Stochastic Models to Optimize Biomanufacturing Operations

By

Tugce Gizem Martagan Akcay

A dissertation submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

(Industrial Engineering)

at the

UNIVERSITY OF WISCONSIN – MADISON

2015

Date of final oral examination: July 27, 2015

The dissertation is approved by the following members of the Final Oral Committee:

Ananth Krishnamurthy, Associate Professor, Industrial and Systems Engineering

Oguzhan Alagoz, Associate Professor, Industrial and Systems Engineering

Gregory DeCroix, Professor, School of Business

Christos Maravelias, Associate Professor, Chemical and Biological Engineering

Dharmaraj Veeramani, Professor, Industrial and Systems Engineering

Abstract

Biomanufacturing methods use live cells to manufacture vaccines and proteins. The use of live cells introduces several operational challenges, including uncertainty in yield and quality, random batch failures, and challenges in meeting specific purity and yield requirements for engineered drugs. In this thesis, we present optimization models to reduce costs and lead times in biomanufacturing operations.

First, we present a stochastic model that balances the risk of batch failures and yield-quality trade-offs to reduce costs in upstream biomanufacturing operations. We develop reliability models for random batch failures, and then provide an infinite horizon Markov decision model to derive the structural properties of the optimal operating policies. We develop stationary policies that closely approximate the optimal value function and are easier to implement in practice.

Second, we analyze a protein purification problem. In this setting, each order denotes an engineered protein having specific production requirements on the yield and quality. We develop a Markov decision model to optimize the pooling decisions for a fixed sequence of chromatography operations. We partition the state space into distinct decision zones that have similar financial characteristics, and then analyze the best starting material and optimal policies that would lead to guaranteed performance outcomes. We present zone-based optimal pooling policies that are easy to implement in practice, and discuss a state aggregation and an action elimination scheme leading to computational advantage in solving realistic industry problems.

Third, we consider the interaction between upstream fermentation and downstream purification operations. We first examine the downstream purification decisions where the joint decision on chromatography techniques and pooling windows are

identified to separate the protein of interest from multiple unwanted impurities. We develop a stochastic optimization model to identify the optimal choice of chromatography techniques and pooling windows at each purification step. Then, we analyze the upstream protein mass decisions, i.e., the best amount of protein to be manufactured through fermentation considering the uncertainty in yield and quality of the downstream purification operations. Based on the financial trade-offs between upstream and downstream operations, insights obtained from the downstream model are used to identify the best decisions for the upstream fermentation operations.

This research provides several contributions to theory and practice. First, the proposed models provide novel functional relationships between yield, quality and costs under various operating policies (optimal and suboptimal) in biomanufacturing operations. Next, we develop Markov decision models that capture both biology-level and manufacturing system level dynamics in a unified framework. We analyze the structural properties, and propose optimal guidelines for industry practices. Further, we develop approximations to solve large data sets. Lastly, we build models that address common challenges and issues encountered in the biomanufacturing industry. Our research findings have been developed in close collaboration with the biomanufacturing industry and have been implemented in practice. To facilitate industry implementation, software prototypes have been developed at Java. Implementation at Aldevron has resulted in 25% reduction in total lead times and 20% reduction in operating costs of protein purification operations on average. Applications of operations research are mostly new to the biotechnology community. We believe that as more companies embrace operations research, it will be an essential part of the protein research and development processes.

Acknowledgements

I would like to express my sincere thanks to my research advisor, Ananth Krishnamurthy. He has provided me continuous support and exciting opportunities to pursue my academic goals. He has also provided a constructive environment where I could peacefully conduct my research work. I greatly appreciate his support and mentorship in every aspects of my academic career over the last five years.

I would like to thank Oguzhan Alagoz, Gregory DeCroix, Christos Maravelias and Raj Veeramani for serving in my committee and providing me valuable feedback and comments. I would also like to thank Rajan Suri for his valuable insights and advise on my career. I would like to express my sincere gratitude to Burak Eksioglu, Sandra Eksioglu, Nihat Kasap and Gurdal Ertek for providing me very exciting opportunities in the earlier phases of my career, and offering me their continuous support since then.

I was very fortunate to have the guidance of our industry collaborators, Peter Leland and Tom Foti. Their feedback and support have been an essential factor for successfully completing this dissertation. I would also like to acknowledge the Madison's bioscience community, BioForward and Wisconsin Economic Development Corporation. This research has been conducted due to the generous support of the National Science Foundation. Further, I would like to thank Material Handling Education Foundation Inc, and Richard and Harriet Fein for their generous scholarships.

My friends, Ashesh, Isil, Merve, and Sanket provided me a wonderful company, and I would like to thank them for all their support and joyful friendship. It was also great to work with Michael Umhoefer, QRM staff and ISyE 641 students. I would like to thank Irem for her valuable friendship and support.

This thesis would not be possible without the continuous support and encouragement of my parents, Haluk and Aysen. They have been my source of inspiration, always holding my hand, and giving me peace. They have done their best to help me achieve my goals and dreams. I would also like to thank my brother and my grand parents for their support. My deepest thanks go to Alp, my best friend and my husband. He never ceased to support and believe in me. I feel very lucky to have him in my life.

To my family.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Biomanufacturing refers to the process of manufacturing therapeutic drugs using live systems, such as, bacteria, insect cells or eukaryotic cells. Recent advances in biomanufacturing, genetics and genomics have led to the development of novel biologic drugs for diseases like cancer, rheumatoid arthritis, and many others. Recombinant proteins and monoclonal antibodies are examples of such novel drugs that are used in the treatment of autoimmune disorders, various types of cancer, and cardiovascular diseases. Since the approval of the first recombinant protein in 1982, the biomanufacturing industry has grown considerably. In the area of biopharmaceuticals (or biologics) alone, more than 5,000 biopharmaceuticals are currently in research and development over the world, and the market analysis reports predict sales of biopharmaceutical drugs to reach beyond $140 billion by 2016 (Beuzekom and Arundel, 2009).

In contrast to the traditional pharmaceuticals that are chemically synthesized, biomanufacturing methods use live systems to produce these drugs. The use of live cells requires biomanufacturing operators to deal with several operational challenges, including batch failures, batch-to-batch variability, uncertainties in yield and quality,

challenges in meeting specific purity and yield requirements specified by the end use or application, and their impacts on costs and lead times. In this study, we develop stochastic optimization models that address these challenges, and show that the manufacturing policies obtained from our optimization models can be powerful in helping biomanufacturers to cope with uncertainties and trade-offs in their operations. We develop optimization models to analyze, optimize and control biomanufacturing systems, improve biomanufacturing efficiency, shorten lead times and reduce costs.

The rest of the chapter is organized as follows. We first provide an overview of the biomanufacturing operations in Section 1.1. Next, we introduce the operational challenges encountered in upstream and downstream operations, and then highlight the research questions in Section 1.2. We specify the approach adopted to address these questions in Section 1.3, and then conclude with the potential societal impacts of the research outcomes in Section 1.4.

1.1 Overview of Biomanufacturing Operations

Biomanufacturing typically consists of two major steps, upstream bioreactor operations, and downstream purification operations (see Figure 1.1). In the upstream bioreactor operations, the cell cultures produce the proteins of interest through fermentation. The fermentation process usually consists of distinct metabolic phases and is carried out in one or more bioreactors (McNeil and Harvey, 2008; Maiti et al., 2009). Bioreactors are typically stainless steel vessels that provide a controlled environment for cells to grow and produce the proteins of interest. At the appropriate time, the proteins of interest along with the byproducts (called as impurities) are harvested from the bioreactor. The yield and throughput of upstream bioreactor

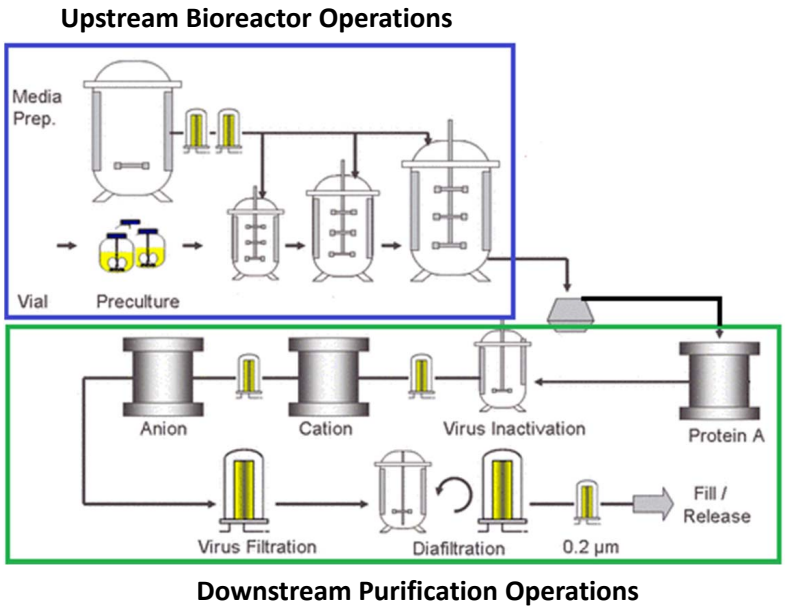

Figure 1.1: Typical manufacturing stages in biomanufacturing operations

operations is a function of the phases dictated by the cell physiology and the operating environment in the bioreactor. Depending on the characteristics of the protein of interest and impurities, the downstream operations could consist of several purification steps, such as, centrifugation, chromatography, ultrafiltration, and diafiltration. Among these, chromatography is a key technology, and most bio-separation processes contain at least one chromatographic step. A typical chromatography step uses a column packed with resins consisting of porous beads. The resins differ in their affinity to the protein and the impurities. Using multiple iterations through the chromatography column, this difference in affinity is exploited to separate the protein of interest from all other unwanted impurities. Downstream purification operations ensure that the final batch meets a predetermined purity requirement specified by the end use or application (i.e., the minimum acceptable quality standard that should be attained before delivering the final batch to the customer).

At each step of the manufacturing operations, the protein yield, throughput and quality of individual production runs are subject to significant variability due to

time varying nature of the cell lines, complex nature of the underlying biological and chemical reactions, stochastic failures due to process uncertainty and contamination. These dynamics at each individual stage makes optimization of biomanufacturing operations extremely challenging. Researchers and practitioners spend significant effort and costs in understanding the complex biology and chemistry to develop these products, and the sophisticated process monitoring and control necessary to determine a reliable manufacturing process that provides high yield and throughput.

Research and development efforts to arrive at stable manufacturing processes that guarantee predicable quality and yields are critical for the biomanufacturing industry. Process repeatability and control of quality variations are critical requirements of Current Good Manufacturing Practices (cGMP) and Process Analytical Technology (PAT) guidelines established by the Food and Drug Administration (FDA) for various levels of approvals. Addressing this need, we develop stochastic optimization models that capture both the biological and chemical characteristics of the underlying processes as well as the operational challenges including batch failures, uncertainties in yield and quality, and batch-to-batch variability. Our study shows that costs, quality, throughput and lead times in biomanufacturing operations could be improved through the effective application of the operations research tools and methodologies. By developing models and analyses for optimization and control of biomanufacturing systems and supply chains, this research leads to a knowledge base that would enable significant improvements in yield and cost reductions in biomanufacturing operations.

1.2 Operational Challenges, Opportunities and Research Questions

Modeling and optimizing biomanufacturing operations require specialized models that are capable of capturing unique trade-offs and challenges in biomanufacturing operations. A critical assessment of the literature and discussions with industry revealed several open issues related to modeling and optimizing biomanufacturing operations. In this section, we provide the biomanufacturing challenges typically encountered in practice, and then highlight the research questions addressed in this study.

Upstream Fermentation Operations

Biomanufacturing operations involve several operational challenges in upstream fermentation processes. First, the use of live cells in the fermentation processes introduce randomness and variability in yield and quality of the batch. The time varying nature of the cell lines implies that the cell culture goes through multiple metabolic phases, and the yield is a complex function of metabolic phases and time-varying parameters, such as, titer, biomass concentration, number of viable cells, growth rate and product formation rate (Narhi and Nordstrom, 2005).

Furthermore, the upstream fermentation system is very vulnerable to various factors ranging from batch failures due to inadequate cell density at inoculation, contamination, cell mutations, and random shock due to equipment failures, operator mistakes or other disruptions in the manufacturing environment. Penalty costs associated with batch failures could be very high in the biomanufacturing industry

due to expensive labor and materials.

Additionally, the critical decision to terminate the fermentation process and harvest the proteins are typically made on the basis of cell physiology and metabolic phases (namely, lag phase, exponential growth phase, stationary phase, deceleration phase, and death phase) in the biomanufacturing practices (Maiti et al., 2009; Gnoth et al., 2007). However, such policies ignore the impact of secondary growth of unwanted byproducts, batch-to-batch variability, uncertainty in yield and quality, and disturbances in the manufacturing environment (due to errors in media formulation, inadequate cell density at inoculation, contamination, mutation or equipment failures) while making harvesting decisions. If the batch obtained from the upstream fermentation contains excess amount of unwanted impurities (byproducts), then this would directly increase the operating costs of subsequent purification operations.

These challenges lead to the following research questions:

RQ 1: How can we develop comprehensive models to determine the optimal harvesting policies in the upstream operations to maximize yield and minimize costs?

RQ 2: Does harvesting at higher amounts of antibody always generate higher profit, or does the resulting amount of metabolic wastes and failure risks lead to costs that justify compromising on yield? When do risks and costs of quality outweigh additional revenue expected from higher yields?

RQ 3: How can we assess the robustness of bioreactor operations to random shocks, failures, and variability? How inefficient are the current harvesting policies? What are the critical degrees of failure risks that present opportunities for improvement, and what is the managerial importance of the harvesting policies that impose predefined limits on unwanted byproducts?

Downstream Purification Operations

Downstream purification operations involve several operational challenges and trade-offs. First, the number of purification steps and overall downstream costs (cost of resins, buffers, operating costs) are strongly impacted by the random yield and batch quality obtained from the upstream process as well as the limitations and randomness in the performance of subsequent chromatography operations and the requirements on the final yield and quality specified by the end use or application.

Secondly, the downstream operations consists of several purification steps and the costs of downstream purification operations can be prohibitive. For instance, the cost of raw materials (resins) used in the purification processes could be up to $4-5M (Farid, 2009). Similarly, biomanufacturing firms incur large penalty costs associated with yield shortages or quality failures. For example, customers might not purchase the batch of protein if it does not conform to the purity requirement.

Third, the decision maker performing the purification operations needs to identify the best choice of equipment (called as the chromatography technique) and the operating policy (called as the pooling window) to satisfy specific customer requirements on yield and quality. However, identifying the best operating policies are challenging since each chromatography technique and pooling window result in different yield and quality outcomes. The downstream purification operations consist of multiple steps in series, and hence the choice of chromatography technique and pooling window at one step affects the performance of purification in subsequent steps. Due to the large number of available chromatography techniques and pooling windows, the problem size could be very large (with millions of combinatorial choices

for purification policies) in most industry settings.

Additionally, individual upstream and downstream purification steps have their own yields, and the overall throughput and yield losses are function of the number of purification steps, and the yield and quality in each individual step (Farid, 2009). Due to the high penalty costs associated with yield shortages and quality failures, biomanufacturing companies need both the optimal policies and also other robust policies that deliver guaranteed performance.

Furthermore, the biomanufacturing industry also needs decision support tools that take into consideration the complex interaction between the upstream fermentation and downstream purification operations. For example, the final purity and yield of a batch depend on several inter-linked factors, such as, the condition of the starting material obtained from upstream, the limitations in the purification outcomes of the available chromatography techniques, production requirements on yield and purity specified by the end use or application, expensive operating costs, and high penalty costs associated with yield shortages and quality failures. Furthermore, upstream decisions that identify the amount of protein to be manufactured in the upstream fermentation processes strongly depend on the performance of the downstream purification operations due to the strict purity and yield requirements specified by the client. Therefore, a unified framework that considers these interactions and financial trade-offs between upstream fermentation and downstream purification decisions is required to reduce the overall costs and lead times in biomanufacturing operations. To the best of our knowledge, these interactions have not been adequately studied in a unifying stochastic optimization framework for biomanufacturing decisions.

These operational challenges lead to the following research questions:

RQ 4: How can we develop comprehensive models to determine both the optimal chromatography technique and the pooling window at each purification step to maximize profitability? For example, what are the optimal pooling windows given the choice of chromatoraphy techniques? What is the best choice of chromatography technique and pooling window to meet purity and yield requirements?

RQ 5: Can the biomanufacturing company determine in advance whether the purity and yield requirements specified by the customer are achievable with the starting material provided by the customer? If achievable, how confident can the biomanufacturing company be about meeting these specific requirements on purity and yield? Are there any performance guarantees? How can profitability be maximized?

RQ 6: How easy or tricky is the purification process likely to be, based on the starting material and limitations of the available purification techniques? What is the impact of the separation capabilities of the available chromatography techniques on costs and revenues?

RQ 7: How would the insights and policies be different, if the biomanufacturing company adopts a risk averse analysis, instead of a risk neutral probabilistic approach? How would the insights be different if the biomanufacturing company use deterministic models rather than stochastic models to identify the best chromatography techniques?

RQ 8: Can we develop models that capture the interaction between the upstream protein amount decisions and downstream purification decisions? What are the best upstream protein amounts that achieve the desired production requirements despite the process uncertainties in the downstream purification operations?

Implications on Biomanufacturing Supply Chain Contracts

The orders for engineered proteins are often placed by large pharmaceutical companies to small, specialized contract biomaufacturers. Although the large pharmaceutical companies often outsource the manufacture of engineered proteins, the manufacturing protocol is typically not fully specified at the time of order placement, because these proteins are uniquely engineered for research and development studies. The contract biomanufacturer therefore undertakes high risk of failure due to uncertainties in the process outcomes. If the contract biomanufacturer agrees to accept an order, then they often perform several initial test runs (called as the scouting experiments) to determine if and how the protein of interest can be manufactured to meet customer specifications. After successful scouting experiments, larger scale production runs are performed. Even during the production runs, the contract biomanufacturer incurs a risk of failure due to complexity and uncertainty in biomanufacturing operations. In Chapter 6, we discuss the potential ways in which the optimization models developed in this thesis could be incorporated in contract design.

1.3 Research Approach and Contributions

Our proposed approach provides and interdisciplinary framework that integrates knowledge related to the (i) biology and chemistry of the biomanufacturing processes, (ii) engineering knowledge about manufacturing systems and production planning, and the (iii) mathematical foundations of uncertainty theory, stochastic control and optimization. The production setting, modeling assumptions and managerial questions analyzed in this study are fairly general and validated through industry feedback. The research outcomes provide insights and guidelines that improve batch quality and yield of upstream bioreactor processes, and reduce penalty costs asso-

ciated with yield shortages and quality failures in downstream operations. We also compare the performance of the optimal policies for both upstream and downstream operations with the current industry practices, and underscore the importance of the operations research methodologies necessary to obtain improvements in biomanufacturing operations. In this section, we first elaborate on the collaboration team, and then discuss the analytical approach and contributions.

1.3.1 Collaboration Team

Due to the complex nature of the biomanufacturing operations, this research was conducted through multidisciplinary research team including academicians (Christos Maravelias from the Department of Chemical and Biological Engineering, Brian Flager from the Department of Chemical and Biological Engineering, Derek Hei from Waisman Biomanufacturing Center), several biomanufacturing companies, government agencies (Wisconsin Economic Development Corporation), and non-profit organizations (BioForward). We adopted a multidisciplinary effort that integrates the knowledge from operations management, stochastic modeling, and biological engineering. Our research has produced data sets and solutions to support common operational challenges encountered in the biomanufacturing industry. We use stochastic optimization tools and models to optimize profitability of the biomanufacturing operations, reduce batch failures, and improve lead times. Our collaboration with the local biomanufacturing industry includes: (1) Biomanufacturing working group sessions (BioWGS), and (2) One-to-one research projects carried out with industry partners.

Biomanufacturing working group sessions provide high engagement with local biomanufacturing companies. The optimization models and insights proposed in this research project have been formulated and validated by a large group of bioman-

ufacturers. For example, Cellular Dynamics International, Epicenter, Functional Biosciences, Gilson, Imbed Bio, Invitrogen, Life Technologies, Mirus Bio, Primorigen Biosciences, Semba Biosciences, ThermoFisher have been participating in these working group sessions. The main objectives of these working group sessions are to (1) to establish a platform to identify community issues and research questions relevant to the biomanufacturing industry, (2) validate the models and insights with a group of industry partners, (3) gather industry data, (4) implement the proposed optimization models and assess their impacts on business practices. Our research outcomes have already been recognized by media and non-profit agencies, including Xconomy (a business news forum in the U.S.) (Engel, 2014), BioForward (a leading Life Science community) (BioForward, 2014), and the Wisconsin Economic Development Corporation (WEDC, 2014). To facilitate industry implementations, software prototypes with friendly user interfaces have been developed. Feedback from our industry partners suggests that our insights are likely to transform the biomanufacturing practices.

These working group sessions have been complemented by one-to-one research projects that involve deeper level of engagement with our industry partners and collaborators. These includes weekly visits to Aldevron in order to analyze their operational challenges, collect data, build credible models that have high impact on practice, and validate the insights. We have been developing decision support tools and models to hedge against uncertainties in biomanufacturing operations and reduce operating costs.

1.3.2 Analytic Approach

To evaluate the trade-offs related to the batch quality, stochastic yield, fermentation operating strategies and the total costs, we develop models based on Markov decision processes. A key characteristic of these models is that they integrate the dynamics

of the protein growth with stochastic models for growth of unwanted byproducts and associated failures. These models are then used to derive the optimal harvesting policies for fermentation process (RQ1). Subsequently, we conduct several numerical analysis using fermentation data available in the literature to identify the optimal harvesting time based on the risks, cost, and yeild/quality trade-offs involved in the fermentation processes (RQ 2). We test the impact of the bioreactor reliability on the optimal costs and policies, and also assess the sensitivity of system performance to sub-optimal harvesting policies under various bioreactor reliability settings (RQ 3). These are summarized in Chapter 3.

Next, we develop stochastic optimization models to determine optimal purification strategies that maximize profitability, i.e, the best selection of the pooling windows given a sequence of chromatography techniques (see Chapter 4), and the best choice of both chromatography techniques and the corresponding pooling windows at each purification step (see Chapter 5)(RQ 4 and RQ 5). We develop optimization models to identify the optimal purification strategies that consider both the specific requirements on purity and yield as well as the limitations in the available chromatography techniques in Chapter 4. Then, we expand this model in Chapter 5 to include the choice of alternative chromatography techniques and multiple impurity types (RQ 4 and RQ 5). We use structural properties of the optimal policies and costs to derive practical guidelines for the purification decisions in Chapter 4 and 5. These guidelines quantify risks and costs in purification operations (RQ 6). We study not only the optimal policies but also guaranteed performance measures to achieve the purity and yield requirements specified by the end use or application. Through numerical analysis, we also investigate the impact of randomness on the chromatography technique selection problem (RQ 7). We build a stochastic optimization model that considers the interlinked nature of the upstream protein amount decisions and

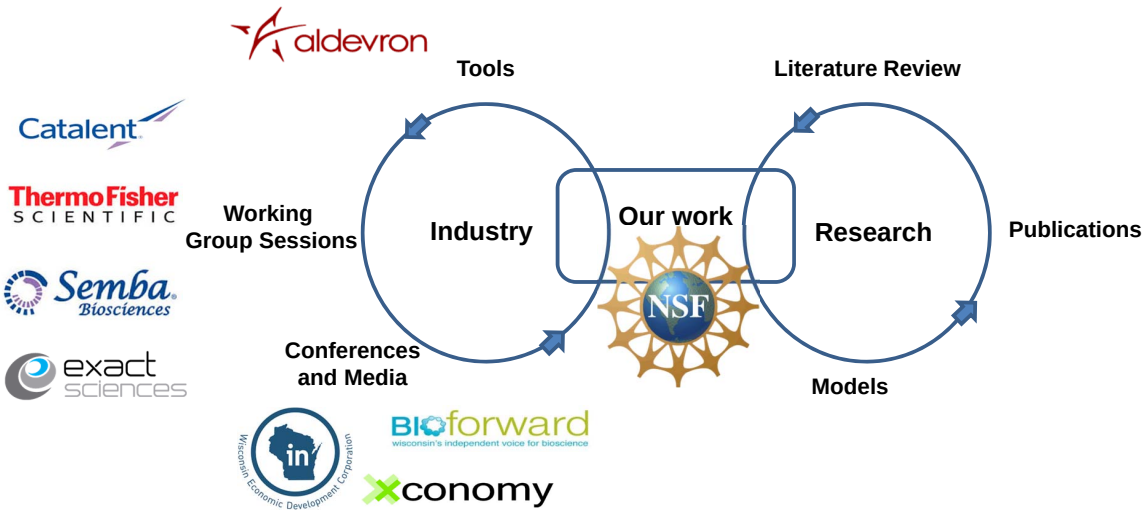

Figure 1.2: Our research approach

downstream purification decisions, and optimize the upstream decisions considering the randomness in downstream purification outcomes (RQ 8).

This study makes several contributions to both theory and practice. We provide models that provide functional relationships between yield, quality and costs under various operating policies (optimal and suboptimal policies typically used in practice). We develop Markov decision models that capture both the biology-level and manufacturing-systems level dynamics in a unified framework to optimize several biomanufacturing decisions in upstream and downstream operations. We analyze the structural properties of the Markov decision models, and propose optimal guidelines for upstream and downstream operating policies. We develop approximations to solve the Markov Decision models using realistic data sets obtained from the literature and industry; and validate our models and insights through discussions with a large group of biomanufacturers at the working group sessions and conferences.

Figure 1.2 summarizes our research approach. Our approach provides a collaborative framework integrating theory and practice. In terms of the theoretical research,

we start with performing a review of the literature and then build novel mathematical models leading to research publications. At the same time, our models and publications are enriched with industry input. We work closely with our industry partners (i.e., Aldevron) to develop specific tools to support biomanufacturing operations. We also share and implement our findings with a larger industry group through working group sessions and conference presentations. These outputs, in turn, enrich our research findings, our models and publications through building credible models by means of industry support, and also quantifying the impacts of our research outcomes on industry practices. To facilitate industry implementations, we developed a user-friendly software at Java. This software has been already in use for daily operations at Aldevron. The implementation of the purification optimization models at Aldevron resulted in 25% reduction in total lead times and 20% reduction in operating costs associated with protein purification on average. Applications of operations research are mostly new to the biotechnology community. We believe that as more companies like Aldevron embrace operations research, it will be an essential part of the protein research and development processes.

1.4 Potential Impacts on Society

Biopharmaceuticals can have significant impacts on well-being of patients by extending life and/or improving the quality of life. Market analysis reports predict sales of biopharmaceuticals to reach beyond $140 billion by 2016. Today, more than $5,000$ biopharmaceuticals are in development over the world (Beuzekom and Arundel, 2009). However, the industry is fast becoming the victim of its own success, since an emerging deficit in biomanufacturing capacity threatens to restrict the development and commercialization of these drugs (Dove, 2002). As biomanufacturing continues to advance, most companies are finding it vital to address biomanufacturing costs,

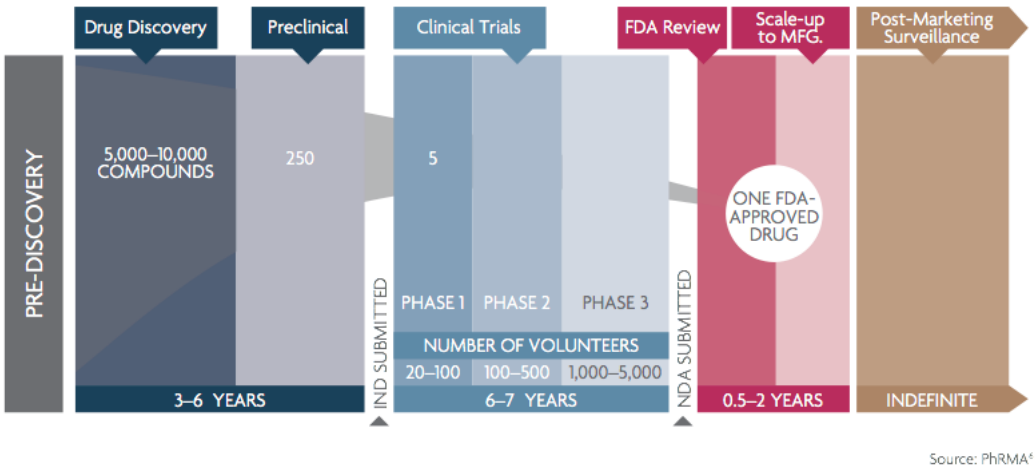

Figure 1.3: Typical timeline in biopharmaceutical research and development

capacity and time lines at multiple levels (Langer, 2009).

Figure 1.3 shows the typical timeline in a biopharmaceutical drug development project. As the figure illustrates, $3-6$ years are required for discovery and preclinical phase, whereas another $6-7$ years are spent for clinical trials. Additionally, according to Forbes, the average cost of bringing a new drug to market is \$1.3 billion (Herper, 2012). Contract biomanufacturing companies are involved in all stages of the research and development (i.e., several dedicated small R&D firms conduct research at discovery and preclinical phase, but also provide support for clinical trials and scale up). Our research outcomes would benefit not only the contract biomanufacturing companies but also the large pharmaceutical companies since they often outsource their research and development projects to contract biomanufacturers. Therefore, we believe that helping contract biomanufacturing companies with reducing their manufacturing costs and timelines could have significant impacts on increasing the speed for drug discovery and also reduce R&D expenditures of the whole supply chain. As Langer (2009) indicates, methods for lowering the costs of manufacturing, increasing the speed of scale-up for clinical testing, and decreasing the cost of product

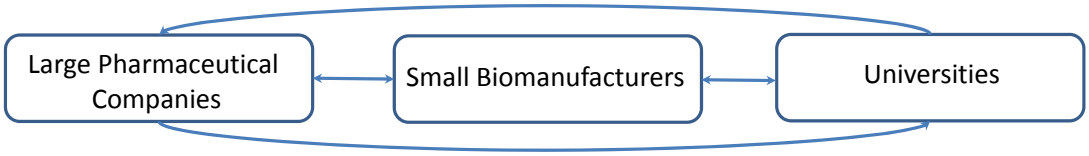

Figure 1.4: Entities in biomanufacturing supply chains

development are all important for factors to speed up the time-to-market.

The outcomes of this research are applicable across different biomanufacturing firms. Especially, most of our industry partners are small-sized companies with less than 50 employees, and have strategic importance in the biopharmaceutical supply chains. Market analysis shows that 75% of biomanufacturing R&D firms are small companies (Beuzekom and Arundel, 2009). Additionally, 85% of new drugs were developed by a small dedicated biomanufacturing firm that was later acquired by a large pharmaceutical company (Beuzekom and Arundel, 2009). Therefore, building optimization models that could reduce their costs and lead times could have significant societal impact by increasing the speed for drug discovery over the long run and the budget allocated for research and development. For example, throughout our collaboration, we have developed several tools to evaluate business risks, improve capacity planning decisions, and reduce lead times and costs. The implementation of the optimization models at Aldevron has resulted in 25% reduction in total lead times and 20% reduction in operating costs of protein purification operations on average. This has played a vital role as they have grown 3-fold during the duration of the collaboration. Applications of operations research are mostly new to both community. We believe that as more companies like Aldevron embrace operations research, regulatory authorities might mandate the use of such methodologies in protein research and development.

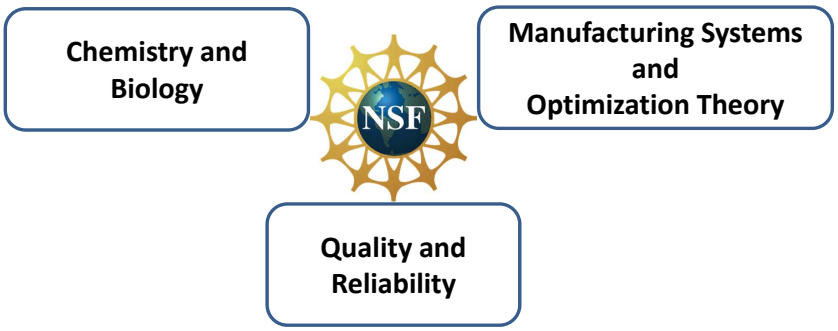

Figure 1.5: Interdisciplinary research approach for biomanufacturing operations

Figure 1.4 represents the main entities and information flow in biomanufacturing supply chains, and also positions the interaction between large pharmaceutical companies, small scale biomanufacturing companies, and universities. As the Figure 1.4 indicates, the large pharmaceutical companies often outsource some of their research and development projects to small biomanufacturers (contract biomanufacturers) to mitigate their risks and reduce their costs. Therefore, the research outcomes developed by the contract biomanufacturing companies are vital for R&D efforts conducted by large pharmaceutical companies. On the other hand, researchers at the universities develop therapies that might commercialize through the contract biomanufacturing companies. Furthermore, if such therapies are shown to be effective and promising, then these contract biomanufacturing companies are quite often acquired by the large pharmaceutical companies. Therefore, the information flow and supply chain dynamics shown in Figure 1.4 illustrate that contract biomanufacturers are critical entities; and hence reducing their costs and lead times could benefit the whole biomanufacturing supply chain by speeding up the time-to-market of therapeutics.

Studies in the biomanufacturing literature spend significant efforts in understanding the complex biology and chemistry behind these new drugs, but they do not typically evaluate the system-level performance of this complex manufacturing setting to address operational challenges and issues. To address this need, we develop

stochastic optimization models that provide a holistic approach in capturing the risks and trade-offs associated with the biomanufacturing operations. We adopt an interdisciplinary research approach that combines the knowledge from chemistry and biology, manufacturing systems and optimization theory, and quality and reliability engineering, as illustrated in Figure 1.5. This interdisciplinary approach provides a unifying framework for manufacturing system challenges and the underlying biology and chemistry behind biomanufacturing operations. Our research has been conducted due to the generous support provided by the National Science Foundation under the grant CMMI 1334933.

Chapter 2

Literature Review

In this chapter, we present an overview of the related literature on the upstream bioreactor and downstream purification operations, and highlight our contributions.

2.1 Related Work in Upstream Operations

In upstream bioreactor operations, there are two research streams that are closely related to this study: (i) modeling and control of fermentation systems (Section 2.1.1), and (ii) reliability modeling and optimization of systems subject to multiple dependent competing failure modes (Section 2.1.2).

2.1.1 Modeling and Control of Fermentation

The fermentation literature investigates the biology and chemistry behind bioreactor dynamics, and focuses on characterizing the best manufacturing techniques and protocol for a specific biologics of interest. The literature on fermentation systems includes both deterministic and stochastic models to optimize fermentation systems. In this section, we provide discussion on both deterministic and stochastic models.

Deterministic Models for Fermentation Systems: Deterministic models include kinetic process models for cell growth and product formation, such as, Monod-type equations and input-output models. For example, a large number of studies provide kinetic models consisting of nonlinear ordinary differential equations to predict the cell behavior and capture the inter-dependency between biomass concentration, growth rates, antibody secretion rates and substrate formation rates during fermentation (Patel et al., 2000; Jang and Barford, 2000; Liu et al., 2003; Tsao et al., 2004; Xing et al., 2010). Kinetic models are typically estimated from empirical fermentation data, and can be integrated in open-loop or close-loop fermentation control models to determine the prescribed recipe that describes the best manufacturing protocol for a specific biologics (Kawohl et al., 2007; Radhakrishnan et al., 1999). The literature on deterministic models of fermentation system often aims to develop control and optimization models to identify the best feeding strategies that either maximize the yield or minimize the difference between the desired process trajectory and the actual trajectory. For example, Luus (1993b) develops an optimal feeding policy to maximize the yield obtained from a batch; whereas Jenzsch et al. (2006) develop feeding strategies to control the specific growth rate at the desired set-point profile.

Stochastic Models for Fermentation Systems: Stochastic modeling and control methodologies have been studied in the chemical and biological engineering literature to model uncertainties in cell growth and product formation. For example, Kawohl et al. (2007) present model predictive control mechanisms using stochastic nonlinear ordinary differential equations and the extended Kalman filter approach. Gnoth et al. (2007) study the variability in biomass concentration and develop feedback control strategies to maintain the biomass concentration at its desired set-point profile. Only a few studies in the literature develop stochastic optimiza-

tion models to capture uncertainties in cell dynamics and optimize the yield and quality obtained from fermentation. For example, Luus (1993a) adopts a dynamic programming approach to identify the optimal feeding strategy that maximizes ethanol concentration subject to predefined constraints on batch volume and feed rates. Saucedo and Karim (1997, 1998) present an MDP formulation to optimize feeding policies in a fed-batch fermentation process. The authors provide an infinite horizon, total discounted cost MDP model to maximize ethanol concentration and reduce costs associated with feeding. The authors present a case study in ethanol production but do not investigate the structural properties of the optimal value function. Peroni et al. (2005) present an Approximate Dynamic Programming approach to maximize the yield and minimize the process time in a fed-batch fermentation systems. The authors develop total discounted profit-to-go function to optimize feed rates and identify optimal harvesting time. The profit-to-go function considers the profit associated with the product concentration accumulating inside the batch. Zero or negative profit-to-go indicate the harvesting decision. The authors encounter the curse of dimentionality because of continuum feeding rates and biomass concentration, and hence develop approximations to estimate the optimal value function.

Research Gaps: Stochastic models in the fermentation literature often do not consider the risks and costs associated with metabolic byproducts, random shocks and batch failures. Furthermore, we observe that studies that build stochastic models for fermentation operations often provide the model formulations and case studies, but do not analytically investigate the structural properties of the optimal value functions and operating policies. Furthermore, simulation studies have been used to understand the impact of batch failures and stochastic yield on quality and throughput (Petrides and Siletti, 2004; Saraph, 2003, 2004). However, these studies

evaluate system level dynamics, and do not optimize cell-level harvesting policies.

To conclude, we observe that the fermentation literature concentrates on cell-level models to understand the underlying biology and chemistry of these complex processes, but often ignores the high-level operational challenges encountered in the biomanufacturing industry. Fermentation studies typically focus on the dynamics of biological processes to determine ways to maximize the yield or minimize the deviations from the set-point profiles. However, high-level operational challenges that are also critical in biological decision making include the parallel accumulation of both yield and unwanted byproducts inside the batch, cost-quality trade-offs between upstream and downstream operations, random disturbances and batch failures (see Section 1). Our contribution in this study is to address this gap through an integrated stochastic model that captures both cell-level dynamics and high-level operational challenges. We develop realistic stochastic models to capture random disturbances on cell growth, failure modes, and the accumulation of both yield and unwanted metabolic wastes. We present a Markov decision model that evaluates various cost-quality trade-offs to identify the optimal fermentation time. The proposed MDP model identifies the optimal condition-based bioreactor harvesting policies that takes into consideration the risks and costs associated with yield, quality and batch failures. Furthermore, we analytically derive the structural properties of the value function and the optimal harvesting policies.

2.1.2 Reliability Modeling and Optimization

Reliability modeling and optimization of systems subject to multiple competing failure modes is another research stream that is closely related to this work. Failure processes involving shocks models or degradation models has been extensively studied in the reliability literature. Two shock models that are closely related to this work

are the extreme shock model, where the failure occurs when the magnitude of a shock exceeds a predefined threshold, and the cumulative shock models where the failure occurs when the cumulative damage (e.g., degradation) exceeds a predefined threshold (Barlow and Proschan, 1965; Nakagawa, 2007).

Shock and Failure Models: Neuts and Bhattacharjee (1981) provide one of the first papers where the survival function for a shock model is explicitly calculated, and a matrix-analtytic method is used to model a shock and wear process. Subsequently, several studies analyze extreme and cumulative shock models to investigate system failure and limiting average availability. For example, Klutke and Yang (2002) analyze a system with hidden failures, where the system degrades at a constant rate, and shocks cause additional degradation. The authors use regenerative arguments to analyze the limiting average availability of the system. Similarly, Wang et al. (2011) model systems where each shock result in a sudden increase in failure rate, and also analyze systems where each shock result in a random increase in the degradation path. The authors develop reliability models to represent the survival behavior of the considered systems. Kharoufeh and Cox (2005) model a degrading system where the rate of degradation is governed by a random environment. The authors derive system lifetime distribution and limiting average availability, and use real sensor data to estimate full and residual lifetime distributions. Similarly, reliability analysis of systems involving multiple catastrophic and degradation failure processes has received a considerable amount of attention in the literature (Huang and Askin, 2003; Li and Pham, 2005; Ye et al., 2011). Optimal time-based or condition-based replacement policies in systems with multiple failure modes have been considerably studied (Rangan et al., 2006; Huynh et al., 2011; Liu et al., 2013b).

Multiple Dependent Failure Processes: We observe that relatively fewer research have been carried out on multiple dependent failure processes. Sheui and Griffith (2002) study shocks that can lead to either minor failures removed by minimal repair or a catastrophic failure. They use renewal theory to develop an extended block replacement policy based and recommend repair or replacement decisions based on the number of shocks since the last replacement. Similarly, Montoro-Cazorla and Pérez-Ocón (2010) analyze shock arrivals following the phase-type distribution based on the number of accumulated shocks. In this setting, shocks deteriorate the system or can lead to a catastrophic failure. The authors study maintenance and replacement policies where corrective repair occurs when the system receives a prefixed number of N shocks, and a preventive repair occurs when the system has undergone $k < N$ nonfatal shocks. Subsequently, Montoro-Cazorla and Pérez-Ocón (2011) extend this analysis to study general distributions of repair times and random number of shocks before a fatal failure. Furthermore, Yu et al. (2014) study a maintenance problem of degrading systems under extreme shock but also consider the impact of procurement lead time while making order-replacement decisions.

There are a few closely related studies to our work that analyze multiple dependent shock failure processes. For example, Rafiee et al. (2013) analyze two failure processes caused by the same random shock process: soft failure due to continuous degradation of the system, and hard failure due to catastrophic impact of shocks. Four different shock patterns that could accelerate the degradation rate of the system are analyzed. The authors provide reliability models for different shock patterns but do not investigate maintenance and replacement decisions. Similarly, Peng et al. (2010) analyze the most related reliability setting to our work. The authors develop reliability models to represent soft failures and hard failures generated through the same shock process, and develop optimal inspection and maintenance policies to

reduce costs. The authors assume fixed costs of replacement, inspection and penalty cost of downtime; and use the renewal theory to identify the optimal inspection and maintenance policy to minimize the expected total cost. Our work differs from Peng et al. (2010) in that we model non-stationary degradation rates and survival probabilities that are mainly driven by the process time, and consider a more generic framework to model shock arrivals. Furthermore, we perform Markov decision analysis to identify the optimal condition-based replacement policies, where the costs and rewards are not fixed but a function of the system state.

Research Gaps: Complementing this vast literature on reliability modeling and analysis, our study considers the effects of the time elapsed since the last shock and its impact in terms of sudden and progressive failures of batches, however, our failure model and plausible remedial actions are different than the studies mentioned above due to the specialized application domain. Aforementioned studies consider either extreme shock models or cumulative shock models. Furthermore, studies involving multiple failure processes often assume independent failures where any of the failure process would cause the system to fail. However, upstream bioreactor operations have multiple dependent failure processes. In the upstream biomanufacturing setting considered, each shock arrival does not necessarily imply a sudden failure due to catastrophic shocks, but could accelerate the accumulation of waste metabolites over time leading to progressive failure. Furthermore, we complement the existing literature on multiple dependent failure processes by developing a generic framework that models non-stationary degradation rates and survival probabilities that are mainly driven by the process time, and use general distribution for shock arrivals.

Despite the vast majority of reliability studies that examine shock/wear models and replacement policies, relatively few studies have considered MDP and reliability

methodologies in a unified framework to analyze systems with complex failure dynamics (Kurt and Kharoufeh, 2010; Elwany et al., 2011; Ulukus et al., 2012). Unique features of the biologic systems also introduce interesting trade-offs between quality and yield in the optimal policy. For example, most classical condition-based replacement models aim to minimize costs associated with inspection, replacements, and failures caused by degradation and shocks (Gottlieb, 1982; Boland and Proschan, 1983; Özekici, 1988). However, accumulation of antibody concentration provides a compelling incentive to run the fermentation for longer duration despite the parallel growth of unwanted metabolic wastes, which has not been encountered in the classical cost-driven condition-based replacement models.

2.2 Related Work in Downstream Operations

Our study in downstream purification operations is related with a broader class of the dynamic programming literature for sequential decision making problems. For example, Bertsekas and Rhodes (1971); Puterman (1994); Bertsekas (2012) provide an excellent overview on the dynamic programming and sequential decision making under uncertainty. However, specific application of these stochastic optimization methodologies in the context of protein purification is limited in the existing operations research literature. Therefore, our literature review mainly focuses on the optimization models for the chromatography operations available in the chemical and biological engineering literature. We classify the relevant literature into two main streams: (i) Models for optimizing the pooling windows (Section 2.2.1), and (ii) Models for optimizing the selection and sequencing of the chromatography techniques (Section 2.2.2).

2.2.1 Optimal Selection of the Pooling Windows

Analysis of the Chromatography Output: Several studies analyze the purification performance of the chromatography operations to identify the yield and purity expected to be obtained using a specific chromatography technique. The decision on the pooling windows determines the fraction of the protein of interest is collected and the corresponding purity obtained at the end of a chromatography step as a result of a given pooling window. To identify the best pooling window decision, the existing literature focuses on characterizing the chromatography output to determine the quality and yield trade-offs associated with each pooling windows. For example, Ngiam et al. (2003) present a framework that captures the trade-off between the purity and yield in chromatography operations to determine the best pooling window strategy needed to meet the specific production requirements on the final purity and yield. They use chromatography data to build fractionation diagrams and also provide diagrams that represent the maximum purification factor versus the yield associated with each pooling strategy. The fractionation diagram denotes the relative change in the mass of protein of interest based on the total mass eluted; whereas the maximum purification factor versus yield diagram represents the trade-off between the purity and yield obtained by the end of a specific chromatography technique (Ngiam et al., 2001). Several other studies analyze the performance of chromatography techniques to model and quantify the purity and yield obtained by the end of a chromatography step (Vasquez-Alvarez et al., 2001; Salisbury et al., 2006).

Optimization Models: Another popular approach adopted in the literature is to optimize the pooling windows using mixed integer linear programming models (MILP). For example, Polykarpou et al. (2011a, 2012b) provide a MILP model that minimizes the number of purification steps and identify the best pooling window decision along with the best selection of the chromatography techniques to achieve

the minimum purity requirement specified by the end user. The proposed models consider the retention time and deviation factors in each chromatography technique to identify the amount of protein and impurity obtained as a function of the pooling windows selected at a chromatography technique. Similarly, Polykarpou et al. (2012a) build a MILP model to identify the best pooling window and chromatography technique, and uses a piece-wise approximation technique to linearize the optimization model.

