as a success or a failure) by the time the available budget is fully consumed, and in a majority of the cases, does allow the decision maker to reach to a final classification before the trial reaches its maximum duration. Although the aggressive testing schedule offers a slightly shorter time-to-market than the optimal testing schedule in the single interim analysis setting, it leads to a correct drug classification only 67.80% of the time, thus, offering a much lower expected profit.

In terms of the expected profit, the optimal testing schedules outperform each of their fixed counterparts in both sequential settings (Table 4.6). We observe the largest improvements in the setting with two interim analyses. In fact, the benefits from learning and updating the decision maker's estimate on the drug's success probability accordingly are most evident in this setting, as the exploration versus exploitation trade-off is taken into account in the resource allocation decision, and this decision can be re-optimized during each of the interim analyses based on the newly acquired information.

4.4.5 Fixed and Optimal Testing Schedules under Optimistic and Pessimistic Priors

In this section, we expand the previous analysis to optimistic and pessimistic priors, which respectively correspond to overshooting the prior ($\hat{p}_E > p_E$) and undershooting the prior ($\hat{p}_E < p_E$).

Table 4.7: Drug misclassification risk (%) for the different trial settings under optimistic and pessimistic priors (* = fraction of time where the final trial results are inconclusive)

Setting	$p_E = 0.56$				$p_E = 0.76$			
	Optimal	Aggressive	Neutral	Conservative	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	0.37	–	0.37/14.63*	–	0.00	–	0.00	–
1 interim analysis	0.40	0.07/14.13*	0.00/30.63*	0.90/13.30*	0.00	0.00	0.13	0.00
2 interim analyses	0.60	0.00/55.00*	0.00/99.20*	0.07/71.80*	0.00	2.90	1.20	0.07

Table 4.8: Expected time-to-market (in number of periods) for the different trial settings under optimistic and pessimistic priors (NA=not applicable, as the drug is futile)

Setting	$p_E = 0.56$				$p_E = 0.76$			
	Optimal	Aggressive	Neutral	Conservative	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	NA	–	NA	–	3.00	–	3.00	–
1 interim analysis	NA	NA	NA	NA	1.50	1.50	1.73	2.93
2 interim analyses	NA	NA	NA	NA	1.00	1.15	1.98	2.63

Table 4.9: Expected profit (in \$ million) for the different trial settings under optimistic and pessimistic priors (* = assuming the resources are reallocated to an effective drug)

Setting	$p_E = 0.56^*$				$p_E = 0.76$			
	Optimal	Aggressive	Neutral	Conservative	Optimal	Aggressive	Neutral	Conservative
Fixed sample size	976.33	–	830.00	–	980.00	–	980.00	–
1 interim analysis	1,051.91	912.07	712.90	838.00	1,064.31	1,066.19	1,055.12	984.34
2 interim analyses	1,071.98	483.11	-11.53	261.33	1,093.48	1,056.59	1,027.10	1,002.56

Across all trial settings, fixed testing schedules translate into high inconclusive testing risks when overshooting the prior ($p_E = 0.56$ in Table 4.7). Unlike drug misclassification, inconclusive testing does not involve false positives in this case, but renders the invested resources useless, because, without a definitive classification, the drug cannot be submitted for FDA approval, nor can the resources be reallocated to a different, more promising drug. The risk of inconclusive testing can be as high as 99.20%, when using a conservative testing schedule in a sequential trial with two interim analyses. This represents the firm's least desirable scenario, as in 99% of the time the available budget is exhausted and the trial reaches its maximum duration, yet no actionable information on the drug's efficacy is available to decision makers. The optimal testing schedules perform significantly better than this and other fixed schedules across all trial settings. When undershooting the prior ($p_E = 0.76$ in Table 4.7), the risk of misclassifying the drug as futile is inherently low, however, unlike the optimal schedules, the fixed testing schedules suffer from misclassification risk in the setting with two interim analyses.

The expected time-to-market is not applicable in the optimistic prior scenario, as the exper-

imental drug's true efficacy does not meet the minimum requirement. In the pessimistic prior scenario and across all trial settings, the optimal testing schedules involve the shortest possible times-to-market, and outperform aggressive, neutral, and conservative testing schedules (Table 4.8). In the sequential adaptive settings, for both fixed and optimal testing schedules, the expected times-to-market under perfect information (see Table 4.5) are longer than under a pessimistic prior ($p_E = 0.76$ in Table 4.8), because the drug's true success probability is larger in the second case, allowing for an earlier classification and an expedited product launch.

Moving on to expected profits, we first discuss the pessimistic prior scenario ($p_E = 0.76$ in Table 4.9). In the fixed sample size setting, fixed and adaptive schedules perform equivalently. However, in the three-period sequential adaptive setting, as a result of low misclassification rates and short times-to-market, the optimal testing schedule outperforms all three fixed testing schedules. In the two-period setting, the optimal testing schedule outperforms the neutral and conservative schedules. The aggressive testing schedule offers an expected profit that is 0.18% higher than the one offered by an optimal testing schedule, because under the aggressive schedule, the number of patients recruited in the first period of the trial is slightly lower than that under an optimal schedule, and both aggressive and optimal schedules perform equally well in terms of misclassification risk and time-to-market. However, given that the decision maker estimates that the experimental drug's efficacy is barely sufficient to meet the minimum treatment difference required, selecting an aggressive testing schedule is unlikely.

As previously mentioned, the figures corresponding to the optimistic prior scenario in Table 4.9 represent the expected profit from the commercialization of a second, effective drug, which has been allocated the resources that were initially going to be invested in the development of a drug that proved to be ineffective. In this scenario, the optimal testing schedules greatly outperform the fixed testing schedules in all trial settings, and most significantly, in the two-interim analysis trial setting. Through a prompt identification of the ineffective drug, optimal testing schedules allow the firm to quickly reallocate the resources to a more promising drug candidate. The promising drug candidate's development can be started immediately, and upon a correct classification, be submitted for FDA approval and offer a high reward. One particular strength of optimal testing schedules is their ability to contribute to a prompt, conclusive, and correct drug classification. If an ineffective drug is under development, it is of vital importance to classify it as futile and terminate

testing as soon as possible, to stop development spending, reallocate resources to a different drug candidate, and prevent trial participants from receiving an ineffective treatment.

4.5 Discussion and Insights

Sequential adaptive clinical trials require decision makers to determine: 1) the number of patients to recruit and observe each period, and 2) stopping criteria to terminate the trial for futility or benefit reasons. In our numerical study, we examine optimal testing schedules for both fixed sample size- and sequential adaptive-parallel group clinical trials, which optimize the patient recruitment decision by incorporating financial constraints and the information about the experimental drug that has been collected thus far. We pair the triangular test's stopping mechanism, which is the best known and widely implemented design among all the sequential methodologies, with fixed testing schedules, reflecting different attitudes towards risk, and compare them with the optimal testing schedules. Given that the optimal testing schedule corresponds to a threshold policy, which implicitly includes a lower stopping boundary or a futility threshold, we fully address the stopping criteria decision by deriving a benefit threshold or success criterion from the triangular test's upper boundary.

The triangular test provides the decision maker with stopping criteria, but does not provide her with a methodology to determine a testing schedule. Towards this end, we examine aggressive, neutral, and conservative testing schedules, which can be chosen depending on the firm's selection criteria. Due to a lack of prescriptive methodologies for resource allocation, the firm might make a decision based on its strategic objectives, financial status, competitor actions, prior beliefs on the experimental drug's efficacy or on the drug's likelihood of getting FDA approval, attitude towards risk, among others. Given that some of these factors may or may not be indicative of the drug's true efficacy, which is unknown to the decision maker, the testing schedule chosen may translate into sub-optimal strategies. For instance, our numerical study shows that choosing a neutral testing schedule in a sequential trial with two interim analyses, based on an optimistic prior expectation on the experimental drug's efficacy, translates into sunk costs for the firm (see Table 4.9), as the firm is unable to correct its prior beliefs and correctly classify the drug as futile, exhausting its resources on an ineffective drug whose efficacy was estimated to be higher. An optimistic prior may also lead

the firm to choose an aggressive testing schedule over a neutral one; while the outcomes of this action are not as negative, they still lead to significantly lower gains than an optimal testing schedule.

The optimal resource allocation strategy guarantees a conclusive classification of the experimental drug by the end of the planning horizon (i.e., by the time the trial reaches its maximum duration, or by the time the full sample size has been tested). In fact, the drug may be classified as successful or futile *before* the full sample size is observed, which represents an additional benefit to both the firm and the patients.

The triangular test's stopping criteria respond to statistical significance and power requirements, but do not take financial constraints into account, nor incorporate newly acquired information on the drug's success probability. Given the uncertainty surrounding the experimental drug's true efficacy, testing schedules need to be flexible in order to be modified based on interim trial results. The optimal resource allocation strategy provides both a testing schedule and a futility stopping criterion, while incorporating financial considerations and learning, where the latter occurs as new information is utilized to adjust the decision maker's estimate on the drug's success probability. The optimal resource allocation strategy can be combined with statistical stopping methodologies such as the triangular test, in order to fulfill statistical significance and power requirements for classifying the experimental drug. This augmented methodology can then be utilized in sequential adaptive clinical trials to address the two main decisions, of determining the number of patients to be recruited and observed in each period and the trial's stopping criteria, offering gains to the trial's sponsor, the trial participants, and the patients in need of the new therapy.

4.6 Conclusions and Directions for Future Research

We examine the performance of optimal adaptive testing schedules, as measured by drug misclassification risk, time-to-market, and the firm's profit, all of which are relevant to drug R&D. We compare the performance of optimal, adaptive testing schedules against that of fixed testing schedules paired with the triangular test's stopping boundaries; the triangular test represents a widely implemented stopping mechanism in sequential clinical trials. To this end, we first expand the stochastic dynamic programming model studied in Rojas-Cordova and Bish [92] to examine a parallel group

trial, and then optimize the firm's resource allocation strategy each period, which translates into an optimal testing schedule. Through a numerical study, which uses publicly available data for drug R&D revenue and costs, we compare the performance of the optimal resource allocation strategy with fixed testing schedules, which represent the different risk attitudes the decision maker may have. Our results suggest that sequential trials that utilize fixed testing schedules, whether aggressive, neutral, or conservative, perform poorly under perfect and imperfect information settings.