Specialized models that capture the complex trade-offs and dynamics involved in multi-step chromatography operations have been studied in the literature. For example, Salisbury et al. (2006) provide a graphical methodology to identify the best operating decisions by taking into consideration the trade-offs between yield, purity and productivity in a two-step chromatography setting. The proposed graphical model identifies the optimal set for the best operating conditions in the first chromatography step that ensures that the material obtained by the end of the second step meets the desired specifications. Similarly, Huuk et al. (2014) provide a process flow optimization approach and identify the separation performance and the pooling windows in a multi-step chromatography setting. The authors use a case study of two consecutive ion exchange chromatography to demonstrate the benefits of the proposed model and optimization approach. Gao and Engell (2005) study an iterative optimization strategy to optimize the set points of a batch chromatography in the presence of set-point perturbations. The authors also provide a simulation study to illustrate the proposed set-point optimization approach in batch chromatography. A comprehensive overview of the other available model-based techniques to optimize and control chromatography operations is provided by Engell and Toumi (2005).

Research Gaps: The optimization of the downstream purification operations involves the models for optimal pooling windows and also optimal chromatography techniques. Therefore, we discuss the research gaps for downstream purification operations in Section 2.2.2, after we introduce the literature on the optimal selection of the chromatography techniques.

2.2.2 Optimal Selection of the Chromatography Techniques

Optimal Sequencing of the Chromatography Techniques: Several studies build deterministic optimization models to determine the optimal selection and sequencing of the chromatography techniques in a multi-step purification setting. For example, Vasquez-Alvarez et al. (2001) and Vasquez-Alvarez and Pinto (2004) develop mixed integer linear programming models to minimize the number of purification steps while achieving a predetermined purity level specified by the end use. These models also provide operating policies that maximize the purity for a given number of purification steps. The authors use physicochemical data associated with chromatography operations to identify the best selection and sequencing of the chromatography steps. Similarly, a mixed integer nonlinear programming model was developed by Lienqueo et al. (2009) to identify the optimal selection of the peptide purification tags. The objective of the study is to maximize the recovery of the protein of interest while minimizing the total costs associated with the purification steps. Nfor et al. (2013) study a framework for the optimization, evaluation and the rational elimination of the least feasible policies to minimize the number of purification steps while meeting the predetermined purity requirement.

Optimal Chromatography Techniques and Pooling Windows: There are a few studies available in the literature that focus on both optimizing the selection of the chromatography techniques and the pooling window at each purification step

simultaneously. Similarly, Polykarpou et al. (2012a,b) present a mixed integer linear programming model to optimize the best chromatography technique and the best pooling window at each technique to minimize the number of purification steps while attaining the pre-defined purity requirement. The authors use a case study consisting of 13 contaminants and 21 candidate steps to demonstrate the use of the model and quantify the computational efficiency of the proposed approach.

Optimal Process Design: Another stream in the literature focuses on the design decisions associated with the chromatography techniques to identify the optimal selection and sequencing of the purification steps and strategies. For example, Liu et al. (2013a) provide a mixed integer nonlinear programming model to optimize design decisions related with the cost-effective chromatography sizing strategies, and provide an industry case example to optimize design decisions in multi-column steps. Similarly, a mixed integer nonlinear programming model was developed by Liu et al. (2014) to optimize the sequencing and sizing of the chromatography operations. The authors optimize chromatography design decisions, such as, the number of columns, column diameter, bed height, and the number of cycles per batch. On the other hand, there are several simulation-optimization studies in the literature that provide stochastic models for chromatography operations and capture the randomness in purification operations to minimize costs. For example, Zhou et al. (2005) provide a framework consisting of mathematical modeling, computer simulation and optimization to quantify process trade-offs and assess the performance of the available operating strategies. Similarly, Nfor et al. (2009) present a novel approach on purification process development that uses biothermodynamics, high throughput experimentation, and simulation tools in order to provide a process understanding on biopharmaceutical manufacturing and respond quickly to quality and market demands. Furthermore, several simulation models are developed to

characterize the complex interaction between the upstream fermentation and several downstream purification operations (Saraph, 2001, 2003, 2004; Petrides and Siletti, 2004; Chhatre et al., 2007; Brunet et al., 2012). These simulation models are used to identify the best design decisions related with upstream fermentation, downstream chromatography and other biomanufacturing operations, including process improvement analysis, such as, process de-bottlenecking, throughput analysis and scheduling.

Research Gaps: Studies described in Section 2.2 often use the chromatography data as input for optimization models to achieve the desired purity level with the minimum number of chromatography steps while minimizing costs. However, the existing optimization studies focus on only minimizing operating costs, and do not capture the financial implications of risks and penalty costs incurred when the specific customers requirements are not achieved. In this study, we provide a unified framework that captures the financial risks, purity and yield requirements as well as the limitations in the available chromatography techniques and penalty costs to optimize upstream and downstream operating decisions. We derive guidelines that quantify risks and costs, and provide performance guarantees for achieving customer requirements on purity and yield. To our knowledge, such guidelines and performance guarantees have not been studied yet in the biomanufacturing literature. The proposed framework in Chapter 5 also considers the interlinked nature of the upstream and downstream operations to make the best decisions on the amount of protein that needs to be obtained at upstream operations by considering the randomness in the downstream purification outcomes. Several simulation studies that capture the interaction between different biomanufacturing steps are available in the literature, but there is a significant room for improvement for stochastic optimization models that mathematically capture the complexity of biomanufacturing systems along with

financial trade-offs and business risks.

To conclude, the research gaps identified and discussed in this section lead to the research questions in Chapter 1. In order to address the opportunities for improvement identified in our literature review, we develop stochastic optimization models that consider reliability issues and yield/quality trade-offs in Chapter 3. We develop a Markov decision model to identify the best pooling strategies for engineered proteins based on specific yield and purity requirements in Chapter 4. Furthermore, we develop an optimization framework to identify the best selection of the chromatography techniques and pooling windows, and also link these downstream purification decisions with upstream protein production decisions in Chapter 5.

Chapter 3

Harvesting Time Optimization Problem

3.1 Introduction

A typical biomanufacturing process is comprised of upstream operations where viable cells produce biologics of interest and downstream operations where the biologics are purified (See Figure 3.1). Upstream operations involve fermentation carried out in bioreactors where viable cells are mixed in a suitable media. The primary output from the fermentation process includes antibodies (or proteins) along with other metabolic wastes as byproducts. The batch obtained from upstream is often stocked in low temperature ($-80\,^{\circ}\mathrm{C}$) refrigerators. Based on specific customer orders, samples are drawn from this stock and purified through a series of downstream operations using

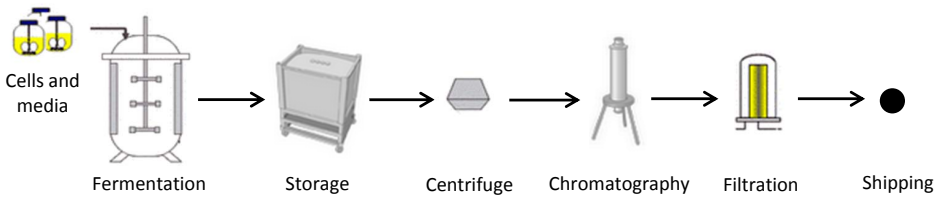

Figure 3.1: Typical biomanufacturing operations

centrifuge, chromatography and filtration. The objective of downstream purification is to separate the desired biologics from unwanted metabolic wastes, and ensure the final batch quality.

$$\text{Viable cells} + \text{Media} \xrightarrow{\text{Fermentation}} \text{Metabolic wastes} + \text{Antibodies of interest}$$

In practice, the critical decision to harvest the batch is typically made on the basis of cell physiology and metabolic phases, and the belief is that high yields is always better (Luus, 1993a; Saucedo and Karim, 1997; McNeil and Harvey, 2008). However, does harvesting at higher amounts of antibody always generate higher profit, or does the resulting amount of metabolic wastes and failure risks lead to costs that might warrant to compromise on yield? When do risks and costs of quality outweigh additional revenue expected from higher yields? Further, what are the critical degrees of failure risks that present business value for improving current harvesting practices, and what is the managerial importance of the predefined control limits for toxic byproducts?

In this chapter, we answer these questions through an integrated stochastic model that captures both cell-level dynamics and system-level tradeoffs. Our work makes several contributions: (i) We develop models to capture random disturbances on upstream biomanufacturing operations. We consider multiple dependent failure processes, and model the simultaneous growth of both yield and unwanted byproducts inside the same batch. Prior work do not consider the risks and costs associated with metabolic byproducts, antibody yield, random shocks and batch failures in a unified framework. (ii) We evaluate various yield/quality tradeoffs using a Markov decision model, and identify the optimal condition-based bioreactor harvesting policies. We incorporate unique characteristics of biomanufacturing operations, and analytically derive the structural properties of the optimal harvesting policy. In particular,

accumulation of the desired antibodies provides an interesting incentive to run the fermentation for longer duration, despite the parallel deterioration of the culture environment. The resulting yield and quality tradeoff is not encountered in the classical cost-driven condition-based replacement models. (iii) Our model integrates two streams of work (fermentation modeling and control, and reliability modeling and optimization) providing a multidisciplinary approach to this important problem. (iv) We demonstrate the use of the model by studying IgG_1 antibody production. We compare the performance of the optimal harvesting policy with popular policies typically used in practice. The analysis suggests that trying to maximize the desired antibody yield is not necessarily optimal from the profit perspective, which would be counter-intuitive to the most popular yield-driven harvesting policies. We leverage insights from optimal policies to develop smart stationary policies that are easier to implement in practice.

The remainder of the chapter is organized as follows. Section 3.2 introduces the operational challenges in biomanufacturing, and provides the background for formulating the bioreactor harvesting problem analyzed in this chapter. The mathematical model is formulated in Section 3.3, and structural properties of the optimal harvesting policy are analyzed in Section 3.4. Numerical studies in Section 3.5 evaluate the sensitivity of the value function to bioreactor reliability levels, and compares the performance of the optimal policy with other harvesting policies typically used in practice. Section 3.6 provides concluding remarks.

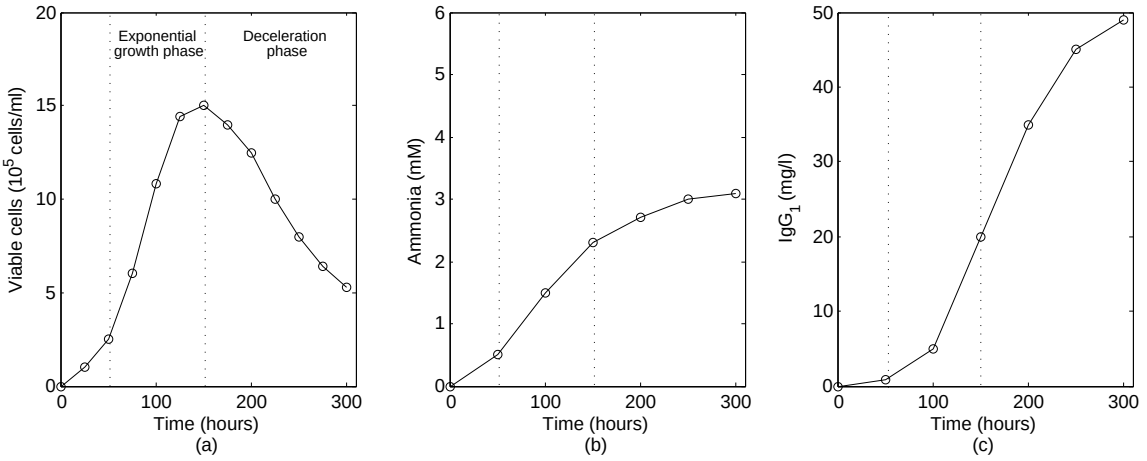

Figure 3.2: Expected amount of viable cells, ammonia and $\mathrm{IgG_1}$ over time (Ozturk et al., 1992)

3.2 Background on Biomanufacturing Operations

3.2.1 Bioreactor Dynamics

Figure 3.2 illustrates the dynamics of a fermentation process using the state information typically available to the decision maker. The figure presents the expected amount of viable cell density, ammonia and $\mathrm{IgG_1}$ antibody levels over time (Ozturk et al., 1992). The time evolution of the viable cell density in Figure 3.2(a) shows that cells undergo several physiological phases during fermentation. For example, the cell culture is observed to enter an *exponential growth phase* at time $t = 50$ hr, and a *deceleration phase* at time $t = 150$ hr. During the deceleration phase, viable cell density drops significantly from 15×10^5 cells/ml to 5×10^5 cells/ml. The cells enter a *death phase* at time $t = 300$ hr which is the final harvesting time for fermentation.

In parallel with physiological phases, cell cultures typically produce toxic metabolic wastes during fermentation. Ammonia is an example of such metabolic wastes accumulating inside the batch, as shown in Figure 3.2(b). Accumulation of metabolic wastes such as ammonia is often an indicator of a deteriorating culture environment

(Tsao et al., 2004). Excessive formation of these byproducts could lead to cytotoxic effects and poor batch quality (Newland et al., 1994; Yang et al., 2000). Current practices typically aim to minimize the formation of waste metabolites during the fermentation process, and maintain its concentration below a predefined threshold (Yang et al., 2000). We use the term *baseline-metabolite profile* to denote the expected concentration of waste metabolites during the fermentation process. The baseline-metabolite profile represents the expected amount of waste metabolites accumulated over time due to inherent cell physiology (See Figure 3.2(b)). However, the actual evolution of waste metabolites is influenced by other extraneous factors (such as disturbances due to cellular activities and environmental conditions) as outlined in Section 3.2.3.

IgG_1 antibody concentration in Figure 3.2(c) represents the yield obtained from the batch. In this example, the cell culture starts producing the antibodies of interest towards the middle of exponential growth phase ($t = 100$ hr), and achieves the highest antibody productivity during the deceleration phase. For some cell cultures, studies provide empirical evidence that the evolution of metabolic wastes and antibody secretion rates are two independent processes evolving over time (Omasa et al., 1992; Ludemann et al., 1994; Tsao et al., 2004; Xing et al., 2010). For example, IgG_1 production rate in Figure 3.2 was reported to be independent of the ammonia levels.

3.2.2 Quality Specifications and Control

Biomanufacturing operations need to abide by Food and Drug Administration (FDA) approved *manufacturing protocols* to guarantee the final product quality. The manufacturing protocol represents all the manufacturing methods, procedures, process parameters and their corresponding specification limits through the course of the manufacturing process. The protocol requires that critical process parameters

lie within their acceptable ranges to ensure product quality (Rathore and Winkle, 2009; Shivhare and McCreath, 2010). For this purpose, continuous monitoring and control of several physico-chemical process parameters, such as pH, temperature and ammonia has become a routine practice (Tsao et al., 2004). However, due to limitations of available sensor technologies, estimation of the biological parameters still need to done through offline measurements carried out at specific instants during the course of the reaction. Assessing the state of the reaction allows the operator to estimate the risk of batch failures and determine if the reaction should be terminated (harvesting available output) or continued to proceed as per the prescribed manufacturing protocol. A *batch failure* occurs when the critical process variables fall outside acceptable limits defined by the manufacturing protocol. Maintaining the critical process parameters within their acceptable limits could be challenging due variability in fermentation systems. For example, Gnoth et al. (2007) operate 12 batches using the same control strategy, and report high variations in product concentrations at each successive run. In upstream biomanufacturing, this batch-to-batch variability is mainly attributable to stochastic, non-linear growth pattern of cells.

In this chapter, we address the bioreactor operator's problem of maximizing profit while still operating within the approved manufacturing protocol. Depending on the state of the reaction, the protocol includes a prescribed set of activities to maintain the desired process trajectory, such as, feeding the cells, adding fresh media, adjusting temperature and pH, etc. In addition to undertaking actions prescribed in the protocol, she can also adjust harvest times for a batch (thereby terminating the fermentation) to minimize failure costs.

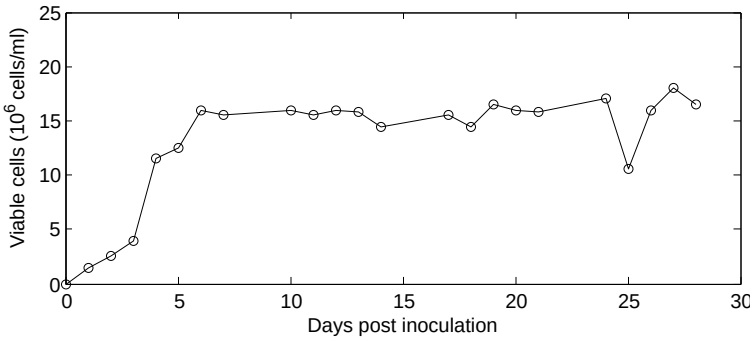

Figure 3.3: Fermentation with a random shock on day 25 (Yang et al., 2000)

3.2.3 Reliability Issues

Although bioreactors are highly-controlled environments, process observations indicate that fermentation systems are subject to random shocks due to equipment failure and/or biology-induced factors. These shocks cause sudden and sharp deviations in viable cell density. For example, Figure 3.3 shows the fermentation data for the clinical production of a monoclonal antibody (Yang et al., 2000). The batch is controlled to maintain a viable cell density of $15 \pm 3 \times 10^6$ cells/ml. On day 25, an unexpected sharp drop in the viable cell density incurs due to shocks that were caused by fouling. Random shocks might cause severe damages leading to either *sudden failure* or *progressive failure* in fermentation systems.

Sudden failure occurs when a batch deteriorates as a result of a shock, and healthy cells die massively and abruptly. Examples of sudden failure could be viral contamination and cell mutation. A *progressive failure* occurs when the amount of byproducts accumulated inside the batch exceeds a deterministic, predefined threshold level. Progressive failures represent failure to meet quality specifications, i.e. if the amount of waste metabolites exceeds the predefined threshold level, the batch becomes impure and toxic due to excessive byproducts. Examples of progressive failure include cytoplasmic acidification and glycosylation. In case of batch failure,

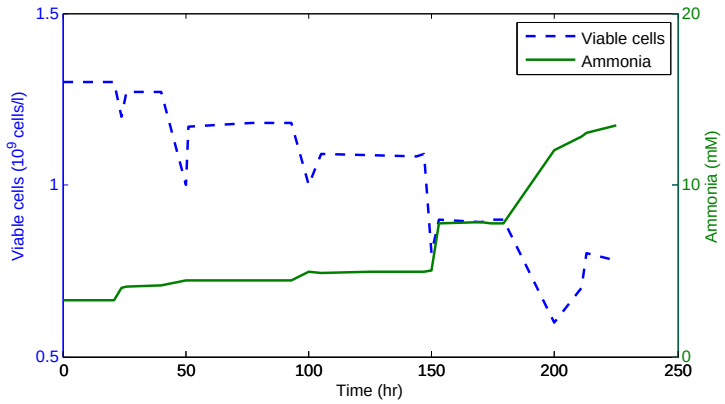

Figure 3.4: Step changes in ammonia due to shocks in viable cells (Newland et al., 1994)

the batch is discarded, and cannot be processed further in downstream.

Even though a shock may not cause a sudden failure, it could still trigger a *progressive failure* by causing step increases in the amount of metabolic wastes. For example, Figure 3.4 illustrates such random increments in ammonia levels due to shocks in viable cell density. Studies have shown that about 80% of all batch failures are due to contamination, operator error, equipment failure and failure to meet specifications (Langer, 2008). Sudden and progressive failures defined in our problem setting account for these most common failure modes. The risks in biomanufacturing operations are typically considered as *low frequency, high-impact risks* (CBI, 2010), as failures could cost up to $1 Million per occurrence in a large scale facility (Langer, 2008). The high cost of failures are mainly due to expensive resources, and sterilization activities during which the production line could stop until the batch is securely discarded and the facility is sterilized.

3.3 Model Formulation

We present a mathematical model for fermentation systems. The objective is to identify the best harvesting policies to maximize total discounted profit. We develop

a discrete time, continuous state space, infinite horizon Markov Decision Process (MDP) model. Section 3.3.1 presents the notation used to build the MDP model. The evolution of monoclonal antibodies is modeled in Section 3.3.2. A bioreactor reliability framework is presented in Section 3.3.3 to model batch failures and metabolic waste concentrations. Models developed in Sections 3.3.2 and 3.3.3 are then integrated in the optimality equation at Section 3.3.4 to identify the optimal bioreactor harvesting policy.

3.3.1 Decision Epochs, States, Actions and Rewards

Decision epochs The bioreactor operator performs measurements at each decision epoch $\boldsymbol{T} = \{t : 0, \tau, 2\tau, 3\tau, \ldots, T\}$ for $\tau > 0$. The time interval τ between any two successive decision epochs is called as a period. The length τ of a period is determined by process-specific characteristics of cell lines, media, and bioreactor operating module; and could range from minutes to days. Decision epochs denote the age of a fermentation process. The maximum age of a batch is bounded by T due to limitations imposed by cell viability, nutrient deficiencies and growth inhibitors (see Figure 3.2).

State space The state of the fermentation is denoted by the ordered triplet (n_t, w_t, m_t) on finite state space $\boldsymbol{N} \times \boldsymbol{W} \times \boldsymbol{M}$, where $\boldsymbol{N} = \{n : 0, \tau, 2\tau, 3\tau, \ldots, T\}$, $\boldsymbol{W} \equiv \left[0, \bar{W}\right) \cup \{\Delta\}$, and $\boldsymbol{M} \equiv \left[0, \bar{M}\right]$. The state $n_t \in \{0, \tau, 2\tau, 3\tau, \ldots, T\}$ for $\tau > 0$ represents the discrete time periods elapsed since the last shock. The state $n_t \in \boldsymbol{N}$ determines the arrival rate of shocks, as described in Section 3.3.3. We let $w_t \in \left[0, \bar{W}\right)$ represent the amount of waste metabolites in the batch at a given time $t \in \boldsymbol{T}$. We assume that fermentation starts at time zero with no waste metabolites inside the bioreactor, $w_0 = 0$. The state space $\boldsymbol{W}$ is bounded by a deterministic, predefined threshold value $\left\{\bar{W}\right\}$ in accordance with federal standards on batch

quality. The state $\{\Delta\}$ denotes a batch failure state. For notational convenience, we define $\boldsymbol{S} \equiv \boldsymbol{W} \setminus \{\Delta\}$. Let $m_t \in \left[0, \bar{M}\right]$ denote the concentration of monoclonal antibodies (mAbs) at time $t \in \boldsymbol{T}$. We assume that fermentation starts at time zero with no mAbs, $m_0 = 0$. Note that mAbs accumulate inside the batch until the batch is harvested or abandoned due to batch failure. The maximum achievable mAbs concentration in a batch culture is bounded by $\bar{M}$ due to limitations in cell viability, specific growth rate and antibody production rate, etc. (McNeil and Harvey, 2008).

Action space The action space $\boldsymbol{A} = \{C, H\}$ consists of two actions. Action $\{C\}$ represents the action of continuing the fermentation according to the established manufacturing protocol. This could include a prescribed set of activities that influence the fermentation process (such as, adding fresh media, adding cells, feeding cells, clearing issues with fouling and clogging, adjusting pH and temperature levels, etc.). These corrective actions in the protocol aim to keep the fermentation dynamics in its desired trajectory. We use $\{C\}$ to denote the set of all prescribed actions in the protocol for that state to continue to operate the fermentation. However, the operator can decide to harvest the batch in order to avoid future batch failures. Action $\{H\}$ represents this decision to terminate the upstream fermentation process for that batch.

Let $a_t\left(n_t, w_t, m_t\right)$ be the action taken at the decision epoch t and state (n_t, w_t, m_t). Failed batches at any decision epoch $t \in \boldsymbol{T}$ are immediately harvested, i.e. $a_t\left(n_t, \Delta, m_t\right) = H$ for all $n_t \in \boldsymbol{N}$ and $m_t \in \boldsymbol{M}$. Similarly, a batch is harvested if it reaches either the maximum fermentation time, $a_T\left(n_T, w_T, m_T\right) = H$, or the maximum mAb concentration, $a_t\left(n_t, w_t, \bar{M}\right) = H$, or the specification limit for the waste metabolites, $a_t\left(n_t, \bar{W}, m_t\right) = H$, for all $t \in \boldsymbol{T}$ and $w_t \in \boldsymbol{S}$. Harvested batches are immediately replaced with the new ones, resulting in $t = n = 0$ and

$w_0 = m_0 = 0$.

Fermentation operating costs The function $r_c(w_t, m_t) = K_c + f(w_t, m_t)$ represents costs of operating the fermentation during one period when the batch is in state $(n_t, w_t, m_t) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$, and is independent of n_t and t. K_c denotes fixed costs of operating the bioreactor during one period, i.e., process monitoring (running and maintaining the sensors, analytics), equipment and labor costs, clean room charges (sterilizing and maintaining the clean room). $f(w_t, m_t)$ represents variable costs of operating the fermentation at state (n_t, w_t, m_t). Variable costs are typically associated with raw material costs (i.e., cost of media, buffers, feeds, and cell lines), inspection costs and quality control (online and offline sampling, analytics) (Farid, 2007; Rathore et al., 2004; Werner, 2004). The function $f(w_t, m_t)$ could have fairly general cost structures.

Cost of batch failure Penalty cost associated with batch failure is denoted by $r(\Delta)$ for all $(n_t, w_t, m_t) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$. We assume that failed batches are discarded to ensure safety. We model $r(\Delta)$ as a predefined parameter independent of (n_t, w_t, m_t) and t. $r(\Delta)$ includes costs associated with sterilization efforts, and costs of initial cell lines, inoculation, buffers, labor and opportunity costs (Langer, 2008). Failure cost can be significantly high due to sterilization efforts, resource costs and other opportunity costs. For example, in case of viral contamination, the production line can stop until the batch is securely disposed and the facility is sterilized. This could translate into various hidden costs, such as, loss of reputation, disruptions in the production plan, and re-scheduling efforts.

Rewards obtained from harvesting The reward function $r_h(w_t, m_t) = r(m_t) - g(w_t)$ captures both the revenue $r(m_t)$ obtained from the batch when it is har-

vested at yield m_t, and the purification costs $g(w_t)$ related with the concentration of waste metabolites w_t. The purification costs include raw material costs (costs of resins and buffers), equipment costs (costs of running a chromatographic cycle, filtration), labor costs (setting up and monitoring the chromatographic runs), cost of quality assurance and control activities (high performance liquid chromatography, labor costs, analytics, documentation), and clean room charges. Among these cost components, raw material costs (especially the cost of resins used in chromatographic runs) tend to dominate the total cost of downstream operations (Farid, 2008).

The amount of waste metabolites represents the batch quality obtained from upstream, and is considered as one of the main drivers of the purification workload in downstream operations. For example, batches with high levels of impurities would require multiple chromatographic runs and inspection steps to achieve the required levels of purity (Kalyanpur, 2002; Werner, 2004; Farid, 2007). This would translate into higher raw material costs, equipment cost, labor and inspection costs in downstream. Therefore, the reward function $r_h(w_t, m_t)$ captures both the revenue obtained from the yield and downstream purification costs related with waste metabolites, and is independent of n_t.

3.3.2 Evolution of Monoclonal Antibodies

The amount of monoclonal antibodies inside the batch increases over time as a consequence of cell physiology. We define $\{M_t, t \in \boldsymbol{T}\}$ as a stochastic process that represents the evolution of mAbs. If the bioreactor operator decides to continue to operate the fermentation according to the prescribed manufacturing protocol at time t where $0 \leq t < T$, then a random amount of mAbs x_t accumulates by the next decision epoch $t + 1$. Therefore, the amount of mAb at time $t + 1$ can be modeled using the

additive function in Equation 3.1:

$$m_{t+1} = \begin{cases} 0 & \text{if the batch is harvested,} \\ m_t + x_t & \text{if the batch is not harvested at time } t. \end{cases} \tag{3.1}$$

where m_t represents the amount of mAb at time t, and the random variable X_t denotes the random amount of mAbs produced during period $[t, t+1)$. We model X_t as a continuous random variable whose probability density function is $f_t^x(\cdot)$ with a general distribution, and support in the interval $\left[x_t^L, x_t^U\right]$ for all $t \in \boldsymbol{T}$. Random increments in the amount of mAb at any time t is bounded such that $0 \leq x_t^L$ and $x_t^U \leq \bar{M}$ as a consequence of the limited cell productivity, substrate limitations and inhibition (Jenzsch et al., 2006). Probability distribution function $f_t^x(\cdot)$ is typically estimated based on fermentation data (Saucedo and Karim, 1997; Ozturk et al., 1992).

3.3.3 Bioreactor Reliability Modeling

In this section, we present a bioreactor reliability model that is used to predict the evolution of waste metabolites and failure risks on a finite, irreducible discrete time Markov chain. We first model the amount of waste metabolites accumulating in the batch due to shocks and cell physiology. Next, we model the shock arrival process, and finally describe the bioreactor reliability function.

Evolution of waste metabolites: Let $\{W_t, t \in \boldsymbol{T}\}$ be a non-stationary stochastic process that represents the evolution of waste metabolites. The total amount of waste metabolites w_t accumulated by time t is a function of two random quantities, B_t and E_t. If the bioreactor operator does not harvest the batch at time t, then the amount

of waste metabolites w_{t+1} at time $t+1$ can be represented using Equation 3.2:

$$w_{t+1} = \begin{cases} w_t + b_t + e_t & \text{if there is a shock during } [t, t+1)\,, \\ w_t + b_t & \text{if there are no shocks during } [t, t+1)\,. \end{cases} \tag{3.2}$$

The random variable B_t represents the additional amount of metabolic wastes secreted due to the baseline metabolic activities during period $[t, t+1)$, and the random variable E_t denotes the amount of metabolic waste generated as a result of the shocks (see Section 3.2.3). B_t is modeled with general distribution $f_t^b(\cdot)$ and support $[0, \infty)$. The baseline-metabolite profile represents the expected trajectory of the concentration of waste metabolites, and is associated with basal metabolic activities during cell growth and reproduction. Similarly, the random variable E_t represents the amount of metabolic waste generated as a result of a shock during period $[t, t+1)$. We assume that E_t is a random variable with general distribution $f_t^e(\cdot)$ and support $[0, \infty)$ for all $t \in \boldsymbol{T}$. Probability distribution functions $f_t^b(\cdot)$ and $f_t^e(\cdot)$ could be estimated based on fermentation dynamics (Omasa et al., 1992; Ozturk et al., 1992; Yang et al., 2000).

Shock arrival process and the evolution of n: The amount of waste metabolites W_t is a function of the shocks on the fermentation system. Let N be a random variable denoting the time between two consecutive shocks, and ρ be a vector of parameters of the unspecified distribution of the time between two shocks. We define $\zeta(n, \rho) =: P(N = n | N \geq n)$ as the probability of a shock in the next time epoch, given that the time elapsed since the last shock is n and the bioreactor operator decides to continue the fermentation process. We note that the function $\zeta(n, \rho)$ could have a fairly general structure. This modeling flexibility is necessary because each bioprocess has its unique failure characteristics and hazard rates. These characteristics are commonly determined by the cell lines, buffers, media, bioreactor operating modules, etc. For example, mammalian cell cultures are likely to have more frequent shock

arrivals than other types of cell lines; whereas single use expression systems have less likelihood of incurring sudden failures such as cross-contamination. As a special case, we let $\rho = (p, \theta)$ and define $\zeta(n, \rho) = 1 - (1-p)^{(n+1)^\theta - n^\theta}$, which is also known as the discrete Weibull distribution (Nakagawa and Osaki, 1975). Characteristics of the discrete Weibull distribution described in Property 3.3.1 allows the flexibility in modeling shocks with increasing ($\theta > 1$), decreasing ($\theta < 1$), or constant ($\theta = 1$) hazard rates for shock arrivals.

Property 3.3.1. *The function $\zeta(n, \rho) = 1-(1-p)^{(n+1)^\theta - n^\theta}$ is an increasing function in n for $\theta > 1$ and a decreasing function in n for $0 < \theta \leq 1$. When $\theta = 1$, shock arrival probability becomes independent of n and reduces to $p \in (0, 1)$, such that, the distribution of N can be represented with geometric distribution.*

Bioreactor reliability function $\boldsymbol{R_t(n, w)}$: We model sudden and progressive failures that may occur due to shocks and their impact on metabolic wastes. Shocks have two possible impacts: *(i)* a shock can either cause a sudden failure due to massive cell deaths and mutation, *(ii)* a shock can accelerate the increase in the amount of waste metabolites to result in progressive failure eventually. We model these two impacts using the bioreactor reliability function. The bioreactor reliability function represents the probability of no batch failures until the next decision epoch, if the bioreactor operator decides to continue the fermentation process. For notational convenience, we suppress the subscript t on state (n_t, w_t, m_t) hereafter, and let $R_t(n, w)$ denote the probability that a batch survives both sudden and progressive failures at state $(n, w) \in \boldsymbol{N} \times \boldsymbol{S}$ during the time period $[t, t+1)$. Then,

$$
\begin{aligned}
R_t(n,w) &= \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w} \left[\int_0^{\bar{W}-w-b} f_t^e(e)\,\mathrm{d}e \right] f_t^b(b)\,\mathrm{d}b \\
&\quad +(1-\zeta(n,\rho)) \int_0^{\bar{W}-w} f_t^b(b)\,\mathrm{d}b \qquad (3.3) \\
&= \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\,\mathrm{d}z + (1-\zeta(n,\rho)) \int_0^{\bar{W}-w} f_t^b(b)\,\mathrm{d}b \qquad (3.4)
\end{aligned}
$$

where, α_t represents the probability that a shock does not cause a sudden failure during the period $[t, t+1)$. For notational convenience, Equation 3.4 uses the convolution $Z_t = E_t + B_t$ with the density function $f_t^{e+b}(z)\mathrm{d}z = f_t^e * f_t^b$. The random variable Z_t models the random increments in the amount of metabolic wastes due to combined impacts of the baseline-metabolites B_t, and random shocks E_t. The bioreactor reliability function establishes the relationship between the risk of sudden failure $\{\alpha_t\}$ and the probability of progressive failure $\left\{w_t : w_t \in [\bar{W}, \infty)\right\}$, through the same shock process $\{\zeta(n,\rho)\}$. The reliability function defined in Equation 3.4 also provides the basis for modeling the evolution of waste metabolites through time.

Property 3.3.2. *The bioreactor reliability function $R_t(n,w)$ is a decreasing function in n for $\theta > 1$, and an increasing function in n for $0 < \theta \leq 1$. When $\theta = 1$, the reliability function becomes independent of n, and can be represented with a constant number across all $n \in \boldsymbol{N}$.*

Property 3.3.2 presents the monotonicity characteristics of $R_t(n,w)$ in n when $\zeta(n,\rho)$ has discrete Weibull distribution. However, insights from Property 3.3.2 is fairly generic and would hold for other shock processes with increasing ($\theta > 1$), decreasing ($\theta < 1$), or constant ($\theta = 1$) hazard rates. Characteristics of the bioreactor reliability function presented in Property 3.3.2 allows us to analyze the structural properties of the optimal policy in Section 3.4.

3.3.4 The Optimality Equation

Let $\mathcal{V}_t(n, w, m)$ be the expected total discounted profit when the batch is in state $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{W} \times \boldsymbol{M}$ at time $t \in \boldsymbol{T}$. The objective is to identify the optimal condition-based harvesting policy to maximize the value function $\mathcal{V}_t(n, w, m)$. The value function considers both the revenue obtained from yield of mAbs and the costs associated with purifying metabolic wastes. For all $t \in \boldsymbol{T} \setminus \{T\}$ and $m \in \boldsymbol{M}$, the optimality equation is expressed as follows (Puterman, 1994):

$$\mathcal{V}_t(n, w, m)$$

$$= \begin{cases} \max\left\{r_h(w, m) + \mathcal{V}_0(0, 0, 0),\ -r_c(w, m) + \beta\mathcal{C}_t(n, w, m)\right\} & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0(0, 0, 0) & \text{if } w = \Delta. \end{cases} \tag{3.5}$$

where, $\mathcal{C}_t(n, w, m)$

$$\begin{aligned} &= \left[1 - R_t(n, w)\right]\left[-r(\Delta) + \mathcal{V}_0(0, 0, 0)\right] \\ &+ \zeta(n, \rho)\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x) \mathcal{V}_{t+1}(0, w + z, m + x)\,\mathrm{d}x\,\mathrm{d}z \\ &+ \left[1 - \zeta(n, \rho)\right] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x) \mathcal{V}_{t+1}(n + 1, w + b, m + x)\,\mathrm{d}x\,\mathrm{d}b. \end{aligned} \tag{3.6}$$

The discount factor is represented by $0 < \beta < 1$. The function $\mathcal{C}_t(n, w, m)$ denotes the expected rewards obtained from the action of continuing the fermentation at time t and state (n, w, m). We note that state transitions are modeled through $\mathcal{C}_t(n, w, m)$. More specifically, the function $\mathcal{C}_t(n, w, m)$ integrates the evolution of both antibodies (in Equation 3.1) and waste metabolites (in Equation 3.2) into single period transition probabilities (in Equation 3.6). For example, the first summand in the right hand side of Equation 3.6 corresponds to the event of batch failure (progressive or sudden batch failure) during the period $[t, t + 1)$. Note that a failed

batch is immediately replaced with a new one leading to $\mathcal{V}_0(0,0,0)$. The second term in Equation 3.6 corresponds to the following series of events: the batch is not harvested at time t, the fermentation system encounters a shock arrival during the period $[t, t+1)$, and the batch survives (sudden and progressive failures) during that period. Similarly, third term in Equation 3.6 corresponds to the event that no shocks arrives during $[t, t+1)$ and the batch survives progressive failure.

Boundary conditions for the optimality equation are as follows. The maximum age T of fermentation implies $a_T(n,w,m) = H$ with $\mathcal{V}_T(n,w,m) = r_h(w,m) + \mathcal{V}_0(0,0,0)$. Similarly, the maximum allowable metabolic waste concentration $\bar{W}$ and the maximum achievable mAb concentration $\bar{M}$ imply $a_t(n,\bar{W},m) = H$ and $a_t(n,w,\bar{M}) = H$, respectively. We consider infinite horizon bioreactor harvesting decisions, where failed or harvested batches are immediately replaced with the new ones, leading to $t = n = 0$ and $w_0 = m_0 = 0$.

3.4 Structural Properties and Optimal Policy

In this section, we first present modeling assumptions on operating costs and rewards, and then analyze structural properties of the optimal bioreactor harvesting policies. The proofs of all results in this section are presented in the Appendix. An important note is that we use a discretization scheme enabling us to use induction on the iterates of value iteration algorithm as a proof technique.

3.4.1 Modeling Assumptions

Assumption 3.4.1. *$r_c(w,m)$ is nondecreasing in (w,m) for all $w \in \boldsymbol{S}, m \in \boldsymbol{M}$, and independent of $(t,n) \in \boldsymbol{T} \times \boldsymbol{N}$.*

Assumption 3.4.2. *$r_h(w, m)$ is nonincreasing in $w \in \boldsymbol{S}$, nondecreasing in $m \in \boldsymbol{M}$, and independent of $(t, n) \in \boldsymbol{T} \times \boldsymbol{N}$. Further, $r_h(w, 0) = 0$ for all $w \in \boldsymbol{S}$, so that batches with no mAbs does not yield any rewards.*

Assumption 3.4.3. *Batch failures are high impact events, implying $r(\Delta) > r_h(0, \bar{M})$, and $r(\Delta) > r_h(w, m) \geq r_c(w, m)$.*

Process observations have shown that batches with higher impurities and antibody densities challenge process monitoring, inspection and corrective/preventive actions (Rathore et al., 2004; Werner, 2004). Therefore, Assumption 3.4.1 implies that single period bioreactor operating costs do not decrease as the concentrations of both waste metabolites and mAbs increase. Assumption 3.4.2 implies that higher yields result in higher revenues, so that $r_h(w, m)$ is nondecreasing in $m \in \boldsymbol{M}$. The rewards obtained from harvesting $r_h(w, m)$ also considers downstream purification costs associated with waste metabolites w (i.e., cost of resins, labor, buffers and filtration). Purification costs are related with downstream operations that separate mAbs from metabolic wastes to meet desired quality and safety standards. Batches with higher metabolic wastes challenge downstream purification operations by increasing the purification workload (Kalyanpur, 2002; Werner, 2004; Farid, 2007). Therefore, the reward function $r_h(w, m)$ is modeled as nonincreasing in $w \in \boldsymbol{W}$, nondecreasing in $m \in \boldsymbol{M}$, and independent of (t, n). We note that $r_h(w, m) \geq 0$ and $r_h(w, 0) = 0$ for all $w \in \boldsymbol{S}$ and $m \in \boldsymbol{M}$. Assumption 3.4.3 holds as we model biomanufacturing systems with low frequency, high impact failures (Langer, 2008; CBI, 2010), implying that $r(\Delta) > r_h(0, \bar{M})$ and $r(\Delta) > r_h(w, m) \geq r_c(w, m)$ for all $w \in \boldsymbol{S}$ and $m \in \boldsymbol{M}$. All costs and rewards are assumed to be finite and bounded.

3.4.2 Structural Analysis

In this section, we analyze structural characteristics of the value function, and provide optimal condition-based harvesting policies in Theorems (3.4.1)-(3.4.3) which are easy to implement in practice.

Property 3.4.1. *$V_t(n, w, m) \geq 0$ for all $t \in \boldsymbol{T}$ and $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$.*

Proof See Appendix.

The value function in Equation (5.7) maximizes the expected benefits obtained from mAbs while minimizing operating costs associated with upstream, downstream, and batch failures. Property 3.4.1 shows that expected value function never becomes negative despite the penalty and operating costs. Property 3.4.1 sets the business case for operating the bioreactor. It would either imply not running the bioreactor at $t = 0$, or harvesting the bioreactor at later stages of fermentation when the accrued operational costs, penalty cost and risks associated with batch failures outweigh expected revenue obtained from mAbs.

Proposition 3.4.1. *The value function $\mathcal{V}_t(n, w, m)$ is nonincreasing in $n \in \boldsymbol{N}$, for all $t \in \boldsymbol{T}$ and $(w, m) \in \boldsymbol{W} \times \boldsymbol{M}$, if $\zeta(n, \rho)$ is nondecreasing in n and*

$$\int_0^{\bar{W}-w} f_t^b(b) r_h(w+b, m)\, \mathrm{d}b \geq \alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z) r_h(w, \bar{M})\, \mathrm{d}z \tag{3.7}$$

Proof See Appendix.

Proposition 3.4.1 shows the monotonicity of the value function in n, and provides sufficiency condition under which $\mathcal{V}_t(n, w, m)$ is nonincreasing in n. Proposition 3.4.1 implies that expected profit obtained from the batch never increases as the time elapsed since the last shock increases. The condition imposes a constraint on the rates at which rewards obtained from harvesting change as a function of w and m. In particular, consider a boundary condition where $w + b$ leads to the upperbound $\bar{W}$.

In this special case, the condition in (3.7) implies:

$$\frac{r_h(\bar{W}, m)}{r_h(w, \bar{M})} \geq \frac{\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\,\mathrm{d}z}{\int_0^{\bar{W}-w} f_t^{b}(b)\,\mathrm{d}b}. \tag{3.8}$$

We note that both right and left-hand side of Equation (3.8) represent a ratio. The right-hand side of Equation (3.8) captures the survival probability under a shock arrival $\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\,\mathrm{d}z$, and no shock arrivals $\int_0^{\bar{W}-w} f_t^{b}(b)\,\mathrm{d}b$. Consider next the expression on left-hand side. Since $r_h(w, m)$ is nonincreasing in w and nondecreasing in m (Assumption 3.4.2), the left-hand side denotes a ratio and positions the optimistic and pessimistic rewards at state (w, m). Note that $r_h(w, \bar{M})$ corresponds to the rewards of harvesting when mAbs concentration reaches the maximum level $\bar{M}$ and the metabolic waste concentration remains constant at w. $r_h(\bar{W}, m)$ corresponds to the rewards from harvesting when cells do not secrete any additional mAbs, but the waste concentration reaches its specification limit $\bar{W}$. Therefore, Equations (3.7) and (3.8) establish the relation between the risk of failure and the rates at which rewards change as a function of w and m. An interesting observation is the fact that $\zeta(n, \rho)$ does not appear in these equation despite dictating the shock arrival process.

Theorem 3.4.1. *There exists an optimal threshold level $n^* \in \boldsymbol{N}$ for all $t \in \boldsymbol{T}$, $w \in \boldsymbol{S}$ and $m \in \boldsymbol{M}$, such that the optimal decision is to harvest the bioreactor if and only if $n \geq n^*$, $\zeta(n, \rho)$ is nondecreasing in n, and the condition in (3.7) holds.*

Proof See Appendix.

Theorem 3.4.1 shows that when $\zeta(n, \rho)$ is nondecreasing in n (i.e., $\theta > 1$ in discrete Weibull), then the optimal policy is of control-limit type based on the time elapsed since the last shock. This implies that it is always optimal to harvest the fermentation when time elapsed since the last shock reaches or exceeds the threshold value n^*, and to continue to operate the fermentation according to the manufacturing protocol till the next decision epoch otherwise. Sufficiency conditions under which there exists a

control limit type policy are useful in practice because these control limits policies provide guidelines that are easy to implement in practice.

Proposition 3.4.2. *The value function $\mathcal{V}_t(n, w, m)$ is nonincreasing in $w \in \boldsymbol{S}$ for all $t \in \boldsymbol{T}$, $n \in \boldsymbol{N}$ and $m \in \boldsymbol{M}$.*

Proof See Appendix.

Proposition 3.4.2 shows that the expected value function is nonincreasing in w, such that the total expected profit obtained from the batch never increases as waste metabolites accumulate inside the batch. Monotonicity of the expected value function in w is utilized to identify the structure of the optimal policy based on the concentration of waste metabolites.

Theorem 3.4.2. *Let $w^+ > w^- \geq 0$ such that $w^+, w^- \in \boldsymbol{W}$. For all $n \in \boldsymbol{N}$, $t \in \boldsymbol{T}$, and $m \in \boldsymbol{M}$, there exists an optimal threshold level $w^* \in \boldsymbol{W} \setminus \{\Delta\}$ such that the optimal decision is to harvest the bioreactor if and only if $w \geq w^*$ and the condition in (3.9) holds:*

$$
\begin{aligned}
& r_h(w^-, m) - r_h(w^+, m) \\
& \leq \beta \left[R_t(n, w^-) - R_t(n, w^+)\right] r(\Delta) \\
& + \zeta(n, \rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z) \left[r_h(w^- + z, m) - r_h(w^-, \bar{M})\right] \mathrm{d}z \\
& + \left[1 - \zeta(n, \rho)\right] \beta \int_0^{\bar{W}-w^-} f_t^{b}(b) \left[r_h(w^- + b, m) - r_h(w^-, \bar{M})\right] \mathrm{d}b
\end{aligned}
\tag{3.9}
$$

where w^* is the waste metabolite-based control limit.

Proof See Appendix.

Theorem 3.4.2 represents sufficiency conditions for the existence of a waste metabolite-based optimal control limit policy, such that it is always optimal to harvest the batch when the amount of metabolic wastes reaches or exceeds the threshold level w^*,

and continue to operate the bioreactor according to the manufacturing protocol till the next decision epoch otherwise. The condition identified in Theorem 3.4.2 captures the risks and costs associated with the amount of metabolic wastes. More specifically, the term $r_h(w^-, m) - r_h(w^+, m)$ on the left-hand side of inequality (3.9) represents the additional benefits obtained from harvesting the batch at lower levels of metabolic wastes w^-, instead of harvesting at w^+, when the mAb concentration is m in both cases. The term $[R_t(n, w^-) - R_t(n, w^+)]\, r(\Delta)$ denotes the additional risk and cost of batch failure when the concentration of waste metabolites increases from w^- to w^+. The second and third terms in right hand side of Equation (3.9) consider two extreme cases for rewards when the fermentation operates at state (n, w^-, m) and the batch does not fail. The second term corresponds to a case where the amount of mAbs reaches the maximum level $\bar{M}$ while the amount of metabolic waste remains at w^-, as denoted by $\zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z) \left[r_h(w^-, \bar{M})\right] \mathrm{d}z + [1-\zeta(n,\rho)]\,\beta \int_0^{\bar{W}-w^-} f_t^b(b) \left[r_h(w^-, \bar{M})\right] \mathrm{d}b$. The third term corresponds to a case where cells do not secrete any additional mAbs, but the amount of metabolic waste increases, as denoted by $\zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z) \left[r_h(w^- + z, m)\right] \mathrm{d}z + [1-\zeta(n,\rho)]\,\beta \int_0^{\bar{W}-w^-} f_t^b(b) \left[r_h(w^- + b, m)\right] \mathrm{d}b$. Therefore, condition (3.9) specified in Theorem 3.4.2 imposes constraints on the rates at which risks and rewards change as a function of the amount of waste metabolites and mAbs.

Proposition 3.4.3. *The value function $\mathcal{V}_t(n, w, m)$ is nondecreasing in $m \in \boldsymbol{M}$ for all $t \in \boldsymbol{T}$, $n \in \boldsymbol{N}$ and $w \in \boldsymbol{W}$.*

Proof See Appendix.

Proposition 3.4.3 shows that expected profit obtained from the batch never decreases with mAb concentration. The result follows since higher mAb yields lead to higher revenues (Assumption 3.4.2). Proposition 3.4.3 is used to investigate structure of the optimal policy with respect to the amount of mAbs.

Theorem 3.4.3. *For all $n \in \boldsymbol{N}$, $t \in \boldsymbol{T}$, and $w \in \boldsymbol{S}$, there exists an optimal threshold level $m^* \in \boldsymbol{M}$ such that the optimal decision is to harvest the bioreactor if and only if $m \geq m^*$ and the condition in (3.10) holds:*

$$
\begin{aligned}
& r_h(w, m^+) - r_h(w, m^-) \\
& \geq \beta\, \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z) \left[r_h(w+z, \bar{M}) - r_h(w+z, m^-) \right] \mathrm{d}x\, \mathrm{d}z \\
& + \beta\, [1 - \zeta(n,\rho)] \int_0^{\bar{W}-w} f_t^{b}(b) \left[r_h(w+b, \bar{M}) - r_h(w+b, m^-) \right] \mathrm{d}x\, \mathrm{d}b
\end{aligned}
\tag{3.10}
$$

for all $n \in \boldsymbol{N}$, $t \in \boldsymbol{T}$, $w \in \boldsymbol{S}$, $\{m^+, m^-, m^*\} \in \boldsymbol{M}$, where m^* is the mAb-based control limit.

Proof See Appendix.

Theorem 3.4.3 identifies sufficiency conditions under which there exists a mAb-based optimal control-limit type policy, such that it is always optimal to harvest the bioreactor when the amount of mAbs reaches or exceeds the threshold level m^*, and to continue to operate the bioreactor according to the manufacturing protocol until the next decision epoch, otherwise. The condition in Equation (3.10) captures the tradeoff between benefits obtained from harvesting at higher levels of mAbs and additional purification costs associated with higher levels of waste metabolites. More specifically, the term $r_h(w, m^+) - r_h(w, m^-)$ denotes the incremental benefits obtained from higher mAb concentrations while the waste concentration remains constant at w. The right-hand side in Equation (3.10) represents the highest achievable reward at state (w, m^-), such that the fermentation reaches the maximum amount of mAbs $\bar{M}$ without any increments in the amount of waste metabolites w. Overall, the condition captures the tradeoff between the incentive of achieving the highest mAb yield, and the risks of incurring higher purification and failure costs.

Structural analysis obtained in this section provide insights for the optimal condition-based harvesting policies based on the amount of mAbs, waste metabolites and the time elapsed since the last shock. We note that optimal control limits are monotonic in n, ammonia levels w, and antibody levels m at a given time t; but not monotonic in t due to nonlinear, non-stationary behavior of the cells and fermentation kinetics. Next, we conduct a numerical analysis on IgG_1 antibody production as an illustrative example in Section 3.5. We derive optimal condition-based harvesting policies for IgG_1 production, and assess the performance of simple harvesting policies typically used in practice.

3.5 Effect of Bioreactor Reliability: Optimal Policy and Current Practice

In this section, we model IgG_1 antibody production to demonstrate optimal condition-based harvesting policies. We conduct numerical analysis to (i) compare the performance of the optimal harvesting policy with simple policies typically used in practice, and (ii) analyze sensitivity of the optimal control limits and value function to various degrees of bioreactor reliability levels at constant hazard rate and increasing hazard rates for shock arrivals. We consider production of IgG_1 for our study (Ozturk et al., 1992). In this production setting, we investigate low-frequency high-impact shocks with constant hazard rate for shock arrivals (Section 3.5.1), and then evaluate systems with increasing hazard rates for shock arrivals (Section 3.5.2).

The cost parameters for this study were determined using a two-step approach. Initial estimates of the cost parameters were obtained through a literature survey (Rathore et al., 2004; Werner, 2004; Farid, 2007; Kelley, 2009). Subsequently, we organized a series of working group sessions with several local biomanufacturing

companies (BioWGS, 2014) to validate our model formulation, parameter choices, and verify how the model/insights compare with industry practices. Based on these insights, we consider linear rewards with $r_h(w,m) = 10m - w$, $r_c(w,m) = 2 + w$ to account for the purification costs and upstream operating costs, respectively. These parameters are representative of industry practices for the manufacture of IgG_1 antibodies. However, we made some assumptions on the cost of batch failure based on discussions at the working group sessions. We learned that penalty cost of failure includes various components that vary from company to company. For example, the production line in the facility can completely stop until the clean room is securely sterilized after a viral batch contamination. This could lead to various opportunity costs (i.e., loss of reputation, re-scheduling efforts, additional general and administrative expenses, loss of production during sterilization periods, high resource costs, etc.). In our numeric analysis, we assume $r(\Delta) = 880$ for the penalty cost, which is approximately 8-10 times greater than the highest achievable revenue per batch and subsequently (in Section 3.5.1) analyze the sensitivity of the performance of the system for values in the range $r(\Delta) = \{440, 880, 1760\}$.

The specification limit for progressive failure is $\bar{W} = 6$ mM of ammonia concentration. In this example, the typical practice suggests to maintain the ammonia concentration below $w = 4mM$, or harvest the batch otherwise. The threshold value of $4mM$ ammonia level has been identified because it was observed that cell specific growth rate decreases by 50% at $4mM$ of ammonia levels (Ozturk et al., 1992). Although sensitivity of the specific growth rate to ammonia could be different for each cell line, setting 50% growth inhibition level as a threshold is a typical practice in fermentation literature. Cells are observed to enter the exponential growth phase at time $t = 50$ hr, and the deceleration phase at time $t = 150$ hr. During the deceleration phase, viable cell density drops significantly from $1.5e^6$ cells/ml to $5e^5$ cells/ml.

Hence, the maximum fermentation time is modeled as $T = 300$ hrs with period $\tau = 50$ hrs as shown in Figure 3.2. Each of the decision epoch is roughly associated with a physiological phase: beginning of the exponential growth phase $t = 50$ hr, middle of the exponential growth phase $t = 100$ hr, end of the exponential growth phase $t = 150$ hr, deceleration phase $t = 200$ hr, end of the deceleration phase $t = 250$ hr, and the death phase $t = 300$ hr. Time between two decision epochs τ is defined based on the current practice which typically involves measurements in the intervals of 2 to 3 days (BioWGS, 2014). Nevertheless, we conduct numerical analysis to investigate the sensitivity to other values of τ, i.e., $\tau = \{10, 25, 50\}$ hours, and observe that optimal policies and managerial insights remain the same at finer levels of discretization.

All computations are carried out using a discount factor of $\beta = \{0.98, 0.99, 0.999\}$. We use discount factors that are close to one because they best reflect the medium-term planning horizon which is most often applicable in the upstream biomanufacturing setting. Numerical analysis reported in this section is based on $\beta = 0.98$. By using such discount factor, we ensure that costs are relatively similar over the short term but far distant costs are less valuable. Although 50 hours seem to be a long time interval, we believe that it best represents a medium term planning horizon on the overall time frame required for a project, which is typically several weeks or months in the context of biomanufacturing operations.

We also replicate all numerical computations using an average cost rate objective in order to ensure that managerial insight derived from the discounted cost model are not attributed to the potential impacts of the discount factor, but rather due to the bioreactor dynamics and trade-offs. We observe that managerial insights obtained from both cost objectives are very similar in our problem context. A discussion on the average cost formulation is provided in the Appendix. To be consistent with

Sections 3.3 and 3.4, figures and tables reported in this section are mainly based on the discounted cost formulation. Where appropriate, we highlight the similarities and differences between the performance of these two cost objectives based on the numerical analysis. Note that we consider discrete Weibull distribution to model shock arrivals, $\zeta(n,\rho) = 1-(1-p)^{(n+1)^\theta - n^\theta}$, where $\rho = (p,\theta)$ are the known parameters of the discrete Weibull distribution.

3.5.1 Bioreactors with Constant Hazard Rate for Shock Arrivals

We first illustrate a bioreactor reliability setting with a geometrically distributed shock arrival process. An example of such bioreactor systems could be continuous perfusion bioreactors, where the culture media is refreshed on a regular basis. Table 3.1 presents a sensitivity analysis to compare the performance of optimal harvesting policy with simple policies typically used in practice. The first column in Table 3.1 denotes bioreactor reliability settings considered for the sensitivity analysis. Reliability settings in Table 3.1 are all identical except the probability of surviving sudden failures α_t, and probability of shock arrival p. We assume constant α_t for all $t \in \boldsymbol{T}$ and denote it by α. Note that $\theta = 1$ represents memoryless shock arrivals where $\zeta(n,\rho)$ is independent of n. Note that Table 3.1 reports the performance of both the discounted cost and average cost objectives.

The first scenario in Table 3.1 corresponds to the maximum reliability setting considered in this chapter with parameters $\alpha = 0.98$, $p = 0.001$. It represents the maximum control on bioreactor with nominal probability of shock arrivals and risk of sudden failures. We estimate that parameters $\alpha = 0.8$ and $p = 0.01$ correspond to an average reliability level typically encountered in this problem setting, based on data from Langer (2008). Hence, we refer to the scenario $\alpha = 0.8$ and $p = 0.01$ as

Table 3.1: Optimal value function $\mathcal{V}^{\Pi^*}$(normalized) and percentage improvements δ^{Π} for $\theta = 1$

	Discounted cost objective			Average cost objective		
Scenarios	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}
Maximum reliability	1.00	7%	6%	1.00	5%	4%
Baseline reliability	0.86	9%	7%	0.86	8%	6%
Less reliable systems with (p)						
(0.02)	0.72	14%	11%	0.70	10%	6%
(0.03)	0.59	23%	17%	0.56	14%	9%
(0.04)	0.44	37%	26%	0.39	17%	10%
(0.05)	0.31	89%	58%	0.27	48%	26%
(0.06)	0.19	>100%	>100%	0.17	>100%	>100%
(0.08)	0.00	-	-	0.00	-	-

baseline reliability setting through the rest of the chapter. Shock arrivals are related with unpredictable cell behaviors and external factors that can not be directly controlled. Therefore, we model different bioreactor reliability levels through parameters $\zeta(n, \rho)$, α and θ. Table 3.1 evaluates fermentation systems that are subject to more frequent shocks than the baseline reliability setting, with parameters $\alpha = 0.8$ and $p \in \{0.02, 0.03, 0.04, 0.05, 0.06, 0.08\}$. The scenario with parameters $\alpha = 0.8$ and $p = 0.06$ corresponds to the least reliable setting with a positive value function.

We evaluate two bioreactor harvesting policies, *yield maximizing policy* Π_1, and *moderately yield aggressive policy* Π_2, which are most typically used in practice. Policy Π_1 suggests to harvest either at the final fermentation time $t = 300$ hr, or at the antibody level 50 mg/l (maximum achievable), or at 4 mM of ammonia level. The motivation for the yield maximizing policy Π_1 is to maximize antibody yield at the time of harvest. Achieving the highest possible yield from the bioreactor has been one of the most popular objectives in practice and research (Luus, 1993a; Saucedo and Karim, 1997; Yang et al., 2000; Jenzsch et al., 2006; Farid, 2007; Kawohl et al., 2007). Whereas the moderately yield aggressive policy Π_2 is less aggressive in antibody yield, and suggests to harvest at either time $t = 300$ hr, or at the antibody level

of $40 mg/l$ (80% of the maximum achievable antibody level), or at 4 mM of ammonia. Note that policies Π_1 and Π_2 differ in terms of their antibody-based control limits. The main motivation driving both of these policies is to achieve high antibody yields while maintaining ammonia level below 4 mM.

Impact of Reliability on Costs

The second column in Table 3.1 presents the optimal value function $\mathcal{V}^{\Pi^*}$ at time $t = 0$ and state $(0, 0, 0)$ under the discounted cost objective, where Π^* denotes the optimal harvesting policy. Similarly, the fifth column in Table 3.1 presents the optimal value function $\mathcal{V}^{\Pi^*}$ based on the average cost objective. All $\mathcal{V}^{\Pi^*}$ in the table are normalized based on $\mathcal{V}^{\Pi^*}$ at the most reliable setting ($\alpha = 0.98$, $p = 0.001$). First, we focus on the managerial insights obtained from the discounted cost model. We observe that $\mathcal{V}^{\Pi^*}$ increasingly decreases as the bioreactor becomes less reliable. For example, consider the baseline reliability setting with $\alpha = 0.8$ and $p = 0.01$. Increasing p from 0.01 to 0.02 decreases the optimal value function from 0.86 to 0.72, resulting in 16% reduction. Whereas increasing p from 0.05 to 0.06 results in 38% reduction in $\mathcal{V}^{\Pi^*}$. Note that $\mathcal{V}^{\Pi^*}$ at $p = 0.08$ is zero, indicating that low reliability levels could negate the business case for manufacturing antibodies. In this specific example, Proposition 3.4.1 indicates that it is optimal to not run the fermentation if the bioreactor reliability is equal to or less than ($\alpha = 0.8$, $p = 0.08$). Insights from the table also show that cell-level shocks have significant impacts on profitability and should not be ignored in the operational decisions. For example, we observe 14% reduction in $\mathcal{V}^{\Pi^*}$ when the bioreactor reliability drops from the most reliable setting to the baseline reliability; whereas this reduction could reach up to 81 % when the reliability level drops to the least reliable setting with a positive value function ($\alpha = 0.08, p = 0.06$). Next, we evaluate the performance of the average cost objective. We observe that insights obtained from the average cost objective are

similar to the ones obtained using the discounted cost formulation. For example, we obtain 14% reduction in the optimal value function when the bioreactor reliability drops from the most reliable setting to the baseline reliability; whereas this reduction could reach up to 83% when the reliability level drops to the least reliable setting with a positive value function ($\alpha = 0.08, p = 0.06$). Therefore, we observe that managerial insights obtained from the discounted cost model are not attributable to the impact of the discount factor but rather related with the bioreactor dynamics and trade-offs.

Table 3.1 also presents the percentage improvements δ^{Π_1} and δ^{Π_2} in the discounted cost value function, that could be achieved at current practice, through implementing the optimal policy in replacement of the yield maximizing policy Π_1, and the moderately yield-aggressive policy Π_2, respectively. We observe that typical practices perform 10% below the optimal policy at high and baseline levels of reliability. This implies that a typical biomanufacturing facility that focuses on maximizing yield while maintaining ammonia concentration below $4mM$ (50% cell death threshold) would on average perform only 7-9 % worse off than the optimal policy. However, the limitations of current policies become more pronounced with increase in p. We observe that the optimal policy could be at least 25% better off than the simple harvesting policies at moderate/low levels of reliability. This implies that harvesting policies aiming to achieve high yield could significantly hurt business profitability under high risks environments, mainly due to quality constraints, purification costs and failure costs. Similarly, we observe that moderately yield-aggressive policy Π_2 outperforms the yield maximizing policy Π_1 in all considered settings; despite the fact that Π_2 is a less aggressive policy in yield but has the same ammonia-based control limits as Π_1. Especially, benefits obtained from adopting a less aggressive policy Π_2 instead of the yield maximizing policy Π_1 is more pronounced at low levels of reliability. How-

ever, we observe that the optimal policy always outperforms both of Π_1 and Π_2. The managerial insights obtained from the average cost model are similar. In Table 3.1, the percentage improvements reported by the average cost objective are slightly less than the ones obtained from the discounted cost objective. For example, we see that current practices perform $4-8\%$ below the optimal policy at high and baseline levels of reliability, whereas the percentage improvement at lower reliability levels are up to 48%. The main difference in the percentage improvements under the average cost and discounted cost models are attributable to the impact of discount factor that penalizes short-term failures. However, insights from both of these cost objectives indicate that substantial improvements in the profitability could be achieved through the optimal harvesting policies.

Impact of Reliability on Policy Structure

We show the structure of the optimal harvesting policy. Figure 3.5 presents the optimal ammonia-based and antibody-based control limits at the most reliable ($\alpha = 0.98, p = 0.001$), baseline reliability ($\alpha = 0.8, p = 0.01$) and least reliable ($\alpha = 0.8, p = 0.06$) settings using the discounted cost objective. In Figure 3.5, we demonstrate the optimal control limits at the beginning of exponential growth phase ($t = 50$ hr), later stage of the exponential growth phase ($t = 100$ hr), earlier stage of the deceleration phase ($t = 200$ hr) and the later stage of the deceleration phase ($t = 250$ hr). Note that the cell culture of interest has a non-growth related antibody production behavior so that cells secrete antibodies of interest mainly during the deceleration phase. The region below the control limit line represents the action of continuing to operate the fermentation, and the region above the control limit denotes harvesting.

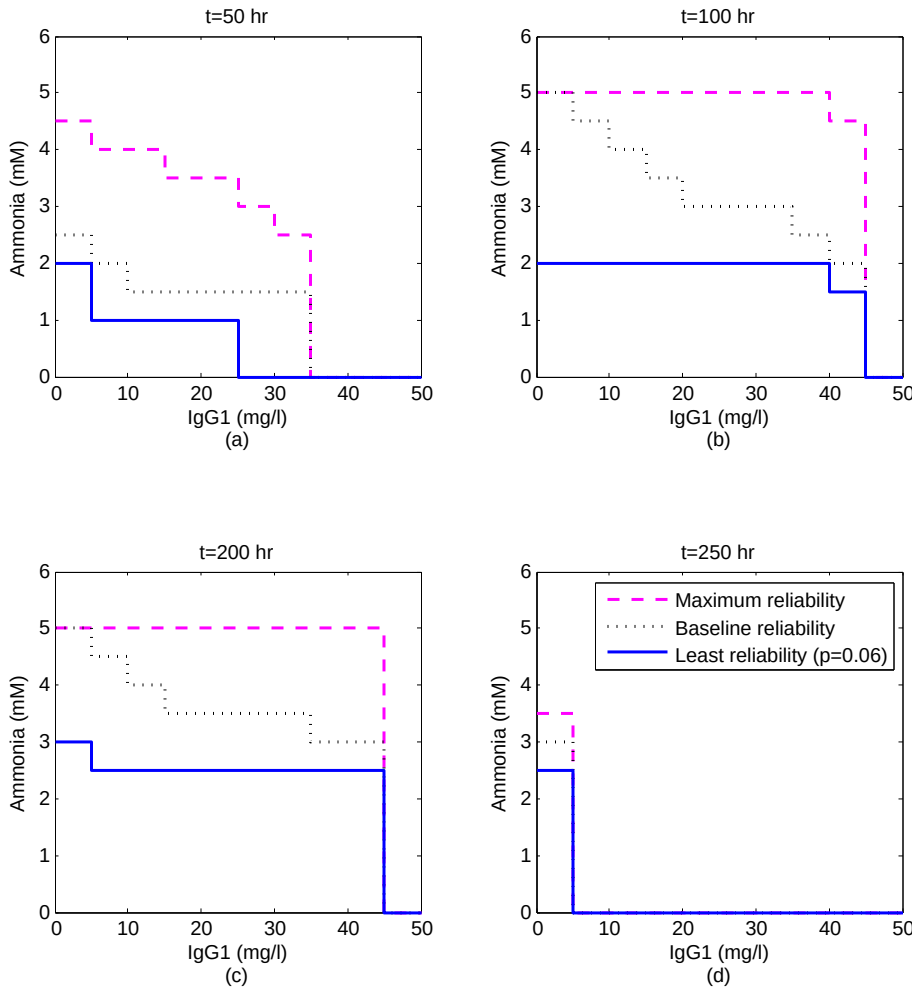

Figure 3.5: Optimal control limits throughout fermentation time t

We observe that, unlike the yield maximizing policy Π_1 and moderately yield-aggressive policy Π_2 with constant ammonia-based and antibody-based control limits over time, the optimal policy dynamically adjusts control limits based on the state and age of fermentation (see Figure 3.5). It is interesting to observe that the optimal ammonia-based control limit did not frequently adopt $4mM$ of ammonia threshold (corresponding to 50% reduction in growth rate) typically used in practice. Optimal ammonia-based control limit is observed to be less conservative than the typical practice at high levels of reliability (i.e., $5mM$ at the most reliable setting and during the deceleration phase $t = 200$, see Figure 3.5(c)), but more conservative at lower

reliability levels (i.e., less than $4mM$ at the baseline reliability setting, time $t = 200$ and antibody level greater than or equal to $15mg/l$; and less than $3mM$ at the least reliable setting and time $t = 200$, Figure 3.5(c)). Optimal control limits obtained from the average cost objective have the same trends and managerial insights, i.e., optimal ammonia-based control limits are observed to be less conservative than the typical practice at high reliability (with a limit of 5mM on the ammonia amount in the most reliable setting at $t = 200$ hours).

Smart Stationary Policies

The optimal control limits in Figure 3.5 are non-stationary because of non-linear and non-stationary nature of antibody and ammonia production. However, harvesting policies typically used in practice, Π_1 and Π_2, are stationary policies with control limits imposed on the fermentation time. Since stationary policies are easier to implement in practice, we exploit the structure of the optimal policies and propose smart stationary policies that can yield close to optimal performance. We define two smart stationary policies, Π_3 and Π_4. Smart stationary policy Π_3 suggests to harvest either at time $t = 250$ hr, or at $50\, mg/l$ antibody level (maximum achievable), or at $4\, mM$ of ammonia level. Smart stationary policy Π_4 harvests either at time $t = 250$ hr, or at $40\, mg/l$ antibody level (80% of the maximum achievable antibody level), or at $4\, mM$ of ammonia. Note that smart stationary policy Π_3 and the yield maximizing policy Π_1 (or similarly Π_2 and Π_4) are the same except that the stationary policy Π_3 (or similarly Π_4) has incentive to harvest earlier ($t = 250$) than the final fermentation time ($t = 300$). Finally, we define another stationary policy Π_5, which is a yield-focused policy ignoring ammonia-based control limits. The yield-focused policy Π_5 suggests to harvest either at time $t = 250$ hr or at $40\, mg/l$ antibody level (80% of the maximum achievable antibody level). Note that the smart stationary policy Π_4 and the yield-focused policy Π_5 are the same except that Π_5

Table 3.2: Performance of the smart stationary policies (for $\theta = 1$)

Scenarios	δ^{Π_3}	δ^{Π_4}	δ^{Π_5}
Maximum reliability	1%	1%	2%
Baseline reliability	1%	1%	17%
Less reliable systems with (p)			
(0.02)	1%	1%	39%
(0.03)	1%	1%	>100%
(0.04)	2%	2%	-
(0.05)	4%	4%	-
(0.06)	22%	22%	-
(0.08)	-	-	-

does not have any ammonia-based control limits. Table 3.2 shows the performance of these simple stationary policies.

The results reveal several managerial insights. First, we observe that the smart stationary policies Π_3 and Π_4 have the same performance, despite the fact that Π_2 outperforms Π_1 in Table 3.1. This would be mainly attributable to both expected bioreactor kinetics over time (Figure 3.2) and the upper-bound on final fermentation time ($t = 250$) suggested in smart stationary policies Π_3 and Π_4. It is interesting to observe that stationary policies Π_3 and Π_4 provide a very good approximation of the optimal value function (within 1% to 4% of the optimal value), except the least reliable setting. In this example, smart stationary policies Π_3 and Π_4 demonstrate how we can use MDP to develop efficient stationary harvesting policies that might be easier to implement. The performance of the smart stationary policies using the average cost objective is identical to that obtained from the discounted cost model, except that the smart stationary polices provide a better approximation of the optimal policy under the least reliable setting ($\alpha = 0.80, p = 0.06$) with $\delta^{\Pi_3} = \delta^{\Pi_4} = 14\%$.

Next, we observe that the yield-focused policy Π_5 under-performs the smart stationary policy Π_4, which highlights the significance of ammonia-based control limits.

This implies that cost of ignoring by-products could be significantly high, especially at low or baseline level reliability systems. Interestingly, Π_5 performs only 2% worse off than the optimal policy at the most reliable setting, but the performance of Π_5 drops rapidly as fermentation becomes less reliable. The dash '-'symbol in table represents a negative value function, implying that the corresponding policy is not a proper strategy to run the business. For example, ignoring by-products at low level reliability systems would result in significant losses in this production setting. The performance of the yield-focused policy is almost identical to that obtained from the discounted cost model.

Harvesting earlier than the final fermentation time could be counter-intuitive when we only consider the expected fermentation kinetics. For example, ammonia levels in Figure 3.2 is expected to remain almost the same during the time interval $t \in (250, 300)$; whereas a $5mg/l$ increase in antibody mass is expected by time $t = 300$. Optimal policy suggests to harvest at $t = 250$ despite the feasibility of achieving an additional yield of $5mg/l$ without compromising from the quality. This would mainly be attributable to costs and risks associated with high levels of ammonia.

Sensitivity Analysis in Penalty Cost

The penalty cost of batch failure includes components that could vary from company to company (due to differences in re-scheduling efforts, administrative expenses, high resource costs required for sterilization, loss of production, and other opportunity costs). In this section, we assess the impact of penalty cost on the optimal value function and evaluate the performance of the popular harvesting policies under different penalty costs. For all numerical analysis in this chapter, we consider a value of $r(\Delta) = 880$, which is $8 - 10$ times higher than the revenue of a batch. However in this section, we consider values in the range $r(\Delta) = \{440, 880, 1760\}$, to asses two

Table 3.3: Impact of penalty cost $r(\Delta)$ on $\mathcal{V}^{\Pi^*}$ and percentage improvements δ^{Π} (for $\theta = 1$)

	$r(\Delta) = 440$			$r(\Delta) = 880$			$r(\Delta) = 1760$		
Scenarios	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}
Maximum reliability	1.005	6%	4%	1.000	7%	6%	0.995	7%	6%
Baseline reliability	0.93	7%	5%	0.86	9%	7%	0.74	15%	11%
Less reliable with (p)									
(0.02)	0.84	9%	7%	0.72	14%	11%	0.48	40%	28%
(0.03)	0.76	10%	7%	0.59	23%	17%	0.25	>100%	>100%
(0.04)	0.68	13%	9%	0.44	37%	26%	0.02	>100%	>100%
(0.05)	0.60	16%	10%	0.31	89%	58%	0.00	-	-
(0.06)	0.51	19%	13%	0.19	>100%	>100%	-	-	-
(0.08)	0.35	33%	21%	0.00	-	-	-	-	-

extreme cases of the value used in the rest of the chapter.

Table 3.3 compares the optimal value function and the performance of popular harvesting policies under different penalty costs. We observe that the setting with maximum reliability is not significantly sensitive to the penalty cost. However, the same is not true for the setting with baseline reliability. In this setting, it is interesting to observe that the popular policy, which aims to maximize the yield, could be $4 - 7\%$ suboptimal even under the lowest penalty cost considered. We also observe that, moderately yield aggressive policy Π_2 outperforms the yield maximizing policy Π_1 in all scenarios considered. Also, for a given reliability setting, the benefit obtained from the optimal policy becomes significant as the penalty cost increases. Analysis using the average cost objective reveals similar managerial insights as the discounted cost objective.

3.5.2 Bioreactors with Increasing Hazard Rate for Shock Arrivals

In this section, we assess bioreactor reliability models with increasing failure rate (IFR) for shock arrivals, and investigate the impact of IFR shock arrivals on optimal value function and control limits. IFR shock arrivals are often found in mammalian

cell cultures or fed-batch fermentation systems due to risks in fouling, contamination, mutation or stability issues. We model IFR shock arrivals using the discrete Weibull model $\zeta(n, \rho)$ described in Section 3.3.3.

Table 3.4 represents a sensitivity analysis on impacts of IFR shock arrivals at various levels of bioreactor reliability. For this purpose, we introduce IFR shock arrivals into 4 scenarios: systems with the maximum ($\alpha = 0.98, p = 0.001$), baseline ($\alpha = 0.8, p = 0.01$), low ($\alpha = 0.8, p = 0.03$), and the least ($\alpha = 0.8, p = 0.06$) reliability settings. The degree of IFR behavior is modeled using parameter $\theta \in \{1.2, 1.4, 1.6, 1.8\}$ as shown in the second column in Table 3.4. Note that higher values of θ in $\zeta(n, \rho) = 1 - (1 - p)^{(n+1)^\theta - n^\theta}$ imply higher probability of shock arrival for a given n. Third column in Table 3.4 represents the maximum likelihood of a shock arrival at the corresponding reliability setting. Note that $n \leq t$, and the 5^{th} decision epoch is the last decision epoch ($t = 250$ hr). Hence, probability $\zeta(5, p)$ denoted in the third column is the probability of a shock during the last period, given that no shocks have been observed in the system since $t = 0$. We also note that the minimum probability of a shock is $\zeta(0, p) = p$.

Impact of Reliability on Cost

In Table 3.4, the optimal value functions $\mathcal{V}^{\Pi^*}$ are normalized based on the most reliable setting in geometrically distributed shock arrival process ($\alpha = 0.98$, $p = 0.001$, $\theta = 1$). Harvesting strategies Π_i for $i \in \{1, \ldots, 5\}$ remain the same as Section 3.5.1. Columns $5 - 9$ represent percentage improvements δ^{Π_i} obtained in the value function when the optimal policy is adopted instead of simple sub-optimal policies Π_i typically used in practice. We observe that managerial insights obtained from Section 3.5.1 are all valid in IFR shock arrival setting as well. We see that moderately yield-aggressive policy Π_2 outperforms the yield maximizing policy Π_1 despite being

Table 3.4: Optimal value function and performance of alternative sub-optimal policies (for $\theta > 1$)

Scenario	θ	$\zeta(5,p)$	$\mathcal{V}^{\Pi^*}$	δ^{Π_1}	δ^{Π_2}	δ^{Π_3}	δ^{Π_4}	δ^{Π_5}
Maximum reliability ($\alpha = 0.98$, $p = 0.001$)	1.0	0.001	1.000	7%	6%	1%	1%	2%
	1.2	0.002	0.998	6%	5%	1%	1%	3%
	1.4	0.003	0.995	7%	5%	1%	1%	4%
	1.6	0.005	0.990	7%	5%	1%	1%	6%
	1.8	0.008	0.983	8%	6%	1%	1%	8%
Baseline reliability ($\alpha = 0.8$, $p = 0.01$)	1.2	0.017	0.791	13%	9%	1%	1%	35%
	1.4	0.027	0.692	21%	16%	1%	1%	75%
	1.6	0.044	0.550	44%	30%	1%	1%	$> 100\%$
	1.8	0.068	0.358	$> 100\%$	$> 100\%$	3%	3%	-
Low reliability ($\alpha = 0.8$, $p = 0.03$)	1.2	0.050	0.381	71%	47%	2%	2%	-
	1.4	0.081	0.137	-	-	57%	56%	-
	1.6	0.127	0.000	-	-	-	-	-
Least reliability ($\alpha = 0.8$, $p = 0.06$)	1.2	0.099	0.000	-	-	-	-	-

less aggressive in antibody yield, and both Π_1 and Π_2 under-perform the optimal policy especially at baseline and low levels of reliability. Difference between the smart stationary policy Π_4 and the yield-focused policy Π_5 emphasizes the importance of ammonia-based control limits. This shows that failure to incorporate batch quality issues at upstream harvesting decisions could significantly increase purification costs at downstream. Smart stationary policies Π_3 and Π_4 are observed to successfully approximate the optimal value function at low and baseline levels of reliability. However, the smart stationary policies have limitations at low levels of bioreactor reliability. This underscores the importance of the optimal condition-based harvesting policies as the bioreactor becomes less reliable. Analysis using the average cost objective reveals similar managerial insights. For example, we observe 4% to 7% improvement in the value function when the optimal policy is adopted instead of Π_1 and Π_2 at the maximum reliability setting; whereas 7% to 26% improvement is observed in the baseline reliability setting for $\theta = \{1.2, 1.4, 1.6\}$, and the improvement exceeds 60% for $\theta = 1.8$. Since the impact of the discount factor is eliminated in the average cost objective, we observe that the policies Π_1 and Π_2 perform slightly better under the

discounted cost objective. However, optimal policies in the average cost model still provide substantial improvements even under the maximum reliability setting, and the potential benefits of the optimal policies are critical as the fermentation system becomes less reliable. Smart stationary policies remain within 1% of the optimal value function and hence provide successful approximation under the average cost setting.

We observe that $\mathcal{V}^{\Pi^*}$ increasingly decreases as the system becomes less reliable due to IFR shock arrivals. The optimal value function drops faster at low reliability systems as θ increases. This implies that $\mathcal{V}^{\Pi^*}$ is more susceptible to IFR behavior at low reliability levels compared to high reliability settings. Furthermore, we observe that θ might have significant impacts on the optimal value function at all levels of reliability. For example, consider the baseline reliability setting with $\alpha = 0.8$ and $p = 0.01$. For $\theta = 1$, the optimal value function is $\mathcal{V}^{\Pi^*} = 0.86$ as indicated in Table 3.1. We observe that when θ increases from 1 to 1.2, the maximum probability of a shock arrival only increases from $\zeta(5, p) = 0.010$ to $\zeta(5, p) = 0.017$; however, the value function drops considerably from 0.86 to 0.791. The impact of IFR behavior becomes more severe as the failure rate increases. For example, as θ increases from 1.2 to 1.6, the probability of shock arrival increase from $\zeta(5, p) = 0.017$ to $\zeta(5, p) = 0.044$, and hence the value function drops significantly from 0.791 to 0.50. This implies that the value function could be highly susceptible to the failure rate, and it is critical to consider the failure behavior of fermentation systems while making harvesting decisions. Analysis under the average cost objective indicates similar behavior of the value function.

Impact of Reliability on Policy Structure

Next, we analyze the impact of IFR shocks on ammonia and antibody-based optimal control limits. For this purpose, we select a high reliability (with $\theta = 1.2$) and a

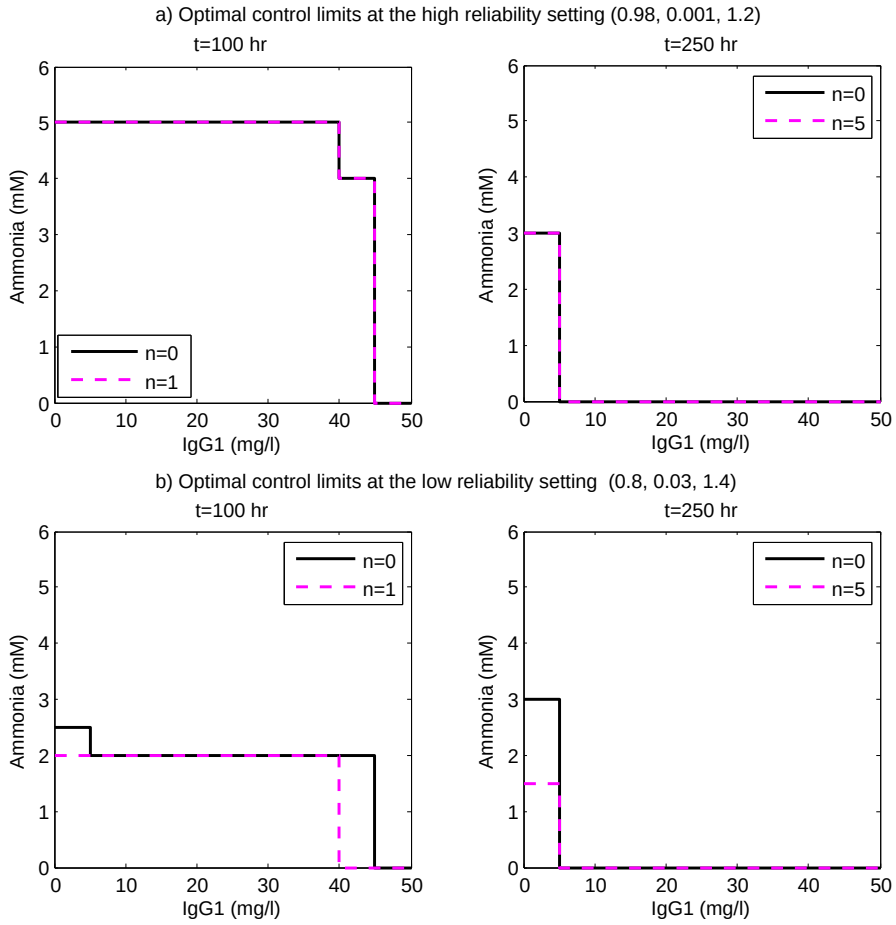

Figure 3.6: Optimal control limits as a function of n at high and low levels of reliability

low reliability ($\theta = 1.4$) setting from Table 3.4, and evaluate how optimal control limits change as a function of n. Figure 3.6 represents ammonia-based and antibody-based optimal control limits at time $t = 100$ (i.e., later stage of the exponential growth phase) and $t = 250$ (i.e., ideal harvesting time suggested by the optimal policy). Note that Figure 3.6 represents ammonia-based and antibody-based optimal control limits at various levels of n for a given time t, ammonia concentration w and antibody level m. It is interesting to observe that optimal control limits are not very sensitive to n. For example, control limits in Figure 3.6 a) are constant in n for each time $t = 100$ and $t = 250$ hrs. This observation could be counter-intuitive since higher levels of n imply higher probability of shock arrivals. In this example, the main reason for constant

control limits over n is mainly attributable to low probabilities of shock arrival at the most reliable setting. However, in low reliability setting, the control limits in Figure 3.6 b) become slightly more conservative as time elapsed since the last shock increases. We also note that optimal control limits are monotonic in n, ammonia levels w, and antibody levels m at a given time t; but not monotonic in t due to non-stationary process kinetics. Analysis of the optimal policies using the average cost objective reveals similar managerial insights. We observe that optimal policies shown in Figure 3.6 are identical under these two objectives, except the policy suggested for the high reliability setting at time $t = 100$ hours. In that setting, the optimal policy suggested by the discounted cost objective has a stationary control limit of 5 mM ammonia when IgG_1 amount is less than 40 mg/l; whereas the average cost objective suggests a non-stationary control limit on ammonia which varies between 4.5 mM and 2.5 mM. However, managerial insights obtained from both of these cost objectives indicate that control limits shown in Figure 3.6 are not very sensitive to n mainly due to the low probability of shock arrivals.

3.6 Conclusions

Biomanufacturing methods use live systems such as bacterial or mammalian cells to manufacture the biologics of interest. Use of live systems introduce several challenges in modeling and optimization of biomanufacturing operations due to batch-to-batch variability, random evolution of fermentation dynamics, parallel growth of unwanted metabolic wastes and desired antibodies inside the same batch, and multiple dependent failure processes caused by random disturbances in cellular activities and fermentation dynamics.

We incorporate biology-induced randomness associated with fermentation dynamics into a cell-level stochastic model to capture the evolution of unwanted metabolic wastes and desired antibodies, and model multiple dependent failure processes in upstream operations. The cell-level reliability model is then integrated in a Markov decision process formulation to analyze system-level decisions on batch harvesting policies. The model studies the impact of cell-level fermentation dynamics and batch failures on system-level financial trade-offs between upstream and downstream operations.

We analyze the structural characteristics of the value function, and derive optimum condition-based harvesting policies for upstream biomanufacturing operations. The optimal policy provides simple guidelines that are easy to implement in practice. Through numerical studies, we evaluate the impact of bioreactor reliability on the optimal control limits and profitability. We compare the performance of the optimal policy for IgG_1 antibody production with popular harvesting policies used in practice. Our study shows that a typical biomanufacturing facility that focuses on maximizing IgG_1 antibody yield could on average improve their profitability by 4-9 %, using the optimal condition-based harvesting policies. The benefits obtained from the optimal policies becomes very significant (more than 40% improvement) at fermentation systems with high risks of batch failures. The case study considered in the numerical analysis shows that trying to maximize the yield is not necessarily optimal from the profit perspective. Although counter-intuitive, the optimal policy has an incentive to harvest the fermentation earlier than the popular practice even when the concentration of metabolic byproducts are expected to remain constant and stable. This would mainly be attributed to financial trade-offs between batch quality, antibody yields and batch failures. MDP model captures system-level trade-offs that are not apparent in the process kinetic models representing cell dynamics alone.

We also use insights obtained from the optimal policy to propose smart stationary policies that could closely approximate the optimal value function.