In contrast, optimal testing schedules, which incorporate budgetary constraints and are flexible to be adjusted based on interim results, offer lower misclassification risks, shorter times-to-market, and higher expected profits. The triangular test includes an upper and a lower stopping boundary (which include benefit and futility thresholds, respectively), and respond to specific statistical significance and power requirements, but does not prescribe the fractions of the full sample size that need to be analyzed in each period (i.e., a testing schedule). We propose to augment the triangular test's upper stopping boundary with an optimal adaptive testing schedule. The combination of a success criterion determined based on statistical significance and power with an optimal and flexible testing schedule is able to offer gains to both pharmaceutical firms and patients, by shortening the drug development time line and reducing R&D costs. Our insights can also be helpful outside the pharmaceutical R&D context, where other problems of sequential decision-making under uncertainty are present.

Our research makes a contribution by quantifying the benefits of an optimal testing schedule for parallel group sequential adaptive clinical trials. We generate important insights on the value of flexible testing schedules that respond to financial constraints as well as newly acquired information on the experimental drug's efficacy, and on how they compare to different types of fixed testing schedules in terms of drug misclassification risk, expected time-to-market, and expected profit. We also make a contribution to the research stream on sequential investment problems, and to healthcare policy literature on parallel group clinical trials.

A limitation of our model is the assumption of a non-limiting patient recruitment (i.e., response collection) rate, which implies that the firm is able to recruit and observe as many patients as needed each period. The analysis of optimal testing schedules under a constrained patient recruitment rate is an interesting future research direction. Our analysis can be expanded to study optimal testing schedules paired with stopping mechanisms derived under different statistical approaches, such as

the repeated significance test approach [3]. Relaxing the assumption of one response per recruited patient is another interesting research opportunity, as some patients may drop out during the study, be unable to be assessed, or do not provide a response. Finally, considering multiple drugs during resource allocation decision-making is an important direction for future research, because of the interdependencies between investment prospects.

Chapter 5

Conclusions and Future Research Directions

This chapter is divided into three parts. The first part provides a summary of the research and results reported in this dissertation, the second part offers a brief description of the potential broader impact of this research, and the third part describes a series of research directions that can be pursued in the future.

5.1 Summary

Sequential adaptive clinical trials promise significant benefits to both pharmaceutical firms and patients, given the constantly increasing costs of drug R&D. These new trial designs require decision makers to determine a resource allocation (i.e., a testing schedule) and to choose when to stop the trial for benefit and futility reasons. These decisions did not have to be addressed in fixed duration (traditional) clinical trials, where a patient sample size was defined *a priori* based on statistical significance and power, and early trial termination was not possible. Decision makers thus face new resource allocation challenges of high complexity, because of the uncertainty on the drug's true efficacy, financial constraints, and the urgency to commercialize the drug as early as possible due to a time-decreasing revenue.

In Chapter 2, we examine early trial termination under the triangular test—the most widely implemented sequential methodology based on straight line boundaries. We build a simulation

model of a Phase 3 sequential trial with an experimental and a control arm, and focus on drug misclassification risk, which can translate into false positives and false negatives. In contrast to existing literature in the medical field focusing on the effects and severity of false positives, we show that false negatives can be more likely to occur. Although the triangular test's stopping boundaries are determined based on chosen statistical significance and power levels, the likelihood of mistakenly classifying an effective drug as futile may exceed the targeted type II error. Because of the detrimental effects of false negatives for both the trial's sponsor and the patients, we suggest investigators to choose stopping boundaries able to detect small but beneficial treatment differences, whenever the characteristics of the drug and the nature of the targeted disease permit.

In light of the results of our analysis of drug misclassification under the triangular test's stopping boundaries, we derive an optimal resource allocation strategy in Chapter 3. We accomplish this with a stochastic dynamic programming model that incorporates Bayesian updates on the drug's probability of success. In addition to identifying structural properties of optimal resource allocation strategies, we carry out a numerical study to compare fixed duration with sequential adaptive trials. We quantify the impact of interim analyses on drug misclassification risk, expected time-to-market, and the firm's profit. Our results suggest that fixed duration trials are significantly more sensitive to pessimistic and optimistic priors on the drug's probability of success than sequential adaptive trials, because they do not allow for learning, thus, for the correction of the decision maker's initial biases. Unlike fixed duration trials, sequential adaptive trials suffer from very low drug misclassification risks, and offer short times-to-market, even with imperfect information. This benefits the firm, which can realize a high payoff and minimize R&D costs, and the patients, who can have expedited access to an effective therapy.

We extend the scope of our stochastic dynamic programming model in Chapter 4, by applying it to a parallel group (i.e., two-armed) sequential adaptive trial. We derive optimal testing schedules under perfect information, optimistic, and pessimistic priors, and compare their performance with that of three different types of fixed testing schedules, which represent different attitudes towards risk that the decision maker may have. We find that sequential trials utilizing fixed testing schedules, whether aggressive, neutral, or conservative, are highly sensitive to imperfect information on the drug's success probability, and likely to suffer from high drug misclassification risks. Fixed testing schedules paired with the triangular test's stopping rules do not guarantee a conclusive

classification of the experimental drug by the time the trial reaches its maximum duration (i.e., by the time the available budget is exhausted and the full patient sample size has been analyzed). In contrast, optimal testing schedules incorporate learning and financial constraints, and offer lower misclassification risks, shorter times-to-market, and higher expected profits, even when paired with more conservative success criteria than the triangular test's benefit thresholds. Since the triangular test and other stopping methodologies do not prescribe the fraction of patient sample size that needs to be tested in each period (i.e., a testing schedule), we propose to augment the triangular test's upper stopping boundary with the proposed optimal adaptive testing schedule.

This dissertation makes a contribution to multiple research streams, such as sequential exploration and sequential investment, dynamic programming and information acquisition, and health-care policy analysis on parallel group clinical trials. This research contains a thorough analysis of both the benefits and the complexity of sequential adaptive clinical trials from a resource allocation decision-making perspective, and provides different stakeholders with insights on the impact of interim analyses on a series of drug R&D-relevant performance measures. To address the lack of prescriptive methodologies for the determination of testing schedules in parallel group sequential trials, this work proposes to augment a well-known and widely implemented sequential methodology with an optimal resource allocation strategy, in order to utilize newly acquired information on the experimental drug's efficacy and adjust patient recruitment throughout the trial.

5.2 Broader Impact

New medications play a vital role in public health, and have a large impact on the quality of human life. The currently long development lead times in drug R&D prevent patients from having prompt access to novel therapies, which can potentially cure or alleviate the symptoms of severe diseases. This research can have a positive impact on the society, by providing decision makers with a better understanding of the optimal testing schedules for adaptive clinical trials, thus, making it possible for patients to have expedited access to effective drugs and therapies.

Adaptive clinical trials promise to increase drug R&D productivity by reducing development costs and shortening development lead times, which in turn, poses new challenges to decision makers. This dissertation analyzes the resource allocation decision problem within different set-

tings and from different perspectives, and proposes a new methodology that addresses the financial investment challenges presented by trial settings, in which early termination, for efficacy or lack thereof, is possible. The literature offers a wide variety of statistical methodologies that address one of these challenges, which consists of the selection of specific criteria to determine whether to stop or continue the trial at different interim points throughout the trial's duration. The main objective of these methodologies is to preserve certain statistical qualities; however, these methodologies are neither based on financial considerations, such as budget availability or potential revenue, nor do they incorporate the information that is collected about the drug's efficacy during the trial.

The methodology we propose addresses the two main challenges presented by sequential adaptive clinical trials: the selection of stopping criteria, and the determination of the number of patients that need to be analyzed at each interim point. Our methodology can be combined with existing statistical methodologies to not only preserve statistical significance and power, but also, to enable the decision maker to adjust her initial beliefs about the drug's efficacy based on the data collected thus far, and to incorporate financial considerations as factors that inform her decision to allocate resources to the drug's development. The efficacy of a drug undergoing development is unknown, thus, decision makers have to rely on expert opinions, past data on similar drugs, and other sources to define whether or not to invest resources. Adaptive designs utilize the information on the drug at hand, while it is being developed, to alter the course of the trial. We then propose to take advantage of these settings in order to improve the way financial decisions are made, and to either expedite the submission of an approval application, or stop testing, based on an accurate evaluation of the drug's efficacy.

By applying an enhanced resource allocation methodology, such as the one we propose, pharmaceutical firms can optimally allocate funds to promising drugs, and promptly withdraw funding from treatments that do not show sufficient evidence of efficacy. This not only entails higher financial gains to the firms, but also represents important benefits to the trial's participants and patients in general: in case the drug under development is ineffective, trial participants will promptly stop receiving this treatment, and if the experimental drug is effective, it will be submitted for regulatory approval much sooner than under a non-adaptive setting, and upon approval, be made available to patients. This dissertation sheds light on the benefits of sequential adaptive clinical trials, provides decision makers with a methodology to allocate resources efficiently within this new setting, and

evaluates the new methodology's performance focusing on measures that are relevant to both the pharmaceutical firms and the patients. With these contributions, our goal is to enable pharmaceutical firms to further explore and apply adaptive designs, and regulatory bodies to define drug evaluation and approval mechanisms that contemplate the new flexibilities and potential of these adaptive designs.

5.3 Future Research Directions

This research can be expanded in various directions. The stochastic dynamic programming model can be extended and applied to other problems with endogenous uncertainty, as well as to other sequential investment decision problems. For example, a wide range of healthcare applications can be approached as sequential decision-making problems with exogenous and endogenous uncertainty, and can benefit from optimization-based approaches, similar to the one proposed in this dissertation.