3.7 Appendix: Proofs

Proof of Property 3.4.1. Consider the value function for all $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T} \setminus T$,

$$\mathcal{V}_t(n, w, m) = \max\{r_h(w, m) + \mathcal{V}_0(0, 0, 0),\ -r_c(w, m) + \beta\, \mathcal{C}_t(n, w, m)\} \tag{3.11}$$

where $\mathcal{C}_t(n, w, m)$ is defined in Equation 3.6. We use Equation 3.11 to analyze two cases: *(i)* If $-r_c(w, m) + \beta\, \mathcal{C}_t(n, w, m) \geq 0$, then $\mathcal{V}_t(n, w, m) \geq 0$ for all $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$. *(ii)* When $-r_c(w, m) + \beta \mathcal{C}_t(n, w, m) < 0$, then $\mathcal{V}_t(n, w, m) \geq 0$ if and only if $r_h(w, m) + \mathcal{V}_0(0, 0, 0) \geq 0$. Hence, the proof reduces to showing that $r_h(w, m) + \mathcal{V}_0(0, 0, 0) \geq 0$ all $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$. Equation 3.11 is revisited to define $\mathcal{V}_0(0, 0, 0) = \max\{r_h(0, 0) + \mathcal{V}_0(0, 0, 0),\ -r_c(0, 0) + \beta\, \mathcal{C}_0(0, 0, 0)\}$. Now consider the case where $-r_c(0, 0) + \beta\, \mathcal{C}_0(0, 0, 0) < 0$, then the state $(0, 0, 0)$ at $t = 0$ becomes an absorbing state with zero rewards, since $r_h(w, 0) = 0$ for all $w \in \boldsymbol{S}$ by Assumption 3.4.2. This implies that, if $-r_c(0, 0) + \beta\, \mathcal{C}_0(0, 0, 0) < 0$, then $V_0(0, 0, 0) = 0$ with $a_0(0, 0, 0) = H$. Therefore, inequality $V_0(0, 0, 0) \geq 0$ always holds. Hence $r_h(w, m) + \mathcal{V}_0(0, 0, 0) \geq 0$ for all $(w, m) \in \boldsymbol{S} \times \boldsymbol{M}$, since $r_h(w, m) \geq 0$ and $r_h(w, 0) = 0$ for all $w \in \boldsymbol{S}$ and $m \in \boldsymbol{M}$. Similarly, the boundary conditions for both the maximum age and mAb concentration imply that $a_T(n, w, m) = a_t(n, w, \bar{M}) = H$ with rewards $r_h(w, m) + V_0(0, 0, 0) \geq 0$ and $r_h(w, \bar{M}) + V_0(0, 0, 0) \geq 0$ respectively. Hence, $\mathcal{V}_t(n, w, m) \geq 0$ for all $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$ from which the result follows. □

Proof of Proposition 3.4.1. The proof is done by induction on the iterates of value iteration algorithm. For $n \in \boldsymbol{N}$ and $m \in \boldsymbol{M}$, we define the initial value for $i = 0$ without loss of generality as:

$$\mathcal{V}_t^0(n,w,m) = \begin{cases} 0, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta), & \text{if } w = \Delta. \end{cases}$$

Next, assume that $\mathcal{V}_t^i(n,w,m)$ is nonincreasing in n, then

$\mathcal{V}_t^{i+1}(n,w,m)$

$$= \begin{cases} \max\left\{r_h(w,m) + \mathcal{V}_0^i(0,0,0), -r_c(w,m) + \beta\, \mathcal{C}_t^i(n,w,m)\right\}, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0^i(0,0,0), & \text{if } w = \Delta. \end{cases} \tag{3.12}$$

We note that neither the cost of harvesting, nor the cost of continuing depends on n in Equation 3.12. To complete the proof, we need to show that $\mathcal{C}_t^i(n,w,m)$ is nonincreasing in n for all $t \in \boldsymbol{T}$, $w \in \boldsymbol{W}$ and $m \in \boldsymbol{M}$. Consider two different values for the time elapsed since the last shock, $n+1$ and n, such that $n+1 > n$. We start with writing the difference between the expected rewards of continuing the fermentation at n and $n+1$ for any time $t \in \boldsymbol{T}$ and state $(w,m) \in \boldsymbol{S} \times \boldsymbol{M}$:

$$\mathcal{C}_t^i(n,w,m) - \mathcal{C}_t^i(n+1,w,m)$$

$$\begin{aligned}
&= [R_t(n+1,w,m) - R_t(n,w,m)]\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right] \\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x) \mathcal{V}_{t+1}^i(0,w+z,m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x) \mathcal{V}_{t+1}^i(n+1,w+b,m+x)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.13) \\
&- \zeta(n+1,p)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x) \mathcal{V}_{t+1}^i(0,w+z,m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&- [1-\zeta(n+1,p)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x) \mathcal{V}_{t+1}^i(n+2,w+b,m+x)\,\mathrm{d}x\,\mathrm{d}b \\
&\geq \zeta(n+1,p)\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z \\
&+ [1-\zeta(n+1,p)] \int_0^{\bar{W}-w} f_t^b(b)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b \\
&- \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z \\
&- [1-\zeta(n,\rho)] \int_0^{\bar{W}-w} f_t^b(b)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b \\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{\epsilon+b}(z) f_t^x(x) \mathcal{V}_{t+1}^i(0,w+z,m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x) \mathcal{V}_{t+1}^i(n+2,w+b,m+x)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.14) \\
&- \zeta(n+1,p)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x) \mathcal{V}_{t+1}^i(0,w+z,m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&- [1-\zeta(n+1,p)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x) \mathcal{V}_{t+1}^i(n+2,w+b,m+x)\,\mathrm{d}x\,\mathrm{d}b \\
&\geq [\zeta(n+1,p) - \zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p) - \zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b)\left[-r(\Delta) + \mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p) - \zeta(n,\rho)] \\
&\quad \alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x) \mathcal{V}_{t+1}^i(0,w+z,m+x)\,\mathrm{d}x\,\mathrm{d}z \qquad (3.15) \\
&+ [\zeta(n+1,p) - \zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) \mathcal{V}_{t+1}^i(n+2,w+b,m+x)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.16)
\end{aligned}$$

$$
\begin{aligned}
&\geq [\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w,\bar{M})\,\mathrm{d}x\,\mathrm{d}z \qquad (3.17)\\
&+ [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\mathcal{V}_{t+1}^i(n+2,w+b,m)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.18)\\
&\geq [\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&- [\zeta(n+1,p)-\zeta(n,\rho)] \\
&\quad \alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[r_h(w,\bar{M})+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}z \qquad (3.19)\\
&+ [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[r_h(w+b,m)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}x\,\mathrm{d}b \qquad (3.20)\\
&= [\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[-r(\Delta)\right]\,\mathrm{d}x\,\mathrm{d}b \qquad (3.21)\\
&- [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[-r(\Delta)\right]\,\mathrm{d}x\,\mathrm{d}b \qquad (3.22)\\
&- [\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[r_h(w,\bar{M})\right]\,\mathrm{d}x\,\mathrm{d}z \\
&+ [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[r_h(w+b,m)\right]\,\mathrm{d}x\,\mathrm{d}b \\
&\geq -[\zeta(n+1,p)-\zeta(n,\rho)]\,\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\left[r_h(w,\bar{M})\right]\,\mathrm{d}x\,\mathrm{d}z \qquad (3.23)\\
&+ [\zeta(n+1,p)-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b)\left[r_h(w+b,m)\right]\,\mathrm{d}x\,\mathrm{d}b \qquad (3.24)\\
&\geq 0 \qquad (3.25)
\end{aligned}
$$

Property 3.3.2 indicates that $R_t(n,w)$ is a decreasing function in n for $\theta > 1$. Therefore, we get $[R_t(n+1,w,m)-R_t(n,w,m)] \leq 0$ and $[\zeta(n+1,p)-\zeta(n,\rho)] \geq 0$. Equation (3.14) follows from Equation (3.13), the induction hypothesis and Property

3.4.1. Equations (3.17) and (3.18) are obtained from Equations (3.15) and (3.16) respectively, and follow from the monotonicity of the value function in w and m (Proposition 3.4.2 and Proposition 3.4.3). The boundary condition $a_t(0, w, \bar{X}) = H$ applies to Equation (3.17) resulting in (3.19) with reward $\left[r_h(w, \bar{M}) + \mathcal{V}_0^i(0, 0, 0)\right]$. Equation (3.20) follows from (3.18) because $\mathcal{V}_{t+1}^i(n+2, w+b, m) \geq [r_h(w + b, m) + \mathcal{V}_0^i(0, 0, 0)]$ by definition of the value function in (5.7). We observe that sum of the Equations (3.21) and (3.22) is greater than or equal to zero since $\int_0^{\bar{W}-w} f_t^b(b)\,\mathrm{d}b \geq \alpha_t \int_0^{\bar{W}-w} f_t^{e+b}\,\mathrm{d}z$. Equations (3.23) and (3.24) reduce to inequality (3.25) under the conditions specified in (3.7). Similarly, the boundary conditions for both the maximum age and mAb concentration imply $a_T(n, w, m) = a_t(n, w, \bar{M}) = H$ with rewards $r_h(w, m)+V_0(0, 0, 0) \geq 0$ and $r_h(w, \bar{M}) + V_0(0, 0, 0) \geq 0$, which are independent of n. Note that since $\mathcal{C}_t^i(n, w, m)$ is nonincreasing in n by the inequality (3.25), then the value function $\mathcal{V}_t(n, w, m)$ is also nonincreasing in $n \in \boldsymbol{N}$ for all $t \in \boldsymbol{T}$ and $(w, m) \in \boldsymbol{W} \times \boldsymbol{M}$. Hence, the result follows from the convergence of value iteration algorithm. □

Proof of Theorem 3.4.1. Assume that the result holds and consider the inequality,

$$r_h(w, m) + \mathcal{V}_0(0, 0, 0) \quad \geq \quad r_c(w, m) + \beta \mathcal{C}_t(n^*, w, m) \tag{3.26}$$

The left-hand side of the inequality (3.26) is constant in n. On the other hand, insights obtained from the proof of Proposition 3.4.1 indicate that $\mathcal{C}_t(n, w, m)$ is nonincreasing in n. Hence, the right hand side of inequality (3.26) is nonincreasing in n. Therefore, inequality (3.26) holds for any $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$ such that $n \geq n^*$. In other words, given that the optimal decision is to harvest at state (n^*, w, m), the optimal decision for any $n \geq n^*$ is also to harvest at $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$. Therefore, the optimal replacement policy is of control limit type policy with control limit n^*. □

Proof of Proposition 3.4.2. The proof is done by induction on the steps of the value iteration algorithm. Let $\mathcal{V}_t^i(n, w, m)$ denote the value function at i^{th} iteration of the value iteration algorithm. We start by defining the initial value for $i = 0$. Without loss of generality, we assume:

$$\mathcal{V}_t^0(n, w, m) = \begin{cases} 0, & w \in \left[0, \bar{W}\right), \\ -r(\Delta), & w = \Delta. \end{cases}$$

Note that $\mathcal{V}_t^0(n, w, m)$ is nonincreasing in w. Next, we assume that $\mathcal{V}_t^i(n, w, m)$ is nonincreasing in w. Then, we obtain from Equation 5.7,

$$\mathcal{V}_t^{i+1}(n, w, m)$$

$$= \begin{cases} \max\left\{r_h(w, m) + \mathcal{V}_0^i(0, 0, 0),\ -r_c(w, m) + \beta\, \mathcal{C}_t^i(n, w, m)\right\}, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0^i(0, 0, 0), & \text{if } w = \Delta. \end{cases} \tag{3.27}$$

Since $r_h(w, m)$ and $-r_c(w, m)$ are nonincreasing in w, it is sufficient to show that $\mathcal{C}_t^i(n, w, m)$ is also nonincreasing in w for all t, n and m. Next, we use a discretization scheme on state space $\boldsymbol{W}$ enabling us to use induction on the iterates of the value iteration algorithm as a proof technique. Consider two arbitrary metabolic waste concentrations w^+ and w^-, such that $w^+ \geq w^- \in \boldsymbol{W} \setminus \{\Delta\}$. To complete the proof, we use the following steps show that $\mathcal{C}_t^i(n, w^-, m) - \mathcal{C}_t^i(n, w^+, m) \geq 0$ for any $t \in \boldsymbol{T}$:

$$\mathcal{C}_t^i(n,w^-,m) - \mathcal{C}_t^i(n,w^+,m)$$

$$
\begin{aligned}
&= R_t(n,w^+)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right] - R_t(n,w^-)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right] && (3.28)\\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w^-+z,m+x)\,\mathrm{d}x\,\mathrm{d}z && (3.29)\\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w^-+b,m+x)\,\mathrm{d}x\,\mathrm{d}b && (3.30)\\
&- \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w^++z,m+x)\,\mathrm{d}x\,\mathrm{d}z &&\\
&- [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w^++b,m+x)\,\mathrm{d}x\,\mathrm{d}b &&\\
&= \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+} f_t^{e+b}(z)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z && (3.31)\\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+} f_t^b(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b && (3.32)\\
&- \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^-} f_t^{e+b}(z)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z && (3.33)\\
&- [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^-} f_t^b(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b && (3.34)\\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w^-+z,m+x)\,\mathrm{d}x\,\mathrm{d}z && (3.35)\\
&+ \zeta(n,\rho)\alpha_t \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w^-+z,m+x)\,\mathrm{d}x\,\mathrm{d}z && (3.36)\\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w^-+b,m+x)\,\mathrm{d}x\,\mathrm{d}b && (3.37)\\
&+ [1-\zeta(n,\rho)] \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w^-+b,m+x)\,\mathrm{d}x\,\mathrm{d}b && (3.38)\\
&- \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w^++z,m+x)\,\mathrm{d}x\,\mathrm{d}z &&\\
&- [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w^++b,m+x)\,\mathrm{d}x\,\mathrm{d}b &&
\end{aligned}
$$

$$
\begin{aligned}
&\geq -\zeta(n,\rho)\alpha_t \int_{\bar{W}-w^+}^{\bar{W}-w^-} f_t^{e+b}(z)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z \\
&- [1-\zeta(n,\rho)] \int_{\bar{W}-w^+}^{\bar{W}-w^-} f_t^{b}(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b \\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^{x}(x)\mathcal{V}_{t+1}^i(0,w^+ + z, m+x)\,\mathrm{d}x\,\mathrm{d}z \qquad (3.39)\\
&+ \zeta(n,\rho)\alpha_t \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^{x}(x)\mathcal{V}_{t+1}^i(0,w^- + z, m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{b}(b) f_t^{x}(x)\mathcal{V}_{t+1}^i(n+1,w^+ + b, m+x)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.40)\\
&+ [1-\zeta(n,\rho)] \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{b}(b) f_t^{x}(x)\mathcal{V}_{t+1}^i(n+1,w^- + b, m+x)\,\mathrm{d}x\,\mathrm{d}b \\
&- \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^{x}(x)\mathcal{V}_{t+1}^i(0,w^+ + z, m+x)\,\mathrm{d}x\,\mathrm{d}z \\
&- [1-\zeta(n,\rho)] \int_0^{\bar{W}-w^+}\int_0^{\bar{M}} f_t^{b}(b) f_t^{x}(x)\mathcal{V}_{t+1}^i(n+1,w^+ + b, m+x)\,\mathrm{d}x\,\mathrm{d}b \\
&\geq -\zeta(n,\rho)\alpha_t \int_{\bar{W}-w^+}^{\bar{W}-w^-} f_t^{e+b}(z)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}z \\
&- [1-\zeta(n,\rho)] \int_{\bar{W}-w^+}^{\bar{W}-w^-} f_t^{b}(b)\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right]\,\mathrm{d}b \\
&+ \zeta(n,\rho)\alpha_t \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^{x}(x)\mathcal{V}_{t+1}^i(0,w^+ + z, m+x)\,\mathrm{d}x\,\mathrm{d}z \qquad (3.41)\\
&+ [1-\zeta(n,\rho)] \int_{\bar{W}-w^+}^{\bar{W}-w^-}\int_0^{\bar{M}} f_t^{b}(b) f_t^{x}(x)\mathcal{V}_{t+1}^i(n+1,w^+ + b, m+x)\,\mathrm{d}x\,\mathrm{d}b \qquad (3.42)\\
&\geq 0 \qquad (3.43)
\end{aligned}
$$

Equations (3.31)-(3.34) are obtained by rewriting the reliability function in Equation 3.28 using the definition in Equation 3.4. We split Equations (3.29)-(3.30) into smaller regions using the fact that $\int_0^{\bar{W}-w^-} V(z)\,\mathrm{d}z = \int_0^{\bar{W}-w^+} V(z)\,\mathrm{d}z + \int_{\bar{W}-w^+}^{\bar{W}-w^-} V(z)\,\mathrm{d}z$, and hence get Equations (3.35) to (3.38). Equations (3.39) and (3.40) are obtained using the induction hypothesis and Property 3.4.1. Similarly, Equations (3.41) and (3.42) are obtained by the induction hypothesis and Property

3.4.1. In Equations (3.41) and (3.42), we note that $\Delta = \left\{w : w \in \left[\bar{W}, \infty\right)\right\}$ with reward $[-r(\Delta) + \mathcal{V}_0^i(0,0,0)]$, from which the inequality (3.43) is established. Similarly, the boundary conditions for both the maximum age and mAb concentration imply $a_T(n,w,m) = a_t(n,w,\bar{M}) = H$ with rewards $r_h(w,m) + V_0(0,0,0) \geq 0$ and $r_h(w,\bar{M}) + V_0(0,0,0) \geq 0$ which are nonincreasing in w by Assumption 3.4.2. Hence, the proof follows from the convergence of value iteration algorithm. Notice that because all the terms on the right hand side of Equation 3.27 are nonincreasing in w, $\mathcal{V}_t^{i+1}(n,w,m)$ is also nonincreasing in w for all $t \in \boldsymbol{T}$, and $(n,w,m) \in \boldsymbol{N} \times \boldsymbol{W} \times \boldsymbol{M}$ from which the result follows.

□

Proof of Theorem 3.4.2. Proof is done by contradiction. Consider two arbitrary waste metabolite concentrations, $w^+, w^- \in \boldsymbol{W} \setminus \{\Delta\}$ such that $w^+ \geq w^-$. Theorem 3.4.2 indicates that if $a_t(n,w^-,m) = H$ then $a_t(n,w^+,m) = H$. Assume by contradiction that $a_t(n,w^-,m) = H$ then $a_t(n,w^+,m) = C$. The contradiction hypothesis implies that

$$r_h(w^-,m) + \mathcal{V}_0(0,0,0) \geq -r_c(w^-,m) + \beta\mathcal{C}_t(n,w^-,m) \tag{3.44}$$

$$r_h(w^+,m) + \mathcal{V}_0(0,0,0) \leq -r_c(w^+,m) + \beta\mathcal{C}_t(n,w^+,m) \tag{3.45}$$

Equations (3.44) and (3.45) lead to the following inequality:

$r_h(w^-, m) - r_h(w^+, m)$

$$
\begin{aligned}
&\geq \beta\mathcal{C}_t(n, w^-, m) - \beta\mathcal{C}_t(n, w^+, m) && (3.46)\\
&\geq \beta R_t(n, w^+)\left[-r(\Delta) + \mathcal{V}_0(0,0,0)\right] - \beta R_t(n, w^-)\left[-r(\Delta) + \mathcal{V}_0(0,0,0)\right]\\
&+ \zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z)\mathcal{V}_{t+1}(0, w^- + z, m)\,\mathrm{d}z && (3.47)\\
&+ \left[1-\zeta(n,\rho)\right]\beta \int_0^{\bar{W}-w^-} f_t^{b}(b)\mathcal{V}_{t+1}(n+1, w^- + b, m)\,\mathrm{d}b && (3.48)\\
&- \zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^+} f_t^{e+b}(z)\mathcal{V}_{t+1}(0, w^+ + z, \bar{M})\,\mathrm{d}z && (3.49)\\
&- \left[1-\zeta(n,\rho)\right]\beta \int_0^{\bar{W}-w^+} f_t^{b}(b)\mathcal{V}_{t+1}(n+1, w^+ + b, \bar{M})\,\mathrm{d}b && (3.50)\\
&\geq \beta\left[R_t(n, w^-) - R_t(n, w^+)\right] r(\Delta)\\
&+ \zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z) r_h(w^- + z, m)\,\mathrm{d}z && (3.51)\\
&+ \left[1-\zeta(n,\rho)\right]\beta \int_0^{\bar{W}-w^-} f_t^{b}(b) r_h(w^- + b, m)\,\mathrm{d}b && (3.52)\\
&- \zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^+} f_t^{e+b}(z) r_h(w^-, \bar{M})\,\mathrm{d}z && (3.53)\\
&- \left[1-\zeta(n,\rho)\right]\beta \int_0^{\bar{W}-w^+} f_t^{b}(b) r_h(w^-, \bar{M})\,\mathrm{d}b && (3.54)\\
&\geq \beta\left[R_t(n, w^-) - R_t(n, w^+)\right] r(\Delta)\\
&+ \zeta(n,\rho)\alpha_t\beta \int_0^{\bar{W}-w^-} f_t^{e+b}(z)\left[r_h(w^- + z, m) - r_h(w^-, \bar{M})\right]\,\mathrm{d}z\\
&+ \left[1-\zeta(n,\rho)\right]\beta \int_0^{\bar{W}-w^-} f_t^{b}(b)\left[r_h(w^- + b, m) - r_h(w^-, \bar{M})\right]\,\mathrm{d}b
\end{aligned}
$$

which contradicts the condition in (3.9). We note that assumption $r(\Delta) > r_h(0, \bar{M})$ on modeling high impact batch failure supports the condition in (3.9). Equation (3.46) follows from Assumption 3.4.1, such that $r_c(w^+, m) - r_c(w^-, m) \geq 0$. Equations (3.47) to (3.50) are obtained from nondecreasing behavior of the value function in m based on Proposition 3.4.3. We note that costs and rewards are independent of t and n by Assumptions (3.4.1) and (3.4.2). Hence, Equations (3.51) and (3.52) follow from (5.7), such that $\mathcal{V}_t(n, w, m) \geq r_h(n, w, m) + \mathcal{V}_0(0,0,0)$ for all

(n,w,m) $\in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$. Equations (3.53) and (3.54) are obtained from the boundary condition $a_t(n, w, \bar{M}) = H$ for all $(n, w) \in \boldsymbol{N} \times \boldsymbol{S}$ and $t \in \boldsymbol{T}$. □

Proof of Proposition 3.4.3. The proof is done by induction on the iterates of value iteration algorithm. We define the initial value for $i = 0$ as

$$\mathcal{V}_t^0(n, w, m) = 0 \text{ for all } m \in \boldsymbol{M}.$$

Next, consider two arbitrary mAb concentration, m^+, $m^- \in \boldsymbol{M}$, such that $m^+ \geq m^-$. Assume that $\mathcal{V}_t^i(n, w, m)$ is nondecreasing in m, then

$$\begin{aligned}
&\mathcal{V}_t^{i+1}(n, w, m^+) \\
&= \begin{cases} \max\left\{r_h(w, m^+) + \mathcal{V}_0^i(0,0,0), -r_c(w, m^+) + \beta\, \mathcal{C}_t^i(n, w, m^+)\right\}, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0^i(0,0,0), & \text{if } w = \Delta. \end{cases} \\
&\geq \begin{cases} \max\left\{r_h(w, m^-) + \mathcal{V}_0^i(0,0,0), -r_c(w, m^-) + \beta\, \mathcal{C}_t^i(n, w, m^+)\right\}, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0^i(0,0,0), & \text{if } w = \Delta. \end{cases} \qquad (3.55) \\
&\geq \begin{cases} \max\left\{r_h(w, m^-) + \mathcal{V}_0^i(0,0,0), -r_c(w, m^-) + \beta\, \mathcal{C}_t^i(n, w, m^-)\right\}, & \text{if } w \in \left[0, \bar{W}\right), \\ -r(\Delta) + \mathcal{V}_0^i(0,0,0), & \text{if } w = \Delta. \end{cases} \qquad (3.56) \\
&= \mathcal{V}_t^{i+1}(n, w, m^-) \qquad (3.57)
\end{aligned}$$

Inequality (3.55) follows from Assumptions (3.4.1) and (3.4.2). Inequality (3.56) is obtained from the induction hypothesis, as shown in (3.59). More specifically, consider $\mathcal{C}_t^i(n, w, m^+)$ for all $n \in \boldsymbol{N}$, $t \in \boldsymbol{T}$, and $w \in \boldsymbol{W}$. We first note the boundary conditions, $a_T(n, w, m) = a_t(n, w, \bar{M}) = H$, with $\mathcal{V}_T(n, w, m) = r_h(w, m) + \mathcal{V}_0(0,0,0)$ and $\mathcal{V}_t(n, w, \bar{M}) = r_h(w, \bar{M}) + \mathcal{V}_0(0,0,0)$, which are nondecreasing in m by Assumption (3.4.2). Next, we consider $\mathcal{C}_t^i(n, w, m^+)$ for all $t \in \boldsymbol{T} \setminus \{T\}$,

$$
\begin{aligned}
\mathcal{C}_t^i(n,w,m^+) &= [1-R_t(n,w)]\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right] \\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w+z,m^+ + x)\,\mathrm{d}x\,\mathrm{d}z \\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w+b,m^+ + x)\,\mathrm{d}x\,\mathrm{d}b \\
&\geq [1-R_t(n,w)]\left[-r(\Delta)+\mathcal{V}_0^i(0,0,0)\right] \qquad (3.58) \\
&+ \zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) f_t^x(x)\mathcal{V}_{t+1}^i(0,w+z,m^- + x)\,\mathrm{d}x\,\mathrm{d}z \\
&+ [1-\zeta(n,\rho)] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^b(b) f_t^x(x)\mathcal{V}_{t+1}^i(n+1,w+b,m^- + x)\,\mathrm{d}x\,\mathrm{d}b \\
&= \mathcal{C}_t^i(n,w,m^-) \qquad (3.59)
\end{aligned}
$$

We note that inequality (3.58) is obtained from the induction hypothesis. The proof follows from the convergence of value iteration algorithm. □

Proof of Theorem 3.4.3. Proof is done by contradiction. Consider two arbitrary mAb concentrations, $m^+, m^- \in \boldsymbol{M}$ such that $m^+ \geq m^-$. Theorem 3.4.3 indicates that if $a_t(n,w,m^-) = H$ then $a_t(n,w,m^+) = H$. Assume by contradiction that $a_t(n,w,m^-) = H$ and $a_t(n,w,m^+) = C$. The contradiction hypothesis implies that

$$
\begin{aligned}
r_h(w,m^-) + \mathcal{V}_0(0,0,0) &\geq -r_c(w,m^-) + \beta\mathcal{C}_t(n,w,m^-) \qquad (3.60) \\
r_h(w,m^+) + \mathcal{V}_0(0,0,0) &\leq -r_c(w,m^+) + \beta\mathcal{C}_t(n,w,m^+) \qquad (3.61)
\end{aligned}
$$

Equations (3.60) and (3.61) lead to the following inequality:

$r_h(w, m^-) - r_h(w, m^+)$

$$
\begin{aligned}
&\geq \beta\zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) q_t(x) \mathcal{V}_{t+1}(0, w+z, m^-)\,\mathrm{d}x\,\mathrm{d}z && (3.62)\\
&+ \beta\left[1-\zeta(n,\rho)\right] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b) q_t(x) \mathcal{V}_{t+1}(n+1, w+b, m^-)\,\mathrm{d}x\,\mathrm{d}b && (3.63)\\
&- \beta\zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{e+b}(z) q_t(x) \mathcal{V}_{t+1}(0, w+z, \bar{M})\,\mathrm{d}x\,\mathrm{d}z && (3.64)\\
&- \beta\left[1-\zeta(n,\rho)\right] \int_0^{\bar{W}-w}\int_0^{\bar{M}} f_t^{b}(b) q_t(x) \mathcal{V}_{t+1}(n+1, w+b, \bar{M})\,\mathrm{d}x\,\mathrm{d}b && (3.65)\\
&\geq \beta\zeta(n,\rho)\alpha_t \int_0^{\bar{W}-w} f_t^{e+b}(z)\left[r_h(w+z, m^-) - r_h(w+z, \bar{M})\right]\,\mathrm{d}z && (3.66)\\
&+ \beta\left[1-\zeta(n,\rho)\right] \int_0^{\bar{W}-w} f_t^{b}(b)\left[r_h(w+b, m^-) - r_h(w+b, \bar{M})\right]\,\mathrm{d}b && (3.67)
\end{aligned}
$$

which contradicts the condition in (3.10). Note that inequality (3.62) follows from Assumption (3.4.1), $r_c(w, m^+) - r_c(w, m^-) \geq 0$. Equations (3.63) to (3.65) follow from Property 3.4.3. We note that rewards are independent of t and n by Assumptions (3.4.1) and (3.4.2). Hence, Equations (3.66) and (3.67) are obtained from the boundary condition $a_t(n, w, \bar{M}) = H$ and the fact that $\mathcal{V}_t(n, w, m) \geq r_h(n, w, m) + \mathcal{V}_0(0, 0, 0)$ for all $(n, w, m) \in \boldsymbol{N} \times \boldsymbol{S} \times \boldsymbol{M}$ and $t \in \boldsymbol{T}$, from which the result follows. □

Chapter 4

Pooling Window Optimization Problem

4.1 Introduction

In this chapter, we focus on the manufacturing processes for engineered proteins as part of the pharmaceutical research and development efforts. These proteins are often engineered for a specific end use or application. For example, a pharmaceutical company could subcontract the manufacturing of a recombinant protein to a smaller biomanufacturing firm as part of its research and development efforts. Manufacturing of these proteins at the biomanufacturing firm would then involve specialized upstream fermentation operations followed by several downstream purification operations. In this chapter, we focus on the downstream purification operations. Purification of engineered proteins are challenging for several reasons. For example, individual proteins have unique chemical and physical properties, and their end use sets constraints on the production methods to satisfy rigorous approval processes. Further, an order has an associated yield requirement (i.e., the desired amount of the protein of interest) and a stringent purity requirement (i.e., the minimum acceptable

quality). The customer typically would not purchase the batch of proteins if it fails to meet the purity requirement. However, they might be willing to compromise on yield at a penalty cost as long as the purity requirement is met.

Table 4.1 presents a typical workflow to purify an engineered protein. Upon the receipt of an order, the scientist at the biomanufacturing firm starts performing small scale scouting runs. Scouting runs represent a set of experiments that helps the scientist collect data about the purification attributes of this protein on several alternative techniques (referred to as *chromatography techniques*). Once the performance of the protein with respect to the available chromatography techniques have been identified, the scientist performs validation runs. The role of the validation runs is to mitigate risks and quantify the yield and purity to be expected in the subsequent production runs at larger scale. For this purpose, the scientist conducts several what-if experiments to explore the performance of alternative operating policies that could potentially achieve the specific requirements on yield and purity. Once the best operating policy is identified, the production run is performed at larger scale to achieve the end product that satisfies the specific production requirements. The overall process often takes 3 to 5 weeks due to the experimental nature of the purification operations. Further, the scouting and validation runs could be as expensive as the production runs themselves. While the scouting and production runs are inevitable for engineered proteins, we believe that the intermediate validation runs present a significant opportunity for reducing lead times and costs through application of the operations research techniques. One of the main objectives in this study is to investigate whether the information obtained from scouting runs can directly be fed into an optimization model to identify the optimal purification polices that can be used in production runs, thereby reducing costs and lowering lead times. As shown in Table 4.1, reducing the time spent in the validation runs via an optimization

Table 4.1: Current and proposed workflow for purification development

Current workflow	Scouting runs 1 week, \$3x cost	$\rightarrow$	Validation runs 1-2 weeks, \$3x cost	$\rightarrow$	Production run 1-2 weeks, \$4x cost
Proposed workflow	Scouting runs 1 week, \$3x cost	$\rightarrow$	Optimization Model $\leq$ 1 day, $\leq$ \$0.1x cost	$\rightarrow$	Production run 1-2 weeks, \$4x cost

model could improve the total cost and lead time up to 33% while also freeing up the associated capacity.

Protein purification operations involve several operational challenges, such as, yield and quality trade-offs, randomness in the starting material, expensive labor and equipment costs, and large penalty costs when the production requirements are not satisfied. Variability in the starting material along with the limitations in chromatography techniques impose significant challenges in meeting the predetermined requirements on purity and yield. For example, if the starting material does not have enough protein and/or has excess amount of impurities, then the specific requirements on the final purity and yield might never be satisfied, even though the scientist takes the optimal courses of purification actions. In such circumstances, committing to the purification order could substantially hurt both the client and the biomanufacturing firm. As pointed out by our industry collaborator, Tom Foti, the Vice President of Aldevron, predicting the failures "earlier than later" is critical.

In this chapter, we provide an optimization framework that quantifies the risks and costs in protein purification operations to answer the following questions: (i) For a given starting material, can the biomanufacturing company determine whether the final purity and yield requirements specified by the customer are achievable at all? If achievable, can we establish performance guarantees for these specific requirements? If not achievable, can we develop guidelines on the starting material to predict the batch failures in advance? *(ii)* How easy or complex is the purification process likely

to be, based on the starting material and limitations in the available purification techniques? How can the total profit be maximized for each purification order? *(iii)* How would the insights and policies be different, if the biomanufacturing company adopts a conservative worst case approach for hedging against uncertainties and large failure costs? By answering these questions using an optimization framework, we believe that biomanufacturing firms can significantly improve their profitability and reduce their lead times for purification of engineered proteins.

To answer these questions, we analyze the protein purification problem using the dynamic programming approach. Our contributions are as follows: First, we investigate the structural properties of the state space, and partition the state space into *decision zones* having similar financial characteristics. More specifically, the decision zones provide a rigorous and formal assessment of the starting material, manufacturing capabilities and business risks at the beginning of each purification run. Next, we propose a *zone-based decision making* approach which is particularly useful in practice since it provides optimal policies based on the condition of the starting material. Insights from the structural analysis are then used to develop a state aggregation and an action elimination scheme that leads to computational advantage in solving realistic industry problems.

A key aspect of our work is that we not only provide *optimal purification policies* using stochastic optimization, but also provide *guaranteed performance* using a worst-case analysis approach to generate the decision zones. We adopt this strategy because of the randomness, high operating costs, and penalty costs involved in industry practices. Biomanufacturing companies often need guaranteed performance measures to ensure profitability and customer satisfaction. Our models provide practical guidelines to evaluate the profitability and failure risk of a starting material

provided by a customer. To our knowledge, such guaranteed performance measures have not been investigated yet in the context of biomanufacturing.

This research is an outcome of an ongoing multi-year collaboration with Aldevron (2013-2015). Aldevron (www.aldevron.com) is a contract biomanufacturing firm specializing in a variety of services including plasmid DNA, protein production services and antibody development. At Aldevron's daily operations, the optimization framework has been in use for all R&D protein purification projects since October 2014. The implementation has resulted in an average of 25% reduction in the total lead times and 20% reduction in operating costs required for protein purification operations, as discussed in Sections 4.7-4.8. Our research outcomes have also been shared and validated with a larger biomanufacturing community (BioWGS, 2014). Through industry implementation, we observe that the optimization framework has the potential for significantly scaling back if not eliminating the validation runs. Our study is one of the first attempts to apply operations research concepts to purification of engineered proteins, and integrates the knowledge from biological engineering and stochastic modeling to derive guidelines that improve industry practices.

The remainder of the chapter is organized as follows. Section 4.2 provides a background on purification operations and introduces the trade-offs and challenges. We develop a mathematical model in Section 5.2, and analyze its structural properties in Sections 4.4-4.6. We demonstrate the implementation of the optimization model in Sections 4.7-4.8, and provide concluding remarks in Section 9.

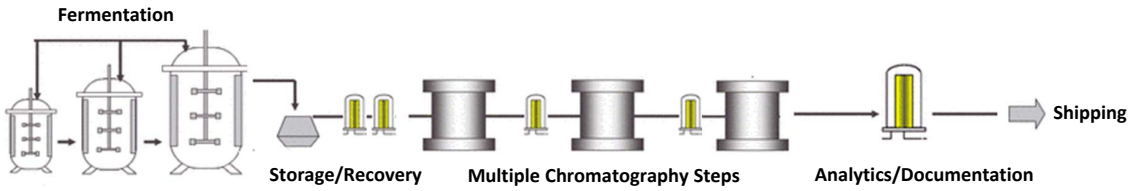

Figure 4.1: Typical manufacturing stages in biomanufacturing

4.2 Background in Protein Purification

A typical biomanufacturing process consists of upstream fermentation operations where bacteria or eukaryotic cells produce the proteins of interest, and downstream purification operations where these proteins are purified through multiple chromatography steps (See Figure 4.1). The primary output of fermentation is a batch mixture that includes the protein of interest and significant amounts of other unwanted impurities derived from the host cell or fermentation medium. After fermentation, this batch must be purified using multiple chromatography steps (typically, $2-6$ steps) based on specific production requirements. The objective of each chromatography operation is to separate the protein of interest from unwanted impurities to achieve the desired purity. In this chapter, we focus on optimizing protein purification decisions related to chromatography operations. We first provide a brief background on chromatography operations, and then introduce the process challenges and trade-offs.

4.2.1 Chromatography Operations

Chromatography is one of the most common but also most challenging operations in biomanufacturing (Farid, 2008; Polykarpou et al., 2011b). The objective of chromatography operations is to separate the protein of interest from unwanted impurities to meet a pre-determined *purity requirement* specified by the end use or application. *Purity* represents the fraction of the proteins of interest in a batch based on the total amount of both host-cell proteins and impurities. Purity requirement is defined by the end use or application of the purified protein. For example, a protein

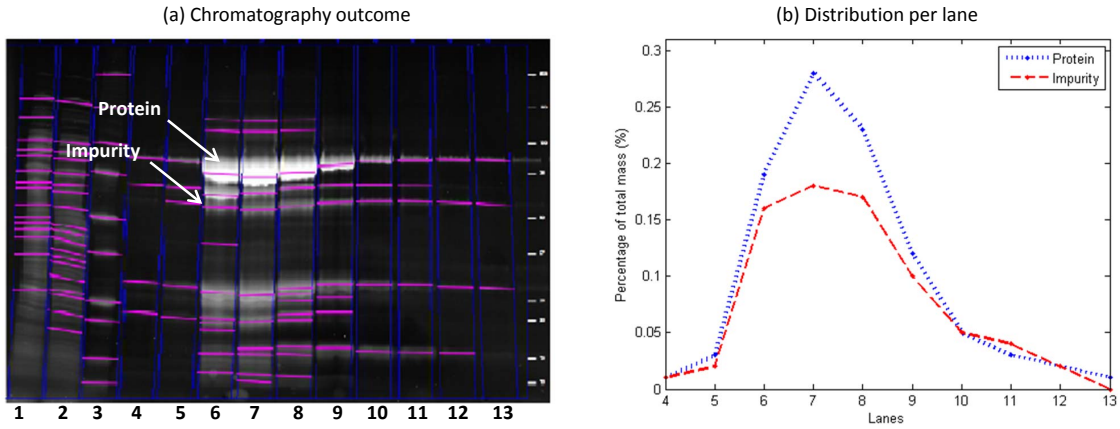

Figure 4.2: An example of chromatography output

to be used as a drug must be highly pure (i.e., 99.9% purity), whereas a protein used for a feed study could have lower purity requirement (i.e, 85% purity).

Chromatography operations are performed in a cylindrical column that is packed with special resins that bind to either the protein of interest or impurities. Chromatography techniques rely on the difference in physico-chemical characteristics between the proteins and impurities to separate one from other, i.e., difference in molecular weight, shape, charge, hydrophobicity, and affinity for a ligand. For example, gel filtration chromatography separates the target protein from impurities based on differences in size and shape, whereas ion-exchange chromatography relies on the difference in electric charges. A typical purification process could involve 2-6 chromatography steps using different techniques, and each step could take 8 hours or more, depending on the process conditions.

Yield and Purity Trade-offs

Figure 4.2 (a) presents an example of chromatography output. This example uses the differential affinity of proteins to divalent metal ions as the separation principal. In Figure 4.2 (a), we can observe distinct columns on the x-axis called as *lanes*. Each

lane can be thought as equivalent to a discrete time interval (i.e., close to 1 minute in practice). Each lane is comprised of some fraction of the total amount of the protein of interest, as well as a fraction of the total amount of the unwanted impurities. The y-axis in Figure 4.2 (a) represents the molecular mass of the target protein and impurities associated with each lane. Figure 4.2(b) plots the fraction of protein and impurity for each lane when the chromatography process described in Figure 4.2 (a) is used.

The scientist performing the purification must decide which lane to *'pool'*. For example, if she pools lanes 6-10, then she collects a large fraction of the protein along with a large fraction of the impurity. On the other hand, if she pools lanes 9-10, she compromises on the yield (i.e., collects a smaller fraction of the protein), but improves the purity (i.e, collects a smaller fraction of the impurity). This illustrates one of the main trade-offs related with yield and purity of a chromatography step. In this example, the scientist can choose to pool any consecutive lanes between lanes 4-13. For instance, lanes 4 to 11, lanes 6 to 10, lanes 7 to 9, are examples of some alternative candidate *pooling windows*. Depending on the outcome of a chromatography step, the scientist actually makes decisions regarding the chromatography technique *and* the pooling window for each of the subsequent steps. In fact, identifying the sequence of chromatography techniques itself is a separate optimization problem. However, we consider purification settings where this sequence is pre-determined based on the scouting runs, and only focus on the pooling window selection problem at each chromatography step.

Challenges in Decision Making

The main challenges in choosing pooling windows to optimize the purification process are summarized as follows. (1) *Yield and quality trade-offs.* Each order is associated

with predetermined protein yield and purity requirements specified by the end use or application. If the final batch does not meet either of the yield and purity requirements, the company incurs large penalty costs associated with yield shortages and quality failures. The scientist might have to compromise on the protein yield to achieve the desired purity. (2) *Engineered proteins.* Each order is unique in the sense that the scientist re-engineers and manufactures each order (protein-impurity mixture) for the first time. This requires evaluating and optimizing each order independently, unlike mass production. (3) *Uncertainty.* Protein and impurity amounts at the beginning of each chromatography step involves uncertainty. Also, each chromatography step has limited capability in terms of separating proteins from impurities. (4) *Interlinked decisions.* Purification involves multiple chromatography steps in series. The output of each step affects the input for subsequent steps and ultimately to possibility for successfully attaining the yield and purity requirements specified by the end use. (5) *Starting batch.* The starting material is typically manufactured through fermentation, and the scientist involved in purification might have limited control over it. Fermentation operations often use bacteria or eukaryotic cells, to manufacture the starting material. The use of these cells introduces variability and uncertainty in the amount of proteins and impurities obtained from the fermentation operation. These in turn affect the subsequent purification decisions. (6) *Problem size.* The problem involves large state and action spaces, challenging the decision making in practice. For example, the state space is typically in terms of milligrams of proteins, whereas the action space increases exponentially in the number of purification steps.

4.3 The Model

We formulate a finite horizon Markov decision model for purification decisions.

Decision epochs: A decision epoch represents the beginning of a chromatography step t, $\mathcal{T} = \{t : 1, \ldots, T-1\}$, where $T-1$ denotes the last chromatography step. We let T represent the terminal period. At time T, the batch is either shipped to the customer or scrapped.

States: The state space is defined as $\mathcal{P} \times \mathcal{I} \cup \{\Delta\}$. The state $p_t \in \mathcal{P}$, $p_t \in [0, p_1]$, denotes the amount of protein of interest available in the batch at the beginning of t^{th} chromatography step. Similarly, $i_t \in \mathcal{I}$, $i_t \in [0, p_1]$ denotes the amount of impurity at the beginning of t^{th} chromatography step. Note that the starting material $(p_1, i_1) \in \mathcal{P} \times \mathcal{I}$ represents the protein and impurity amounts obtained from fermentation operations. Typically, both of the protein and impurity amounts are in milligrams. The batch has the maximum possible amount of protein and impurity (p_1, i_1) at the beginning of the first chromatography step. The stopping state $\{\Delta\}$ is an absorbing state with no rewards, and represents a batch which is either shipped to the customer or scrapped.

Actions: The action space is defined as $\mathcal{A}_t = \mathcal{W}_t \cup \{S\}$. Let $a_t(p_t, i_t)$ denote the action selected at the beginning of chromatography step $t \in \mathcal{T}$ at state (p_t, i_t). The action $w_t \in \mathcal{W}_t$ denotes the pooling window w_t corresponding to the chromatography step $t \in \mathcal{T}$. Let $\mathcal{L}_t$ denote an ordered set of lanes available at each chromatography step t, where $\mathcal{L}_t = \{1, 2, \ldots, L_t\}$. Then, a pooling window w_t corresponds to a subset of consecutive lanes from the set $\mathcal{L}_t$, where the set of all possible pooling windows at a chromatography step $t \in \mathcal{T}$ is $\mathcal{W}_t = \big\{(i, \ldots, j) \in \mathcal{L}_t : \; j = i + k, \; i = \{1, \ldots, L_t\}, k = \{0, 1, \ldots, L_t - i\}\big\}$. The total number of possible pooling windows at each chromatography step $t \in \mathcal{T}$ is denoted by N_t. Note that N_t is large but finite and bounded. The action $\{S\}$ represents the action of stopping the purification process. Once the purification stops, the batch is

either shipped or scrapped. The operator can decide to stop the purification at the beginning of any chromatography step $t \in \mathcal{T}$. Note that, at the terminal period T, the only available action is to stop, $a_T(p_T, i_T) = S$ for all $(p_T, i_T) \in \mathcal{P} \times \mathcal{I}$. Similarly, $a_t(\Delta) = S$ for all $t \in \mathcal{T} \cup \{T\}$.

Transitions: The transition probabilities are defined based on the mathematical models for chromatography operations (Vasquez-Alvarez et al., 2001; Salisbury et al., 2006; Polykarpou et al., 2011b). We adopt these models to identify the amount of proteins and impurities that remain in the batch after completion of the chromatography step $t \in \mathcal{T}$. At each chromatography step $t \in \mathcal{T}$, a random fraction $\{\Psi_t|w_t\}$ of the impurity i_t is carried over the next step $t+1$, implying that the remaining amount of impurity was eliminated through the chromatography step t. Therefore,

$$i_{t+1} = (\psi_t|w_t)i_t. \tag{4.1}$$

The random fraction $\{\Psi_t|w_t\}$ has distribution $g_t(\cdot|w_t)$ with finite support $[\psi_t^l|w_t, \psi_t^u|w_t]$ for all $w_t \in \mathcal{W}_t$, $t \in \mathcal{T}$. It represents a random fraction of the amount of impurities i_t that remains inside the batch at the beginning of $(t+1)^{th}$ step, given that there are i_t units of impurities at the beginning of chromatography step t, and the pooling window w_t is selected for the chromatography step t.

Similarly, at each chromatography step $t \in \mathcal{T}$, a random fraction $\{\Theta_t|w_t\}$ of the protein of interest is carried over the next step $t+1$, implying that the remaining amount of the protein was eliminated during the chromatography step t. Therefore,

$$p_{t+1} = (\theta_t|w_t)p_t. \tag{4.2}$$

$\{\Theta_t|w_t\}$ represents the random fraction of protein p_t that remains inside the batch at the beginning of $(t+1)^{th}$ step, given that there are p_t units of protein at

the beginning of chromatography step t, and the pooling window w_t is selected. The random fraction $\{\Theta_t|w_t\}$ has distribution $f_t(\cdot|w_t)$ with finite support $[\theta_t^l|w_t, \theta_t^u|w_t]$ for all $w_t \in \mathcal{W}_t$, $t \in \mathcal{T}$.

The probability density functions $f_t(\cdot|w_t)$ and $g_t(\cdot|w_t)$ can be different for each chromatography step t, depending on physico-chemical characteristics of the proteins, impurities, and specific chromatography technique chosen at each step. Regardless, the finite support for each $f_t(\cdot|w_t)$ and $g_t(\cdot|w_t)$ can be determined based on the physico-chemical characteristics of the molecules and chromatography techniques, described in Section 4.2. We assume that Θ_t and Ψ_t are independent based on the fact that proteins of interest and impurities have distinct physical and chemical characteristics (Vasquez-Alvarez et al., 2001; Polykarpou et al., 2011b). Available chromatography techniques mainly differ in terms of how they exploit these unique characteristics to separate proteins from impurities.

One of the key performance measures for a chromatography technique is its *purification capability* under the pooling window w_t. The purification capability is defined in terms of the best and worst possible amounts of proteins and impurities that might remain in the batch at the end of the chromatography step t. Note that for each chromatography step t and pooling window w_t, $\{\Theta_t|w_t\}$ and $\{\Psi_t|w_t\}$ have finite support $(\theta_t^l|w_t, \theta_t^u|w_t)$ and $(\psi_t^l|w_t, \psi_t^u|w_t)$, respectively. This implies that $(\theta^u, \psi^l|w_t)$ and $(\theta^l, \psi^u|w_t)$ denote the best and worst purification capabilities when w_t is chosen at the chromatography step t. Then, let $(\bar{\theta}_t, \bar{\psi}_t|w_t)$ define the expected purification capability of the chromatography step t using window w_t. This purification capability associated with each window w_t at the chromatography step t is used to identify a stochastic ordering of actions for the structural analysis in Section 4.6. The system transitions to state $\{\Delta\}$ if the purification process is abandoned at any step $t \in \mathcal{T}$,

or at the end of planning horizon T.

Purity Requirement and Costs: The quality of a batch at chromatography step t is measured in terms of its *purity*, defined as $\gamma_t = \frac{p_t}{p_t+i_t}$ for $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ and $t \in \mathcal{T} \cup \{T\}$. Batch purity is a critical performance measure, and a minimum purity level γ_d is often a part of the requirements specified by the end use or application. If a batch does not meet this minimum purity requirement, i.e., $\gamma_t < \gamma_d$, then the customer would not purchase that batch. Only batches that achieve the predefined quality standard $\gamma_t \geq \gamma_d$ are shipped to the customer. The desired purity level could range from 85% to 99.9% based on specific requirements of each order. Since biomanufacturing firms do not have financial incentives to attain purity levels higher than the minimum requirement, the goal is often to meet the purity requirement γ_d.

After the completion of a purification project, if the batch does not meet the minimum purity requirement, a penalty cost of quality failure c_f is incurred. Penalty cost c_f could vary from company to company, and could include penalties associated with project delays, loss of reputation, and its impact on future orders. Operating costs of a chromatography step t is denoted by c_t, and include raw material costs (resins and buffers), equipment and labor costs, and quality control costs (HPLC, analytics, documentation). Operating costs could be different at each step t based on the type of resin, buffer, column, and other specifications of the chromatography techniques used at each step (Farid, 2008).

Yield Requirement and Revenue: At the completion of purification, the revenue obtained from a batch depends on both its purity and yield. If the batch meets the minimum purity requirement $\gamma_t \geq \gamma_d$, then the revenue obtained from per unit of protein is r; otherwise $r = 0$. Each order has a predetermined *yield requirement* p_d

specified by the end use or application, such that, the total revenue expected to obtain from the purification project is $\mathbb{1}_{\gamma_t \geq \gamma_d} r(p_d)$, where $\mathbb{1}(\cdot)$ is the indicator function. Further, $\mathbb{1}_{\gamma_t \geq \gamma_d} r(p_t) = \mathbb{1}_{\gamma_t \geq \gamma_d} r(p_d)$ for $p_t \geq p_d$, i.e, customers do not pay for proteins manufactured in excess of their yield requirement. However, if the batch yields less protein than the customer requirement p_d, the biomanufacturing company incurs a yield penalty cost $c_\ell(p_d - p_t)^+$, where c_ℓ denotes the shortage cost per unit of protein short. Note that both the penalty cost and revenues obtained from the purification are linear in protein amounts. If purification stops at step $t \in \mathcal{T} \cup \{T\}$ and $\mathbb{1}_{\gamma_t \geq \gamma_d, p_t < p_d}$, the total profit is given by $r(p_t) - c_\ell(p_d - p_t)^+$. Also note that $r(p_t) - c_\ell(p_d - p_t)^+$ can be negative. In practice, the typical yield requirement could be as low as 25%-40% of the amount of proteins p_1 initially available in the batch, reflecting the manufacturing challenge involved in achieving purity levels of 85%-99.9%. Therefore, the final reward $r_S(p_t, i_t)$ resulting from stopping the purification process at state $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ is as follows:

$$r_S(p_t, i_t) = \begin{cases} -c_f & \text{if } \gamma_t < \gamma_d, \\ r(p_d) & \text{if } \gamma_t \geq \gamma_d \text{ and } p_t \geq p_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } \gamma_t \geq \gamma_d \text{ and } p_t < p_d, \end{cases} \tag{4.3}$$

for $t \in \mathcal{T}$ when $a_t(p_t, i_t) = S$, and for $t = T$. Note that $r \leq c_\ell \leq c_f$. The stopping state $\{\Delta\}$ has no rewards $r_S(\Delta) = 0$.

The Value Function: We formulate a finite horizon non-discounted Markov decision model with the following value function $\mathcal{V}_t(p_t, i_t)$ for all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$:

$$\mathcal{V}_t(p_t, i_t) = \max_{\{w_t, S\} \in A_t} \Big\{ r_S(p_t, i_t), -c_t + \mathop{\mathbb{E}}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \Big\}, \text{for } t = \{1, \ldots, T-1\} \tag{4.4}$$

$$\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T), \tag{4.5}$$

where the expectation is taken with respect to the probability distribution of θ_t and ψ_t, i.e.,

$$\underset{\theta_t,\psi_t|w_t}{\mathbb{E}} \mathcal{V}_{t+1}(p_t\theta_t, \psi_t i_t) = \int_{\psi_t^l|w_t}^{\psi_t^u|w_t} \int_{\theta_t^l|w_t}^{\theta_t^u|w_t} f_t(\theta_t|w_t) g_t(\psi_t|w_t) \mathcal{V}_{k+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi \tag{4.6}$$

Note that $\mathcal{V}_t(\Delta) = 0$ for $t \in \mathcal{T}$. Let π_t^* denote the optimal pooling policy from step $t \in \mathcal{T}$ until the final time T. If w_t^* maximizes the right hand side of Equation (4.4) for each (p_t, i_t) and t, the policy $\pi_1^* = \{w_1^*, \ldots, w_T^*\}$ is optimal (Puterman, 1994).

Note that purity and yield requirements are not imposed as hard constraints in the mathematical model, but through cost penalties. In the biomanufacturing industry, customers understand the challenges involved in the manufacture of engineered proteins, and recognize that the manufacturer might not necessarily be able to meet both the yield and purity requirements. However, severe penalty costs associated with failure to meet the yield and purity requirements makes this a challenging business environment. We do not consider the discount factor in our formulation because the purification operations typically consists of $2-6$ chromatography steps, which represent a short-term planning horizon (1 to 7 days) compared to the overall protein manufacturing lead time (i.e., $7-8$ weeks on average). In this setting, discounting the value function could lead to a bias in decision making. Further, since the motivating industry setting involves contract biomanufacturers where each batch represents an engineered protein uniquely made for a specific customer demand, a finite horizon optimization model for each batch is reasonable.

4.4 Structural Analysis of the State Space and Bounds

In this section, we investigate the structural properties of the state space and provide guidelines to quantify risks and costs associated with chromatography operations. We partition the state space into nonempty sets called as zones (namely *failure* zone in Section 4.4.1, *target* zone in Section 4.4.2 and *risk* zone in Section 4.4.3). To do so, we first establish some important structural properties of the value function in Proposition 4.4.1.

Proposition 4.4.1. *The value function $\mathcal{V}_t(p_t, i_t)$ is nondecreasing in $p_t \in \mathcal{P}$ for a given $i_t \in \mathcal{I}$, and nonincreasing in $i_t \in \mathcal{I}$ for a given $p_t \in \mathcal{P}$, for all $t \in \mathcal{T} \cup \{T\}$.*

Proof See Appendix.

Monotonicity of the value function in Proposition 4.4.1 implies that the optimal profit obtained from a batch never decreases as the protein amount increases, and never increases as the impurity amount increases. Note that Proposition 4.4.1 holds for any probability density functions $f_t(\cdot)\mathrm{d}\theta$ and $g_t(\cdot)\mathrm{d}\psi$ as long as they are well behaved (i.e., finite moments). In subsequent sections, we use the monotonicity of the value function to identify several structural properties of the state space.

4.4.1 Failure Zones

We define the failure zone $\mathbb{F}_t$ which characterizes the minimum amount of purity and yield required for a batch at the beginning of each chromatography step $t \in \mathcal{T}$, i.e., the biomanufacturer has no financial incentives to perform the purification, if the batch does not meet these minimum requirements at the start of each chromatogra-

phy step $t \in \mathcal{T}$.

Theorem 4.4.1. *[Failure Zone] The optimal policy has the property that for some* (p'_t, i'_t)*, where* $\frac{p'_t}{(p'_t+i'_t)} < \gamma_d$ *for* $t \in \mathcal{T}$*, the optimal action is to abandon the purification* $a^*_t(p'_t, i'_t) = S$ *for all states in* $\mathbb{F}_t = \{(p_t, i_t) \in \mathcal{P} \times \mathcal{I} : p_t \leq p'_t \text{ and } i_t \geq i'_t\}$*.* $\mathbb{F}_t$ *is called as the failure zone at the chromatography step* $t \in \mathcal{T}$*.*

Proof See Appendix.

Theorem 4.4.1 states that there exists a set of $(p_t, i_t) \in \mathbb{F}_t$, called as the failure zone $\mathbb{F}_t$, where the optimal action is to abandon the purification operation and scrap the batch, $\mathbb{1}_{(p_t,i_t)\in\mathbb{F}_t} a^*(p_t, i_t) = S$, for $t \in \mathcal{T}$. Note that Theorem 4.4.1 does not require the specific knowledge of the probability density functions $f_t(\cdot)\mathrm{d}\theta$ and $g_t(\cdot)\mathrm{d}\psi$; and only uses the monotonic behavior that follows from Equations (4.1)-(4.2), i.e., the state $(p_{t+1}, i_{t+1}) \in \mathcal{P} \times \mathcal{I}$ at step $t+1 \in \mathcal{T}$ is a non-increasing function of the state $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ at step $t \in \mathcal{T}$. Figure 4.3 illustrates an example of the failure zone using industry data for a chromatography step. Next, Proposition 4.4.2 characterizes the failure zone $\mathbb{F}_t$ at step $t \in \mathcal{T}$, in terms of the costs and purification capabilities of the remaining steps.

Proposition 4.4.2. *A batch state* $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ *with* $\frac{p_t}{(p_t+i_t)} < \gamma_d$ *belongs to the failure zone* $\mathbb{F}_t$ *at the chromatography step* $t \in \mathcal{T}$*, if either of the following conditions hold:*

(i) $i_t > p_t \dfrac{1-\gamma_d}{\gamma_d} \displaystyle\prod_{w_j} \frac{(\theta^l_j|w_j)}{(\psi^u_j|w_j)}$

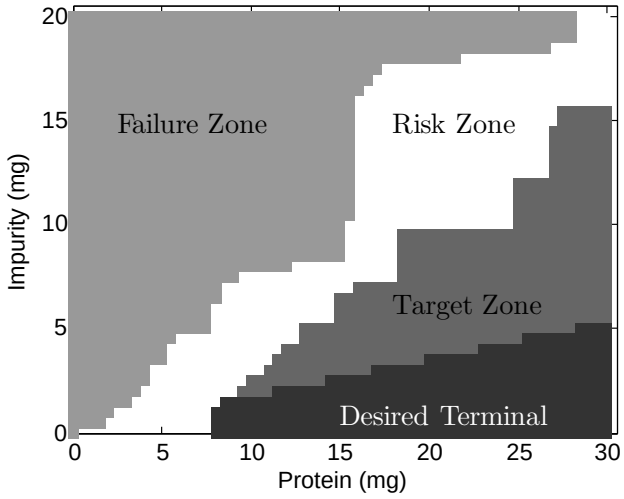

Figure 4.3: An example of the zones for a chromatography step using industry data

for all $\pi_t = (w_t, w_{t+1}, \ldots, w_{T-1})$, *and* $j = \{t, \ldots, T-1\}$,

(ii) $p_t \leq \dfrac{-c_f + c_t + c_l\big(p_d - p_t \prod_{w_j}(\theta_j^l|w_j)\big)^+}{r \prod_{w_j}(\theta_j^l|w_j)}$ *and* $i_t \leq p_t \dfrac{1-\gamma_d}{\gamma_d} \prod_{w_j} \dfrac{(\theta_j^l|w_j)}{(\psi_j^u|w_j)}$

for all $\pi_t = (w_t, w_{t+1}, \ldots, w_{T-1})$, *and* $j = \{t, \ldots, T-1\}$, $\{t, j\} \in \mathcal{T}$.

Proof See Appendix.

Condition (i) in Proposition 4.4.2 represents the case where the purity and yield requirements lie outside the purification capability of all the purification strategies $w_j \in \mathcal{W}_j$ available in the remaining steps $j = t, \ldots, T-1$. Condition (ii) corresponds to the case where none of the purification strategies w_j available in the remaining steps $j = t, \ldots, T-1$ provide adequate financial incentives for continuing the purification process. Note that Proposition 4.4.2 defines $\mathbb{F}_t$ using the worst possible realization of the purification outcomes $(\theta_t^l, \psi_t^u|w_t)$ across all pooling windows w_t at all chromatography steps $t \in \mathcal{T}$. This yields a worst-case classification of states in $\mathbb{F}_t$ and provides a guaranteed performance measure to hedge against failures and large penalties.

4.4.2 Target Zones

We characterize a particular subset in the state space, called as the *target zone* $\mathbb{T}_t$ at the start of a chromatography step $t \in \mathcal{T}$, where both yield and purity requirements specified by the customers can be eventually achieved at the end of the planning horizon T, if $(p_t, i_t) \in \mathbb{T}_t$ and the optimal pooling actions are taken at each step $t, \ldots, T-1$ (See Section 4.6 for a discussion on the optimal actions). Such guaranteed performance measures are critical in a biomanufacturing setting, because it helps to hedge against uncertainties and failures, and also better manage customer expectations. Furthermore, the target zone $\mathbb{T}_t$ provides a measure of the difficulty involved in meeting customer requirements on final purity and yield from the state (p_t, i_t) at the chromatography step $t \in \mathcal{T}$. To characterize the target zone $\mathbb{T}_t$ at each step t, we use recursion based on the worst possible outcome corresponding to an action w_t at each chromatography step $t \in \mathcal{T}$. We first define the terminal zone $\mathbb{S}$ of the purification operations in Definition 4.4.1.

Definition 4.4.1. *The terminal zone* $\mathbb{S}$ *corresponds to the set of protein and impurity states that meet both yield and purity requirements specified by the end use or application, i.e.,*

$$\mathbb{S} = \left\{ (p_t, i_t) : \; p_t \geq p_d, \; i_t \leq \frac{1-\gamma_d}{\gamma_d} p_t \right\} \textit{ for } (p_t, i_t) \in \mathcal{P} \times \mathcal{I}, t \in \mathcal{T}. \tag{4.7}$$

It follows that, if the batch is in the terminal zone at the beginning of a chromatography step $t \in \mathcal{T}$, i.e., $(p_t, i_t) \in \mathbb{S}$, then the purification can be stopped at the step t, and the batch can be shipped to the customer.

Next, the terminal zone $\mathbb{S}$ is used to define the *target zone* $\mathbb{T}_t$ for each chromatography step $t \in \mathcal{T}$. If a batch (p_t, i_t) belongs to the target zone $\mathbb{T}_t$ at the beginning of the

chromatography step t, then there *exists* a sequence of actions that will guarantee that both yield and purity requirements can be achieved by the end of the planning horizon T. Clearly, based on the terminal zone $\mathbb{S}$ in Definition 4.4.1, the target zone $\mathbb{T}_T$ at the end of the planning horizon T is defined as $\mathbb{T}_T = \Big\{(p_T, i_T) :\ p_T \geq p_d,\ i_T \leq \frac{1-\gamma_d}{\gamma_d} p_T\Big\}$. Next, we define the operator J_t^w as follows,

$$J_t^w(p_{t+1} \geq x, i_{t+1} \leq y\, p_{t+1}) :$$

$$(p_{t+1} \geq x, i_{t+1} \leq y\, p_{t+1}) \rightarrow (p_t \geq \frac{x}{(\theta_t^l|w)}, i_t \leq y p_t \frac{(\theta_t^l|w)}{(\psi_t^u|w)}). \tag{4.8}$$

The input to the operator $J_t^w(\cdot)$ are the bounds on p_{t+1} and i_{t+1} at the beginning of $(t+1)^{th}$ step, and the operator scales its input to provide output corresponding to the bounds on p_t and i_t at the start of t^{th} step for a particular choice of window $w \in \mathcal{W}_t$. The operator $J_t^w(\cdot)$ uses the worst-case purification capability of the window $w \in \mathcal{W}_t$ to determine these bounds. Using the operator $J_t^w(\cdot)$, Proposition 4.4.3 characterizes the target zones $\mathbb{T}_t$ for each chromatography step $t \in \mathcal{T}$.

Proposition 4.4.3. *The target zone $\mathbb{T}_t$ at the beginning of the chromatography step t is defined as*

$$\mathbb{T}_T = \Big\{(p_T, i_T) :\ p_T \geq p_d,\ i_T \leq \frac{1-\gamma_d}{\gamma_d} p_T\Big\}, \tag{4.9}$$

$$\mathbb{T}_t = \bigcup_{w \in \mathcal{W}_t} J_t^w(\mathbb{T}_{t+1}), \tag{4.10}$$

for the pooling windows $w \in \mathcal{W}_t$ and steps $t \in \mathcal{T}$, where $J_t^w(\cdot)$ is defined in Equation (4.8).

Proof See Appendix.

The target zone $\mathbb{T}_t$ in Proposition 4.4.3 is obtained recursively using the worst case purification capability $(\theta_t^l, \psi_t^u|w_t)$ for each available pooling window w_t in each of the remaining chromatography steps $t \in \mathcal{T}$. It identifies all states (p_t, i_t) that

can lead to the terminal zone $\mathbb{S}$ by the end of the planning horizon T. Figure 4.3 demonstrates an example of the target zone for a chromatography step using industry data. The following characteristics of the target zones follow from Proposition 4.4.3, and provide important managerial implications (Bertsekas and Rhodes, 1971):

(i) At the beginning of t^{th} chromatography step, if the starting state of the batch (p_t, i_t) belongs to the target zone $\mathbb{T}_t$, then the scientist can always guarantee that there exists at least one purification strategy that leads to the terminal zone $\mathbb{S}$ by the time T.

(ii) The target zone provides bounds $(\hat{p}_t,\ \hat{i}_t\}$ on the starting state at step t, such that, $\mathbb{T}_t = \{(p_t, i_t) :\ p_t \geq \hat{p}_t,\ i_t \leq \hat{i}_t\}$ for $t \in \mathcal{T}$.

The characteristics of the target zones (item (i) and (ii) listed above) have practical implications for managing purification operations. For example, item (i) indicates that target zones provide performance guarantees in terms of achieving the purity and yield requirements since Proposition 4.4.3 uses the worst-case outcomes in each chromatography step. Despite limitations in the chromatography techniques, such performance guarantees about the ability to meet specific customer requirements provide a competitive advantage to the biomanufacturing firms by ensuring customer satisfaction. Item (ii) states that the target zone has a threshold type structure, and hence can be easily interpreted and implemented in practice.

4.4.3 Risk Zones and Bounds on the Value Function

As a direct consequence of the target zones $\mathbb{T}_t$ and failure zones $\mathbb{F}_t$, we define the risk zone, $\mathbb{R}_t = \left\{(p_t, i_t) |\, p'_t < p_t < \hat{p}_t,\ \hat{i}_t < i_t < i'_t, \text{ where } (p'_t, i'_t) \in \mathbb{F}_t, \text{ and } (\hat{p}_t, \hat{i}_t) \in \mathbb{T}_t\right\}$ for all $t \in \mathcal{T}$. The risk zone includes all states (p_t, i_t) that are neither in the target zone $\mathbb{T}_t$, nor in the failure zone $\mathbb{F}_t$ at the beginning of the step $t \in \mathcal{T}$. This subset of the

state space is a measure of the financial risk associated with purifying a particular batch. For example, a batch that is in the risk zone $\mathbb{R}_t$ at the start of the step t could either achieve the final purity and yield requirements at time T or fail to do so leading to large penalties associated with shortage costs or quality failures. Next, we characterize the bounds on the optimal value function $\mathcal{V}_t^*(p_t, i_t)$ based on the zones at each step $t \in \mathcal{T}$ as follows:

$$\mathcal{V}_t^*(p_t, i_t) = -c_f \text{ for all } (p_t, i_t) \in \mathbb{F}_t, t \in \mathcal{T}. \tag{4.11}$$

$$\sum_{j=t}^{T-1} -c_j + r(p_d) \leq \mathcal{V}_t^*(p_t, i_t) \leq r(p_d) \text{ for all } (p_t, i_t) \in \mathbb{T}_t, t \in \mathcal{T}. \tag{4.12}$$

$$-c_f \leq \mathcal{V}_t^*(p_t, i_t) \leq \sum_{j=t}^{T-1} -c_j + r(p_d) \text{ for all } (p_t, i_t) \in \mathbb{R}_t, t \in \mathcal{T}. \tag{4.13}$$

Note that Equation (4.11) for the failure zone is a direct consequence of Theorem 4.4.1. Similarly, the cost bounds on the target zone in inequality (4.12) follow from Proposition 4.4.3 and the stopping cost structure in Equation (4.3). The cost bounds on the risk zone in inequality (4.13) follow from the monotonicity of the value function in Proposition 4.4.1. Since the state space is continuous and has high dimension, these bounds provide managerial insights to quantify the risks and costs for states within each zone. For example, any batch state in the failure zone will result in large penalty cost $-c_f$; whereas a batch state in the target zone can lead to a large reward up to $r(p_d)$. On the other hand, a batch in the risk zone can lead to financial losses due to shortage costs. Insights from the bounds provide basis for an aggregation scheme in Section 4.5 and are used in the structural analysis of the optimal purification policies in Section 4.6.

4.5 State Aggregation, Action Elimination and Ordering Scheme

We use insights from the structural analysis of the state space to construct a state aggregation and an action elimination procedure for the Markov decision model. Recall that the state space is continuous, and the size of the action space increases exponentially in the number of purification steps. Therefore, a state aggregation and action elimination procedure could provide computational advantage solving complex industry problems. Additionally, we define a stochastic ordering scheme for the pooling windows $w_t \in \mathcal{W}$ at step $t \in \mathcal{T}$ based on the quality and yield trade-offs associated with each chromatography step t.

4.5.1 State Aggregation

The state aggregation scheme groups certain nonempty subset of the original system states into a single aggregate state. We first define an aggregate state called as the failure state d_t at the chromatography step $t \in \mathcal{T}$, and characterize the aggregation scheme for the failure state d_t in Proposition 4.5.1.

Proposition 4.5.1. *All batch states $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ that are an element of the failure zone $(p_t, i_t) \in \mathbb{F}_t$ can be grouped and viewed as a single state called as the failure state d_t.*

Proof See Appendix.

Proposition 4.5.1 indicates that all original system states that are in the failure zone $\mathbb{F}_t$ can be grouped and viewed as a single state, the failure state d_t at the chromatography step $t \in \mathcal{T}$. Hence, in the original problem, the failure state is an

absorbing state with reward $r(d_t) = -c_f$.

Note that the bounds on the value function derived in Section 4.4.3 indicate that the optimal value $\mathcal{V}_t^*(p_t, i_t)$ is constant over the partition $\left\{\mathbb{1}_{p_t \leq p_t', i_t \geq i_t'} | (p_t', i_t') \in \mathbb{F}_t\right\}$ of the original state space $\mathcal{P} \times \mathcal{I}$ at each $t \in \mathcal{T} \cup \{T\}$. More specifically, we have $\mathcal{V}_t^*(p_t, i_t) = -c_f$ for all $(p_t, i_t) \in \{\mathbb{1}_{p_t \leq p_t', i_t \geq i_t'} | (p_t', i_t') \in \mathbb{F}_t\}$. Since all protein-impurity pair that satisfy Proposition 4.5.1 are already an element of the failure zone $\mathbb{F}_t$, the aggregation scheme proposed for the purification problem is exact, in the sense that the aggregate state d_t encompasses subsets of the original system states that have equal costs (Bertsekas, 2012).

4.5.2 Action Elimination and Ordering

In this section, we first provide an action elimination procedure in Proposition 4.5.2, which is then used to create a stochastic ordering scheme for the pooling windows $w_t \in \mathcal{W}_t$ at each step $t \in \mathcal{T}$.

Proposition 4.5.2. *Let w_t^i and w_t^j be two distinct pooling windows at the chromatography step $t \in \mathcal{T}$, such that, $F_t(\Theta|w_t^i) \geq_{st} F_t(\Theta|w_t^j)$, $G_t(\Psi|w_t^i) \leq_{st} G_t(\Psi|w_t^j)$, and $(\theta_t^l|w_t^i) < (\theta_t^l|w_t^j)$, $(\theta_t^u|w_t^i) < (\theta_t^u|w_t^j)$, and $(\psi_t^l|w_t^i) > (\psi_t^l|w_t^j)$, $(\psi_t^u|w_t^i) > (\psi_t^u|w_t^j)$. Then,*

(i) $\mathcal{V}_t(p_t\theta_t, \psi_t i_t|w_t^i) < \mathcal{V}_t(p_t\theta_t, \psi_t i_t|w_t^j)$ for $t \in \mathcal{T}$, and for all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$.

(ii) The pooling window w_t^i is said to be strictly dominated by the pooling window w_t^j at step $t \in \mathcal{T}$, such that, $a_t^*(p_t, i_t) \neq w_t^i$ as a direct result of part (i), for $t \in \mathcal{T}$ and for all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$.

Proof See Appendix.

Conditions in Proposition 4.5.2 ensures that the pooling window w_t^i leads to lower amount in protein and higher amount in impurity compared to the pooling window w_t^j, given that both w_t^i and w_t^j have the same starting condition (p_t, i_t) at the chromatography step t. Proposition 4.5.2 indicates that the pooling window w_t^i is expected to result in lower profits than the pooling window w_t^j for all (p_t, i_t) at the chromatography step t, $\mathcal{V}_t(p_t\theta_t, \psi_t i_t | w_t^i) < \mathcal{V}_t(p_t\theta_t, \psi_t i_t | w_t^j)$. As a direct outcome, the pooling window w_t^i can be eliminated from the set of actions $\mathcal{W}_t$ at the chromatography step t. Let $\hat{\mathcal{W}}_t$ denote the set of actions at step $t \in \mathcal{T}$ obtained after executing this action elimination procedure, i.e., $\hat{\mathcal{W}}_t \subset \mathcal{W}_t$ at step $t \in \mathcal{T}$.

Figure 4.4 shows an example of a strictly dominated pooling window using industry data described in more detail in Section 4.8. Consider two pooling windows w^i and w^j with the following characteristics: The window w^i pools the lanes 7 to 11, and its purification capability is $(\bar{\theta} = 0.71, \bar{\psi} = 0.53 | w^i)$ with the bounds $(\theta^l = 0.64, \psi^l = 0.47 | w^i)$ and $(\theta^u = 0.78, \psi^u = 0.58 | w^i)$. The window w^j corresponds to the lanes 5 to 8 with the purification capability $(\bar{\theta} = 0.73, \bar{\psi} = 0.52 | w^j)$, and the bounds $(\theta^l = 0.65, \psi^l = 0.46 | w^j)$ and $(\theta^u = 0.8, \psi^u = 0.51 | w^j)$. Also, we note that there exists a stochastic dominance in the probability distributions of these two pooling windows, i.e., $F_t(\Theta | w^i) \geq_{st} F_t(\Theta | w^j)$, $G_t(\Psi | w^i) \leq_{st} G_t(\Psi | w^j)$. Therefore, the conditions in Proposition 4.5.2 are satisfied, and pooling the lanes 5 to 8 is better off than pooling the lanes 7 to 11. Hence, the pooling window w^i is strictly dominated by w^j for this specific chromatography step.

Next, we provide a stochastic ordering mechanism for the pooling windows $w_t \in \hat{\mathcal{W}}_t$ at each step $t \in \mathcal{T}$. Note that there are N_t possible windows at the chromatography step $t \in \mathcal{T}$. Let the index n denote the position of the window w_t^n in our ordering scheme, i.e., the pooling window w_t^n is the n^{th} pooling window

among N_t possible windows that are stochastically ordered for the step $t \in \mathcal{T}$. Subsequently, Property 1 and Assumption 4.5.1 provide necessary conditions for a stochastic ordering of all pooling windows $w_t^n \in \hat{\mathcal{W}}_t$ at the chromatography step $t \in \mathcal{T}$.

Property 4.5.1 (Property 1.). $(\theta_t^l|w_t^{n-1}) < (\theta_t^l|w_t^n) < (\theta_t^l|w_t^{n+1})$, $(\theta_t^u|w_t^{n-1}) < (\theta_t^u|w_t^n) < (\theta_t^u|w_t^{n+1})$, *and* $(\psi_t^l|w_t^{n-1}) < (\psi_t^l|w_t^n) < (\psi_t^l|w_t^{n+1})$, $(\psi_t^u|w_t^{n-1}) < (\psi_t^u|w_t^n) < (\psi_t^u|w_t^{n+1})$ *for all* $\{w_t^{n-1}, w_t^n, w_t^{n+1}\} \in \hat{\mathcal{W}}_t$, *and* $t \in \mathcal{T}$.

Assumption 4.5.1. $F_t(\Theta|w_t^{n-1}) \geq_{st} F_t(\Theta|w_t^n) \geq_{st} F_t(\Theta|w_t^{n+1})$ *and* $G_t(\Psi|w_t^{n-1}) \geq_{st} G_t(\Psi|w_t^n) \geq_{st} G_t(\Psi|w_t^{n+1})$ *for all* $\{w_t^{n-1}, w_t^n, w_t^{n+1}\} \in \hat{\mathcal{W}}_t$, *and* $t \in \mathcal{T}$.

Property 1 indicates that the bounds $({\theta^l}_t, \psi_t^l|w_t^n)$ and $({\theta^u}_t, \psi_t^u|w_t^n)$ of the pooling window w_t^n increase in the index n, for each chromatography step $t \in \mathcal{T}$. Characteristics in Property 1 are typically found in the chemical engineering literature to analyze the pooling windows and generate fractionation diagrams (Ngiam et al., 2001, 2003). Assumption 4.5.1 ensures that for a given chromatography step $t \in \mathcal{T}$, both the fraction of protein $(\theta_t|w_t^n)$ and the fraction of impurity $(\psi_t|w_t^n)$ associated with n^{th} pooling window w_t^n increases as its rank index n increases. This behavior illustrates the basic yield-quality trade-off encountered in protein purification, as discussed in Section 4.2.

Based on the discussion above, we order the pooling windows $w_t^n \in \mathcal{W}_t$ for each chromatography step $t \in \mathcal{T}$ using the following procedure: The pooling windows $w_t^n \in \mathcal{W}_t$ are first listed in ascending order based on $(\bar{\theta}_t|w_t^n)$. Next, strictly dominated windows are eliminated based on Proposition 4.5.2 to obtain the set of feasible actions $\hat{\mathcal{W}}_t$. Ties are broken randomly. Property 1 and Assumption 4.5.1 imply that the windows w_t^n are listed in ascending order in terms of $(\bar{\psi}_t|w_t^n)$. The same procedure is performed for each chromatography step t independently. This procedure leads to

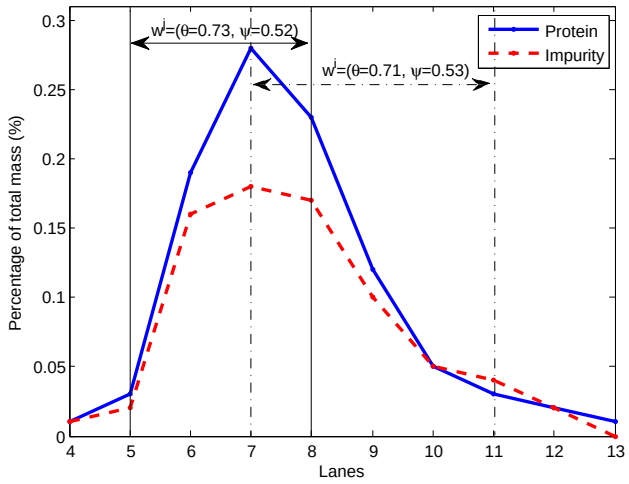

Figure 4.4: Example of a dominated pooling window (based on purification data from Aldevron)

an ascending ordering of the pooling actions in terms of both protein and impurity fractions $(\bar{\theta}_t, \bar{\psi}_t | w_t)$ as well as $(\theta_t^l, \psi_t^l | w_t)$ and $(\theta_t^u, \psi_t^u | w_t)$ associated with all pooling windows $w_t \in \hat{\mathcal{W}}_t$ at each chromatography step $t \in \mathcal{T}$. If the pooling windows satisfy Property 1 and Assumption 4.5.1, we can make the following observations after executing the action elimination and ordering procedure:

(i) a pooling window $w_t^i \in \hat{\mathcal{W}}_t$ that has the least fraction of protein $(\bar{\theta}_t | w_t^i) \leq (\bar{\theta}_t | w_t^j)$, among all the pooling windows $w_t^j \in \hat{\mathcal{W}}_t \setminus \{w_t^i\}$ at the step $t \in \mathcal{T}$, also has the least fraction of impurity $(\bar{\psi}_t | w_t^i) \leq (\bar{\psi}_t | w_t^j)$ at the chromatography step $t \in \mathcal{T}$.

(ii) a pooling window $w_t^j \in \hat{\mathcal{W}}_t$ that has the most fraction of protein $(\bar{\theta}_t | w_t^j) \geq (\bar{\theta}_t | w_t^i)$ among all the pooling windows $w_t^i \in \hat{\mathcal{W}}_t \setminus \{w_t^j\}$, also has the most fraction of impurity $(\bar{\psi}_t | w_t^j) \geq (\bar{\psi}_t | w_t^i)$ at the chromatography step $t \in \mathcal{T}$.

(iii) all pooling windows $w_t^n \in \hat{\mathcal{W}}_t$ can be sorted in ascending order based on their purification capability $(\bar{\theta}_t, \bar{\psi}_t | w_t^n)$, i.e., $(\bar{\theta}_t | w_t^{n-1}) \leq (\bar{\theta}_t | w_t^n)$ and $(\bar{\psi}_t | w_t^{n-1}) \leq (\bar{\psi}_t | w_t^n)$, for all $\{n, n+1\} \in N_t$ at each chromatography step $t \in \mathcal{T}$.

Note that observations $(i) - (iii)$ above are direct consequence of Proposition 4.5.2 and the trade-off between purity and yield associated with the pooling windows.

Next, we fix a pooling window $w_t^n \in \hat{\mathcal{W}}_t$ at chromatography step $t \in \mathcal{T}$. Any pooling window $w_t^j \in \hat{\mathcal{W}}_t$ having larger index than n, $j > n$, is referred to as *larger window*. For example, if the pooling window w_t^j is larger than w_t^n, it has higher protein and impurity fractions then n^{th} window, i.e., $(\theta_t, \psi_t | w_t^j) > (\theta_t, \psi_t | w_t^n)$. Similarly, any pooling window $w_t^j \in \hat{\mathcal{W}}_t$ with index smaller than n, $j < n$, is referred to as *smaller window*, i.e., $(\theta_t, \psi_t | w_t^j) < (\theta_t, \psi_t | w_t^n)$. Hence, the smallest window resulting from this ordering scheme gets the least fraction of protein but also the least fraction of impurity; whereas the largest window collects the highest fraction of both protein and impurity.

4.6 Structural Analysis of the Optimal Policy

In this section, we focus on identifying the structural properties of the optimal policies by exploiting the structural properties of the state space as defined in Section 4.4.

4.6.1 Optimal Policies in the Failure Zone and Risk Zone

Recall that, if the batch state is in the failure zone $(p_t, i_t) \in \mathbb{F}_t$ at step $t \in \mathcal{T}$, then Theorem 4.4.1 indicates that the optimal policy is to stop the purification, $a_t^*(p_t, i_t \mid p_t, i_t \in \mathbb{F}_t) = S$. In this section, we analyze structural properties of the optimal policy when the batch is in the risk zone $\mathbb{R}_t$. To do so, we first define the set of all protein-impurity pairs $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ that meet the final purity requirement specified by the customer, called as the purity set $\mathbb{P} = \left\{(p_t, i_t) : i_t \leq p_t \frac{1-\gamma_d}{\gamma_d}\right\}$. Note that $\mathbb{S} \subset \mathbb{P}$, as the terminal zone $\mathbb{S}$ meets both the purity and yield requirements, whereas $\mathbb{P}$ only meets the purity requirement. Note that, if $(p_t, i_t) \in \mathbb{S}$ or $(p_t, i_t) \in \mathbb{P}$, then the optimal action is to stop at the beginning of the chromatography step $t \in \mathcal{T}$. In Proposition 4.6.1, we first define the *effective purity set* $\mathbb{P}_t^e$ at chromatography step $t \in \mathcal{T}$ that corresponds to all protein-impurity pairs that can lead to the purity

set $\mathbb{P}$ by time T.

Proposition 4.6.1. *The effective purity set at the beginning of the chromatography step t is*

$$\mathbb{P}_T^e = \Big\{(p_T, i_T) : \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T \Big\}, \tag{4.14}$$

$$\mathbb{P}_t^e = \bigcup_{w \in \mathcal{W}_t} K_t^w(\mathbb{P}_{t+1}^e), \tag{4.15}$$

for pooling windows $w \in \mathcal{W}_t$ and $\{t, t+1\} \in \mathcal{T}$, where $K_t^w(\cdot)$ is defined as:

$$K_t^w(y\, p_{t+1} \geq i_{t+1}) : (y\, p_{t+1} \geq i_{t+1}) \rightarrow (y p_t \frac{(\theta_t^u|w)}{(\psi_t^l|w)} \geq i_t) \tag{4.16}$$

for any positive real numbers $y \in R^+$, the pooling window $w \in \mathcal{W}_t$, and $\{t, t+1\} \in \mathcal{T}$.

Proof See Appendix.

Note that if a state (p_t, i_t) belongs to the effective purity set $\mathbb{P}_t^e$ at a step $t \in \mathcal{T}$, then there exists at least one purification policy $\pi_t = \{w_t, w_{t+1}, \ldots, w_{T-1}\}$ that could achieve the desired purity levels specified by the end use or application. Note that Proposition 4.6.1 uses the best-case realizations of the purification capabilities $(\theta_t^l, \psi_t^u|w_t)$. This ensures that the set $\mathbb{P}_t^e$ only includes the states (p_t, i_t) from which the final purity requirement can definitely be achieved by the time T. Proposition 4.6.1 is used to identify the structural characteristics of the optimal purification policy in Theorem 4.6.1.

Theorem 4.6.1. *[Risk Zone] The optimal action $a_t^*(p_t, i_t)$ at state $(p_t, i_t) \in \mathbb{R}_t$ and chromatography step $t \in \mathcal{T}$ has the property that, $a_t^*(p_t, i_t) = \big\{w_t^* \in \mathcal{W}_t : (p_{t+1}, i_{t+1}|\, p_t, i_t, w_t^*) \in \mathbb{P}_{t+1}^e\big\}$ for $\{t, t+1\} \in \mathcal{T}$, and for all $(p_t, i_t) \in \mathbb{R}_t$ with $\gamma_t < \gamma_d$.*

Proof See Appendix.

Theorem 4.6.1 indicates that when the batch state is in the risk zone $\mathbb{R}_t$ at step t, then the optimal policy selects the pooling windows in such away as to keep the batch state (p_{t+1}, i_{t+1}) within the effective purity set $\mathbb{P}^e_{t+1}$ of the next decision epoch $t+1 \in \mathcal{T}$. Theorem 4.6.1 provides guidelines to choose the best candidates for pooling windows at the risk zone. We note that the purification example in Section 4.8 illustrates the lack of threshold-type optimal policies using industry data. However, the guidelines obtained from Theorem 4.6.1 can significantly help the scientist evaluate and understand which pooling windows are good or bad choices for a chromatography step based on financial risks.

4.6.2 Optimal Policies in the Target Zone

We explore the optimal purification policies when the starting state of the batch is in the target zone, $(p_t, i_t) \in \mathbb{T}_t$, at the beginning of the chromatography step $t \in \mathcal{T}$. We break this analysis into two cases: In Case 1, yield shortages are not allowed, i.e., the biomanufacturing firm is committed to meet both purity and yield requirements. However, in Case 2, yield shortages are permitted even though the batch state is in the target zone, i.e., the biomanufacturing firm might meet the purity but not the yield requirement at the expense of incurring a penalty cost. First, we define the problem of reachability of a target set (Bertsekas and Rhodes, 1971), and then use the characteristics of the reachability problem to identify the optimal policies in Case 1 and 2.

Definition 4.6.1. *The target set $\mathbb{T}_T$ is said to be reachable at time T, from state (p_t, i_t) and chromatography step t, if there exists at least one purification policy $\pi_t = (w_t, \cdots w_{T-1})$, such that, the batch state $(p_{t+1}, i_{t+1}) = (\theta_t p_t, \psi_t i_t | w_t)$ is contained in $\mathbb{T}_T$*

at time T for all possible sequence of purification capabilities at the chromatography steps $\{t, t+1, \ldots, T-1\}$.

Definition 4.6.1 indicates that both the yield and purity requirements are said to be reachable from state (p_t, i_t) and step t, only if there exists a purification policy that attains these minimum requirements by time T, despite incurring the worst possible purification capabilities in all chromatography steps. As a direct consequence of Definition 4.6.1 and Proposition 4.4.3, the reachability problem has the following characteristics (Bertsekas and Rhodes, 1971).

Property 4.6.1 (Property 2.)**.** *The target zone $\mathbb{T}_T$ is reachable at time T from all points of the target zone $\mathbb{T}_t$ defined in Proposition 4.4.3, for $t = \{1, \ldots, T-1\}$.*

Property 4.6.2 (Property 3.)**.** *If the target zone $\mathbb{T}_{T-1}$ is reachable at the chromatography step $T-1$, from state $(p_t, i_t) \in \mathbb{T}_t$ and chromatography step $t \in \mathcal{T}$, then the target zone $\mathbb{T}_T$ is reachable at time T from all points of the target zone $(p_t, i_t) \in \mathbb{T}_t$ at the chromatography step $t \in \mathcal{T}$.*

Property 2 implies that the final yield and purity requirements can be attained by time T as long as the batch state at the chromatography step t is an element of the target zone $\mathbb{T}_t$ defined in Proposition 4.4.3. Property 3 states that the reachability problem from the chromatography step t to the end of the planning horizon T, can be reduced to the reachability problem from the chromatography step t to the chromatography step $T-1$. Therefore, if the batch state (p_t, i_t) is in the target zone $\mathbb{T}_t$ at step $t \in \mathcal{T}$, then there is a sequence of actions such that the subsequent states $(p_{t+1}, i_{t+1}), \ldots, (p_T, i_T)$ are always in the target zones $\mathbb{T}_{t+1}, \ldots, \mathbb{T}_T$ regardless of the disturbances in the chromatography steps.

Optimal Policy for Case 1 (Yield shortage not allowed): We first investigate a special case of the problem, where the scientist has to perform chromatography

steps in such a way as to satisfy both yield and purity requirements at the end of the planning horizon T, when the starting state at the chromatography step $t \in \mathcal{T}$ is within its target zone, $(p_t, i_t) \in \mathbb{T}_t$. Then, the problem is equivalent to the problem of reachability of a target set, as stated in Definition 4.6.1. Characteristics of the reachability problem in Properties (2)-(3) are used in deriving the optimal policies for Case 1. Theorem 4.6.2 provides the optimal purification policies when the batch state is in the target zone, $(p_t, i_t) \in \mathbb{T}_t$, at chromatography step $t \in \mathcal{T}$.

Theorem 4.6.2. [Target Zone, Case 1] When $(p_t, i_t) \in \mathbb{T}_t$, the optimal policy has the following characteristics: $a_t^*(p_t, i_t) = \left\{w_t^* \in \mathcal{W}_t : \{(\theta_t^l p_t, \psi_t^u i_{t+1} | w_t^*) \in \mathbb{T}_t \mid (p_t, i_t) \in \mathbb{T}_t\}\right\}$ for $\{t, t+1\} \in \mathcal{T}$, for all $(p_t, i_t) \in \mathbb{T}_t$ with $\gamma_t < \gamma_d$.

Proof See Appendix.

Theorem 4.6.2 provide guidelines to select the optimal pooling window w_t at a chromatography step $t \in \mathcal{T}$. Theorem 4.6.2 indicates that the optimal action at step t will perform the purification in such a way as to stay within the target zone $\mathbb{T}_{t+1}$ of the next chromatography step $t+1 \in \mathcal{T}$, when the starting batch state is in the target zone at the beginning of a chromatography step t, $(p_t, i_t) \in \mathbb{T}_t$. Note that the optimal policy is a direct implication of the characteristics of the target zones in Property (2)-(3). Recursive application of Theorem 4.6.2 to all remaining chromatography steps indicates that, if the batch state is $(p_t, i_t) \in \mathbb{T}_t$, then the optimal policy is to select the pooling windows in a way as to ensure that the subsequent states $(p_{t+1}, i_{t+1}), \ldots, (p_{T-1}, i_{T-1})$ lie within their respective target zones $\mathbb{T}_{t+1}, \ldots, \mathbb{T}_{T-1}$ in all remaining chromatography steps $t+1, \ldots, T-1$. Note that, if the initial state of the batch is in its target zone, then the definition of the target zones in Proposition 4.4.3 ensures that there exists at least one optimal policy that satisfies Theorem 4.6.2. Also note that the optimal policy is to stop when the state

(p_t, i_t) is an element of the terminal zone, $(p_t, i_t) \in \mathbb{S}$, at a chromatography step $t \in \mathcal{T}$.

Optimal Policies for Case 2 (Yield shortage allowed): We define a new reachability problem by allowing yield shortage, i.e., $p_T \leq p_d$, despite the batch state (p_t, i_t) being in the target zone $\mathbb{T}_t$ at a chromatography step $t \in \mathcal{T}$. The optimal policies in Case 2 consider the trade-offs between operating costs and shortage costs. Although compromising on yield might not be ideal, especially when it is know that the initial batch (p_1, i_1) is in the target zone $\mathbb{T}_1$; allowing shortage costs could help the biomanufacturing firm reduce the number of purification steps by compromising on yield. In some cases, this could increase their profitability despite shortage costs. To analyze the optimal policies in Case 2, we relax the yield requirement from Case 1. Then, the structural analysis becomes similar to Section 4.6.1, except that, Theorem 4.6.3 provides optimal policies with guaranteed performance for achieving the purity requirement despite allowing shortage costs.

Theorem 4.6.3. [Target Zone, Case 2] For all $(p_t, i_t) \in \mathbb{T}_t$, such that, $\gamma_t < \gamma_d$ at the chromatography step $t \in \mathcal{T}$, the optimal policy has the characteristic that $a_t^*(p_t, i_t) = \left\{ w_t^* \in \mathcal{W}_t : (p_{t+1}, i_{t+1} |\, p_t, i_t, w_t) \in \mathbb{P}_{t+1}^e \right\}$, where the effective purity set $\mathbb{P}_t^e$ is defined in Proposition 4.6.1, and uses the operator $L_t^w(y\, p_{t+1} \geq i_{t+1}) : (y\, p_{t+1} \geq i_{t+1}) \rightarrow (y p_t \frac{(\theta_t^l | w)}{(\psi_t^u | w)} \geq i_t)$ for all $t \in \mathcal{T}$, and $y \in R^+$.

Proof See Appendix.

In Theorem 4.6.3, the optimal policy takes into consideration the trade-off between shortage costs and operating costs. Note that Theorem 4.6.3 provides guaranteed performance on achieving the final purity requirement, since the operator L_t^w in Theorem 4.6.3 takes into consideration the worst-case realization of the purification capabilities. In Case 2, although the biomanufacturing firm has the capability for

meeting both the purity and yield requirements by the time T, the optimal policy can choose to compromise on yield in order to reduce the number of purification steps and operating costs. In practice, in order to maintain good long-term relationships with the customers, the biomanufacturing firm might decide to meet both the demand and purity requirements whenever they have enough purification capabilities – even if this decision might not be the best decision that increases profitability of a particular order. In such cases, the decision maker will proceed with the optimal policy suggested in Theorem 4.6.2, instead of Theorem 4.6.3.

4.6.3 Worst-Case Guarantees versus Probabilistic Approach

Biomanufacturing firms often adopt a worst-case approach while making operating decisions in order to hedge against uncertainties and reduce the risk of incurring large penalty costs. It was for this reason that the analysis in Section 4.4 and Section 4.6 are conducted using worst-case purification capabilities associated with the pooling windows. Alternatively, we can generalize this approach using a probabilistic approach based on the expected purification capabilities and costs. In this section, we re-interpret the zones identified in Section 4.4 and policies in Section 4.6 in the context of the worst-case and probabilistic approaches. Subsequently, we use industry data in Section 4.8 to investigate how different would the insights and policies be different, if the biomanufacturing firm adopts the worst case approach compared to the probabilistic one.

Let $\mathcal{V}_t^w(p_t, i_t)$ and $\mathcal{V}_t^p(p_t, i_t)$ be the value function at step $t \in \mathcal{T}$ and state $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ under the worst case and the probabilistic approach, respectively. The value function $\mathcal{V}_t^w(p_t, i_t)$ is obtained by solving the set of Equations (4.4)-(4.5) using the worst outcomes of the purification capability $(\theta_t^l, \psi_t^u | w_t)$ at each chromatography step $t \in \mathcal{T}$ and pooling window $w_t \in \mathcal{W}_t$; whereas the probabilistic approach

obtains the value function $\mathcal{V}_t^p(p_t, i_t)$ solving the set of Equations (4.4)-(4.5) based on the underlying probability distributions $f_t(\cdot)$ and $g_t(\cdot)$, as discussed in Section 5.2.