In certain clinical trials, patient recruitment and observation is uncertain due to participants dropping out, or therapies not showing evaluable results, therefore, our analyses can be expanded to the study of stochastic patient recruitment and response collection rates. Given the complexities of the pharmaceutical market, competitors' actions, and changes in regulatory requirements, the firm's revenue may be uncertain and could be modeled as stochastic rather than deterministic. Within the sequential adaptive clinical trial context, stochastic dynamic programming can be applied to decisions other than resource allocation; for instance, the allocation of patients with different characteristics to different trial arms, with each arm corresponding to a different therapy, or a different dosage of a given therapy. In these problems, the reward may not be financial, but rather measured in terms of a specific health characteristic.

Determining optimal resource allocation strategies for trials following adaptive designs other than sequential testing (e.g., sample size re-estimation, Phase II/III hybrid designs, among others) offers a promising area of research. The shift towards these novel drug R&D designs will require researchers to examine the nature and complexity of the decisions that pharmaceutical firms, clinical investigators, and policy makers need to face. This area would benefit from interdisciplinary efforts that take into account the different perspectives and needs of the stakeholders involved, in

order to perform a holistic analysis and make a contribution to scientific literature, as well as to pharmaceutical R&D and regulatory practices.

Bibliography

- [1] Ahuja, V. and Birge, J.R. (2016). “Response-adaptive designs for clinical trials: simultaneous learning from multiple patients.” *European Journal of Operational Research*, **248** (2), pp. 619–633.
- [2] American Society of Clinical Oncology (ASCO) (2015). *Phases of Clinical Trials*. <http://www.cancer.net/navigating-cancer-care/how-cancer-treated/clinical-trials/phases-clinical-trials> [Accessed: 2016-09-28].
- [3] Armitage, P., McPherson, C., and Rowe, B. (1969). “Repeated significance tests on accumulating data.” *Journal of the Royal Statistical Society. Series A (General)*, pp. 235–244.
- [4] Barker, A., Sigman, C., Kelloff, G., Hylton, N., Berry, D., and Esserman, L. (2009). “I-SPY 2: An Adaptive Breast Cancer Trial Design in the Setting of Neoadjuvant Chemotherapy.” *Clinical Pharmacology & Therapeutics*, **86** (1), pp. 97–100.
- [5] Bassler, D., Montori, V.M., Briel, M., Glasziou, P., and Guyatt, G. (2008). “Early stopping of randomized clinical trials for overt efficacy is problematic.” *Journal of Clinical Epidemiology*, **61** (3), pp. 241–246.
- [6] Bellissant, E., Benichou, J., and Chastang, C. (1990). “Application of the triangular test to phase II cancer clinical trials.” *Statistics in medicine*, **9** (8), pp. 907–917.
- [7] Berndt, E.R., Nass, D., Kleinrock, M., and Aitken, M. (2015). “Decline in economic returns from new drugs raises questions about sustaining innovations.” *Health Affairs*, **34** (2), pp. 245–252.
- [8] Berry, D.A. (1978). “Modified two-armed bandit strategies for certain clinical trials.” *Journal of the American Statistical Association*, **73** (362), pp. 339–345.
- [9] Berry, D.A. (2005). “Introduction to Bayesian methods III: use and interpretation of Bayesian tools in design and analysis.” *Clinical Trials*, **2** (4), pp. 295–300.
- [10] Berry, D.A. (2011). “Adaptive clinical trials: the promise and the caution.” *Journal of Clinical Oncology*, **29** (6), pp. 606–609.
- [11] Berry, D.A. (2012). “Adaptive clinical trials in oncology.” *Nature Reviews Clinical Oncology*, **9** (4), pp. 199–207.
- [12] Berry, D.A. and Fristedt, B. (1985). *Bandit problems: sequential allocation of experiments (Monographs on Statistics and Applied Probability)*. Springer.
- [13] Berry, S.M., Carlin, B.P., Lee, J.J., and Muller, P. (2010). *Bayesian adaptive methods for clinical trials*. CRC Press.

- [14] Bickel, J.E. and Smith, J.E. (2006). "Optimal sequential exploration: A binary learning model." *Decision Analysis*, **3** (1), pp. 16–32.
- [15] Brown, D.B. and Smith, J.E. (2013). "Optimal sequential exploration: Bandits, clairvoyants, and wildcats." *Operations Research*, **61** (3), pp. 644–665.
- [16] Chow, S.C. and Chang, M. (2008). "Adaptive design methods in clinical trials—a review." *Orphanet Journal of Rare Diseases*, **3** (1), p. 1.
- [17] Christian, B. and Cremaschi, S. (2015). "Heuristic solution approaches to the pharmaceutical R&D pipeline management problem." *Computers & Chemical Engineering*, **74**, pp. 34–47.
- [18] Colvin, M. and Maravelias, C.T. (2008). "A stochastic programming approach for clinical trial planning in new drug development." *Computers & Chemical Engineering*, **32** (11), pp. 2626–2642.
- [19] Colvin, M. and Maravelias, C.T. (2010). "Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming." *European Journal of Operational Research*, **203** (1), pp. 205–215.
- [20] Crowley, J. and Hoering, A. (2012). *Handbook of statistics in clinical oncology*. CRC Press.
- [21] Cyert, R.M., DeGroot, M.H., and Holt, C.A. (1978). "Sequential investment decisions with Bayesian learning." *Management Science*, **24** (7), pp. 712–718.
- [22] Cyert, R.M. and March, J.G. (1963). "A behavioral theory of the firm." *Englewood Cliffs, NJ*, **2**.
- [23] David, E., Tramontin, T., and Zimmel, R. (2009). "Pharmaceutical R&D: the road to positive returns." *Nature Reviews Drug Discovery*, **8** (8), pp. 609–610.
- [24] David, F.S., Bobulsky, S., Schulz, K., and Patel, N. (2015). "Creating value with financially adaptive clinical trials." *Nature Reviews Drug Discovery*, **14** (8), pp. 523–524.
- [25] Décamps, J.P., Mariotti, T., and Villeneuve, S. (2005). "Investment timing under incomplete information." *Mathematics of Operations Research*, **30** (2), pp. 472–500.
- [26] Denrell, J. (2007). "Adaptive learning and risk taking." *Psychological review*, **114** (1), p. 177.
- [27] Denrell, J. and Le Mens, G. (2011). "Seeking positive experiences can produce illusory correlations." *Cognition*, **119** (3), pp. 313–324.
- [28] Denrell, J. and March, J.G. (2001). "Adaptation as information restriction: The hot stove effect." *Organization Science*, **12** (5), pp. 523–538.
- [29] DiMasi, J.A., Grabowski, H.G., and Hansen, R.W. (2016). "Innovation in the pharmaceutical industry: new estimates of R&D costs." *Journal of health economics*, **47**, pp. 20–33.
- [30] DiMasi, J.A., Grabowski, H.G., and Hansen, R.W. (2016). "Innovation in the pharmaceutical industry: new estimates of R&D costs." *Journal of Health Economics*, **47**, pp. 20–33.
- [31] DiMasi, J.A., Hansen, R.W., and Grabowski, H.G. (2003). "The price of innovation: new estimates of drug development costs." *Journal of health economics*, **22** (2), pp. 151–185.
URL <http://www.sciencedirect.com/science/article/pii/S0167629602001261>

- [32] Ding, M., Eliashberg, J., and Stremersch, S. (2013). *Innovation and Marketing in the Pharmaceutical Industry: Emerging Practices, Research, and Policies*, volume 20. Springer Science & Business Media.
- [33] Ellenberg, S.S., DeMets, D.L., and Fleming, T.R. (2010). "Bias and trials stopped early for benefit." *JAMA*, **304** (2), pp. 156–159.
- [34] Enea, G. and Lo Nigro, G. (2011). "A real options based model to select a balanced R&D portfolio." In *15th Annual International Conference on Real Options, Turku, Finland*. Realoptions.org.
- [35] Esserman, L.J. and Woodcock, J. (2011). "Accelerating identification and regulatory approval of investigational cancer drugs." *JAMA*, **306** (23), pp. 2608–2609.
- [36] Ghaffarzadegan, N., Lyneis, J., and Richardson, G.P. (2011). "How small system dynamics models can help the public policy process." *System Dynamics Review*, **27** (1), pp. 22–44.
- [37] Gordon Lan, K., Simon, R., and Halperin, M. (1982). "Stochastically curtailed tests in long term clinical trials." *Sequential Analysis*, **1** (3), pp. 207–219.
- [38] Goss, P.E., Ingle, J.N., Martino, S., Robert, N.J., Muss, H.B., Piccart, M.J., Castiglione, M., Tu, D., Shepherd, L.E., and Pritchard, K.I. (2003). "A randomized trial of letrozole in postmenopausal women after five years of tamoxifen therapy for early-stage breast cancer." *New England Journal of Medicine*, **349** (19), pp. 1793–1802.
- [39] Halliday, R.G., Drasdo, A.L., Lumley, C.E., and Walker, S.R. (1997). "The allocation of resources for R&D in the world's leading pharmaceutical companies." *R&D Management*, **27** (1), pp. 63–77.
- [40] Hassanzadeh, F., Modarres, M., Nemati, H.R., and Amoako-Gyampah, K. (2014). "A robust R&D project portfolio optimization model for pharmaceutical contract research organizations." *International Journal of Production Economics*, **158**, pp. 18–27.
- [41] Haybittle, J. (1971). "Repeated assessment of results in clinical trials of cancer treatment." *The British journal of radiology*, **44** (526), pp. 793–797.
- [42] He, P., Lai, T.L., and Liao, O.Y. (2012). "Futility stopping in clinical trials." *Statistics and its Interface*, **5** (4), pp. 415–423.
- [43] Huber, G.P. (1991). "Organizational learning: The contributing processes and the literatures." *Organization science*, **2** (1), pp. 88–115.
- [44] Hulot, J.S., Cucherat, M., Charlesworth, A., Van Veldhuisen, D.J., Corvol, J.C., Mallet, A., Boissel, J.P., Hampton, J., and Lechat, P. (2003). "Planning and monitoring of placebo-controlled survival trials: comparison of the triangular test with usual interim analyses methods." *British journal of clinical pharmacology*, **55** (3), pp. 299–306.
- [45] I-SPY Trials (2012). *FDA releases draft Guidance for accelerated approval of breast cancer treatments using I-SPY 2 process*. <http://www.ispytrials.org/blog/post/fda-releases-draft-guidance-for-accelerated-approval-of-breast-cancer-treatments-using-i> [Accessed: 2016-09-21].
- [46] Jacob, W.F. and Kwak, Y.H. (2003). "In search of innovative techniques to evaluate pharmaceutical R&D projects." *Technovation*, **23** (4), pp. 291–296.

- [47] Jennison, C. and Turnbull, B.W. (1999). *Group sequential methods with applications to clinical trials*. CRC Press.
- [48] Jitlal, M., Khan, I., Lee, S., and Hackshaw, A. (2012). “Stopping clinical trials early for futility: retrospective analysis of several randomised clinical studies.” *British Journal of Cancer*, **107** (6), pp. 910–917.
- [49] Kaminsky, P. and Yuen, M. (2014). “Production capacity investment with data updates.” *IIE Transactions*, **46** (7), pp. 664–682.
- [50] Kim, K. and Demets, D.L. (1987). “Design and analysis of group sequential tests based on the type I error spending rate function.” *Biometrika*, pp. 149–154.
- [51] Klenke, A., Mattner, L., et al. (2010). “Stochastic ordering of classical discrete distributions.” *Advances in Applied Probability*, **42** (2), pp. 392–410.
- [52] Kolonko, M. and Benzing, H. (1985). “The sequential design of Bernoulli experiments including switching costs.” *Operations Research*, **33** (2), pp. 412–426.
- [53] Korn, E.L., Freidlin, B., and Mooney, M. (2010). “Bias and trials stopped early for benefit.” *JAMA*, **304** (2), pp. 156–159.
- [54] Kornish, L.J. and Keeney, R.L. (2008). “Repeated commit-or-defer decisions with a deadline: The influenza vaccine composition.” *Operations Research*, **56** (3), pp. 527–541.
- [55] Kouvelis, P., Milner, J., and Tian, Z. (2014). “Clinical Trials for New Drug Development: Optimal Investment and Application.” Unpublished working paper.
- [56] Lacey, M.J., Hanna, G.J., Miller, J.D., Foster, T.S., and Russell, M.W. (2014). “Impact of Pharmaceutical Innovation in HIV/AIDS Treatment during the Highly Active Antiretroviral Therapy (HAART) Era in the US, 1987-2010: An Epidemiologic and Cost-impact Modeling Case Study.”
- [57] Lan, K.G. and DeMets, D.L. (1983). “Discrete sequential boundaries for clinical trials.” *Biometrika*, **70** (3), pp. 659–663.
- [58] Larson, R.C., Ghaffarzadegan, N., and Diaz, M.G. (2012). “Magnified effects of changes in NIH research funding levels.” *Service science*, **4** (4), pp. 382–395.
- [59] Levis, A.A. and Papageorgiou, L.G. (2004). “A hierarchical solution approach for multi-site capacity planning under uncertainty in the pharmaceutical industry.” *Computers & Chemical Engineering*, **28** (5), pp. 707–725.
- [60] Macready, W.G. and Wolpert, D.H. (1998). “Bandit problems and the exploration/exploitation trade-off.” *IEEE Transactions on Evolutionary Computation*, **2** (1), pp. 2–22.
- [61] Madani, O., Lizotte, D.J., and Greiner, R. (2004). “The budgeted multi-armed bandit problem.” In *International Conference on Computational Learning Theory*. Springer, pp. 643–645.
- [62] Maravelias, C.T. and Grossmann, I.E. (2001). “Simultaneous planning for new product development and batch manufacturing facilities.” *Industrial & Engineering Chemistry Research*, **40** (26), pp. 6147–6164.

- [63] McCardle, K.F. (1985). "Information acquisition and the adoption of new technology." *Management Science*, **31** (11), pp. 1372–1389.
- [64] McPherson, K. (1982). "On choosing the number of interim analyses in clinical trials." *Statistics in Medicine*, **1** (1), pp. 25–36.
- [65] Mueller, P.S., Montori, V.M., Bassler, D., Koenig, B.A., and Guyatt, G.H. (2007). "Ethical issues in stopping randomized trials early because of apparent benefit." *Annals of Internal Medicine*, **146** (12), pp. 878–881.
- [66] National Cancer Institute (2014). *Surveillance, epidemiology, and end results program*. <https://seer.cancer.gov/statfacts/html/ld/all.html> [Accessed: 2017-04-02].
- [67] National Cancer Institute (NCI) (2015). *NCI Dictionary of Cancer Terms*. <https://www.cancer.gov/publications/dictionaries/cancer-terms?cdrid=45833> [Accessed: 2016-09-28].
- [68] National Center for Health Statistics (2015). "Health, United States, 2014: With Special Feature on Adults Aged 55?64." *Report*.
URL <https://www.cdc.gov/nchs/data/abus/abus14.pdf>
- [69] Nesse, E. (2016). *Clinical Trial Design*. https://ccrod.cancer.gov/confluence/download/attachments/71041052/Clinical_Trial_Design.pdf [Accessed: 2016-09-28].
- [70] Nissen, M. (2016). *Read the Label on Pharma's New Drug Sales*. <https://www.bloomberg.com/gadfly/articles/2016-08-16/big-pharma-new-drug-sales-tell-only-part-of-the-story> [Accessed: 2016-09-28].
- [71] O'Brien, P.C. and Fleming, T.R. (1979). "A multiple testing procedure for clinical trials." *Biometrics*, pp. 549–556.
- [72] Oh, H.C. and Karimi, I. (2004). "Regulatory factors and capacity-expansion planning in global chemical supply chains." *Industrial & Engineering Chemistry Research*, **43** (13), pp. 3364–3380.
- [73] Orloff, J., Douglas, F., Pinheiro, J., Levinson, S., Branson, M., Chaturvedi, P., Ette, E., Gallo, P., Hirsch, G., Mehta, C., et al. (2009). "The future of drug development: advancing clinical trial design." *Nature Reviews Drug Discovery*, **8** (12), pp. 949–957.
- [74] Papageorgiou, L.G., Rotstein, G.E., and Shah, N. (2001). "Strategic supply chain optimization for the pharmaceutical industries." *Industrial & Engineering Chemistry Research*, **40** (1), pp. 275–286.
- [75] Park, H., Clear, T., Rouse, W.B., Basole, R.C., Braunstein, M.L., Brigham, K.L., and Cunningham, L. (2012). "Multilevel simulations of health delivery systems: A prospective tool for policy, strategy, planning, and management." *Service Science*, **4** (3), pp. 253–268.
- [76] Patel, N.R., Ankolekar, S., Antonijevic, Z., and Rajicic, N. (2013). "A mathematical model for maximizing the value of phase 3 drug development portfolios incorporating budget constraints and risk." *Statistics in Medicine*, **32** (10), pp. 1763–1777.
- [77] Paul, S.M., Mytelka, D.S., Dunwiddie, C.T., Persinger, C.C., Munos, B.H., Lindborg, S.R., and Schacht, A.L. (2010). "How to improve R&D productivity: the pharmaceutical industry's grand challenge." *Nature Reviews Drug Discovery*, **9** (3), pp. 203–214.

- [78] Peace, K.E. (1992). *Biopharmaceutical sequential statistical applications*. CRC Press.
- [79] Peto, R., Pike, M., Armitage, P., Breslow, N.E., Cox, D., Howard, S., Mantel, N., McPherson, K., Peto, J., and Smith, P. (1977). "Design and analysis of randomized clinical trials requiring prolonged observation of each patient. II. analysis and examples." *British journal of cancer*, **35** (1), p. 1.
- [80] Pharmaceutical Research and Manufacturers of America (2016). "2016 Biopharmaceutical Research Industry Profile." *Report*.
URL <http://www.phrma.org/report/industry-profile-2016>
- [81] Pharmaceutical Research and Manufacturers of America (2016). "A Decade of Innovation in Chronic Diseases." *Report*.
URL <http://www.phrma.org/report/a-decade-of-innovation-in-chronic-diseases>
- [82] Pharmaceutical Research and Manufacturers of America (PhRMA) (2015). *Biopharmaceutical Industry-Sponsored Clinical Trials: Impact on State Economies*. <http://www.phrma.org/sites/default/files/pdf/biopharmaceutical-industry-sponsored-clinical-trials-impact-on-state-economies.pdf> [Accessed: 2016-09-28].
- [83] Pocock, S., Wang, D., Wilhelmsen, L., and Hennekens, C.H. (2005). "The data monitoring experience in the Candesartan in Heart Failure Assessment of Reduction in Mortality and morbidity (CHARM) program." *American heart journal*, **149** (5), pp. 939–943.
- [84] Pocock, S. and White, I. (1999). "Trials stopped early: too good to be true?" *The Lancet*, **353** (9157), pp. 943–944.
- [85] Pocock, S.J. (1977). "Group sequential methods in the design and analysis of clinical trials." *Biometrika*, **64** (2), pp. 191–199.
- [86] Pocock, S.J. (1982). "Interim analyses for randomized clinical trials: the group sequential approach." *Biometrics*, pp. 153–162.
- [87] Pocock, S.J. (2005). "When (not) to stop a clinical trial for benefit." *JAMA*, **294** (17), pp. 2228–2230.
- [88] President's Council of Advisors on Science and Technology (2012). "Report to the President on Propelling Innovation In Drug Discovery, Development, and Evaluation." *Technical report*, Executive Office of the President. <https://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-fda-final.pdf> [Accessed: 2016-10-12].
- [89] Rajapakse, A., Titchener-Hooker, N.J., and Farid, S.S. (2005). "Modelling of the biopharmaceutical drug development pathway and portfolio management." *Computers & Chemical Engineering*, **29** (6), pp. 1357–1368.
- [90] Rogers, M.J., Gupta, A., and Maranas, C.D. (2002). "Real options based analysis of optimal pharmaceutical research and development portfolios." *Industrial & Engineering Chemistry Research*, **41** (25), pp. 6607–6620.
- [91] Rogers, M.S., Chang, A.M., and Todd, S. (2005). "Using group-sequential analysis to achieve the optimal sample size." *BJOG: An International Journal of Obstetrics & Gynaecology*, **112** (5), pp. 529–533.