Similarly, we let $\mathbb{F}_t^w$ and $\mathbb{T}_t^w$ be the failure zone and target zone under the worst-case approach, respectively. $\mathbb{F}_t^p$ and $\mathbb{T}_t^p$ are the failure and target zone under the probabilistic approach, respectively. Theorem 4.4.1 and the bounds on the value function derived in Section 4.4.3 indicates that $\mathbb{F}_t^w$ corresponds to all (p_t, i_t) with the worst-case value function $\mathcal{V}_t^w(p_t, i_t) = -c_f$. Similarly, Proposition 4.4.2 uses the worst-case purification capabilities to identify $\mathbb{F}_t^w$. On the other hand, using Theorem 4.4.1 and the bounds on the value function derived in Section 4.4.3, we observe that $\mathbb{F}_t^p$ represents all states (p_t, i_t) having an expect value $\mathcal{V}_t^p(p_t, i_t) = -c_f$. Similarly, $\mathbb{F}_t^p$ can be characterized by replacing (θ_t^l, ψ_t^u) with the expected purification capabilities $(\bar{\theta}_t, \bar{\psi}_t)$ for all $t \in \mathcal{T}$ in Proposition 4.4.2. Proposition 4.4.3 defines the target zone under the worst-case approach $\mathbb{T}_t^w$. Based on the bounds derived in Section 4.4.3, the target zone $\mathbb{T}_t^p$ under the probabilistic approach corresponds to all (p_t, i_t) whose expected value lies within $\sum_{j=t}^{T-1} -c_j + r(p_d) \leq \mathcal{V}_t^*(p_t, i_t) \leq r(p_d)$ at each step t. No doubt, the target zone under the probabilistic approach $\mathbb{T}_t^p$ can be characterized using the expected purification outcomes $(\bar{\theta}_t, \bar{\psi}_t)$ instead of the worst case performance (θ_t^l, ψ_t^u) in Proposition 4.4.3; however this would not be useful in providing guaranteed performance for achieving the specific requirements on purity and yield. It is easy to observe that $\mathbb{F}_t^p \subset \mathbb{F}_t^w$, and $\mathbb{T}_t^w \subset \mathbb{T}_t^p$ at each chromatography step $t \in \mathcal{T}$. The guidelines on the optimal policies stated in Theorems 4.4.1-4.6.3 would use the probabilistic zones $\mathbb{F}_t^p$ and $\mathbb{T}_t^p$ for the probabilistic approach, and $\mathbb{F}_t^w$ and $\mathbb{T}_t^w$ for the worst-case analysis.

In the probabilistic approach, note that the aggregation scheme in Proposition 4.5.1 is exact when the failure zone $\mathbb{F}_t^p$ considered in the aggregation scheme is

defined based on the best purification outcomes $(\theta_t^u, \psi_t^l|w_t)$ for each step $t \in \mathcal{T}$ and pooling windows $w_t \in \mathcal{W}_t$. Using the best purification outcomes in the aggregation scheme for the probabilistic approach ensures that we capture sufficient conditions under which the firm has no financial incentives for performing the purification.

4.7 Implementation at Aldevron

In this section, we elaborate on the project timeline, implementation challenges and results obtained at Aldevron.

4.7.1 Timeline

The optimization framework has been constructed, revised, validated, and implemented over a two-year period (2013-2015) through continuous interaction with Aldevron's protein purification team and senior management. Our research collaboration with Aldevron started in February 2013. Through weekly company visits, we observed operational challenges that are typical to the biomanufacturing operations, collected data, validated our models, carried out the implementation, and quantified the savings. The purification optimization model was built during August 2013-February 2014. Data collection and revisions were performed during February-June 2014. Results obtained from the mathematical model were validated during June-September 2014 by various test runs comparing the current practice with the optimal policies. Insights obtained at Aldevron were shared with a broader biomanufacturing community through series of working group sessions (BioWGS, 2014; BioForward, 2014), followed by the actual implementation and use of the model in Aldevron's daily operations since October 2014.

4.7.2 Implementation Results

Two years into collaboration, the optimization model has been currently in use for all R&D protein purification orders. Since the implementation of the optimization framework, Aldevron has realized lead time and cost reductions. On average, the implementation has led to 25% reduction in total lead times and 20% reduction in operating costs involved in R&D protein purification. These lead time and cost savings were mainly due to the following three factors:

1. *Scaling back the validation runs.* Using the information obtained from the scouting experiments to run the optimization model has allowed to scale back the number of validation runs needed prior to full scale production. For the majority of purification projects, the scientists were able to take the process information obtained from scouting runs, and then feed this information directly into the optimization model. In minor instances, the scouting experiments indicates some potential issues with variability and stability of the proteins. In such cases, the scientists kept performing the validation runs to gain further data and process understanding.

2. *Formal assessment of the risks and better understanding of manufacturing capabilities.* The optimization model provides a rigorous and formal assessment of the business risks at the beginning of each purification run. This information is especially critical in communicating the manufacturing challenges with the customer. For example, one of the major challenges in purification operations is the variability in the starting material. Without formal assessment of the manufacturing capabilities and risks, it is very difficult to predict and react to the challenges in attaining the production requirements. The optimization framework provides an improved understanding of the business risks and financial trade-offs involved in protein purification operations. The proposed zone-based decision making approach provides a quick and reliable analysis of the manufacturing capabilities leading to better and easier

communication with the clients. The knowledge on "guaranteed performance" or "guaranteed failure" obtained by the end of scouting runs has been invaluable for both the clients and the biomanufacturing company.

3. *Process economics taken into consideration.* Prior to the use of the optimization framework, potential operating policies were assessed based on historical experience. Given the combinatorial nature of the pooling strategies, it was inevitable for the scientist to take shortcuts to avoid getting overwhelmed with the number of available pooling choices at each step. As a result, the scientists often used to focus on meeting the purity requirement, and did not consider the overall financial implications while making pooling decisions. In contrast, the optimization model provides a formal framework that captures the uncertainties in purification outcomes, financial trade-offs, and the limitations in manufacturing capabilities. As a result, the purification policies suggested by the optimization model are based on the process economics as well as chemical characteristics (i.e., scouting data), and hence has led to lower costs and shorter lead times.

Cost and lead time reductions were determined in two phases: *1. Validation phase (June-September 2014):* During the Summer 2014, we collected scouting data for all engineered purification orders, and then identified the decision zones and optimal operating polices based on this information. However, the optimal policies and decision zones were generated only for validation purposes, and were not implemented in daily practice. In this phase, the scientists kept performing the purification operations based on their expertise. For validation purposes, the policies proposed by the optimization model were compared against the ones adopted by the scientists. This information was used to quantify potential savings (costs and lead times) that could have been achieved if the optimal policies were used instead of the current practice. *2. Implementation phase (since October 2014):* Once the optimization framework was

implemented, savings obtained as a result of the framework were quantified through a policy evaluation mechanism. For each purification project, we collected information about the operating policy that the scientist would have used if the optimization model was not implemented. Then, we used this information to evaluate the performance of that specific policy associated with that specific order (i.e., evaluate the value function for a given policy), and then compared it against the performance of the optimal policy. Since protein purification operations require high costs, long lead times and limited resources, it was not possible to conduct both the optimal policies and other business practices simultaneously in the laboratory for the purpose of quantifying the savings.

4.7.3 Implementation Challenges

Feedback from the broader biomanufacturing community beyond Aldevron has been a core part of the problem definition, analysis and validation. For example, we organized a series of working group sessions with the local biomanufacturing firms during various phases of this research (BioWGS, 2014). The objectives were to understand problem characteristics, validate assumptions, define managerial questions and identify relevant optimization techniques. Our models and insights have also been shared with a larger biomanufacturing community (BioForward, 2014). Application of operations research tools to solve these problems are new to the industry, and the response has been more of cautious enthusiasm. This is mainly due to the fact that biomanufacturing processes are highly regulated, and changing their current practice impacts the regulatory approval process. Feedback from the community is that as more companies embrace the application of operations research models to optimize operations, both biomanufacturing firms and regulatory authorities are likely to view such approaches as being essential for reducing costs and lead times in the research and development. However, operations research implementations at Aldevron have

already started to gain an important visibility in the Wisconsin's bioscience community through BioForward and the Wisconsin Economic Development Corporation (BioForward, 2014; WEDC, 2014).

Other implementation challenges were related to formatting the data required to run the optimization model. The dynamic programming algorithm used in the optimization framework is coded using Java and includes a user-friendly interface. Initially, the data obtained from scouting runs were in the format of gel pictures as shown in Figure 4.2. A special biomanufacturing image processing software was used to convert these gel images into the protein and impurity amounts corresponding to each lane. This information was stored in a table format at MS Excel, and then fed into the Java code to run the optimization model. Although the resulting data was reliable, the overall process of converting the gel images into a data format compatible with our optimization tool was laborious. To overcome this challenge, the scientists adopted another alternative data collection technique that eliminates the process of reading the gel images. This alternative technique allows to directly measure the protein and impurity amounts instead of reading this information from gel images. Although this alternative data collection technique is cheaper and faster, and the only reason why it was not used before is due to the typical industry practices and culture. Special training sessions were conducted to get the buy-in of all purification scientists and also help them in getting familiar with the optimization framework. Overall, the protein purification team has been very satisfied with the way how the tool helped their decisions.

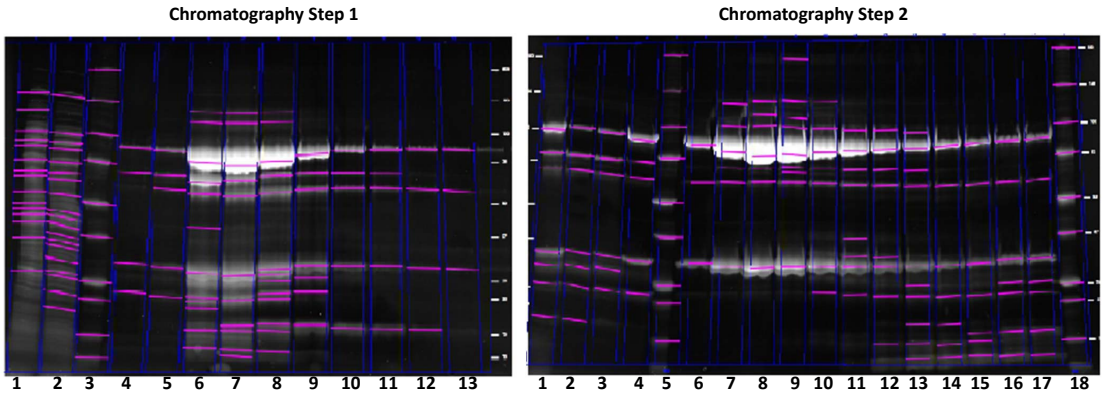

Figure 4.5: Two step chromatography outputs

4.8 An Illustrative Case Study

Since each purification order is custom-engineered and unique, each order has its own operating policies and managerial insights specific to that order. Therefore, we believe that it would not be useful to explain the optimal policies and managerial insights for every single protein considered in the implementation process at Aldevron. Instead, we elaborate on one of the custom-engineered purification orders at Aldevron (Section 4.8.1) and explain the way how the optimization framework was implemented for that actual order. More specifically, we demonstrate the decision zones, identify the optimal policies, and discuss the managerial insights for an actual purification order (Sections 4.8.2-4.8.3). Further, we compare the difference between worst-case and probabilistic approaches in terms of the decision zones and optimal policies (Section 4.8.4), and quantify the computational savings due to action elimination and state aggregation (Section 4.8.5). To protect client confidentiality, actual data and cost information obtained from Aldevron are masked. However, the parameters and assumptions used in this section are typical and valid across the biomanufacturing industry.

4.8.1 Problem Setting and Parameters

The protein of interest considered involved in the implementations are all engineered proteins used for *in vitro* studies in biomanufacturing. In this section, we consider a protein purification problem with two chromatography steps, as shown in Figure 4.5. The first step uses the binding affinities between proteins and metal ions as a separation principle, and the second step uses separation based on electric charge. Figure 4.5 shows that the first step has 10 candidate lanes (starting from lane 4 to 13) leading to 55 candidate pooling windows. The second step has candidate 12 lanes (from lane 6 to 17) leading to 78 candidate pooling windows. Main characteristics of the pooling windows can be summarized as follows. In the first chromatography step, the smallest pooling window w_1^1 has the purification capability of $0.009 \leq (\theta_1|w_1^1) \leq 0.011$ and $0.003 \leq (\psi_1|w_1^1) \leq 0.0036$, and the largest pooling window w_1^{55} has $0.878 \leq (\theta_1|w_1^{55}) \leq 1$ and $0.67 \leq (\psi_1|w_1^{55}) \leq 0.819$. In the second chromatography step, the smallest pooling window w_2^1 corresponds to $0.042 \leq (\theta_2|w_2^1) \leq 0.051$ and $0.003 \leq (\psi_2|w_2^1) \leq 0.004$, and the largest pooling windows w_2^{78} has $0.856 \leq (\theta_2|w_2^{78}) \leq 1$ and $0.6 \leq (\psi_2|w_2^{78}) \leq 0.741$. In both of the chromatography steps t, the purification capabilities are uniformly distributed within 10% of their mean $(\bar{\theta}_t, \bar{\psi}_t)$ for $t = \{1, 2\}$. All pooling windows in both of the chromatography steps satisfy data characteristics and assumptions in Section 4.5.2.

The yield requirement is 8 milligram (mg) of protein with a purity level equal or greater than 85%. The actual cost information obtained from Aldevron is masked for confidentiality purposes, and representative values are used instead. The operating costs of each chromatography step is $c_t = \$15$ for $t = 1, 2$. These include costs associated with labor, materials, equipment, inspection and analytics. Rewards obtained per mg of protein is $r = \$5$, and the shortage cost per mg of protein is $c_l = \$6$.

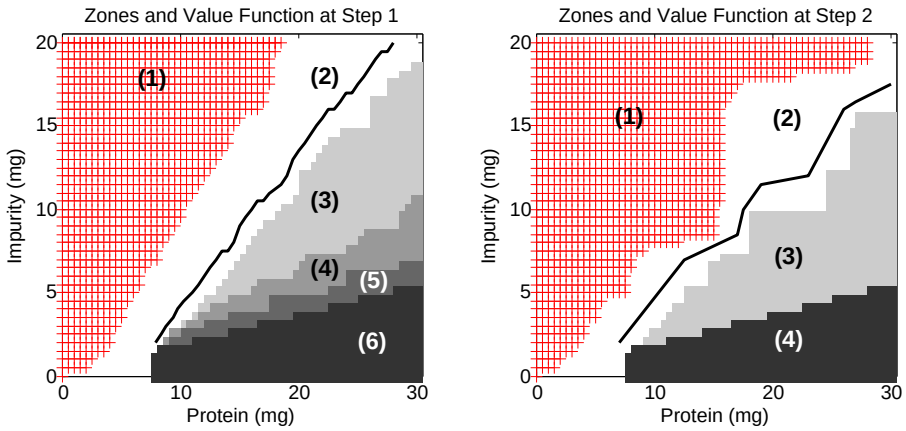

Figure 4.6: Optimal value and decision zones for the first and second step

The penalty cost of failure is $c_f = \$48$, which is equivalent to the maximum possible shortage cost considered in our purification setting.

4.8.2 Decision Zones and Their Financial Implications

We investigate financial implications of a batch condition obtained from fermentation. For this purpose, we analyze the structural properties of the optimal value function, and characterize the failure, risk and target zones for each chromatography step. Figure 4.6 presents the zones and the optimal value function at each chromatography step. The value function associated with the zones in Figure 4.6 are as follows: For the first step, (1) $\mathbb{E}\mathcal{V}_1(p_1, i_1) = -48$, (2) $-48 < \mathbb{E}\mathcal{V}_1(p_1, i_1) < 10$, (3-6) $10 \leq \mathbb{E}\mathcal{V}_1(p_1, i_1) \leq 40$, and the solid line for $\mathbb{E}\mathcal{V}_1(p_1, i_1) = 0$. Optimal value for second step: (1) $\mathbb{E}\mathcal{V}_2(p_2, i_2) = -48$, (2) $-48 < \mathbb{E}\mathcal{V}_2(p_2, i_2) < 25$, (3-4) $25 \leq \mathbb{E}\mathcal{V}_2(p_2, i_2) \leq 40$, and the solid line for $\mathbb{E}\mathcal{V}_2(p_2, i_2) = 0$. Managerial insights derived from Figure 4.6 are discussed below, and summarized in Table 4.2.

In Figure 4.6, the region (1) in both of the chromatography steps corresponds to the failure zone with $\mathbb{E}\mathcal{V}_t(p_t, i_t|(p_t, i_t) \in \mathbb{F}_t) = -48$ for $t = \{1, 2\}$. This region represents protein-impurity states $(p_t, i_t) \in \mathbb{F}_t$, $t = \{1, 2\}$, where the biomanufacturing

firm is better off with abandoning the purification. As expected, the failure zone in the second chromatography step is observed to be larger than the one in the first chromatography step.

Region (2) in Figure 4.6 represents the risk zone $\mathbb{R}_t$ with $-48 < \mathbb{E}\mathcal{V}_1(p_1, i_1) < 10$ for the first step, and $-48 < \mathbb{E}\mathcal{V}_2(p_2, i_2) < 25$ for the second step. If the protein-impurity state is an element of the risk zone, the biomanufacturing firm is expected to incur financial losses due to combined impact of shortage costs and operating costs, even if the final batch meets the purity requirement.

Regions $(3-6)$ at the first chromatography step represent the target zone $\mathbb{T}_t$ where the firm is capable of meeting both the demand and purity requirements of the end use or application. For example, if the starting batch (p_1, i_1) is in the region (3), the firm can expect to achieve the final yield and purity requirements through two chromatography steps using the optimal policies, and yielding $\mathbb{E}\mathcal{V}_1(p_1, i_1) = 10$. However, the optimal policy in region (4) suggests that the firm might be better off compromising on yield to achieve the final purity requirement at the end of the first step, and incurring some penalty cost associated with shortage costs. In this case, the operating costs for the second chromatography step are greater than the penalty cost of shortage incurred. Note that although both the demand and purity requirements could be met in the region (4), it is financially better to abandon the purification at the end of the first step, and incur the shortage penalties, i.e., $10 < \mathbb{E}\mathcal{V}_1(p_1, i_1) < 25$. In practice, intangible costs associated with loss of goodwill may motivate the firm to choose pooling windows that keep the batch state within the target zone of the next step (Theorem 4.6.2), with $\mathbb{E}\mathcal{V}_1(p_1, i_1) = 10$.

Table 4.2: Summary of the insights based on Figure 4.6

	Region	Range of $\mathbb{E}\mathcal{V}_t(p_t, i_t)$	Business Implications
Step 1	(1)	$\mathbb{E}\mathcal{V}_1(p_1, i_1) = -48$	Stop and scrap the batch.
	(2)	$-48 < \mathbb{E}\mathcal{V}_1(p_1, i_1) < 10$	Risk zone with high potential losses. Can meet the purity, but will incur high operating and shortages.
	(3)	$\mathbb{E}\mathcal{V}_1(p_1, i_1) = 10$	Can meet both purity and yield requirements in two steps.
	(4)	$10 < \mathbb{E}\mathcal{V}_1(p_1, i_1) < 25$	Can meet both purity and yield requirements in two steps. However, financially better off with single step, despite shortages.
	(5)	$\mathbb{E}\mathcal{V}_1(p_1, i_1) = 25$	Can meet both purity and yield requirements in one step.
	(6)	$\mathbb{E}\mathcal{V}_1(p_1, i_1) = 40$	Stop. Desired terminal state.
Step 2	(1)	$\mathbb{E}\mathcal{V}_2(p_2, i_2) = -48$	Stop and scrap the batch.
	(2)	$-48 < \mathbb{E}\mathcal{V}_2(p_2, i_2) < 25$	Risk zone with high potential losses. Can meet the purity, but will incur high operating and shortage costs.
	(3)	$\mathbb{E}\mathcal{V}_2(p_2, i_2) = 25$	Can meet both purity and yield requirements in one step.
	(4)	$\mathbb{E}\mathcal{V}_2(p_2, i_2) = 40$	Stop. Desired terminal state.

With a batch starting in the region (5), the firm can expect to achieve the final purity and yield requirements at the end of the first chromatography step, with $\mathbb{E}\mathcal{V}_1(p_1, i_1) = 25$. In this case, a second step chromatography is not required. Region (6) in the first step and region (4) in the second step represent all protein-impurity pairs meeting the specific requirements on purity and yield.

The solid line in Figure 4.6 for both of the chromatography steps corresponds to all state impurity pairs (p_t, i_t) having $\mathbb{E}\mathcal{V}_t(p_t, i_t) = 0$ for $t = 1, 2$. The states to the left of the solid line correspond to a region where the firm should expect financial losses due to combined impact of shortage costs and operating costs, even if the final batch met the purity requirement. Due to the monotonicity of the value function (Proposition 4.4.1), the profit is nondecreasing in protein amounts p_t for a given impurity level i_t. Hence, the solid line in the first chromatography step has an important managerial implication: if the starting condition of a batch is on the left hand side of the solid line, then the firm might prefer to scrap the starting material, re-work in-house or request the provider to send a new starting material.

4.8.3 Optimal Policies and Comparison with Practice

Next, we compare the purification policy used in practice with the optimal policy. We analyze optimal policies for batches with starting state $p_1 \in [25, 30]$ mg, and $i_1 \in [15, 20]$ mg, and compare with policy adopted at Aldevron for a batch having $(p_1, i_1) \in (27.5, 17.5)$. From Figure 4.6, we observe that this starting range $p_1 \in [25, 30]$ mg and $i_1 \in [15, 20]$ mg encompasses the risk zone, the solid line with no profit, and the target zone. Note that the starting state $(27.5, 17.5)$ is in the risk zone for the first chromatography step. Therefore, the firm does not have any performance guarantees for this purification order in terms of yield and purity. We quantify and compare the risks and costs associated with this starting state.

State-dependent optimal policies: Let π^* denote the optimal policy, and $\mathcal{V}_1(p_1, i_1|\pi^*)$ represent the optimal value function. Table 4.3 presents a snapshot of the optimal policy for the first chromatography step. Table 4.3 only displays the optimal policies at selected states (i.e., in the intervals of 2.5 mg) in order to improve the readability. In Table 4.3, $\{S\}$ represents the stopping action. For other actions in Table 4.3, we present the staring lane, end lane, and the corresponding action index. For example, $L6$-8 (21) means pooling from lane 6 to lane 8 (including lane 8), and this action the 21^{st} feasible action out of 36 pooling windows at the first chromatography step. Cells colored in gray represent the target zone based on Figure 4.6, and the entries in bold correspond to the failure zone.

We make the following observations regarding the optimal policy. First, we observe that the optimal action is to stop the purification for states in the failure zone $\mathbb{F}_1$ (i.e., top left of Table 4.3), and for the terminal zone $\mathbb{S}$ (i.e, bottom right of Table 4.3). Second, we see that in the target zone $\mathbb{T}_1$, the optimal policies do not have a threshold-type structure, but they do satisfy the characteristics in

Table 4.3: A snapshot of the optimal pooling policies for selected states (first chromatography step)

Protein (mg)	10	12.5	15	17.5	20	22.5	25	27.5	30
Impurity (mg)									
20	**S**	**S**	**S**	**S**	L4–13 (36)	L7–8 (17)	L7–8 (17)	L6–8 (21)	L7–9 (19)
17.5	**S**	**S**	**S**	L7–8 (17)	L4–10 (29)	L6–9 (24)	L7–8 (17)	L6–8 (21)	L7–10 (20)
15	**S**	**S**	L7–9 (19)	L7–9 (19)	L4–10 (29)	L4–9 (26)	L6–9 (24)	L4–8 (23)	L6–12 (30)
12.5	**S**	**S**	L4–8 (23)	L4–8 (23)	L4–13 (36)	L5–10 (28)	L4–9 (26)	L4–10 (29)	L5–10 (28)
10	**S**	L5–10 (28)	L6–10 (27)	L4–9 (26)	L6–9 (24)	L5–8 (22)	L7–8 (17)	L7–10 (20)	L7 (13)
7.5	L6–8 (21)	L4–9 (26)	L4–9 (26)	L4–9 (26)	L5–10 (28)	L5–8 (22)	L7–9 (19)	L7–8 (17)	L7–8 (17)
5	L4–10 (29)	L5–10 (28)	L5–12 (33)	L7–8 (17)	L6–7 (16)	L6–8 (21)	L7–8 (17)	L7–10 (20)	S
2.5	L4–9 (26)	L4–8 (23)	S	S	S	S	S	S	S
0	S	S	S	S	S	S	S	S	S

Theorem 4.6.2-4.6.3. Third, in the risk zone $\mathbb{R}_1$, we observe a non-increasing trend in the action index as the impurity amount decreases for a given protein amount. For example, for $p_1 = 17.5$, the optimal policy chooses actions of higher order as the impurity amount decreases. However, this monotonic trend is not present for all protein-impurity pairs. For example, for $p_1 = 20$ mg and $i_1 = 20$ mg, the optimal policy suggests to pool lanes $4 - 13$ with the action index 36. When $i_1 = \{15, 17.5\}$, it adopts a smaller window (lanes $4 - 10$ with the action index 29). However, at $i_1 = 12.5$, it switches back to lanes $4 - 13$ with the action index 36. Such deviations in threshold type policies are also observed in the second step. We also observe the lack of threshold type policies as the protein amount increases for a fixed impurity.

For our starting state $(p_1, i_1) = (27.5, 17.5)$, the optimal policy suggests to pool lanes $6 - 8$ in the first chromatography run, and lanes $6 - 9$ in the second chromatography run, $\pi^* = \{\text{Lanes } 6 - 8, \text{Lanes } 6 - 9\}$ with $\mathcal{V}_1(27.5, 17.5|\pi^*) = \$9$.

Comparison with current practice: Based on the scouting and validation experiments, Aldevron decided to pool lanes $6 - 9$ in the first chromatography step, and lanes $7 - 10$ in the second chromatography step. Characteristics of these pooling windows are as follows: $0.747 \leq \Theta_1 \leq 0.913$, $0.545 \leq \Psi_1 \leq 0.666$, and $0.524 \leq \Theta_2 \leq 0.640$, $0.204 \leq \Psi_2 \leq 0.250$. We let $\pi_1'' = \{\text{Lanes } 6 - 9, \text{Lanes } 7 - 10\}$ denote the

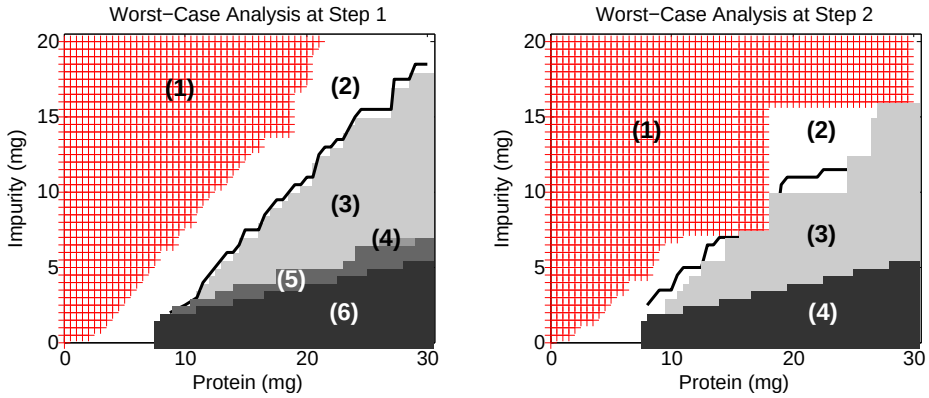

Figure 4.7: Zones for two chromatography steps under the worst-case analysis

pooling policy used in practice, with the value function $\mathcal{V}_1(p_1, i_1|\pi_1'')$. As a result of the policy π_1'', 13.314 mg of protein and 2.212 mg of impurity were obtained at the end of the production run. Therefore, the yield and purity requirements specified by the end use or application (8 mg of protein with $\geq 85\%$ purity) were satisfied for this order. However, when we compare the realization of purification capabilities during the production run against the supports derived at scouting runs, we observe that the realizations were in favor of the biomanufacturing firm for that specific production run (i.e, closer to the mean, with realizations $\theta_1 = 0.832$, $\psi_1 = 0.602$ and $\theta_2 = 0.582$, $\psi_2 = 0.21$). Therefore, we evaluate the performance of the policy π_1'' even though yield and purity requirements were satisfied in our example run. We observe that the value function associated with the current practice is $\mathcal{V}_1(27.5, 17.5|\pi_1'') = \$7.2$, whereas the value function of the optimal policy is $\mathcal{V}_1(27.5, 17.5|\pi_1^*) = \$9$. Therefore, for the stating state $(27.5, 17.5)$, we observe that 25% improvement in the expected profit is achieved through optimization.

4.8.4 Comparison between Worst-Case and Probabilistic Approach

In this section, we evaluate the difference between pessimistic (worst-case) and probabilistic approaches for the two step purification process shown in Figure 4.5. The

Table 4.4: Comparison between the worst-case and probabilistic analysis

Evaluation criteria	**Step 1**	**Step 2**
Difference in zones	**% overlap**	**% overlap**
Failure	85.5%	81.5%
Risk	69.0%	60.1%
Target	94.3%	99.6%
% Difference in value function	**% of total states**	**% of total states**
No difference	65.5%	80.0%
$> 0\%$ and $< 1\%$	24.4%	20.0%
$> \%1$	9.6%	0%
$> \%5$	0.5%	0%
Difference in policies	**% of total states**	**% of total states**
Same window	54%	74%
Smaller window in worst-case	39%	26%
Larger window in worst-case	7%	0%

zones and their value function using worst-case analysis are presented in Figure 4.7. Table 4.4 compares and quantifies the difference between the worst-case (Figure 4.7) and probabilistic approaches (Figure 4.6) in terms of zones, optimal value function and optimal purification policies.

Difference in zones: Figure 4.7 presents the failure zone $\mathbb{F}_t^w$(Region 1) and target zone $\mathbb{T}_t^w$ (Regions 3 to 6) under the worst-case approach for the chromatography steps $t = 1, 2$. When we compare Figure 4.6 with Figure 4.7, we observe that the failure zone under the worst-case approach $\mathbb{F}_t^w$ is larger than the failure zone in the probabilistic case $\mathbb{F}_t^p$, $t = 1, 2$. This is expected since the worst-case analysis takes into consideration the worst possible realization of purification capabilities. However, we observe high degree of overlap between $\mathbb{F}_t^w$ and $\mathbb{F}_t^p$ in this purification project. For example, Table 4.4 shows that 85.5% of the states that belong to $\mathbb{F}_1^w$ also belong to $\mathbb{F}_1^p$; whereas, the degree of overlap between these two zones reduces to 81.5% for the second step purification. We observe that the overlap percentage for the zones is lower in step 2, because the second step is the last chromatography step, and hence

the worst-case approach tends to minimize the failure risks and total costs (operating and failure costs) using a pessimistic scenario for purification capabilities.

Table 4.4 shows that 94.3% of the states in the probabilistic target zone $\mathbb{T}_1^p$ also belong to the worst-case target zone $\mathbb{T}_1^w$ for the first step. The overlap between $\mathbb{T}_2^p$ and $\mathbb{T}_2^w$ increases 99.6% in the second step. Hence, $\mathbb{T}_t^p$ of this purification project is not very sensitive to uncertainty in the purification capabilities for $t = 1, 2$. This behavior could be attributed to several factors, such as, low standard deviation associated with the purification capabilities, and high separation capabilities of both chromatography steps. Therefore, due to high degree of overlap, either of the probabilistic and worst case analysis could be used to characterize the target zone in this purification order.

Difference in the value function: Next, we evaluate the difference in the optimal value function between the worst-case and probabilistic approaches. Let $\mathcal{V}_t^p(p_t, i_t)$ and $\mathcal{V}_t^w(p_t, i_t)$ represent the value function under the probabilistic approach and worst-case approach, respectively, for $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ and $t = \{1, 2\}$. Table 4.4 presents the percentage difference between these two value functions, i.e., $\frac{\mathcal{V}_t^p(p_t,i_t) - \mathcal{V}_t^w(p_t,i_t)}{\mathcal{V}_t^w(p_t,i_t)}$ for $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ and $t = \{1, 2\}$. Table 4.4 indicates that, in 65.5% of all states, there is no difference in the optimal value functions between these two approaches. This observation would be attributed to the high degree of overlap between the target and failure zones of these two approaches. Similarly, we observe that in 89.9% of all states, the difference between the value functions are less than 1%. For a small fraction of the state space (0.5% of all states), the percentage difference in value functions is higher than 5%. These states typically correspond to protein-impurity pairs within the risk zone and having zero profit or less.

Difference in optimal policies: Lastly, we evaluate the difference between the optimal purification policies for the worst-case and probabilistic approaches, as shown in Table 4.4. In the first step, we observe that 54% of the states have the same optimal purification policies suggested by these two approaches. These include states belonging to target and failure zone under both of these two approaches. The main difference in the optimal policy is observed in the risk zone, where, the worst-case approach adopts smaller pooling windows than the probabilistic approach in 39% of the states. These states typically correspond to lower protein and impurity levels (i.e., states with zero profit or losses). Interestingly, the worst-case approach adopts larger windows compared to the probabilistic approach in 7% of the state space at the first chromatography step. This behavior is particularly observed at high protein and impurity levels (i.e., when the starting batch has more than 22 mg of protein, and more than 19 mg of impurity). The overlap between the purification policies suggested by the worst-case and probabilistic approach is higher in the second chromatography step. For example, both of these two approaches suggest the same policies in 74% of the states. In the remaining 26% of the state space, the worst-case approach adopts smaller pooling windows than the probabilistic approach, as expected. Furthermore, we observe that the worst-case approach abandons the purification when the batch purity gets closer to 50% (i.e, $p_2 \leq 17.5, i_2 \geq 16$); whereas the probabilistic approach performs the purification using small pooling windows to eliminate impurities.

4.8.5 Impact of State Aggregation and Action Elimination

Using the state aggregation scheme in Proposition 4.5.1 and the action elimination procedure in Proposition 4.5.2, we obtain significant savings in the computational effort required to obtain solutions to industry sized problems. For example, applying the state aggregation scheme to the purification project presented in Section 4.8.2 has led to grouping 35.5% of the state space into a single aggregate state in the first step,

and similarly 43.5% of the state space in the second step. After eliminating strictly dominated actions at each chromatography step based on Proposition 4.5.2, the total number of pooling windows reduced from 55 to 36 candidate windows in the first step, and from 78 to 20 candidate windows in the second chromatography step. The combined impact of the state aggregation and action elimination procedures resulted in 54% reduction in the CPU time.

4.9 Conclusions

We focus on protein purification operations conducted by biomanufacturers and pharmaceutical companies. Each order represents an engineered protein having purity and yield requirements specified by the end use or application, and biomanufacturing firms often incur high penalty costs when these specific requirements are not achieved. However, achieving of both the purity and yield requirements is often challenging in a typical biomanufacturing setting, since the biomanufacturing firms might have to compromise on the protein yield in order to achieve the desired purity level. Furthermore, the starting material involves significant variability and uncertainty in terms of the protein and impurity amounts which affects the subsequent purification decisions. The problem involves continuous state space and a large action space due to the interlinked nature of the purification decisions. Limitations in the purification capabilities of the available chromatography techniques further challenge the purification decisions. Due to high penalty costs and strict production requirements on purity, biomanufacturing decisions need practical guidelines and guaranteed performance measures to hedge against uncertainties in their operations.

We develop an optimization framework which captures the yield and purity trade-offs, uncertainty in the starting material, limitations in the purification capabilities,

and interlinked decisions involving multiple purification steps for engineered proteins. We analyze the structural properties and establish theoretical results that provide practical guidelines for quantifying the risks and costs, and optimize purification decisions based on the specific production requirements on yield and purity. Our analysis partitions the state space into distinct, nonempty subsets called as the *failure zone*, *risk zone* and *target zone*. These zones provide an analysis of financial trade-offs and business risks based on the condition of the starting material and the limitations in manufacturing capabilities. For each zone, we then provide practical guidelines for purification decisions to maximize the total profitability. Zone-based decision making is particularly practical and easy to implement in most biomanufacturing settings. We also provide guaranteed performance measures using a worst-case analysis, and compare the managerial insights obtained from the probabilistic and worst-case approaches.

The optimization framework has been developed and implemented at Aldevron. Furthermore, the model and managerial insights have been shared and validated with a larger industry group (BioWGS, 2014; BioForward, 2014). Implementation insights at Aldevron indicate an average of 25% reduction in lead times and 20% reduction in operating costs. We believe that our optimization framework provides a rigorous analysis of the risks and financial trade-offs involved in protein purification. Applications of operations research techniques are mostly new to the biomanufacturing community, and most biomanufacturing processes are bound by strict regulatory controls. As more companies like Aldevron embrace operations research and integrate it into practice, regulatory authorities might mandate the use of such approaches to improve biomanufacturing efficiency.

4.10 Appendix: Proofs

Proof of Proposition 4.4.1. We prove the monotonicity of the value function using proof by induction. We first investigate the value function $\mathcal{V}_T(p_T, i_T)$ at stopping time $t = T$. Note that $\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T)$. It is easy to observe that stopping costs $r_S(p_T, i_T)$ in Equation (4.3) are nondecreasing in $p_T \in \mathcal{P}$ for a given $i_T \in \mathcal{I}$; and nonincreasing in $i_T \in \mathcal{I}$ for a given $p_T \in \mathcal{P}$.

Next, we assume by induction hypothesis that $\mathcal{V}_t(p_t, i_t)$ is nondecreasing in $p_t \in \mathcal{P}$ for a given $i_t \in \mathcal{I}$, and for all $t \in \mathcal{T}$. First, we proceed with investigating the monotonicity of the value function in p_t for a given $i_t \in \mathcal{I}$. Let $p_t^- < p_t$, $\{p_t^-, p_t\} \in \mathcal{P}$ for $t \in \mathcal{T}$. By definition of the value function in Equation (4.4), we have, for $i_t \in \mathcal{I}$ and $t \in \mathcal{T}$,

$$\begin{aligned}
\mathcal{V}_t(p_t, i_t) &= \max_{w_t \in W_t} \Big\{ r_S(p_t, i_t), -c_t + \mathop{\mathbb{E}}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t | w_t) \Big\} \\
&\geq \max_{w_t \in W_t} \Big\{ r_S(p_t^-, i_t), -c_t + \mathop{\mathbb{E}}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t | w_t) \Big\} \qquad (4.17) \\
&\geq \max_{w_t \in W_t} \Big\{ r_S(p_t^-, i_t), -c_t + \mathop{\mathbb{E}}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t^-, \psi_t i_t | w_t) \Big\} \qquad (4.18) \\
&= \mathcal{V}_t(p_t^-, i_t) \qquad (4.19)
\end{aligned}$$

where, Equation (4.17) follows from the cost structure in Equation (4.3), and Equation (4.18) is obtained from the induction hypothesis. Proof for monotonicity of the value function in $i_t \in \mathcal{I}$ for a given $p_t \in \mathcal{P}$ and for all $t \in \mathcal{T}$ is entirely analogous. □

Proof of Theorem 4.4.1. It is sufficient to show that if $a^*(p_t', i_t') = S$ then $a^*(p_t, i_t) = S$ for all states in $\mathbb{F}_t = \{(p_t, i_t) \in \mathcal{P} \times \mathcal{I} : \ p_t \leq p_t' \text{ and } i_t \geq i_t'\}$, $t \in \mathcal{T}$.

Note that at time $t = T$, the only available action is to stop with rewards $\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T)$. Next, we fix any $(p_t, i_t) \in \mathbb{F}_t$, and assume by contradic-

tion hypothesis that $a^*(p'_t, i'_t) = S$, but $a^*(p_t, i_t) = w$, for $(p_t \leq p'_t, i_t \geq i'_t)$, where $w \in \mathcal{W}_t$ and $w \neq S$. This implies that,

$$r_S(p'_t, i'_t) > -c_t + \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p'_t, \psi_t i'_t) \mathrm{d}\theta \mathrm{d}\psi \tag{4.20}$$

and

$$-c_t + \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi > r_S(p_t, i_t) \tag{4.21}$$

which together imply

$r_S(p'_t, i'_t) - r_S(p_t, i_t)$

$$\begin{aligned} > & \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p'_t, \psi_t i'_t) \mathrm{d}\theta \mathrm{d}\psi \\ & - \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi. \end{aligned} \tag{4.22}$$

Note that $r_S(p'_t, i'_t) - r_S(p_t, i_t) = 0$ due to stopping cost structure in Equation (4.3). Theorem 4.4.1 defines (p'_t, i'_t) such that $\gamma_d > \frac{p'_t}{(p'_t + i'_t)}$. Hence, $r_S(p'_t, i'_t) = -c_f$, and also $r_S(p_t, i_t) = -c_f$ since $(p_t \leq p'_t, i_t \geq i'_t)$. Therefore, inequality (5.13) indicates that the term on its right hand side is negative. However,

$$\begin{aligned} & \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p'_t, \psi_t i'_t) \mathrm{d}\theta \mathrm{d}\psi \\ & - \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi \qquad (4.23) \\ \geq & \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi \\ & - \int_{\psi^l|w}^{\psi^u|w} \int_{\theta^l|w}^{\theta^u|w} f_t(\theta_t|w) g_t(\psi_t|w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi \qquad (4.24) \\ = & 0. \end{aligned}$$

Therefore, the term on the right hand side of inequality (5.13) is positive, which contradicts the inequality (5.13), and hence the proof follows. Note that Equation (5.14) follows from the monotonicity of the value function in Proposition 4.4.1, and the fact that $\mathbb{E}\mathcal{V}_{t+1}(\theta_t p'_t, \psi_t i'_t)$ is negative by the contradiction hypothesis, and note that $(p_t \leq p'_t, i_t \geq i'_t) \in \mathbb{F}_t$. □

Proof of Proposition 4.4.2. We prove Proposition 4.4.2 by induction. We first focus on condition (i). Let $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ with $\gamma_d > \frac{p_t}{(p_t+i_t)}$ at step $t \in \mathcal{T}$. Assume by induction hypothesis that (p_t, i_t) satisfy the condition (i). Then, at the last chromatography step $T-1$, we have,

$$
\begin{aligned}
&\mathcal{V}_{T-1}(p_{T-1}, i_{T-1}) \\
&= \max_{\{w_{T-1} \in \mathcal{W}_{T-1}\}} \Big\{ r_S(p_{T-1}, i_{T-1}), -c_{T-1} + \mathop{\mathbb{E}}_{\theta_{T-1}, \psi_{T-1} | w_{T-1}} r_S(\theta_{T-1} p_{T-1}, \psi_{T-1} i_{T-1}) \Big\} \\
&= \max \Big\{ -c_f, -c_{T-1} - c_f \Big\} \qquad (4.25) \\
&= -c_f
\end{aligned}
$$

Note that Equation (4.25) follows from the induction hypothesis and the stopping costs structure defined in Equation (4.3).

Similarly, at the chromatography step $t \in \mathcal{T}$, we have,

$$
\begin{aligned}
\mathcal{V}_t(p_t, i_t) &= \max_{\{w_t \in \mathcal{W}_t\}} \Big\{ r_S(p_t, i_t), -c_t + \mathbb{E}\, \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \Big\} \\
&\geq \max \Big\{ -c_f, -c_t - c_f \Big\} \qquad (4.26) \\
&= -c_f
\end{aligned}
$$

where, inequality in Equation (4.26) follows from the induction hypothesis and stopping costs as condition (i) holds. Hence, abandoning the purification at state

(p_t, i_t) leads to less financial losses than continuing the purification under condition (i). The proof for condition (ii) is entirely analogous to that of condition (i). □

Proof of Proposition 4.4.3. We use backward induction. At the end of the planning horizon T, the target zone is by definition:

$$\mathbb{T}_T = \Big\{(p_T, i_T) : \; p_T \geq p_d, \; \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T \Big\}.$$

In order for (p_T, i_T) to be element of $\mathbb{T}_T$ by time T, it is sufficient that, at $T-1$, we have

$$(p_{T-1}, i_{T-1}) \;\; \in \;\; \Big\{ \bigcup_{w \in \mathcal{W}_{T-1}} p_{T-1} \geq \frac{p_d}{(\theta^l_{T-1}|w)}, \frac{1-\gamma_d}{\gamma_d} p_{T-1} \frac{(\theta^l_{T-1}|w)}{(\psi^u_{T-1}|w)} \geq i_{T-1} \Big\},$$

which is equivalent to

$$(p_{T-1}, i_{T-1}) \;\; \in \bigcup_{w \in \mathcal{W}_{T-1}} J^w_{T-1}(\mathbb{T}_T),$$

based on J^w_{T-1} in Equation (4.8). Hence, $\mathbb{T}_{T-1} \equiv \bigcup_{w \in \mathcal{W}_{T-1}} J^w_{T-1}(\mathbb{T}_T)$. Repeated application of the same procedure leads to

$$\mathbb{T}_{T-k} = \bigcup_{w \in \mathcal{W}_{T-k}} J^w_{T-k}(\mathbb{T}_{T-k+1}), \text{for } k = 1, \ldots, T-1.$$

which is equivalent to the Equation (4.10) in Proposition (4.4.3). □

Proof. Proof of Proposition 4.5.1 Let the partition $\left\{I_{p_t \leq p'_t, i_t \geq i'_t}\right\}$ represent all protein-impurity pairs satisfying Proposition 4.5.1. Note that Proposition 4.5.1 has the same conditions as the failure zone in Proposition 4.4.2. Therefore, based on Theorem 4.4.1,

we have $\mathcal{V}_t^*(p_t, i_t) = -c_f$ over the partition $\left\{I_{p_t \leq p'_t, i_t \geq i'_t}\right\}$ specified in Proposition 4.5.1, since $I_{p_t \leq p'_t, i_t \geq i'_t} \in \mathbb{F}_t$ for all chromatography steps $t \in \mathcal{T}$ and pooling windows $w_t \in \mathcal{W}_t$. Hence, the aggregate failure state can be modeled as an absorbing state with reward $r(d_t) = -c_f$, and the aggregation scheme is exact since the failure state d_t encompasses subsets of the original system states that have same costs and transitions.

□

Proof of Proposition 4.5.2: First, we fix any protein-impurity pair $(p_t, i_t) \in \mathcal{P} \times \mathcal{T}$ at the chromatography step $t \in \mathcal{T}$. Let w_t^i and w_t^j be two distinct pooling windows at the chromatography step $t \in \mathcal{T}$, such that, $F_t(\Theta|w_t^i) \geq_{st} F_t(\Theta|w_t^j)$, $G_t(\Psi|w_t^i) \leq_{st} G_t(\Psi|w_t^j)$, and $(\theta_t^l|w_t^i) < (\theta_t^l|w_t^j)$, $(\theta_t^u|w_t^i) < (\theta_t^u|w_t^j)$, and $(\psi_t^l|w_t^i) > (\psi_t^l|w_t^j)$, $(\psi_t^u|w_t^i) > (\psi_t^u|w_t^j)$, as specified in Proposition 4.5.2. Next, we evaluate the optimal value function $\mathcal{V}_t(p_t, i_t|w_t^j)$ at state (p_t, i_t) under the pooling action w_t^j at the chromatography step t:

$$
\begin{aligned}
&\mathcal{V}_t(p_t, i_t|w_t^j) \\
&\quad = -c_t + \int_{\psi^l|w_t^j}^{\psi^u|w_t^j} \int_{\theta^l|w_t^j}^{\theta^u|w_t^j} f_t(\theta_t|w_t^j) g_t(\psi_t|w_t^j) \mathcal{V}_{t+1}(p_t\theta_t, \psi_t i_t|w_t^j) \mathrm{d}\theta \mathrm{d}\psi \\
&\quad > -c_t + \int_{\psi^l|w_t^i}^{\psi^u|w_t^i} \int_{\theta^l|w_t^i}^{\theta^u|w_t^i} f_t(\theta_t|w_t^i) g_t(\psi_t|w_t^i) \mathcal{V}_{t+1}(p_t\theta_t, \psi_t i_t|w_t^i) \mathrm{d}\theta \mathrm{d}\psi \qquad (4.27) \\
&\quad = \mathcal{V}_t(p_t, i_t|w_t^i).
\end{aligned}
$$

Note that Equation (4.27) follows from the conditions in Proposition 4.5.2 and the monotonicity of the value function in Proposition 4.4.1. Hence, for any $(p_t, i_t) \in \mathcal{P} \times \mathcal{T}$ at the chromatography step $t \in \mathcal{T}$, the value function $\mathcal{V}_t(p_t, i_t|w_t^j)$ under the pooling window w_t^j denotes strictly higher profit then the value function $\mathcal{V}_t(p_t, i_t|w_t^i)$ under the pooling window w_t^i. Hence, w_t^i is said to be strictly dominated by w_t^j at the chromatography step $t \in \mathcal{T}$ since $\mathcal{V}_t(p_t, i_t|w_t^j) > \mathcal{V}_t(p_t, i_t|w_t^i)$ and thus $a^*(p_t, i_t) \neq w_t^i$.