- [92] Rojas-Cordova, A. and Bish, E.K. (2017). “Optimal Sequential Resource Allocation: The Case of Sequential Adaptive Clinical Trials.” Under Review.
- [93] SEBILLE, V. and BELLISSANT, E. (2000). “Comparison of four sequential methods allowing for early stopping of comparative clinical trials.” *Clinical Science*, **98** (5), pp. 569–578.
- [94] Smith, J.E. and Ulu, C. (2012). “Technology adoption with uncertain future costs and quality.” *Operations Research*, **60** (2), pp. 262–274.
- [95] Solak, S., Clarke, J.P.B., Johnson, E.L., and Barnes, E.R. (2010). “Optimization of R&D project portfolios under endogenous uncertainty.” *European Journal of Operational Research*, **207** (1), pp. 420–433.
- [96] Subramanian, D., Pekny, J.F., Reklaitis, G.V., and Blau, G.E. (2003). “Simulation-optimization framework for stochastic optimization of R&D pipeline management.” *AIChE Journal*, **49** (1), pp. 96–112.
- [97] Sun, E., Lakdawalla, D., Reyes, C., Goldman, D., Philipson, T., and Jena, A. (2008). “The determinants of recent gains in cancer survival: an analysis of the Surveillance, Epidemiology, and End Results (SEER) database.” *Journal of Clinical Oncology*, **26** (15 suppl), pp. 6616–6616.
- [98] Tobias, J.S., Monson, K., Gupta, N., MacDougall, H., Glaholm, J., Hutchison, I., Kadalayil, L., and Hackshaw, A. (2010). “Chemoradiotherapy for locally advanced head and neck cancer: 10-year follow-up of the UK Head and Neck (UKHAN1) trial.” *The lancet oncology*, **11** (1), pp. 66–74.
- [99] Todd, S., Whitehead, A., Stallard, N., and Whitehead, J. (2001). “Interim analyses and sequential designs in phase III studies.” *British Journal of Clinical Pharmacology*, **51** (5), pp. 394–399.
- [100] UKCCCR Anal Cancer Trial Working Party (1996). “Epidermoid anal cancer: results from the UKCCCR randomised trial of radiotherapy alone versus radiotherapy, 5-fluorouracil, and mitomycin.” *The Lancet*, **348** (9034), pp. 1049–1054.
- [101] Ulu, C. and Smith, J.E. (2009). “Uncertainty, information acquisition, and technology adoption.” *Operations Research*, **57** (3), pp. 740–752.
- [102] U.S. FDA (2010). *Guidance for Industry: Adaptive Design Clinical Trials for Drugs and Biologics*. <http://www.fda.gov/downloads/Drugs/GuidanceComplianceRegulatoryInformation/Guidances/UCM201790.pdf> [Accessed: 2016-12-05].
- [103] U.S. FDA (2011). *Oncologic Drugs Advisory Committee (ODAC) Meeting Report Silver Spring, MD*. <http://www.fda.gov/downloads/AdvisoryCommittees/CommitteesMeetingMaterials/Drugs/OncologicDrugsAdvisoryCommittee/UCM245644.pdf> [Accessed: 2015-05-28].
- [104] U.S. FDA (2014). *Fast Track, Accelerated Approval, and Priority Review*. <http://www.fda.gov/ForPatients/Approvals/Fast/ucm405447.htm> [Accessed: 2015-05-28].
- [105] US Food and Drug Administration (2010). “Guidance for industry: adaptive design clinical trials for drugs and biologics.” *Washington DC, USA: Food and Drug Administration*.
- [106] US Food and Drug Administration (2011). “Oncologic Drugs Advisory Committee (ODAC) Meeting Report.” <http://www.fda.gov/downloads/AdvisoryCommittees/CommitteesMeetingMaterials/Drugs/OncologicDrugsAdvisoryCommittee/UCM245644.pdf> [Accessed: 2015-05-28].

- [107] US Food and Drug Administration (2014). *Fast Track, Accelerated Approval, and Priority Review*. <http://www.fda.gov/ForPatients/Approvals/Fast/ucm405447.htm> [Accessed: 2015-05-28].
- [108] Wald, A. (1973). *Sequential analysis*. Courier Corporation.
- [109] Whitehead, J. (1997). *The design and analysis of sequential clinical trials*. John Wiley & Sons.
- [110] Whitehead, J. (2002). “Sequential methods in clinical trials.” *Sequential Analysis*, **21** (4), pp. 285–308.
- [111] Whitehead, J. (2004). “Stopping clinical trials by design.” *Nature Reviews Drug Discovery*, **3** (11), pp. 973–977.
- [112] Ye, F. (2008). *Design and Analysis of Phase III Clinical Trials*. <https://medschool.vanderbilt.edu/cqs/files/cqs/media/2008Jun2008Fei.pdf> [Accessed: 2016-09-28].
- [113] Zannad, F., Stough, W.G., McMurray, J.J., Remme, W.J., Pitt, B., Borer, J.S., Geller, N.L., and Pocock, S.J. (2012). “When to stop a clinical trial early for benefit: lessons learned and future approaches.” *Circulation: Heart Failure*, **5** (2), pp. 294–302.

Appendix A

Appendix for Chapter 2

A.1 Drug Testing Process

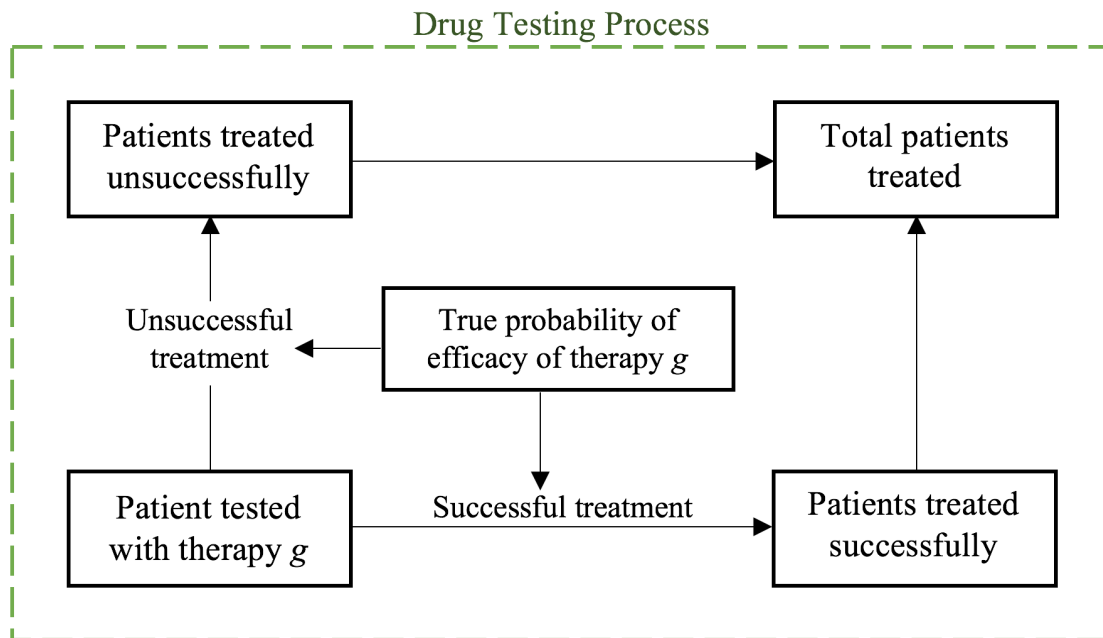


Figure A.1: Conceptual representation of the drug testing process for treatment group g , where $g \in \{E(\text{experimental}), C(\text{control})\}$