□

Proof of Proposition 4.6.1: We use backward induction to identify the effective purity set. At the end of the planning horizon T, the effective purity set is, by definition,

$$\mathbb{P}_T^e = \Big\{(p_T, i_T) : \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T\Big\}.$$

Using the same definition at chromatography step $T-1$, we have,

$$\mathbb{P}_{T-1}^e = \Big\{(p_T, i_T) : \bigcup_{w \in \mathcal{W}_{T-1}} \frac{1-\gamma_d}{\gamma_d} \frac{(\theta_{T-1}^u|w)}{(\psi_{T-1}^l|w)} p_T \geq i_T\Big\},$$

which is equivalent to

$$\mathbb{P}_{T-1}^e = \bigcup_{w \in \mathcal{W}_{T-1}} K_{T-1}^w(\mathbb{P}_T^e).$$

For $k = 1, \ldots, T-1$, the operator $K_{T-k}^w(\cdot)$ is defined as

$K_{T-k}^w(y\, p_{T-k+1} \geq i_{T-k+1}) : (y\, p_{T-k+1} \geq i_{T-k+1}) \rightarrow (y p_{T-k} \frac{(\theta_{T-k}^u|w)}{(\psi_{T-k}^l|w)} \geq i_{T-k})$. Using backward induction, successive application of the same procedure to the chromatography steps $T-k$, $k = 2, \ldots, T-1$ leads to the following recursion:

$$\mathbb{P}_T^e = \Big\{(p_T, i_T) : \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T\Big\},$$

$$\mathbb{P}_{T-k}^e = \bigcup_{w \in \mathcal{W}_{T-k}} K_{T-k}^w(\mathbb{P}_{T-k+1}^e),$$

as stated in Proposition 4.6.1. □

Proof of Theorem 4.6.1: Theorem 4.6.1 identifies the optimal policy for states that are element of the risk zone $(p_t, i_t) \in \mathbb{R}_t$. We note that, all state impurity pairs such

that $a_t^*(p_t, i_t) = S$ at time $t \in \mathcal{T}$ are classified as $(p_t, i_t) \in \mathbb{F}_t$, by definition of the failure zone.

First, we classify the pooling actions in the risk zone into two distinct sets: $\bar{W}_t = \{\bar{w}_t \in \mathcal{W}_t : (\theta_t^u p_t, \psi_t^l i_t | \bar{w}_t) \notin \mathbb{P}_{t+1}^e | (p_t, i_t) \in \mathbb{R}_t, a_t(p_t, i_t) = \bar{w}_t\}$, and $\check{W}_t = \{\check{w}_t \in \mathcal{W}_t : (\theta_t^u p_t, \psi_t^l i_t | \check{w}_t) \in \mathbb{P}_{t+1}^e | (p_t, i_t) \in \mathbb{R}_t, a_t(p_t, i_t) = \check{w}_t\}$ for all $t \in \mathcal{T}$. An example of action type $\bar{w}_t$ could be a pooling window that leads the the failure zone over the next decision epoch; whereas an example of action type $\check{w}_t$ is a pooling window that keeps the system state within the risk zone of the next decision epoch. Hence, our objective function for this sub-problem can be rewritten as:

$\mathcal{V}_t((p_t, i_t) | (p_t, i_t) \in \mathbb{R}_t) =$

$$\max_{w_t \in \{\bar{W}_t \cup \check{W}_t\}} \Big\{ r_S(p_t, i_t), -c_t + \mathbb{E}\, \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t | w_t) \Big\}, \tag{4.28}$$

and

$$\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T). \tag{4.29}$$

At the end of the planning horizon T, we have $\mathcal{V}_T(p_T, i_T) =$

$$r_S(p_T, i_T) = \begin{cases} -c_f & \text{if } (p_T, i_T) \notin \mathbb{P}_T^e, \\ r(p_d) & \text{if } (p_T, i_T) \in \mathbb{P}_T^e \text{ and } p_t \geq p_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } (p_T, i_T) \in \mathbb{P}_T^e \text{ and } p_t < p_d. \end{cases} \tag{4.30}$$

As a result of Equations (4.28)-(4.30), we observe that at time $T-1$, the optimal action is to keep the system state within the effective purity set of the next period, i.e., $a_{T-1}^*(p_{T-1}, i_{T-1}) = \{\check{w}_{T-1} \in \check{\mathcal{W}}_{T-1} : (\theta_{T-1} p_{T-1}, \psi_{T-1} i_{T-1} | \check{w}_{T-1}) \in \mathbb{P}_T^e\}$ for all $(p_{T-1}, i_{T-1}) \in \mathbb{R}_{T-1}$.

Similarly, at time $t \in \mathcal{T} \setminus T-1$, by the definition of the desired purity set in Proposition 4.6.1, we observe that a batch state $\{(p_{t+1}, i_{t+1}) \notin \mathbb{P}_{t+1}^e | (p_t, i_t) \in \mathbb{R}_t\}$ has no chance of meeting the final purity requirement by the time T, even under

the best possible realization of the purification capabilities. Hence, using the cost structure in Equation (4.30), we have $\mathcal{V}_{t+1}(p_{t+1}, i_{t+1}) = r_S(p_{t+1}, i_{t+1}) = -c_f$ for all $\{(p_{t+1}, i_{t+1}) \notin \mathbb{P}^e_{t+1})\}$. As a result, the optimal action at step t has the characteristic that $a^*_t(p_t, i_t) = \left\{\check{w}_t \in \check{\mathcal{W}}_t : (\theta_t p_t, \psi_t i_t |\, \check{w}_t) \in \mathbb{P}^e_{t+1}\right\}$ for $\{t, t+1\} \in \mathcal{T}$ and for all $(p_t, i_t) \in \mathcal{R}_t$. □

Proof of Theorem 4.6.2. Theorem 4.6.2 analyzes the optimal policies for all $(p_t, i_t) \in \mathbb{T}_t$, $t \in \mathcal{T}$ in Case 1. We use backward induction. At the end of the planning horizon T, we have $\mathcal{V}_T(p_T, i_T) =$

$$r_S(p_T, i_T) = \begin{cases} -c_f & \text{if } \gamma_T < \gamma_d, \\ r(p_d) & \text{if } (p_T, i_T) \in \mathbb{T}_T \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } \gamma_T \geq \gamma_d \text{ and } p_t < p_d. \end{cases} \tag{4.31}$$

Hence, at time $T-1$, the optimal pooling action is in such a way as $a^*_{T-1}(p_{T-1}, i_{T-1}) = \left\{w^*_{T-1} \in \mathcal{W}_{T-1} : ((\theta^l_{T-1} p_{T-1}, \psi^u_{T-1} i_{T-1} | w^*_{T-1}) \in \mathbb{T}_T \,|\, (p_{T-1}, i_{T-1}) \in \mathbb{T}_{T-1})\right\}$ for all $(p_{T-1}, i_{T-1}) \in \mathbb{T}_{T-1}$ with $\gamma_{T-1} < \gamma_d$. Note that the structure of the target zones in Proposition 4.4.3 and Property (1)-(3) ensure that there exists at least one such policy. We proceed similarly with time $t \in \mathcal{T}$. Note that the bounds on the value function derived in Section 4.4.3 indicate that (1) $\mathcal{V}^*_t(p_t, i_t) = -c_f$ for all $(p_t, i_t) \in \mathbb{F}_t$, $t \in \mathcal{T}$. (2) $\sum_{j=t}^{T-1} -c_j + r(p_d) \leq \mathcal{V}^*_t(p_t, i_t) \leq r(p_d)$ for all $(p_t, i_t) \in \mathbb{T}_t$, $t \in \mathcal{T}$. (3) $-c_f \leq \mathcal{V}^*_t(p_t, i_t) \leq \sum_{j=t}^{T-1} -c_j + r(p_d)$ for all $(p_t, i_t) \in \mathbb{R}_t$, $t \in \mathcal{T}$. Hence, based on the bounds of the value function, the optimal pooling policy for all $(p_t, i_t) \in \mathbb{T}_t$ is in such as way as $a^*_t(p_t, i_t) = \left\{w^*_t \in \mathcal{W}_t : ((\theta^l_t p_t, \psi^u_t i_{t+1} | w^*_t) \in \mathbb{T}_{t+1} \mid (p_t, i_t) \in \mathbb{T}_t)\right\}$ for $\{t, t+1\} \in \mathcal{T}$, for all $(p_t, i_t) \in \mathbb{T}_t$. Proposition 4.4.3 and Property (1)-(3) ensure that there exists at least one such policy. □

Proof of Theorem 4.6.3. Since Case 2 is relaxing the yield requirement from Case 1, The proof is entirely analogous to that of Theorem 4.6.1, and hence omitted. □

Chapter 5

Simultaneous Optimization of Upstream and Downstream Operations

In this chapter, we focus on the interaction between the upstream fermentation and downstream purification decisions, such that, the scientist identifies the best amount of protein to be manufactured in the upstream fermentation operations, and also determines the best choice of equipment (called as the *chromatography technique*) and the best operating policy (called as the *pooling window*) for the downstream operations. In most industry practices, meeting the specific requirements on the final yield and purity could be challenging due to high operating costs, random yield losses, uncertain process outcomes, and financial trade-offs. For example, drugs that are in the final phases of their research and development (and whose potential end-users are humans) should abide by stringent purity requirements, i.e., the batch shipped to the customer should be free of contaminants and other unwanted impurities. Achieving high purity for a batch could require multiple purification steps with expensive operating costs and failure risks. However, 'a good starting material' that involves

'sufficient' amount of protein can significantly alleviate the risks and challenges in protein purification operations. Increasing the protein mass obtained from the upstream batch is possible and also a popular research topic. However, improving the protein mass could require expensive operating and re-engineering costs, and might not always outweigh its resulting benefits in the downstream purification operations.

We develop a stochastic optimization framework to address common challenges encountered in industry practices, and answer the following research questions: Given the limitations in the purification operations and financial trade-offs, does producing high amount of protein in upstream operations always result in higher profit? What is the 'best' protein mass to start a specific purification order? What are the best choices of chromatography techniques and pooling windows based on specific production requirements? Can we develop exact approximation mechanisms to solve industry size problems? How can biomanufacturing firms overcome the challenges in conforming to high purity requirements?

Our research has been conducted in close collaboration with Aldevron, a contract biomanufacturer specializing in protein manufacturing. Furthermore, our research questions, mathematical model, assumptions, and managerial insights have been validated through a series of working group sessions with the local biomanufacturing industry in Wisconsin (BioWGS, 2014) and have been shared with a larger biotechnology community (BioForward, 2014; Engel, 2014). Our study is one of the first of its nature that combines the knowledge from operations research and chemical engineering to develop a stochastic optimization framework for engineer-to-order protein manufacturing. The proposed optimization framework addresses manufacturing system-level challenges that are often encountered in practice. These challenges include random yield losses, uncertain purity outcomes, financial trade-offs, strict

production requirements, failure risks, and the interlinked nature of the manufacturing steps. Most literature focus on modeling and optimizing biological and chemical parameters but often ignore these manufacturing system challenges. However, these challenges are often encountered in industry practices and lead to several managerial and operational issues in decision making. In this study, we provide a unifying framework that links the underlying biology and chemistry of protein manufacturing processes with manufacturing system challenges to improve profitability. Furthermore, we investigate the structural properties of the optimization problem, and develop optimal polices that are easy to implement in practice. We study a state aggregation scheme that significantly reduces the curse of dimentionality in large problems, and demonstrate the use of the optimization model with a case study at Aldevron.

The remainder of the chapter is organized as follows. We provide a background on biomanufacturing operations and define the problem in Section 5.1. We develop an optimization model in Section 5.2, and analyze its structural properties in Section 5.3. The model formulation is revised using a state aggregation scheme in Section 5.4. An industry case study is presented in Section 5.5, and concluding remarks are provided in Section 5.6.

5.1 Background on Protein Manufacturing and Prior Work

In this section, we discuss the manufacturing processes, production requirements, operational and financial trade-offs involved in protein manufacturing, and provide the problem definition.

5.1.1 Yield and Purity Requirements

Purity is a typical measure of the batch quality in biomanufacturing operations, and represents the fraction of the total amount of the protein of interest based on the total amount of proteins and impurities available in the batch. Several chromatography techniques are used to separate the protein of interest from unwanted impurities in order to meet a predefined *purity requirement* specified by the end use or application. Depending on the end use, the purity requirement could range from 80% to 99.9% for a batch. For example, if the end users are humans, then the batch often needs to be free of unwanted impurities to satisfy regulatory requirements. In this chapter, we are specifically interested in optimizing the purification decisions of drugs whose potential end-users are humans, and hence the final batch should abide by very high purity requirements (i.e., $\geq 99.9\%$ purity).

The yield requirement represents the amount (mass) of the protein of interest that should be obtained by the end of the purification operations. We note that we use the terms 'amount' and 'mass' interchangeably during the rest of the chapter. The yield requirement is typically specified by the customer along with the purity requirement. When the yield requirement is not achieved, the biomanufacturing company incurs penalty costs associated with per unit of protein short. The yield requirement often corresponds to 25% to 40% of the protein of interest available in the starting material, demonstrating the challenges in meeting both the purity and yield requirements simultaneously. Furthermore, the biomanufacturer often encounters a *trade-off between the upstream and downstream operating costs* (Section 5.1.2) as well as a *trade-off between the yield and purity* obtained at each chromatography step in the downstream operations (Section 5.1.3 and Section 5.1.4).

5.1.2 Optimizing Protein Mass in Upstream Operations

The amount of the protein of interest along with unwanted impurities increase simultaneously during the fermentation process. The scientist operating the fermentation does not have the ability to block or interfere with the formation of the impurities since these are natural metabolites. Different types and amounts of impurities that are expected to be obtained by the end of the fermentation operation are known to the scientist. However, the scientist has the ability to control the amount of the protein of interest obtained through the fermentation process. To do so, there are several upstream controls that increase the protein formation during fermentation. These include feeding the cell culture, adjusting the harvesting times, increasing the productivity of cell lines through biological and chemical inferences, re-engineering cells to increase their productivity, etc. Increasing the protein concentration obtained from a batch is a very popular research topic in the biomanufacturing literature, and is referred to as the *problem of increasing the titer*. Although increasing the protein mass obtained from a batch is possible, it may significantly boost the upstream operating costs since these costs are nondecreasing in the protein amount obtained. Consequently, increasing the protein mass increases the upstream costs but also reduces the risk of incurring lost sale penalty costs by the end of the purification operations. On the other hand, increasing the protein mass "more than required" could hurt the profitability of a batch since the customers do not purchase the proteins manufactured in excess of their yield requirement. Identifying the best amount of protein that should be obtained from the upstream batch is challenging since the scientist incurs random yield losses and uncertainty in the batch purity while performing the downstream purification operations. Although higher protein mass increases the upstream costs, it alleviates the downstream purification risks. Due to this trade-off, this scientists needs a framework that identifies the optimal protein mass based on the expected performance of the downstream operations and

the specific production requirements.

Studies on upstream fermentation often focus on developing models for cell growth and product formation (Patel et al., 2000; Tsao et al., 2004; Xing et al., 2010). Several studies investigate optimal control strategies to increase the product formation and protein amount (Saucedo and Karim, 1997; Yang et al., 2000; Peroni et al., 2005; Gnoth et al., 2007; Xing et al., 2010). However, fermentation studies typically focus on cell-level dynamics, and often do not capture risk and cost trade-offs associated with unwanted impurities and limited separation capabilities of purification operations. In this chapter, we do not develop fermentation control policies to increase the product formation or protein amount. Instead, we focus on identifying the optimal amount of protein that should be obtained by the end of the fermentation given that the scientist has multiple control options for increasing the protein mass up to a desired amount. We build a framework that links the upstream protein mass decisions with the downstream purification decisions, failure risks, financial trade-offs and specific production requirements for engineer-to-order proteins.

5.1.3 Optimizing Pooling Decisions for a Given Chromatography Technique

Figure 5.1 (a) shows an example of the outcome obtained from a chromatography operation. Each column in Figure 5.1 (a) is called as a *lane* and corresponds to a discrete time interval which could often be as small as one minute in most industry settings. The y-axis in Figure 5.1 (a) represents the molecular mass of the target protein and impurities associated with each lane. Each lane is comprised of some fraction of the total amount of the protein of interest, as well as some fraction of different types of unwanted impurities. For example, consider the lane 8 in Figure 5.1 (a). In this lane, there are 8 different types of molecules, i.e., one of the molecules in

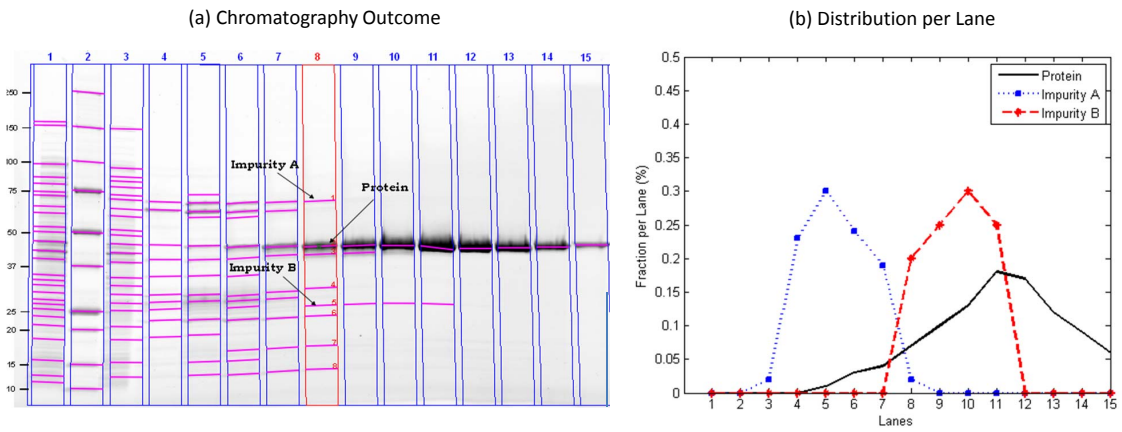

Figure 5.1: An example of chromatography outcome

Figure 5.1 (a) is the protein of interest and the remaining ones are different types of unwanted impurities as indicated by the arrows. A typical practice in the literature is to translate the chromatography outcome in Figure 5.1 (a) into visual diagrams shown in Figure 5.1(b). Figure 5.1(b) plots the fraction of protein and the fraction of a selected impurity type available in each lane when the chromatography technique described in Figure 5.1 (a) is used.

The scientist performing the chromatography operation decides which *chromatography technique* and *pooling window* to use simultaneously. First, we introduce the pooling window decision for a given chromatography technique. Consider the chromatography outcome in Figure 5.1(b). For simplicity, we plot only two different impurity types in Figure 5.1(b), namely Impurity A and B, and use this example to discuss the process trade-offs and pooling decisions. If the scientist pools lanes 3-15 in Figure 5.1(b), then she collects all of the protein along with Impurity A and B. On the other hand, if she pools lanes 9-15, she compromises on the yield (i.e., collects a smaller fraction of the protein) but completely eliminates Impurity A and some fraction of Impurity B. Alternatively, the scientist might decide to significantly compromise on the yield (i.e., lose $\approx 55\%$ of the protein) by pooling lanes 12-15, but this

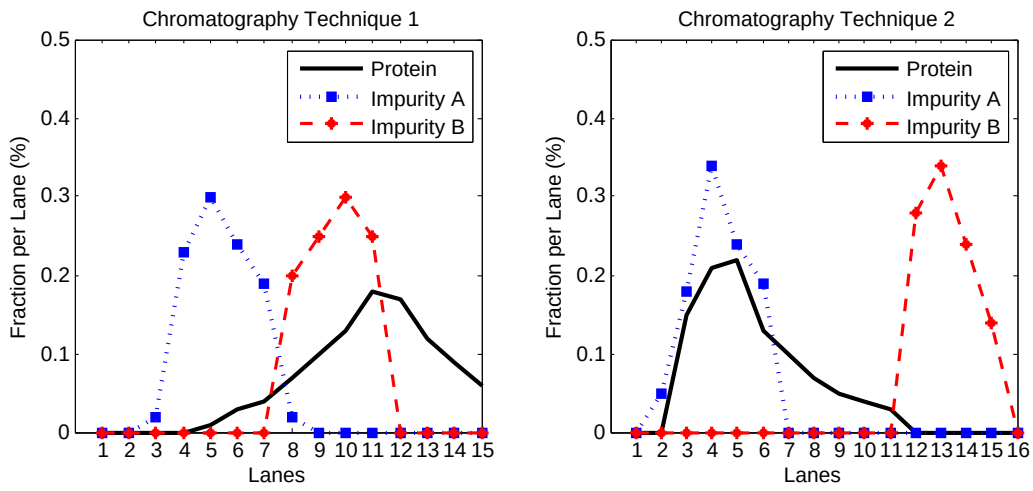

Figure 5.2: Difference in the separation capability of two chromatography techniques

decision helps achieving 100% purity since all unwanted impurities are eliminated. This simple example illustrates the yield and purity trade-off typically encountered in chromatography operations. Pooling window decisions can become more complex when a chromatography technique demonstrates differential affinity to different types of impurities. In practice, depending on the outcome of a chromatography step, the scientist makes decisions regarding the chromatography technique *and* the pooling window for each of the purification steps. This leads to another problem, namely the choice of chromatography technique.

5.1.4 Optimizing the Choice of Chromatography Technique

Figure 5.2 shows the output obtained from two different chromatography techniques based on industry data. Both of the graphs in Figure 5.2 use the same starting material but different techniques to separate the protein of interest from unwanted impurities. The starting material contains the protein of interest along with Impurity A and B. The x-axis in Figure 5.2 represents the lanes, and the y-axis denotes the fraction of the protein and impurities corresponding to each lane. The solid curve in Figure 5.2 represents the protein of interest, and other curves are associated with dif-

ferent types of impurities available in each lane for a given chromatography technique.

Figure 5.2 demonstrates that the choice of chromatography technique strongly influences the relative positions of the protein and impurities and their corresponding amount in each lane. For example, Impurity A is located on the left hand side of the protein of interest under the first chromatography technique in Figure 5.2. However, the same Impurity A overlaps completely with the protein of interest under the second chromatography technique presented in Figure 5.2. It is clear to see that the second technique is not capable of separating Impurity A from the protein of interest. However, it provides better separation outcome for Impurity B compared to the first chromatography technique. The pooling window decision in this setting is now influenced by the choice of chromatography technique. For example, if the scientist needs to completely eliminate Impurity B, then she could pool the lanes 12-15 on the first chromarography technique (compromising on $\approx 55\%$ of the protein of interest) or she could pool the lanes 2-11 on the second chromatography technique without compromising on the protein. The relative locations of the protein and impurities and their corresponding amounts in each lane are complex functions of the physical and chemical properties, and vary for each different chromatography technique. Closed-form expressions to estimate chromatography outcomes based on physico-chemical characteristics are available in the literature (Vasquez-Alvarez et al., 2001; Polykarpou et al., 2011b). In this chapter, we use industry data to capture the expected outcome of a chromatography technique. In practice, a starting material could contain up to 100 different types of impurities to be separated using 5-10 available chromatography techniques. The combinatorial complexity arising from multiple types of impurities and chomatography techniques, and their associated yield and purity trade-offs makes purification decisions complex.

There are several studies in the literature that build deterministic optimization models to identify the optimal selection and sequencing of the chromatography techniques (Vasquez-Alvarez et al., 2001; Vasquez-Alvarez and Pinto, 2004; Lienqueo et al., 2009; Nfor et al., 2013). A common characteristics of the literature is that mixed integer linear programming models are developed to minimize the number of purification steps while achieving a predetermined purity level. There are only few studies in the literature that focus on optimizing both the selection of the chromatography techniques and the pooling window simultaneously. (Polykarpou et al., 2011b, 2012a,b). However, most studies focus on minimizing the number of purification steps and do not capture the manufacturing system challenges, such as, randomness in purification outcomes, business risks and financial trade-offs in engineer-to-order protein manufacturing.

5.1.5 Summary of Challenges in Decision Making

In practice, the decisions on upstream protein mass, downstream chromatography techniques and pooling windows are challenging due to (*i*) *Cost trade-offs between upstream and downstream:* Higher protein mass increases the operating costs in upstream processes but reduces the failure risks and associated costs in downstream operations, (*ii*) *Purity and yield trade-offs in purification:* The biomanufacturing company often needs to compromise on the protein yield to achieve the stringent requirements on the final purity, (*iii*) *Uncertainty in chromatography outcome:* The amount of protein and impurities available in each lane are often subject to variability and uncertainty, (*iv*) *Different types of impurities:* The starting material could contain various types of impurities whose separation performance differs for each available chromatography technique, (*v*) *Interlinked decisions:* Biomanufacturing operations involve multiple manufacturing steps in series where the outcome of one step directly impacts the performance of subsequent steps, and (*vi*) *Large penalty costs:*

The biomanufacturing company incurs large penalty costs associated with failure to meet the specific purity and yield requirements. To address these challenges, we develop a stochastic optimization model that maximizes the total expected profit of an engineer-to-order protein batch.

5.2 Model Formulation

In this section, we provide a Markov decision model that maximizes the total expected profit obtained from a specific order. We decompose the *optimization problem* into two sub-problems: the *upstream protein mass problem* and the *downstream purification problem*.

Decision Epochs: We consider discrete time, finite horizon Markov decision model. The set $\mathcal{T} = \{t : 1, \ldots, T-1\}$ denotes the decision epochs for the downstream purification problem, where each decision epoch $t \in \mathcal{T}$ represents the beginning of a chromatography step. The number of chromatography steps required to achieve the desired purity level is finite and bounded by $T - 1$ due to the limitations in the number of chromatography techniques. Next, we let $t = 0$ denote the decision epoch for the upstream protein mass problem that determines the amount of protein available in the starting material. Hence, the set $\mathcal{T} \cup \{0\}$ denotes all the decision epochs of the optimization problem. The end of the planning horizon is captured by T, such that, the batch is either shipped to the customer or scrapped at a penalty cost at the final time $t = T$.

State Space: First, we focus on the state space for the downstream purification problem. Let $p_t \in \mathcal{P}$ represent the amount (mass) of the protein of interest available in the batch at the beginning of the chromatography step $t \in \mathcal{T}$. Note that $p_t \in [0, p_1]$

for all chromatography steps $t \in \mathcal{T}$ since the amount of protein p_t available at the beginning of each step t is bounded by the starting material p_1. Next, we define the different types of impurities and their corresponding amounts available in the batch. Let $\mathcal{K} = \{k : 1, 2, \ldots, K\}$ be the set of K distinct types of impurities where $K < \infty$. Let $i_{k,t} \in \mathcal{I}_k$ denote the amount of the impurity type $k \in \mathcal{K}$ available in the batch at the beginning of the chromatography step $t \in \mathcal{T}$. Hence, the impurity states $(i_{1,t}, \ldots, i_{K,t}) \in \mathcal{I}_1 \times \ldots \times \mathcal{I}_K$ represent the set of all distinct types of impurities $k \in \mathcal{K}$ and their corresponding amounts $i_{k,t}$ available in the batch at the beginning of the chromatography step $t \in \mathcal{T}$. Note that $i_{k,t} \in [0, i_{k,1}]$ for each impurity type $k \in \mathcal{K}$ and chromatography step $t \in \mathcal{T}$. The amount of protein p_1 and the amount of impurities $i_{k,1}$ for each impurity type $k \in \mathcal{K}$ at the beginning of the first chromatography step is finite and determined by the upstream fermentation operations. We define the state $\{\Delta\}$ as the stopping state, i.e., an absorbing state with zero cost representing the end of the optimization problem where the batch is either shipped or scrapped. Therefore, $(p_t, i_{1,t}, \ldots, i_{K,t}) \cup \{\Delta\}$ is the state of the downstream purification problem for all chromatography steps $t \in \mathcal{T}$. Note that we use the terms 'amount' and 'mass' interchangeably.

Next, we consider the state space for the upstream protein mass problem. The amount of protein p_0 and impurity $i_{k,0}$ for all impurity types $k \in \mathcal{K}$ available in the batch at the beginning of the upstream fermentation operation is represented by the state $(p_0, i_{1,0}, \ldots, i_{K,0})$. It is assumed that there are no proteins or impurities available in the batch at the beginning of the fermentation process. Therefore, the starting state of the upstream protein mass problem is $p_0 = 0$ and $i_{k,0} = 0$ for all impurity types $k \in \mathcal{K}$.

Action Space: We first present the set of actions for the downstream purification problem. Let $\mathcal{C} = \{c : 1, 2, \ldots, C\}$ be the set of available chromatography techniques, and hence the action $c \in \mathcal{C}$ denotes the choice of the chromatography technique c to be used at a given purification step $t \in \mathcal{T}$. The action space $\mathcal{C}$ is finite and countable since there are limited number of chromatography techniques available in a biomanufacturing facility. Next, let $\mathcal{L}_c = \{1, 2, \ldots, L_c\}$ denote an ordered set of lanes available at each chromatography technique $c \in \mathcal{C}$. Note that the maximum number of lanes L_c on a chromatography technique c could be different for each technique $c \in \mathcal{C}$. Then, a pooling window $w_c \in \mathcal{W}_c$ on the chromatography technique c represents a subset of consecutive lanes from the set $\mathcal{L}_c$. More specifically, the set of all possible pooling windows for a chromatography technique c is $\mathcal{W}_c = \{(i, \ldots, j) \in \mathcal{L}_c : j = i + m, i = \{1, \ldots, L_c\}, m = \{0, 1, \ldots, L_c - i\}\} = \{w_c : w_1, w_2, \ldots, w_{N_c}\}$, where N_c is the maximum number of pooling windows available on a chromatography technique c. Note that the set $\mathcal{W}_c = \{w_c : w_1, w_2, \ldots, w_{N_c}\}$ represents the set of all possible pooling windows that can be adopted at the chromatography technique $c \in \mathcal{C}$. The set of pooling windows is finite, countable and bounded by N_c since there are limited number of lanes L_c at each chromatography technique c (See Section 5.1). Let $a_t(p_t, i_{1,t}, \ldots, i_{K,t})$ represent the action taken at the beginning of the purification step $t \in \mathcal{T}$ and state $(p_t, i_{1,t}, \ldots, i_{K,t})$. In the downstream purification problem, the scientist makes the joint decision $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ on both the chromatography technique $c \in \mathcal{C}$ and the pooling window $w_c \in \mathcal{W}_c$ used at the beginning of the purification step t. Additionally, the scientist has the possibility for taking the stopping action $\{S\}$ by either scrapping the batch or shipping it to the customer. We note that the only available action at the stopping state $\{\Delta\}$ is $a_t(\Delta) = S$ for all purification steps $t \in \mathcal{T}$.

In the upstream protein mass problem, the fermentation starts with $p_0 = 0$ and $i_k = 0$ for all $k \in \mathcal{K}$. The scientist has the ability to control the fermentation process

through a set of controls $\mathcal{U}$ that determines the amount of protein p_1 obtained by the fermentation operation. The control $u(p_1) \in \mathcal{U}$ represents a fermentation operating policy that leads to $p_1 \in [0, p_{max}]$ units of protein of interest by the end of the fermentation process. In practice, this operating policy is a combination of actions related to harvesting time, physical and biological parameters, and feeding policies. In this chapter, we do not build a detailed optimization model for the fermentation operating policies. Instead, we consider a high-level fermentation problem where we assume that a one to one mapping exists between the amount of protein p_1 obtained at the end of the fermentation and the control policy $u(p_1) \in \mathcal{U}$. We refer the reader to the fermentation optimization literature (discussed in Section 5.1.2) where detailed control models are developed to establish the relation between the fermentation operating actions $u(p_1) \in \mathcal{U}$ and the resulting protein amount p_1 (Saucedo and Karim, 1997; Yang et al., 2000; Xing et al., 2010). Note that the maximum amount of the protein that could be achieved through any actions used during the fermentation is bounded by p_{max} due to the limitations of the cell culture.

Transitions: In the downstream purification problem, the state transitions define the amount of the protein of interest p_t and the amount $i_{k,t}$ of each impurity type $k \in \mathcal{K}$ that remain in the batch after the chromatography technique $c \in \mathcal{C}$ is performed using the pooling window $w_c \in \mathcal{W}_c$ at a purification step $t \in \mathcal{T}$. First, we model the changes in the amount $i_{k,t}$ after the completion of the chromatography step $t \in \mathcal{T}$. For a given impurity type $k \in \mathcal{K}$, let $\{\Psi_k|c, w_c\}$ be a random fraction of the impurity amount $i_{k,t}$ that remains inside the batch after performing the chromatography technique c using the pooling window w_c. This implies that the remaining amount of the impurity type k was eliminated as a result of the action $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ taken at the beginning of that purification step t. Therefore,

$$i_{k,t+1} = (\psi_k|c, w_c) i_{k,t}. \tag{5.1}$$

Let the random fraction $\{\Psi_k|c, w_c\}$ have the distribution $g_k(\cdot|c, w_c)$ for all $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ and $k \in \mathcal{K}$, with finite support $[\psi_k^l|(c, w_c), \psi_k^u|(c, w_c)]$. Note that Equation (5.1) holds for all impurity types $k \in \mathcal{K}$. Since a chromatography technique exploits the physico-chemical properties of the protein and impurities to separate one from other, the distribution $g_k(\cdot|c, w_c)$ is unique for each impurity type $k \in \mathcal{K}$; and is independent of time $t \in \mathcal{T}$ and impurity amount $i_{k,t}$ for all $k \in \mathcal{K}$. However, the distribution $g_k(\cdot|c, w_c)$ is defined by the choice of chromatography technique c, pooling window w_c and physico-chemical characteristics of the impurity type $k \in \mathcal{K}$. We refer the reader to the chemical engineering literature for details (Vasquez-Alvarez et al., 2001; Polykarpou et al., 2011b).

Similarly, let $\{\Theta|c, w_c\}$ be the random fraction of protein p_t that remains inside the batch at the beginning of $(t+1)^{th}$ step, given that there are p_t units of protein at the beginning of the purification step t, and the action (c, w_c) is performed. This implies that the remaining amount of the protein of interest is eliminated during that chromatography step t. Therefore,

$$p_{t+1} = (\theta|c, w_c) p_t. \tag{5.2}$$

The random fraction $\{\Theta|c, w_c\}$ has distribution $f(\cdot|c, w_c)$ and a finite support $[\theta^l|(c, w_c),\ \theta^u|(c, w_c)]$ for all $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$. Note that the distribution $f(\cdot|c, w_c)$ is a function of the chromatography technique $c \in \mathcal{C}$ and pooling window $w_c \in \mathcal{W}_c$, but is independent of the time $t \in \mathcal{T}$ and impurities $i_{k,t}$ for all $k \in \mathcal{K}$. The finite support of the distributions $f(\cdot|c, w_c)$ and $g_k(\cdot|c, w_c)$ for all $k \in \mathcal{K}$ can be determined either using industry data or based on the physico-chemical characteristics of the protein and

impurities and their response to each chromatography technique c (Vasquez-Alvarez et al., 2001; Polykarpou et al., 2011b). The state transitions as a result of the action $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ are therefore captured by

$$\Big(p_{t+1}, i_{1,t+1}, \ldots, i_{K,t+1}\Big) = \Big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\Big). \tag{5.3}$$

In Equations (5.1)-(5.3), we note that the random fractions $(\Theta, \Psi_1, \ldots, \Psi_K | c, w_c)$ are independent of protein mass p_t, impurity mass i_t, impurity type k and purification step t, but only depend on the chromatography technique c and its pooling window w_c. This is because a chromatography technique separates a specific impurity type k by exploiting its physical and chemical properties, such as, molecular weight, electric charge, hydrophobicity, etc. These physicochemical characteristics are specific to the type of protein and impurities (regardless of their masses) and also independent of other impurities available in the batch and their masses (Vasquez-Alvarez et al., 2001; Polykarpou et al., 2011b). Note that the system transitions to the stopping state $\{\Delta\}$ when $a_t(p_t, i_{1,t}, \ldots, i_{K,t}) = S$ at any $t \in \mathcal{T} \cup \{0\}$. At the final time T, the only available action is to stop the purification, i.e., $a_T(p_T, i_{1,T}, \ldots, i_{K,T}) = S$ for all $p_T \in \mathcal{P}$, $i_{k,T} \in \mathcal{I}_k, k \in \mathcal{K}$.

In the upstream protein mass problem, the scientist has the ability to control and adjust the protein mass $p_1 \in \mathcal{P}$ obtained at the end of the fermentation, but can not interfere with the amounts of various impurities obtained along with the protein at the end of the fermentation. The main reason is because impurities are natural by-products, such as, ammonia and lactate, and hence their formation can not be interfered or prevented due to cellular dynamics. Given that the fermentation starts with $p_0 = 0$ and $i_{k,0} = 0$ for all $k \in \mathcal{K}$, the batch state $(p_1, i_{1,1}, \ldots, i_{K,1})$ obtained at the end of the fermentation using the control $u(p_1) \in \mathcal{U}$ is given by

$$\Big(p_1, i_{1,1}, \ldots, i_{K,1}\Big) = \Big(p_0, i_{1,0}, \ldots, i_{K,0} \mid u(p_1)\Big) \tag{5.4}$$

Purity Requirement and Purification Costs: Cost of running a purification step using the chromatography technique $c \in \mathcal{C}$ is denoted by r_c, and consists of setup costs (i.e., calibration, column preparation and washing), material costs (i.e., resins and buffers), equipment and labor costs. The biomanufacturing company also incurs high penalty costs when the purity and yield requirements specified by the end use are not met. Let γ_d and p_d denote the purity requirement and the yield requirement respectively. *The batch purity* γ_t at state $(p_t, i_{1,t}, \ldots, i_{K,t})$ is a quality measure defined by $\gamma_t = \frac{p_t}{p_t + \sum_k i_{k,t}}$ for any $t \in \mathcal{T} \cup \{T\}$. In practice, if the drug is in the final phase of the clinical trials or the end users are humans, then the purity requirement is often very high with $\gamma_d \geq 99.9\%$. The biomanufacturing company incurs high penalty costs c_f when the batch fails to meet the predefined purity requirement γ_d. The customers are typically large pharmaceutical companies conducting clinical trials, and hence a batch that does not meet the purity requirement can not be further utilized for the research and development efforts at the customer's site. Therefore, the customers often do not purchase the batch when it fails to meet the purity requirement. The penalty cost of quality failure c_f could range from company to company, since it includes penalty costs associated with project delays, loss of reputation, cost of disappointing the customers and its impact on future orders, etc.

Yield Requirement and Revenue: At any manufacturing step $t \in \mathcal{T} \cup \{0, T\}$, the revenue obtained from stopping the batch at the state $(p_t, i_{1,t}, \ldots, i_{K,t})$ is defined as follows:

$$r_s(p_t, i_{1,t}, \ldots, i_{K,t}) = \begin{cases} r(p_d) & \text{if } p_t \geq p_d \text{ and } \gamma_t \geq \gamma_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } p_t < p_d \text{ and } \gamma_t \geq \gamma_d \\ -c_f & \text{otherwise.} \end{cases} \quad (5.5)$$

If the batch meets both of the production requirements (i.e, the case $p_t \geq p_d$ and $\gamma_t \geq \gamma_d$ in Equation 5.5) then the customers purchase only the amount ordered p_d, and do not pay for proteins manufactured in excess of their yield requirement, i.e., $\mathbb{1}_{\gamma_t \geq \gamma_d, p_t \geq p_d} r(p_t) = r(p_d)$ where $\mathbb{1}$ is the indicator function and $r(p_d)$ represents the revenue obtained from p_d units of protein sold. On the other hand, if the batch satisfies the purity requirement but not the yield requirement (i.e., the case $\gamma_t \geq \gamma_d$ and $p_t < p_d$ in Equation 5.5), then the biomanufacturing company obtains a revenue $r(p_t)$ associated with the protein amount p_t but also incurs a penalty cost $c_\ell(p_d - p_t)$ due to yield shortages. Typically, the revenue is characterized per unit of protein delivered, and the yield penalty cost is per unit of protein in short, i.e., $r(p_t) = r \times p_t$ and $c_\ell(p_d - p_t) = c_\ell \times (p_d - p_t)$ for all $0 < p_t \leq p_d$, where r is the revenue per unit of protein sold and c_ℓ is the yield penalty cost per unit of protein in short. If the batch does not conform to the minimum purity requirement (i.e., the case $\gamma_t < \gamma_d$ in Equation 5.5) then no revenue is obtained and the biomanufacturing firm incurs a large penalty cost of failure, c_f. Note that $c_f > r(p_d)$, and depending on the amount $(p_d - p_t)^+$, the yield penalty cost $c_\ell(p_d - p_t)^+$ could be as large as the penalty cost c_f.

Upstream Costs: The cost of upstream fermentation operations is captured by $c_u(p_1)$, and represents the cost of labor, equipment, inspection and raw materials (buffers and cell lines) required to obtain $p_1 \in \mathcal{P}$ units of protein by the end of the fermentation process. We assume that $c_u(p_1)$ is nondecreasing in p_1 since additional

resources are required to increase the protein mass. Although we do not assume any specific form for the upstream cost $c_u(p_1)$, it is often a piece-wise linear function in most industry settings (See Section 5.5.3).

The Value Function: We develop a finite horizon Markov decision model that identifies the best choice of the chromatography technique and its pooling window for the downstream purification problem, and the best choice of the protein amount $p_1 \in \mathcal{P}$ for the upstream problem. The objective is to maximize the total expected profit obtained from a batch. Let $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$ denote the value function for the downstream purification problem when there are p_t units of the protein of interest and $(i_{1,t}, \ldots, i_{K,t})$ units of impurity type $k \in \mathcal{K}$ are present in the batch at the beginning of t^{th} purification step, $t \in \mathcal{T}$. At the end of the planning horizon $t = T$, the value function is

$$\mathcal{V}_T(p_T, i_{1,T}, \ldots, i_{K,T}) = r_s(p_T, i_{1,T} \ldots, i_{K,T}) + \mathcal{V}(\Delta). \tag{5.6}$$

For all $t \in \mathcal{T}$, the value function of the downstream purification problem is

$$\begin{aligned} \mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) \;=\; & \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ -r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\big), \\ & r_s(p_t, i_{1,t} \ldots, i_{K,t}) + \mathcal{V}(\Delta) \Big\}, \end{aligned} \tag{5.7}$$

where, the expectation operator is

$$\begin{aligned} & \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\big) \\ = \; & \int_\theta \int_{\psi_1} \ldots \int_{\psi_K} \mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\big) \times \\ & f(\theta|c, w_c) \times g_1(\psi_1|c, w_c) \times \ldots \times g_K(\psi_K|c, w_c) \mathrm{d}\psi_K \ldots \mathrm{d}\psi_1 \mathrm{d}\theta. \end{aligned} \tag{5.8}$$

Combining the upstream protein mass problem at time $t = 0$ with the downstream purification problem at $t = 1$, the value function $\mathcal{V}_0$ of the optimization model is given by

$$\mathcal{V}_0(p_0, i_{1,0}, \ldots, i_{K,0}) = \max_{u(p_1)\in\mathcal{U}} -c_u(p_1) + \mathcal{V}_1\big(p_1, i_{1,1}, \ldots, i_{K,1}|u(p_1)\big). \tag{5.9}$$

Note that $\{\Delta\}$ is an absorbing state with no costs or rewards, and the starting state for the upstream protein mass problem is $p_0 = 0$ and $i_{k,0} = 0$ for all $k \in \mathcal{K}$.

5.3 Structural Properties

In this section, we first investigate the structural properties of the downstream purification problem at $t \geq 1$ and then use these structural characteristics to generate managerial insights for the upstream protein mass problem at $t = 0$. In the reminder of the chapter, we use a discretization scheme to analyze the structural properties of the optimal value function and policies. All proofs are available in the Appendix. Modeling assumptions on the transition probabilities and stopping rewards are summarized in Assumption 5.3.1 and 5.3.2.

Assumption 5.3.1. *Let $w_c^n \in \mathcal{W}_c$ be the n^{th} pooling window on the chromatography technique $c \in \mathcal{C}$ at any step $t \in \mathcal{T}$. For each technique c, the pooling windows $w_c^n \in \mathcal{W}_c$ can be ordered, such that, $\int_0^j f(\theta|c, w_c^n)\mathrm{d}\theta \leq \int_0^j f(\theta|c, w_c^{n+1})\mathrm{d}\theta$ and $\int_0^m g_k(\psi_k|c, w_c^n)\mathrm{d}\psi_k \leq \int_0^m g_k(\psi_k|c, w_c^{n+1})\mathrm{d}\psi_k$ for all $k \in \mathcal{K}$ on a given technique $c \in \mathcal{C}$, $0 \leq j \leq 1$, $0 \leq m \leq 1$. This implies that $\int_\theta \int_{\psi_1} \ldots \int_{\psi_K} \Big\{ f(\theta|(c, w_c^{n+1}))g_1(\psi_1|(c, w_c^{n+1}))\times \ldots \times g_K(\psi_K|(c, w_c^{n+1})) - f(\theta|(c, w_c^n))g_1(\psi_1|(c, w_c^n)) \times \ldots \times g_K(\psi_K|(c, w_c^n))\Big\}\mathrm{d}\psi_K \ldots \mathrm{d}\psi_1\mathrm{d}\theta > 0$.*

Assumption 5.3.2. *Let $p^+ \in \mathcal{P}$ and $p^- \in \mathcal{P}$, such that, $p^+ > p^-$. Then, $r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1})$*

$$\geq r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^n) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^n)$$

for a given impurity level $(i_1, \ldots, i_K)$ *and chromatography technique* $c \in \mathcal{C}$.