A.2 Detailed Description of Simulation Model

Table A.1: Variables and parameters used in the simulation model

Variables and Parameters	Formula	Units
1: Control outcome	$TR_C = \text{IF THEN ELSE } [J > (1 - p_C), 1, 0],$ $J \sim U[0, 1]$	Patient
2: Desired number of tests to stop phase	$M = 244$	Patient
3: Interim analysis point reached	$IR = \text{IF THEN ELSE } (N = IA, 1, 0)$	Dimensionless
4: Interim analysis recruitment limit	$IA = 122$	Patient
5: Maximum duration reached	$MD = \text{IF THEN ELSE } (N < M, 0, 1)$	Dimensionless
6: Patient assigned to control group	$AC = \text{IF THEN ELSE } (PT = 1, 1 - AE, 0)$	Dimensionless
7: Patient assigned to treatment group	$AE = \text{IF THEN ELSE } (PT = 1 : \text{AND:}$ $A > 0.5, 1, 0), J \sim U[0, 1]$	Dimensionless
8: Patient tested	$PT = \text{IF THEN ELSE } (MD = 0 : \text{AND:}$ $RD = 1, 1, 0)$	Dimensionless
9: Random seed testing	$RS = 89$	Dimensionless
10: Random seed patient assignment	$RA = 46$	Dimensionless
11: Recruitment decision	$RD = \text{IF THEN ELSE } (SB = 1 : \text{OR:}$ $SF = 1, 0, 1)$	Dimensionless
12: Treatment outcome	$TR_E = \text{IF THEN ELSE } [J > (1 - p_E), 1, 0]$	Patient
13: Stopping for benefit	$\frac{d}{dt}SB = \text{IF THEN ELSE } [(IR = 1$ $: \text{AND: } Z \geq BT), 1, 0]$	Dimensionless
14: Stopping for futility	$\frac{d}{dt}SF = \text{IF THEN ELSE } [(IR = 1$ $: \text{AND: } Z \leq BT), 1, 0]$	Dimensionless
15: Successful results control group	$\frac{d}{dt}S_c = \text{IF THEN ELSE } AC = 1, TR_C, 0]$	Patient
16: Successful results treatment group	$\frac{d}{dt}S_e = \text{IF THEN ELSE } AE = 1, TR_E, 0]$	Patient
17: Threshold for benefit	$BT = 5.82 + 0.257V - (0.583 * \text{SQRT}(V))$	Dimensionless
18: Threshold for futility	$FT = -5.82 + 0.772V + (0.583 * \text{SQRT}(V))$	Dimensionless
19: Total sample	$N = N_e + N_c$	Patient
20: Total sample of control group	$N_c = S_c + U_c$	Patient
21: Total sample of treatment group	$N_e = S_e + U_e$	Patient
22: Total number of successes	$S = S_e + S_c$	Patient
23: True probability of success of control	$p_C = 0.5$	Dimensionless
24: True probability of success of treatment	$p_E = 0.7$	Dimensionless
25: Unsuccessful results control group	$\frac{d}{dt}U_c = \text{IF THEN ELSE } AC = 1, 1 - TR_C, 0]$	Patient
26: Unsuccessful results treatment group	$\frac{d}{dt}U_e = \text{IF THEN ELSE } AE = 1, 1 - TR_E, 0]$	Patient
27: Statistic measuring treatment difference	$Z = (N_c S_e - N_e S_c) / N$	Dimensionless
28: Statistic measuring accrued information	$V = [N_c N_e S(N - S)] / N^3$	Dimensionless

Appendix B

Appendix for Chapter 3

Throughout, we use the following relationship between the beta function, $B(\gamma, \zeta)$, and the gamma function, $\Gamma(\cdot)$:

$$B(\gamma, \zeta) = \frac{\Gamma(\gamma) \cdot \Gamma(\zeta)}{\Gamma(\gamma + \zeta)}.$$

Proof of Lemma 1. For $t = 1, 2, 3, \dots, T$, the decision maker models N_t^+ as Binomial(N_t, \hat{p}_t). Then, for some $m \in [1, \dots, T + 1 - t]$, $\sum_{j=1}^m N_{t-1+j}^+ \sim \text{Binomial}(N_{t,m}^{Cum}, \hat{p}_t)$.

The posterior predictive probability of having k successful observations in the first period, given the prior parameters λ_0 and β_0 , can be derived using a Beta-binomial distribution:

$$\begin{aligned} \Pr(\gamma_1 = k | \lambda_0, \beta_0) &= \Pr(N_1^+ = k | \lambda_0, \beta_0), \text{ for } k \in [0, 1, \dots, N_1], \\ &= \binom{N_1}{k} \frac{B(k + \lambda_0, N_1 - k + \beta_0)}{B(\lambda_0, \beta_0)} \\ &= \binom{N_1}{k} \frac{1}{B(\lambda_0, \beta_0)} \int_0^1 p^{\lambda_0 + k - 1} (1 - p)^{\beta_0 + N_1 - k - 1} dp \end{aligned}$$

$$\begin{aligned}
&= \binom{N_1}{k} \frac{\Gamma(\lambda_0 + \beta_0)}{\Gamma(\lambda_0)\Gamma(\beta_0)} \cdot \frac{\Gamma(\lambda_0 + k)\Gamma(\beta_0 + N_1 - k)}{\Gamma(\lambda_0 + \beta_0 + N_1)} \\
&= \binom{N_1}{k} \frac{(\lambda_0 + \beta_0 - 1)!(\lambda_0 + k - 1)!(\beta_0 + N_1 - k - 1)!}{(\lambda_0 - 1)!(\beta_0 - 1)!(\lambda_0 + \beta_0 + N_1 - 1)!} \\
&= \binom{N_1}{k} \frac{\prod_{j=0}^{k-1} (\lambda_0 + j) \prod_{j=0}^{N_1-k-1} (\beta_0 + j)}{\prod_{j=0}^{N_1-1} (\lambda_0 + \beta_0 + j)}.
\end{aligned}$$

Similarly, the posterior predictive probability of having k successful observations given the prior parameters γ_{t-1} and ζ_{t-1} can be derived using a Beta-binomial distribution:

$$\begin{aligned}
\Pr\left(\sum_{j=1}^m N_{t-1+j}^+ = k \mid \gamma_{t-1}, \gamma_{t-2}, \dots, \gamma_0\right) &= \Pr\left(\sum_{j=1}^m N_{t-1+j}^+ = k \mid \gamma_{t-1}\right), \text{ for } k \in [0, 1, \dots, N_{t,m}^{Cum}], m \in [1, \dots, T+1-t], \\
&= \binom{N_{t,m}^{Cum}}{k} \frac{1}{B(\gamma_{t-1}, \zeta_{t-1})} \int_0^1 p^{\gamma_{t-1}+k-1} (1-p)^{\zeta_{t-1}+(N_{t,m}^{Cum}-k)-1} dp \\
&= \binom{N_{t,m}^{Cum}}{k} \frac{\Gamma(\gamma_{t-1} + \zeta_{t-1})}{\Gamma(\gamma_{t-1})\Gamma(\zeta_{t-1})} \cdot \frac{\Gamma(\gamma_{t-1} + k)\Gamma(\zeta_{t-1} + N_{t,m}^{Cum} - k)}{\Gamma(\gamma_{t-1} + \zeta_{t-1} + N_{t,m}^{Cum})} \\
&= \binom{N_{t,m}^{Cum}}{k} \frac{(\gamma_{t-1} + \zeta_{t-1} - 1)!(\gamma_{t-1} + k - 1)!(\zeta_{t-1} + N_{t,m}^{Cum} - k - 1)!}{(\gamma_{t-1} - 1)!(\zeta_{t-1} - 1)!(\gamma_{t-1} + \zeta_{t-1} + N_{t,m}^{Cum} - 1)!} \\
&= \binom{N_{t,m}^{Cum}}{k} \frac{\prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t,m}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t,m}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)}. \quad \blacksquare
\end{aligned}$$

Proof of Lemma 2. $\Pr(\gamma_t \geq \gamma_c(M) \mid \gamma_{t-q} = \gamma) = \Pr(\sum_{j=t-q+1}^t N_j^+ \geq \gamma_c(M) - \gamma \mid \gamma_{t-q} = \gamma)$ is the posterior predictive probability of success, where $\sum_{j=t-q+1}^t N_j^+ = N_{t-q+1,q}^{Cum}$ is the sum of q Beta-binomial random variables, with a total number of observations of $\sum_{j=t-q+1}^t N_j$. Then, using Lemma 1, we can write:

$$\Pr(\gamma_t \geq \gamma_c(M) \mid \gamma_{t-q} = \gamma) = \frac{\sum_{k=(\gamma_c(M)-\gamma)^+}^{N_{t-q+1,q}^{Cum}} \binom{N_{t-q+1,q}^{Cum}}{k} \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+1,q}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)}.$$

To show that $\Pr(\gamma_t \geq \gamma_c(M) \mid \gamma_{t-q} = \gamma)$ is increasing in γ , we evaluate the following:

$$\begin{aligned}
& \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma) - \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma - 1) = \Pr(\gamma_t - \gamma_{t-q} \geq (\gamma_c(M) - \gamma)^+) - \Pr(\gamma_t - \gamma_{t-q} \geq (\gamma_c(M) - \gamma)^+ + 1) \\
&= \frac{\sum_{k=(\gamma_c(M)-\gamma)^+}^{N_{t-q+1,q}^{Cum}} \binom{N_{t-q+1,q}^{Cum}}{k} \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+1,q}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} - \frac{\sum_{k=(\gamma_c(M)-\gamma)^++1}^{N_{t-q+1,q}^{Cum}} \binom{N_{t-q+1,q}^{Cum}}{k} \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+1,q}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} \\
&= \frac{\binom{N_{t-q+1,q}^{Cum}}{\gamma_c(M)-\gamma} \prod_{j=0}^{\gamma_c(M)-\gamma-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+1,q}^{Cum}-\gamma_c(M)+\gamma-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} > 0, \text{ for } \gamma_c(M) > \gamma
\end{aligned}$$

$$\Rightarrow \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma_1) > \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma_2), \forall \gamma_1 > \gamma_2 : \gamma_1, \gamma_2 \in \mathbb{Z}^+, \gamma_1 < \gamma_c(M).$$

To show that $\Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma)$ is increasing in q , we evaluate the following:

$$\begin{aligned}
& \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q} = \gamma) - \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-(q-1)} = \gamma) = \Pr(\gamma_t - \gamma_{t-q} \geq (\gamma_c(M) - \gamma)^+) - \Pr(\gamma_t - \gamma_{t-q+1} \geq (\gamma_c(M) - \gamma)^+) \\
&= \frac{\sum_{k=(\gamma_c(M)-\gamma)^+}^{N_{t-q+1,q}^{Cum}} \binom{N_{t-q+1,q}^{Cum}}{k} \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+1,q}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} - \frac{\sum_{k=(\gamma_c(M)-\gamma)^+}^{N_{t-q+2,q-1}^{Cum}} \binom{N_{t-q+2,q-1}^{Cum}}{k} \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+2,q-1}^{Cum}-k-1} (\zeta_{t-1} + j)}{\prod_{j=0}^{N_{t-q+2,q-1}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} \\
&= \frac{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + j) + \sum_{k=(\gamma_c(M)-\gamma)^+}^{N_{t-q+2,q-1}^{Cum}} \left\{ \prod_{j=0}^{k-1} (\gamma_{t-1} + j) \prod_{j=0}^{N_{t-q+2,q-1}^{Cum}-k-1} (\zeta_{t-1} + j) \left[\binom{N_{t-q+1,q}^{Cum}}{k} - \binom{N_{t-q+2,q-1}^{Cum}}{k} \right] \right\}}{\prod_{j=0}^{N_{t-q+1,q}^{Cum}-1} (\gamma_{t-1} + \zeta_{t-1} + j)} > 0
\end{aligned}$$

$$\Rightarrow \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q_1} = \gamma) > \Pr(\gamma_t \geq \gamma_c(M) | \gamma_{t-q_2} = \gamma), \forall q_1 > q_2 : q_1, q_2 \in \mathbb{Z}^+, q_1 < t,$$

completing the proof. \blacksquare

Proof of Lemma 3. Observe that $V_t(N_t = 0) \equiv V_t(\gamma_{t-1}, R_{t-1})|_{N_t=0} = 0, \forall \gamma_{t-1}, \forall R_{t-1}$.

(i) Given that $N_t^*(\gamma_{t-1}) = 0$, we have that $V_t(N_t = 0, \gamma_{t-1}) \geq V_t(N_t > 0, \gamma_{t-1}), \forall N_t : 0 < N_t < R_{t-1}$.

Additionally, by Lemma 2, $V_t(\gamma_{t-1})$ decreases as γ_{t-1} decreases. Thus, we have, $V_t(N_t = 0, \gamma) = V_t(N_t = 0, \gamma_{t-1}) \geq V_t(N_t > 0, \gamma_{t-1}) \geq V_t(N_t > 0, \gamma), \forall \gamma \leq \gamma_{t-1}$, leading to $N_t^*(\gamma) = 0, \forall \gamma \leq \gamma_{t-1}$.

(ii) Given that $N_t^*(R_{t-1}) = 0$, we have that $V_t(N_t = 0, R_{t-1}) \geq V_t(N_t > 0, R_{t-1}), \forall N_t : 0 < N_t <$

R_{t-1} . Thus, we have that $V_t(N_t, \gamma_{t-1}, R_{t-1}) \leq 0$ and $\Pr(N_t^+ \geq \gamma_c(M) - \gamma_{t-1})|_{R_{t-1}} < \frac{c}{g_T}, \forall N_t : 0 < N_t < R_{t-1}$. Hence, $\Pr(N_t^+ \geq \gamma_c(M) - \gamma_{t-1})|_R < \frac{c}{g_T}, \forall R \leq R_{t-1}, \forall N_t : 0 < N_t < R_{t-1}$. Therefore, $V_t(N_t = 0, R) > V_t(N_t > 0, R), \forall R \leq R_{t-1}$, leading to $N_t^*(R) = 0, \forall R \leq R_{t-1}$. This completes the proof. ■

Proof of Lemma 4. (i) If $N_t^*(\gamma_{t-1}, R_{t-1}) = 0$, then $V_t(\gamma_{t-1}, R_{t-1}) \leq 0$, and $\Pr(N_t^+ \geq \gamma_c(M) - \gamma_{t-1})|_{R_{t-1}} = 0, \forall N_t : 0 < N_t < R_{t-1}$. From the definition of R_t , we have that $R_{t+q} \leq R_{t-1}$, for $q \in \{0, \dots, T-t\}$. Then, the result directly follows from Lemma 3, part (ii).

(ii) If $R_{t-1} < \gamma_c(M) - \gamma_{t-1}$, then $\Pr(N_t^+ \geq \gamma_c(M) - \gamma_{t-1})|_{R_{t-1}} = 0$, and $V_t(\gamma_{t-1}, R_{t-1}) \leq 0, \forall N_t$ in the feasible region. From the definition of R_t , we know that $R_{t+q} \leq R_{t-1}$, for $q \in \{0, \dots, T-t\}$. Then, the result directly follows from Lemma 3, part (ii). This completes the proof. ■

To prove Theorem 1, we first provide some structural properties of the $V_T(\cdot)$ function in Lemma A.

Lemma A. Properties of $V_T(\cdot)$. For a sufficiently large $E[\hat{p}_T]$, the value function at time T , $V_T(\cdot)$, is first convex and then concave in N_T , with a unique maximizer. Specifically, $V_T(\cdot)$ is strictly decreasing with a slope of $-c$ for $N_T \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$. At some value of $N_T \geq \gamma_c(M) - \gamma_{T-1}$, $V_T(\cdot)$ becomes increasing, reaching its unique maximizer, after which it becomes strictly decreasing again with a slope of $-c$. Otherwise, i.e., if $E[\hat{p}_T]$ is not sufficiently large, we have that $V_T(\cdot) \leq 0, \forall N_T$.

Proof of Lemma A. From Eq.(2), we have:

$$V_T(N_T, \gamma_{T-1}, R_{T-1}) = -cN_T + g_T \Pr(N_T^+(N_T) \geq \gamma_c(M) - \gamma_{T-1} | \gamma_{T-1} < \gamma_c(M)).$$

$\forall N_T \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$, we have $\Pr(N_T^+(N_T) \geq \gamma_c(M) - \gamma_{T-1} | \gamma_{T-1} < \gamma_c(M)) = 0 \Rightarrow V_T(N_T, \gamma_{T-1}, R_{T-1}) = -cN_T$. Thus, $V_T(N_T, \gamma_{T-1}, R_{T-1})$ is monotonically decreasing in N_T at a rate of $-c$ for $N_T \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$.

For $N_T \geq \gamma_c(M) - \gamma_{T-1}$, using the law of total probability and Lemma 1, we can write:

$$\begin{aligned}
\Pr(N_T^+(N_T) \geq \gamma_c(M) - \gamma_{T-1} | \gamma_{T-1} < \gamma_c(M)) &= \Pr(N_T^+(N_T) \geq \gamma_c(M) - \gamma_{T-1}) \quad (\text{by Lemma 1}) \\
&= \int_0^1 \Pr(N_T^+(N_T) \geq \gamma_c(M) - \gamma_{T-1} | \hat{p}_T = \tilde{p}) f_{\hat{p}_T}(\tilde{p}) d\tilde{p} \\
&= \int_0^1 \bar{F}(N_T^+(\gamma_c(M) - \gamma_{T-1} - 1; N_T, \tilde{p}) f_{\hat{p}_T}(\tilde{p}) d\tilde{p},
\end{aligned} \tag{B.1}$$

where $\bar{F}(k; N_T, \tilde{p})$ denotes the tail (1 - CDF) of Binomial (N_T, \tilde{p}) at $k \in \mathbb{Z}^+$.

$$\begin{aligned}
\text{Define } \Delta\Pr(N_T, \tilde{p}) &\equiv \bar{F}(\gamma_c(M) - \gamma_{T-1} - 1; N_T + 1, \tilde{p}) - \bar{F}(\gamma_c(M) - \gamma_{T-1} - 1; N_T, \tilde{p}) \\
&= 1 - F(\gamma_c(M) - \gamma_{T-1} - 1; N_T + 1, \tilde{p}) - 1 + F(\gamma_c(M) - \gamma_{T-1} - 1; N_T, \tilde{p}) \\
&= \sum_{k=0}^{\gamma_c(M) - \gamma_{T-1} - 1} \left[\binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} - \binom{N_T + 1}{k} \tilde{p}^k (1 - \tilde{p})^{N_T + 1 - k} \right] \\
&= \sum_{k=0}^{\gamma_c(M) - \gamma_{T-1} - 1} \left[\binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} \left[1 - \frac{(N_T + 1)}{(N_T + 1 - k)} (1 - \tilde{p}) \right] \right] \\
&= \sum_{k=0}^{\gamma_c(M) - \gamma_{T-1} - 1} \left[\binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} \left[\frac{\tilde{p}(N_T + 1) - k}{(N_T + 1 - k)} \right] \right].
\end{aligned}$$

Thus, for $k \leq \tilde{p}(N_T + 1)$, or equivalently, $N_T \geq \frac{k}{\tilde{p}} - 1$, the component with index k ($k = 0, \dots, \gamma_c(M) - \gamma_{T-1} - 1$) in the sum above is non-negative, and otherwise it is negative. However, by usual stochastic ordering of the binomial distribution [51], we know that $\Delta\Pr(N_T, \tilde{p}) \geq 0$, $\forall \tilde{p}(0, 1)$ and $N_T \in \mathbb{Z}^+$. Therefore, for $N_T \geq \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$, all $\gamma_c - \gamma_{T-1}$ components in the sum above are non-negative; and otherwise, some components are non-negative while some components are negative, but the sum is always non-negative.

Obviously, as N_T increases, the number of negative components in the sum decreases. Let β_k denote component k in the sum, $\beta_k \equiv \binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} \left[\frac{\tilde{p}(N_T + 1) - k}{(N_T + 1 - k)} \right]$. $\Delta\Pr(N_T, \tilde{p})$ must contain at least one strictly positive component for $N_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$, and observe that $\beta_0 = \tilde{p}(1 - \tilde{p})^{N_T} > 0$.

Also observe that:

$$\begin{aligned}
\Delta V_T(N_T) &\equiv V_T(N_T + 1, \gamma_{T-1}) - V_T(N_T, \gamma_{T-1}) \\
&= -c + g_T \int_0^1 \Delta \Pr(N_T, \tilde{p}) f_{\tilde{p}_T}(\tilde{p}) d\tilde{p} \\
&= -c + g_T E_{\tilde{p}} [\Delta \Pr(N_T, \tilde{p})]
\end{aligned}$$

Thus, $V_T(\cdot)$ is increasing in N_T if $E_{\tilde{p}} [\Delta \Pr(N_T, \tilde{p})] \geq \frac{c}{g_T}$, and is decreasing in N_T otherwise, where $\frac{c}{g_T} < 1$ by assumption.

$\Delta \Pr(N_T, \tilde{p}) \geq 0$, and contains at least one strictly positive component, $\tilde{p}(1 - \tilde{p})^{N_T}$, for $N_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$, thus $\tilde{p}(1 - \tilde{p})^{N_T} + \sum_{k=1}^{\gamma_c(M) - \gamma_{T-1} - 1} \beta_k \geq 0$. By our assumption of a sufficiently large \tilde{p} , $\exists N'_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$ such that $\tilde{p}(1 - \tilde{p})^{N'_T} \geq \frac{c}{g_T} - \sum_{k=1}^{\gamma_c(M) - \gamma_{T-1} - 1} \beta_k$. Thus, $\Delta \Pr(N_T, \tilde{p}) \geq \frac{c}{g_T}$, $E_{\tilde{p}} [\Delta \Pr(N_T, \tilde{p})] \geq \frac{c}{g_T}$, and $V_T(\cdot)$ starts increasing in N'_T .

Let $\Delta \beta_k$ denote the rate of change of β_k as N_T increases, that is:

$$\begin{aligned}
\Delta \beta_k &\equiv \beta_k(N_T + 1) - \beta_k(N_T) \\
&= \left\{ \binom{N_T + 1}{k} \tilde{p}^k (1 - \tilde{p})^{N_T + 1 - k} \left[\frac{\tilde{p}(N_T + 2) - k}{(N_T + 2 - k)} \right] \right\} - \left\{ \binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} \left[\frac{\tilde{p}(N_T + 1) - k}{(N_T + 1 - k)} \right] \right\} \\
&= \binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k} \left[\frac{-N_T^2 \tilde{p}^2 + N_T \tilde{p}(2k - 3\tilde{p}) + 2\tilde{p}(k - \tilde{p}) + k(1 - k)}{(N_T + 1 - k)(N_T + 2 - k)} \right].
\end{aligned}$$

Thus, β_k is increasing in N_T for $N_T \leq \frac{2k - 3\tilde{p} \pm \sqrt{4k(1 - \tilde{p}) + \tilde{p}^2}}{2\tilde{p}}$.

Some of the non-negative components in $\Delta \Pr(N_T, \tilde{p})$ are strictly decreasing in N_T , while others are first increasing and then decreasing in N_T . Earlier in this proof, we showed that the number of non-negative components in $\Delta \Pr(N_T, \tilde{p})$ increases in N_T for $N_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$. Since $\frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1 > N_T > k, \forall k \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$, and for a sufficiently large \tilde{p} , any given component $\beta_k < 0$ will become non-negative at some $N_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$. Although β_k may later decrease in N_T , it will never fall below zero once it becomes non-negative. Further, from usual stochastic ordering, we know that $\Pr(N_T, \tilde{p})$ is monotone increasing in $N_T, \forall N_T, \forall \tilde{p}$. There-

fore, if $V_T(\cdot)$ starts increasing in N'_T , and $\tilde{p}(1 - \tilde{p})^{N'_T} \geq \frac{c}{g_T} - \sum_{k=1}^{\gamma_c(M) - \gamma_{T-1} - 1} \beta_k$, then $\Delta \Pr(N'_T, \tilde{p})$ and $V_T(\cdot)$ will keep increasing monotonically in N_T for $N'_T \leq N_T < \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} - 1$.

For a large $N_T > \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}} > \gamma_c(M) - \gamma_{T-1} - 1$, all components in $\Delta \Pr(N_T, \tilde{p})$ will be non-negative, as the non-negativity condition will always be met, i.e., $N_T \geq \frac{k}{\tilde{p}} - 1, \forall k \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$. However, as N_T keeps increasing, $\binom{N_T}{k} \tilde{p}^k (1 - \tilde{p})^{N_T - k}$ will decrease and reach zero $\forall k \in [0, \gamma_c(M) - \gamma_{T-1} - 1]$, thus $\beta_k = 0, \forall k \in [0, \gamma_c(M) - \gamma_{T-1} - 1], \Delta \Pr(N_T, \tilde{p}) = 0 < \frac{c}{g_T}$, and $V_T(\cdot)$ will monotonically decrease in N_T for $N_T > \frac{\gamma_c(M) - \gamma_{T-1} - 1}{\tilde{p}}$ with a slope of $-c$. This completes the proof. ■

In order to study the structural properties of an optimal solution in any period t , we first show in Lemma B that there exists a region in which the expected profit to-go function in period t is monotone decreasing in N_t .

Lemma B. There exists some $\tilde{N}_t(\gamma_{t-1})$ such that $E[V_{t+1}^*(N_t)]$ is decreasing in N_t , for $N_t \geq \tilde{N}_t(\gamma_{t-1}, R_{t-1})$.

Proof of Lemma B. The expected profit to-go function, $E[V_{t+1}^*(N_t)] =$

$\sum_{n_t^+ = 0}^{\min N_t, \gamma_c(M) - \gamma_{t-1} - 1} [V_{t+1}^*(\gamma_{t-1} + n_t^+, R_{t-1} - N_t | \gamma_c(M) - \gamma_{t-1}) \cdot \Pr(N_t^+ = n_t^+ | N_t^+ < \gamma_c(M) - \gamma_{t-1})]$, is comprised of a sum of $N_t + 1$ components in Region 1, defined as $\{0 \leq N_t \leq \gamma_c(M) - \gamma_{t-1} - 1\}$, and a sum of $\gamma_c(M) - \gamma_{t-1}$ components in Region 2, defined as $\{\gamma_c(M) - \gamma_{t-1} \leq N_t\}$, with each component having the following structure: $V_{t+1}^* \cdot \frac{\Pr(N_t^+ = n_t^+)}{\sum_{k=0}^{\min(N_t, \gamma_c(M) - \gamma_{t-1} - 1)} \Pr(N_t^+ = k)}$.

In Region 1, the expected profit to-go function behaves in a non-monotone manner as N_t increases; as N_t increases, the number of components inside the sum increases; however, the direction of change for each component in the sum can be either positive or negative, depending on problem parameters.

In Region 2, the expected profit to-go function is comprised of the sum of a constant number of components in N_t . Further, in a subregion $\tilde{N}_t(\gamma_{t-1}) \leq N_t$ of Region 2, the number of components of zero value increases because, in period $t + 1$, the number of successes needed to obtain the revenue g_{t+1} exceeds the leftover budget R_t , while each of the non-zero components inside the sum decreases as N_t increases, because the probability mass function $\Pr(N_t^+ = n_t^+)$ decreases at a higher rate than the CDF $\sum_{k=0}^{\min(N_t, \gamma_c(M) - \gamma_{t-1} - 1)} \Pr(N_t^+ = k)$. Hence, the expected profit to-go

function must decrease in N_t for some $N_t \geq \tilde{N}_t(\gamma_{t-1})$. This completes the proof. ■

Proof of Theorem 1. To simplify the notation, we omit conditioning on the event $\{\gamma_{t-1} < \gamma_c(M)\}$.

From Eq. (2), we have: $N_t^{*A}(\gamma_{t-1}, R_{t-1}) = 0 \Leftrightarrow$

$$\Pr(N_t^+(\gamma_{t-1}, R_{t-1}) \geq \gamma_c(M) - \gamma_{t-1}) \leq \frac{cN_t - [E[V_{t+1}(\gamma_{t-1}, R_{t-1})] \Pr(N_t^+(\gamma_{t-1}, R_{t-1}) < \gamma_c(M) - \gamma_{t-1})]}{g_t}.$$

From Lemma 2, we know that $\Pr(N_t^+(\gamma_{t-1}, R_{t-1}) \geq \gamma_c(M) - \gamma_{t-1})$ is increasing in γ_{t-1} . Therefore, if there exists a $\underline{\gamma}_{t-1}^A$ for which $N_t^{*A}(\underline{\gamma}_{t-1}^A, R_{t-1}) = 0$, then $N_t^{*A}(\gamma_{t-1}, R_{t-1}) = 0$ for all $\gamma_{t-1} \leq \underline{\gamma}_{t-1}^A$, and $N_t^{*A}(\gamma_{t-1}, R_{t-1}) > 0$ for all $\gamma_{t-1} > \underline{\gamma}_{t-1}^A$.

$E[\hat{p}_t]$ is derived from the expectation of the Beta distribution with parameters γ_{t-1} , and ζ_{t-1} .

This completes the proof. ■

Proof of Corollary. The proof follows directly from Theorem 1. ■

Proof of Lemma 5. Part (i) From the proof of Lemma A, $N_{t+1}^{*A}(\gamma_t, R_t) = 0 \Leftrightarrow \Pr(N_{t+1}^+(\gamma_t, R_t) \geq \gamma_c(M)) \leq \frac{cN_{t+1}}{g_{t+1}}$, and from the proof of Lemma B, $E[V_{t+1}(\gamma_{t-1}, R_{t-1})]$ is a sum of non-negative components. Thus,

$$\Pr(\underline{p}_t^A \leq \underline{p}_{t+1}^A(\gamma_{t-1}, R_{t-1})) = 1. \quad \blacksquare$$

Proof of Lemma 5. Part (ii) The proof follows directly from Lemmas 2 and 3, and Theorem 1.

■

Proof of Theorem 2. The proof follows directly from Lemmas 2 and 3, and Theorem 1. ■

Proof of Theorem 3. The proof follows directly from Lemmas 3 and A. ■

Sensitivity analysis on the reduction rate per period for g_t

Table B.1: Optimal first-period solution (N_1^*) for the two-period adaptive trial (the parameters used are specified in Section 5.1)

Revenue reduction rate per period	Prior ($E(\hat{p}_1)$)																		
	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95
1%	5	5	5	5	5	5	5	5	5	200	185	185	180	170	160	150	135	125	115
5%	5	5	5	5	5	5	5	5	5	200	185	185	180	170	160	150	135	125	115
50%	5	5	5	5	5	5	5	5	5	200	195	185	180	175	165	150	145	135	125