Using the index n, Assumption 5.3.1 presents an ordering scheme for the pooling windows $w_c^n \in \mathcal{W}_c$ on a given chromatography technique $c \in \mathcal{C}$. In practice, the first pooling window of this ordering scheme is called as the *yield-aggressive window* and the last window is the *quality-aggressive window* for a given chromatography technique $c \in \mathcal{C}$. Assumption 5.3.1 can be easily validated with industry data. For example, the chromatography outcome in Figure 5.1 suggests that Assumption 5.3.1 often holds in practice. This ordering scheme is used in Theorem 5.3.2 to investigate the structural properties of the optimal pooling windows on a given chromatography technique. For notational convenience, we suppress the subscript t when possible in the reminder of the chapter.

Assumption 5.3.2 corresponds to a broader version of the commonly used supperadditivity property in terms of the states and pooling actions. It implies that the difference in stopping rewards between pooling a larger indexed window w_c^{n+1} and a smaller indexed pooling window w_c^n is higher for a larger protein state $p^+ \in \mathcal{P}$ than a smaller protein state $p^- \in \mathcal{P}$. Although this corresponds to the superadditivity assumption commonly used in the Markov decision literature, Lemma 5.3.1 characterizes under which conditions Assumption 5.3.2 holds in our specific problem setting. The term $r_s(\theta p_t | c, w_c)$ in Lemma 5.3.1 represents the stopping rewards (Equation 5.5) obtained from the following sequence of events: there are p_t units of protein at the purification step $t \in \mathcal{T}$ and the scientist pools the window w_c on the chromatography technique c at the purification step t, and the purification stops at the step $t + 1$. For notational convenience we drop the subscript t in Lemma 5.3.1.

Lemma 5.3.1. *Under the ordering scheme in Assumption 5.3.1 and the stopping cost structure in Equation (5.5), Assumption 5.3.2 always holds for all $t \in \mathcal{T}$ except under the following conditions:*

$$(i) \quad r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) + r(\theta p^-|c, w_c^n) - r_\ell(p_d - \theta p^-|c, w_c^n)$$
$$\leq r(\theta p^-|c, w_c^{n+1}) - r_\ell(p_d - \theta p^-|c, w_c^{n+1}) + r(p_d)$$

$$(ii) \quad r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) - r(\theta p^+|c, w_c^n) + r_\ell(p_d - \theta p^+|c, w_c^n)$$
$$\leq r(\theta p^-|c, w_c^{n+1}) - r_\ell(p_d - \theta p^-|c, w_c^{n+1}) - r(\theta p^-|c, w_c^n) + r_\ell(p_d - \theta p^-|c, w_c^n)$$

(iii) There exists w_c^n and w_c^{n+1}, such that, $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$ but $\frac{(\theta p^-|c, w_c^n)}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^n)} < \gamma_d$ for $\{p^+, p^-\} \in \mathcal{P}$, and $i_k \in \mathcal{I}_k$, $k \in \mathcal{K}$.

Proof See Appendix.

Conditions (i) and (ii) in Lemma 5.3.1 compare the incremental changes in the stopping rewards and lost sale costs when the protein amount p and the ordering index n associated the pooling windows increase. Condition (iii) corresponds the case where higher indexed windows and higher protein amounts in the starting material allow to satisfy the purity requirement but lower indexed windows and lower protein amounts fail to meet the purity requirement when the purification stops at some $t \in \mathcal{T}$. To provide practical insights, we refer to the yield-aggressive and quality-aggressive windows as a special case of this condition. Then, condition (iii) implies that the starting material containing higher protein amount p^+ can meet the purity requirement using the quality-aggressive policy, however, the yield-aggressive policy using less protein amount p^- in the starting material fails to meet the purity requirement when the purification stops at some $t \in \mathcal{T}$. Next, we investigate the

structural properties of the downstream purification problem in Proposition 5.3.1, and then use these characteristics to provide managerial insights for the upstream protein mass problem.

Proposition 5.3.1. (1) *For the downstream purification problem, the value function* $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$ *is nondecreasing in* p_t *and nonincreasing in* $i_{k,t}$ *for all* $k \in \mathcal{K}$ *at* $t \geq 1$.

(2) *At a given impurity level* $(i_{1,t}, \ldots, i_{K,t})$ *and purification step* $t \in \mathcal{T}$*, there exists three protein threshold values,* $\check{p}_t$*,* $\bar{p}_t$ *and* $\hat{p}_t$*, such that,* $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) = -c_f$ *for all* $p_t \leq \check{p}_t$*,* $\mathcal{V}_t(\bar{p}_t, i_{1,t}, \ldots, i_{K,t}) = 0$*, and* $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) = a > 0$ *for all* $p_t \geq \hat{p}_t$ *where* a *is a constant. Note that* $\hat{p}_t \geq \bar{p}_t \geq \check{p}_t$ *for all* $t \in \mathcal{T}$.

Proof See Appendix.

Part (1) of Proposition 5.3.1 indicates that the value function associated with the downstream purification problem never decreases as the protein amount p_t increases, and never increases as the impurity amount $i_{k,t}$ increases for all impurity types $k \in \mathcal{K}$ at any chromatography step $t \in \mathcal{T}$. Note that $\mathcal{T} = \{t : 1, \ldots, T-1\}$. The monotonicity property of the value function is used in part (2) of Proposition 5.3.1 as well as in the subsequent theorems to characterize the structural properties of the optimal policies.

Part (2) of Proposition 5.3.1 reveals managerial insights about the protein amount in the starting material, and leads to important implications for the upstream protein mass problem at $t = 0$. It shows that there exists three protein thresholds values, $\check{p}_t$, $\bar{p}_t$ and $\hat{p}_t$ for a given impurity level $(i_{1,t}, \ldots, i_{K,t})$ at the purification step $t \in \mathcal{T}$. The protein thresholds values $\check{p}_1$, $\bar{p}_1$ and $\hat{p}_1$ associated with the first chromatography step $t = 1$ provide important managerial insights for the upstream protein mass

problem: When the upstream protein mass is $\check{p}_1 < p_1 < \bar{p}_1$, then it implies that the amount of protein obtained from the fermentation process is too less, such that, the biomanufacturing company is expected to incur a financial loss rather than profit. The threshold value $\bar{p}_1$ corresponds to the case where the downstream purification operations result in an expected profit of zero. Lastly, the threshold value $\hat{p}_1$ implies that increasing the upstream protein mass more than $\hat{p}_1$ units does not increase the profitability of purification operations due to the specific yield requirements. Using these insights, Theorem 5.3.1 identifies the batch states $(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t})$ that are expected to result in failure at a downstream purification step $t \in \mathcal{T}$.

Theorem 5.3.1. *The optimal policy has the property that for some* $(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t})$ *where* $\gamma_d > \frac{\check{p}_t}{\check{p}_t + \sum_k i'_{k,t}}$, *it is optimal to stop the purification,* $a_t^*(p_t, i_{1,t}, \ldots, i_{K,t}) = S$, *for all* $p_t \leq \check{p}_t$ *and* $i_{k,t} \geq i'_{k,t}$ *at time* $t \in \mathcal{T}$, $k \in \mathcal{K}$.

Proof See Appendix.

Theorem 5.3.1 shows that there exists a partition on the state space, $p_t \leq \check{p}_t$ and $i_{k,t} \geq i'_{k,t}$ for all $k \in \mathcal{K}$ at $t \in \mathcal{T}$, such that, the biomanufacturing company has no financial incentives for conducting the purification operations when the starting batch state is in this partition. Practically, Theorem 5.3.1 defines some control limits on the protein and impurity amounts to characterize a deficient starting material at the beginning of each chromatography step, i.e., the biomanufacturing company expects to fail to meet the specific customer requirements and incurs large penalty costs when the starting batch state is in this partition. In other words, if failure is inevitable in downstream, it is better to fail earlier than later. Next, Theorem 5.3.2 evaluates the structural properties of the optimal pooling windows for a given chromatography technique.

Theorem 5.3.2. *For all $p_t > \check{p}_t$, if Assumption 5.3.1 and Assumption 5.3.2 hold, then the optimal pooling policy $a_t^*(p_t, i_{1,t}, \ldots, i_{K,t}|c) = w_c^{n*}$ is nondecreasing in p_t for $p_t \geq p_t^*$ at a given impurity level $(i_{1,t}, \ldots, i_{K,t})$ when the chromatography technique $c \in \mathcal{C}$ is used at the purification step $t \in \mathcal{T}$.*

Proof See Appendix.

Theorem 5.3.2 characterizes the structural properties of the optimal pooling windows w_c^{n*} on a chromatography technique $c \in \mathcal{C}$ using Assumptions 5.3.1-5.3.2. Note that Assumption 5.3.1 only provides a stochastic ordering scheme for the pooling windows w_c^n using a chromatography technique c. This assumption holds in practice due to the principles of chromatographic separation described in Section 5.1. However, Assumption 5.3.1 does not guarantee a stochastic ordering for the chromatography techniques. In most industry settings, different chromatography techniques have different affinities to various types of impurities, and hence a stochastic ordering across chromatography techniques is impractical. Therefore, Theorem 5.3.2 focuses on the structural characteristics of the optimal pooling policies for a given chromatography technique c. Theorem 5.3.2 suggests that the scientist adopts policies that are more quality-aggressive as the amount of protein in the starting material increases. As an alternative interpretation of Theorem 5.3.2, we see that the optimal pooling policies at each purification step preserve at least some fraction θ^* of the protein, and the scientist tends to be less concerned in terms of yield losses as the amount of protein in the starting material increases.

Although the downstream purification problem ($t \geq 1$) has some nice structural properties, the value function of the optimization model (including the upstream protein mass decision at $t = 0$ and downstream operations at $t \geq 1$) is not monotonically increasing or decreasing due to the cost trade-offs between upstream and

downstream operations. It is possible to determine some conditions under which the value function of the overall optimization model is monotonic. However, these conditions will be restrictive for industry problems and will not have practical relevance (for example, we use industry data in Section 5.5 to identify the optimal protein mass and the optimal purification policies, and see that the value function $\mathcal{V}_0$ of the optimization problem does not have monotonic properties in most cost settings encountered in practice). However, it is still possible to analytically evaluate the performance of different operating policies for the upstream problem. For example, Theorem 5.3.3 compares and evaluates the performance of popular upstream protein mass decisions, and identifies the conditions under which an upstream policy outweighs its alternatives. To generate managerial insights in Theorem 5.3.3, we use a discretization scheme δ where $\epsilon = p_d/\delta$ represents the specific yield requirement of the customer. In practice, the discretization unit δ often corresponds to a milligram or gram, depending on the specific end use or application.

Theorem 5.3.3. *Let Π_1 and Π_2 be two different upstream operating policies with the corresponding value function, $\mathcal{V}_0^{\Pi_1}$ and $\mathcal{V}_0^{\Pi_2}$, respectively.*

(1) Let Π_1 be the upstream policy $a_0 = u(p_1)$, such that, $p_1 > \hat{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the upstream policy $a_0 = u(\hat{p}_1)$. Then, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$.

(2.a) Let Π_1 be the upstream policy $a_0 = u(p_1)$, such that, $\bar{p}_1 \leq p_1 < \hat{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the policy $a_0 = u(\hat{p}_1)$. Then, $\mathcal{V}_0^{\Pi_1} > \mathcal{V}_0^{\Pi_2}$ if the following condition holds: $c_u(\hat{p}_1) - c_u(p_1) > r(p_d)$.

(2.b) Let Π_1 and Π_2 be the upstream policies identical to part (2.a). Then, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$ if the following condition holds: $c_u(\hat{p}_1) - c_u(p_1) < r(p_d) - r(p_{\epsilon-1}) + c_\ell(p_d - p_{\epsilon-1})$.

(3) Let Π_1 be the upstream policy $a_0 = u(p_1)$, such that, $\check{p}_1 \leq p_1 < \bar{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the upstream policy $a_0 = \bar{p}_1$. Then, $\mathcal{V}_0^{\Pi_1} > \mathcal{V}_0^{\Pi_2}$ if the following condition holds: $c_u(\bar{p}_1) - c_u(p_1) > c_f$.

(4) Let Π_1 *be the upstream policy* $a_0 = u(p_1)$*, such that,* $p_1 \leq \check{p}_1$*,* $p_1 \in \mathcal{P}$*. Let* Π_2 *be the upstream policy* $a_0 = S$*. Then,* $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$*.*

Proof See Appendix.

Part (1) of Theorem 5.3.3 compares two upstream operating policies: Π_1 increases the upstream protein mass more than $\hat{p}_1$ units, whereas Π_2 keeps the upstream protein mass at $\hat{p}_1$ units. It is assumed that the scientists takes the optimal courses of actions for the downstream purification problem. Part (1) shows that the upstream policy Π_2 is always better off than Π_1 due to the specific yield requirements.

Part (2) compares the policy Π_1 that maintains the upstream protein mass in the range $\bar{p}_1 \leq p_1 < \hat{p}_1$ against the policy Π_2 that increases the upstream protein mass up to $\hat{p}_1$ units. Part $(2.a)$ states that the policy Π_1 is better off than Π_2 when the cost of increasing the upstream protein mass is very expensive. For example in practice, it is possible to encounter instances where increasing the protein mass requires excessive re-engineering efforts to improve the productivity of the cell lines. In such cases, increasing the protein mass might not be financially justified compared to the expected revenue from the batch. On the other hand, part $(2.b)$ considers identical polices as part $(2.a)$, but identifies cost configurations where Π_2 is better off than Π_1 in terms of maximizing the total expected profit of the optimization problem. The condition in part $(2.b)$ compares the incremental increases in the upstream operating costs associated with higher protein mass $(\bar{p}_1)$ against the incremental changes in the revenue and lost sale costs associated with lower levels of protein mass $(\bar{p}_1 \leq p_1 < \hat{p}_1)$. Using the discretization scheme δ where $\epsilon = p_d/\delta$, we note that the term $\epsilon - 1$ in part $(2.b)$ corresponds to the protein amount with one unit of lost sales.

Part (3) considers the upstream policy Π_1 where the upstream protein mass is maintained in the range $\check{p}_1 \leq p_1 < \bar{p}_1$ leading to a negative expected profit for the downstream purification problem; whereas the policy Π_2 adopts $\bar{p}_1$ units of protein resulting in zero profit for the downstream purification operations. Part (3) identifies the cost configuration where Π_1 is better off than Π_2 in terms of the expected profit of the optimization problem. The condition in part (3) compares the penalty costs of failures with the cost of increasing the protein mass. In practical context, it represents a cost configuration where the cost of failures are less critical than the cost of efforts required to increase the protein mass. Although not very often, it is possible to encounter such cost settings for proteins in research and development phase. Part (4) shows that it is always better off to abandon the purification and incur large failure costs rather than starting the purification with the protein amount $p_1 \leq \check{p}_1$.

The structural properties and managerial insights derived in this section have been developed based on feedback from our industry partners, and are easy to implement in practice. In Section 5.5, we demonstrate the use of the optimization model and illustrate the structural properties using an industry case study from Aldevron. One of the challenges in solving industry size problems is associated with the curse of dimentionality due to the large state space of the optimization problem. Therefore, we revise the model formulation in Section 5.4 using a state aggregation scheme, and discuss under which conditions the proposed aggregation scheme is exact.

5.4 The Reduced-Dimension MDP Model

A typical industry setting could contain up to 100 different types of impurities with 5-10 candidate chromatography techniques, each having 50-100 pooling window

choices. Although the action space is manageable, the size of the state space increases exponentially in the number of impurities. The state space could easily explode in most industry settings, making the optimization problem challenging to solve and analyze. To address this issue, we use an aggregation scheme to revise the state space, transitions and the value function of the optimization model, and refer to this revised version as the reduced-dimension model. Then, we use the structural insights obtained in Section 5.3 to identify the conditions under which the reduced-dimension model is exact. The state space, transitions, rewards and the value function of the reduced-dimension model are as follows.

State Space: Each impurity state $i_{k,t} \in \mathcal{I}_k$ is a binary variable $i_{k,t} \in \{0, 1\}$ for all $k \in \mathcal{K}$ at time $t \in \mathcal{T} \cup \{T\}$, such that, the state $i_{k,t}$ indicates whether the specific impurity type $k \in \mathcal{K}$ is present in the batch ($i_{k,t} = 1$) or has been completely eliminated ($i_{k,t} = 0$) by the time $t \in \mathcal{T} \cup \{T\}$. Therefore, the impurity state $(i_{1,t}, \ldots, i_{K,t}) \in \mathcal{I}_1 \times \ldots \mathcal{I}_K$ is a 2^K dimensional vector representing which impurities are actually present in the batch and which ones have been completely removed. The state $p_t \in \mathcal{P}$ representing the amount of the protein available in the batch at time $t \in \mathcal{T} \cup \{T\}$ remains the same as Section 5.2. In the reduced-dimension model, the starting state is $(p_1, 1, \ldots, 1)$ and the desired terminal states meeting both the yield and purity requirements are $(p_t, 0, \ldots, 0)$ where $p_t \geq p_d$ for $p_t \in \mathcal{P}$ and $t \in \mathcal{T} \cup \{T\}$.

State Transitions: In the reduced-dimension model, the transitions in the protein state $p_t \in \mathcal{P}$ remain the same as the dynamics in Equation (5.2)-(5.4). However, the transitions associated with the impurity states $(i_{1,t}, \ldots, i_{K,t})$ are simplified using the probability distribution function $P_k(i_{k,t+1}|i_{k,t}, c, w_c)$ for each impurity type $k \in \mathcal{K}$ at time $t \in \mathcal{T} \cup \{T\}$. Note that $\{i_{k,t}, i_{k,t+1}\} \in \{0, 1\}$ for all $\{t, t+1\} \in \mathcal{T} \cup \{T\}$. At each chromatography step $t \in \mathcal{T}$, the function $P_k(i_{k,t+1}|i_{k,t}, c, w_c)$ represents

the probability of achieving the impurity state $i_{k,t+1} \in \{0,1\}$ as a result of the chromatography technique c and the pooling window w_c, given that the state of the impurity type $k \in \mathcal{K}$ before that purification step is $i_{k,t} \in \{0,1\}$. The original model defined in Section 5.2 captures the transition probabilities using the probability distribution functions $g_k(\cdot|c,w_c)d\psi_k$, whereas the transition probabilities in the reduced-dimension model are probability mass functions $P_k(i_{k,t+1}|i_k,c,w_c)$ for each impurity type $k \in \mathcal{K}$, $t \in \mathcal{T}$, $(c,w_c) \in \mathcal{C} \times \mathcal{W}_c$. For each chromatography technique $c \in \mathcal{C}$, the function $P_k(i_{k,t+1}|i_{k,t},c,w_c)$ can be easily determined using the chromatography data (e.g., Figure 5.1). The transition probabilities for the upstream protein mass problem at time $t = 0$ remain the same as Equation (5.4) since the upstream decisions control the amount of protein at the beginning of the first step purification operation.

Since a chromatography technique exploits the difference in the physicochemical properties of each impurity type k as a separation principle, the probability distributions $P_k(i_{k,t+1}|i_{k,t},c,w_c)$ are independently distributed for each impurity type k under the action $(c,w_c) \in \mathcal{C} \times \mathcal{W}_c$. Therefore, the state transitions associated with all impurities $(i_{1,t+1},\ldots,i_{K,t+1}) \in \mathcal{I}_1 \times \ldots \times \mathcal{I}_K$ are captured by the joint probability distribution $P_k(i_{1,t+1}|i_{1,t},c,w_c) \times \ldots \times P(i_{K,t+1}|i_{K,t},c,w_c)$. Note that once an impurity type k is eliminated from the batch, the same impurity k is never regenerated, i.e., $P_k(i_{k,t+1} = 1|i_{k,t} = 0,c,w_c) = 0$ for all $(c,w_c) \in \mathcal{C} \times \mathcal{W}_c$ at $\{t,t+1\} \in \mathcal{T}$. However, the probability $P_k(i_{k,t+1} = 0|i_{k,t} = 1,c,w_c)$ of eliminating a specific impurity type $k \in \mathcal{K}$ is a function of the chromatography technique $c \in \mathcal{C}$, the pooling window $w_c \in \mathcal{W}_c$, and the physicochemical characteristics of the impurity $k \in \mathcal{K}$ as discussed in Section 5.1.

Rewards and the Value Function: The costs and rewards remain the same as Section 5.2. We note that the revenue $r_s(p, i_1, \ldots, i_K)$ obtained from stopping the batch at state $(p, i_1, \ldots, i_K)$ is a special case of Equation (5.5) where $\gamma_d = 100\%$. Therefore, the value function of the downstream purification problem and the upstream protein mass problem remain the same as Equations (5.6)-(5.9) along with the same boundary conditions. We note that the expectation operation in the reduced-dimension model is

$$\mathbb{E}\mathcal{V}_{t+1}\big(p_{t+1}, i_{1,t+1} \ldots, i_{K,t+1} \mid p_t, i_{1,t} \ldots, i_{K,t}, (c, w_c)\big)$$

$$= \int_\theta \sum_{i_{1,t+1}}^{1} \cdots \sum_{i_{K,t+1}}^{1} f(\theta|c, w_c) \times P_1(i_{1,t+1}|i_{1,t}, c, w_c) \times \ldots \times P_K(i_{K,t+1}|i_{K,t}, c, w_c) \times \mathcal{V}_{t+1}\big(\theta p_t, i_{1,t+1}, \ldots, i_{K,t+1}\big)\mathrm{d}\theta$$

Next, Proposition 5.4.1 summarizes the state aggregation scheme used to formulate the reduced-dimension model, and identifies the conditions under which this aggregation scheme is exact.

Proposition 5.4.1. *For each impurity type $k \in \mathcal{K}$, the values of the impurity state $i_{k,t} \in [0, i_{(k,1)}]$ can be aggregated and viewed as a binary variable, $i_{k,t} \in \{0, 1\}$, such that, $i_{k,t} = 0$ represents the case where all amount of the impurity type k has been removed from the batch by the time $t \in \mathcal{T} \cup \{T\}$, and $i_{k,t} = 1$ denotes the case where a positive amount $i_{k,t} \in (0, i_{k,1}]$ of impurity type k is present in the batch at time $t \in \mathcal{T} \cup \{T\}$. This aggregation scheme is exact for 100% purity requirement.*

Proof See Appendix.

Proposition 5.4.1 indicates that the state aggregation scheme is exact for a special instance of the optimization problem where the final purity requirement is $\gamma_d = 100\%$. Proposition 5.4.1 uses the specific characteristics of the transition

probabilities and stopping costs under 100% purity requirement, and indicates that it is sufficient to keep track of which impurity types are present in the batch rather than capturing their corresponding amounts at 100% purity requirement. The idea with the aggregation scheme is simple and intuitive. It relies on the fact that the transition probabilities depend on the type of impurities but not their amounts (i.e., chromatography performs the separation based on the physical and chemical properties of the protein and impurities) and that all impurity types needs to be removed under 100% purity requirement.

Proposition 5.4.1 provides a way to address the curse of dimentionality for a special instance of the optimization problem which is frequently encountered in practice. Excessively high purity requirements are frequently encountered in the biomanufacturing industry, and dealing with these requirements is a significant challenge in practice. It is important to note that when the drugs are used as feed study in the earlier phases of the research and development, the purity requirement could be as low as 85%. However, drugs that are in the later phases of clinical trials, and whose end-users are potentially humans, should abide by very high purity standards. Due to high penalty costs, uncertainties, and manufacturing trade-offs involving both upstream and downstream operations, these excessive purity requirements impose an important layer of challenge in biomanufacturing practices.

All structural characteristics and managerial insights discussed in Section 5.3 are valid for the reduced-dimension model since it is a special case of the optimization model considered in Section 5.2. For example, insights related with the monotonicity of the downstream value function and the protein threshold values in Proposition 5.3.1, the structural characteristics of the pooling windows in Theorem 5.3.2 and the upstream policies considered in Theorem 5.3.3 are not impacted by the aggrega-

tion scheme. However, the threshold values $i'_{k,t}$ associated with each impurity $k \in \mathcal{K}$ in Theorem 5.3.1 lose their managerial implications since the reduced-dimension model does not capture the amount of impurities. Instead, Theorem 5.3.1 can be interpreted such that it is optimal to stop the purification operation when $p_t \leq \check{p}_t$ for a given impurity state $(i_{1,t}, \ldots, i_{k,t})$ using the chromatography technique $c \in \mathcal{C}$ at time $t \in \mathcal{T}$.

5.5 Insights from Industry Case Study

In this section, we demonstrate the application of the optimization model using a case study from Aldevron. We first introduce the problem setting in Section 5.5.1, and then discuss the managerial insights for the downstream purification problem in Section 5.5.2 and the upstream protein mass problem in Section 5.5.3.

5.5.1 Problem Setting

The purification data considered in this case study is obtained from our industry partner, Aldevron. The production requirements are 10 milligrams of protein at 100% purity. The target protein is manufactured for in vitro studies. Scouting runs indicate that the purification process involves 6 candidate chromatography techniques with an average of 80 pooling windows per chromatography technique. The starting material consists of a mixture of 9 different impurities along with the protein of interest. Expected separation outcomes of each available chromatography technique is presented in Figure 5.3. The solid line in Figure 5.3 represents the expected fraction of the protein of interest corresponding to each lane using a specific chromatography technique. The dotted lines are associated with different types of impurities that are available in the starting material. Chromatography techniques differ from each other by their separation outcomes, i.e., the relative locations of impurities and the protein of interest, and their corresponding amount per each lane.

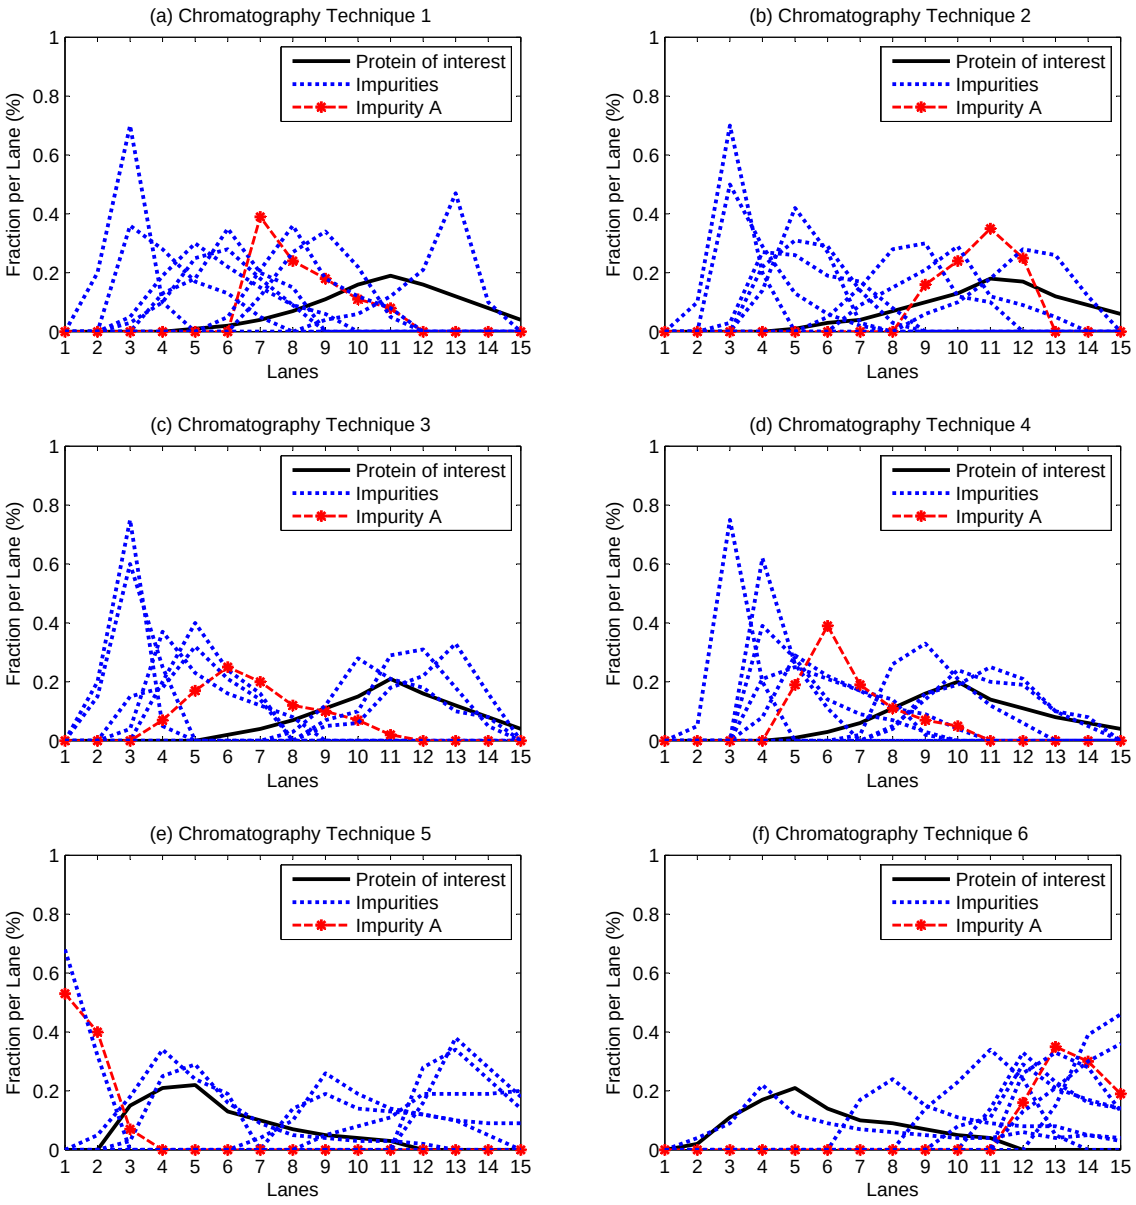

Figure 5.3: Expected separation outcomes of the candidate chromatography techniques

For example, Figure 5.3 tracks a specific impurity type (named as Impurity A) on different chromatography techniques. As the figure shows, the location and amount of Impurity A corresponding to each lane is different for each chromatography technique. Similarly, the separation outcome associated with the protein of interest varies for each technique in Figure 5.3. We note that we mask the information for remaining impurities in Figure 5.3 to protect client confidentiality. However, as Figure 5.3 illustrates, the combinatorial nature of the separation outcomes along with strict quality requirements and high operating costs make this a challenging problem in practice.

Cost and revenue information used in this case study represents industry standards based on feedback from several local biomanufacturing companies (BioWGS, 2014), and is normalized for confidentiality purposes. The normalized values of costs and revenue are as follows: chromatography operating costs are $r_c = \$3$ for each available chromatography technique since they all use similar types of resins and buffers in this case study, failure cost $c_f = \$24$, lost sale cost $c_\ell = \$1.2$ per milligram of protein in short, revenue $r = 1\$$ per milligram of protein produced at 100% purity. Upstream operating costs $c_u(p_1)$ may vary based on the size and number of bioreactors, feeding strategies and harvesting polices adopted to achieve the desired amount of protein. Therefore, we provide a separate discussion on upstream costs and perform sensitivity analysis in Section 5.5.3.

5.5.2 Insights for Purification Decisions

We present the optimal polices and structural insights for the purification problem considered in the case study, and discuss the managerial implications of the analytic results derived in Section 5.3.

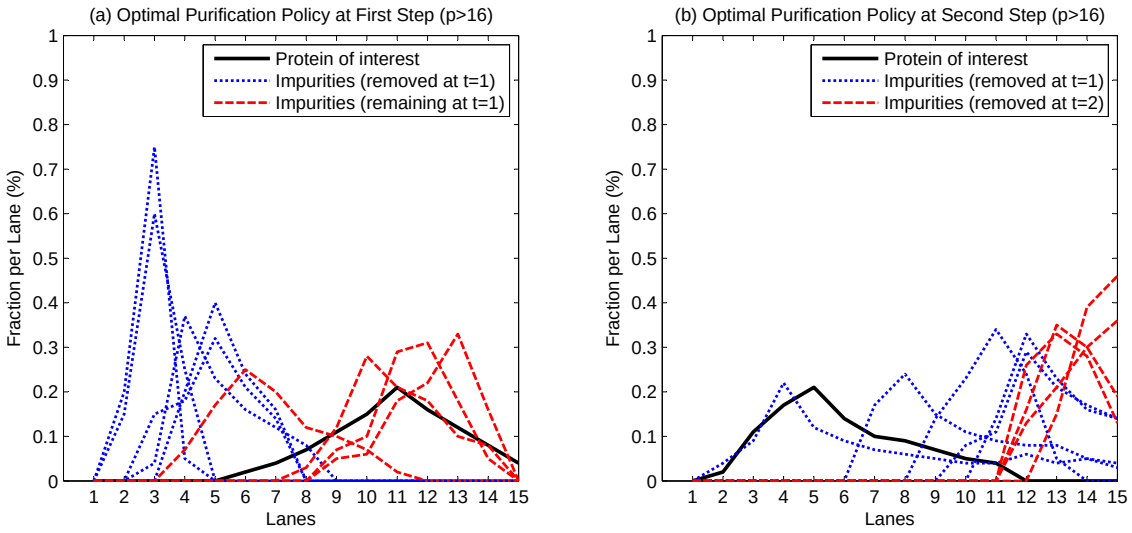

Figure 5.4: Expected outcomes of the optimal policies at each step (for $p_1 > 16$ milligrams)

Overview of the Optimal Policies: Figure 5.4 presents the optimal policies for the downstream purification problem considered in the case study when $p_1 > 16$ milligrams. We note that the optimal policy in Figure 5.4 has two purification steps where the chromatography technique in Figure 5.3 (c) and Figure 5.3 (f) are selected for the first and second step, respectively. Purification policies shown in Figure 5.4 are the optimal ones when the batch obtained from upstream contains $p_1 > 16$ milligrams of protein along with 9 types of impurities. Dotted lines in Figure 5.4 represent the different types of impurities that will be removed using the optimal chromatography technique and pooling window in the first step; whereas dashed lines correspond to the impurities that will be removed using the optimal policy in the second step. The optimal pooling window for the first chromatography step corresponds to the lanes 9 to 15 in Figure 5.4 (a). Then, based on the expected outcome of the first step, we expect to pool the lanes 2 to 11 in the second step, as shown in Figure 5.4 (b). The chromatography technique chosen in the first step is expected to completely eliminate 5 out of 9 different types of impurities but also requires $(1 - \theta) = 12.5\%$ yield losses on average. The chromatography

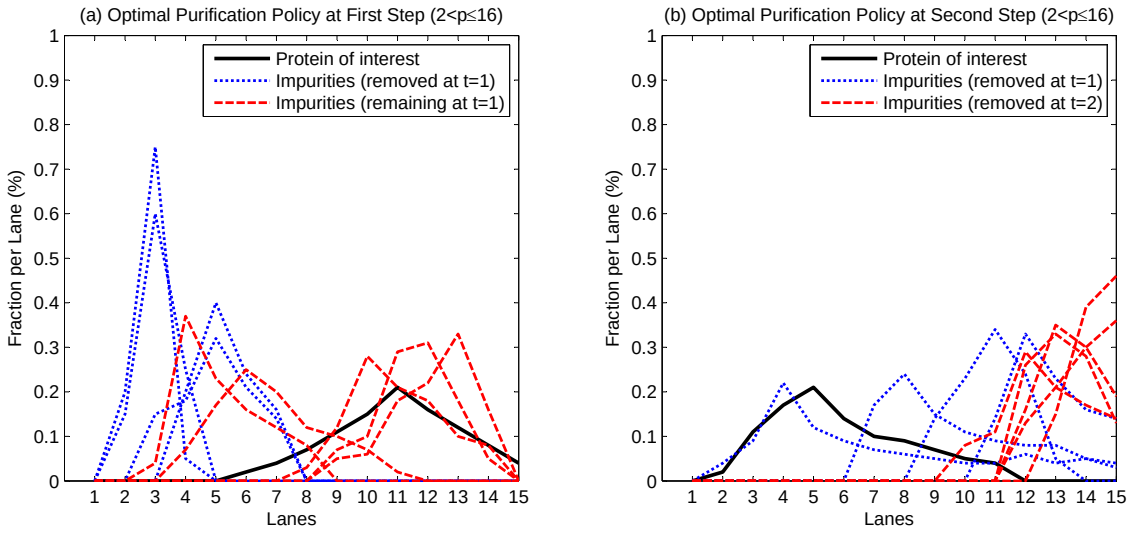

Figure 5.5: Expected outcomes of the optimal policies at each step (for $2 < p_1 \leq 16$ milligrams)

technique selected in the second step is expected to have high separation capability for all remaining impurity types with no significant yield losses expected on average (i.e., $1 - \theta \approx 0\%$). It is interesting to observe that the optimal policy in this example adopts higher yield losses in the first step and lower yield losses in the final step.

The optimal purification policies change as a function of the starting material obtained from upstream operations. For example, Figure 5.5 illustrates the optimal purification policies when the starting material contains $2 < p_1 \leq 16$ milligrams of protein along with 9 different impurities. It is interesting to note that the choice of the chromatography technique remains the same as the ones in Figure 5.4 (although this need not always be the case). However, the optimal pooling window are different for both of the chromatography steps compared to Figure 5.4. In Figure 5.5 (a), the optimal policy pools the lanes 8 to 15 in the first step, eliminating 4 out of 9 impurities. Based on the expected outcome of the first step, we expect that the optimal policy pools the lanes 2 to 9 in the second step as shown in Figure 5.5 (b). Although there is a small difference between the pooling window shown in Figure 5.4

and Figure 5.5, it directly impacts the different types of impurities eliminated at each purification step. For example, the specific impurity starting at lane 2 and ending at lane 9 in Figure 5.5 (a) remains in the batch by the end of the first step when $2 < p_1 \leq 16$ milligrams; although that impurity is completely removed during the first step when $p_1 > 16$ milligrams, as shown Figure 5.4 (a). In this specific case study, we observe that the optimal purification policies always choose the same chromatography technique as the ones in Figure 5.4 but varies in terms of the pooling polices when the protein amount p_1 in the starting material changes but the starting impurity types remain constant. Based on this insight, we next analyze the value function and optimal pooling policies for different values of p_1.

Expected Profit and Optimal Thresholds: Figure 5.6 (a) plots the optimal value function of the downstream purification problem based on the protein amount p_1 involved in the starting material. Note that the value function $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{9,1})$ in Figure 5.6 represents the expected profit at the beginning of the first chromatography step given that the batch obtained from the upstream fermentation operations contains all 9 different types of impurities shown in Figure 5.3. The value function $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{9,1})$ of the downstream purification problem in Figure 5.6 is nondecreasing in the protein amount p_1. The critical protein-thresholds for the starting material are: $\check{p}_1 = 2$ milligrams with $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{9,1}) = -24$ for all $p_1 \leq \check{p}_1 = 2$, $\bar{p}_1 = 15.5$ milligrams with $\mathcal{V}_1(15.5, i_{1,1}, \ldots, i_{9,1}) = 0$, and $\hat{p}_1 = 23$ milligrams with $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{9,1}) = \$14$ for all $p_1 \geq \hat{p}_1 = 23$. As the value function indicates, the biomanufacturing company expects to incur losses in downstream purification operations when the starting material involves less than 15.5 milligrams of protein. In practice, depending on the upstream operating costs, the biomanufacturing firm could reduce or even avoid these losses through optimizing the upstream protein mass decisions (See Section 5.5.3). On the other hand, the downstream purification

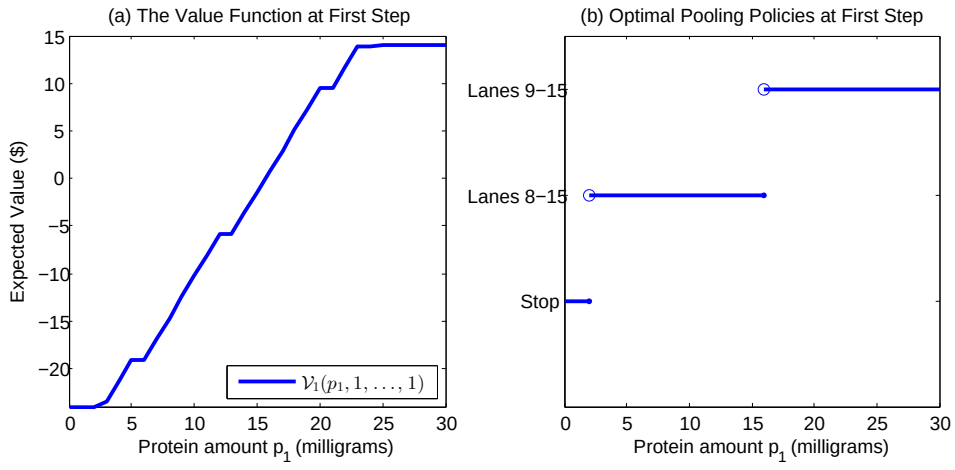

Figure 5.6: Expected value and the optimal pooling policies in the first step as a function of p_1 when $i_{k,1} = 1$ for all $k \in \mathcal{K}$

problem achieves its maximum profit when the starting material obtained from upstream contains more than 23 milligrams of protein. Note that the demand requirement is $p_d = 10$ milligrams but $\hat{p}_1 \geq 23$ milligrams, illustrating the challenge in achieving the purity requirement without significantly compromising on yield.

Figure 5.6 (b) presents the optimal pooling windows as a function of the protein amount p_1 when the chromatography technique indicated in Figure 5.4 (or Figure 5.5) are used. Figure 5.6 (b) shows that if the starting material contains $p_1 \leq 2$ milligrams of protein, then it is optimal to stop the purification despite incurring a large penalty cost for purity failures. On the other hand, if the starting material contains $2 < p_1 \leq 16$ milligrams of protein, then the optimal policy suggest to pool the lanes 8-15 in the first purification step; whereas it is optimal to pool the lanes 9-15 if the starting material contains $p_1 > 16$ milligrams of protein. Figure 5.4 and Figure 5.5 compares the specific impurities that are expected to remain in the batch by the end of the first purification step when these two different pooling polices are adopted. The main trade-off between these two pooling policies involves whether to completely eliminate the specific impurity starting in lane 2 and ending in lane

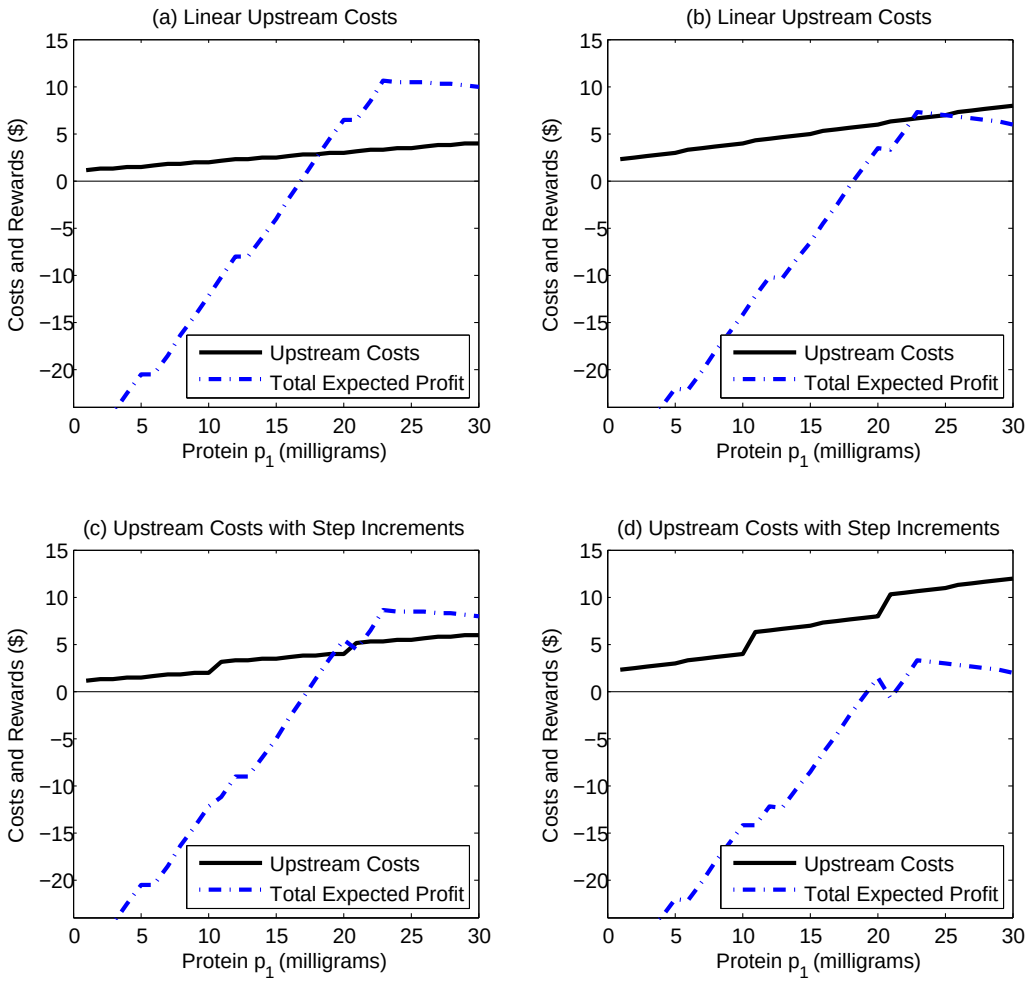

Figure 5.7: Total expected profit as a function of the upstream operating costs $c_u(p_1)$

9 by the end of the first chromatography step or not. The optimal pooling policy shown in Figure 5.6 (b) demonstrates that the optimal policy is likely to tend towards quality-aggressive policies (i.e., pooling the lanes 8-15 vs. lanes 9-15) as the protein amount in the starting material increases for a given chromatography technique. We note that this finding aligns with the structural characteristics of the optimal pooling policies in Theorem 5.3.2.

5.5.3 Insights for Upstream Protein Mass Decisions

We investigate the optimal amount of protein that should be manufactured in the upstream fermentation operations by taking into consideration the interlinked nature of the upstream protein mass and downstream purification decisions. Upstream operating costs often vary as a function of several parameters, such as, the specific type and number of bioreactors used in upstream processes, feeding strategies, cell re-engineering efforts, labor and materials used, harvesting policies, etc. Based on the feedback from our industry partners, we consider four different cost structures for the upstream operations (see Figure 5.7), and investigate the impact of these alternative strategies on the total expected profit.

Figure 5.7 considers four different cost settings for fermentation operations as a function of the protein amount p_1 manufactured in the upstream processes. These upstream cost settings are defined based on industry feedback. Figure 5.7 also presents the total expected profit and the optimal protein mass decisions. The solid lines in Figure 5.7 correspond to the upstream operating costs, and the dashed lines represent the total expected profit of the optimization problem based on both upstream and downstream operations. Figure 5.7 (a) and Figure 5.7 (b) assume a linear cost structure for the fermentation operations as a function of the protein amount p_1. This cost structure corresponds to the case where a single large bioreactor is used to manufacture p_1 units of protein. In such cases, the biomanufacturing company incurs a total operating cost $c_u(p_1) = mp_1 + n$ where m represents the variable costs (i.e., costs of cell culture, buffers, media, process monitoring and control. etc.) and n is the fixed cost of operating the bioreactor (i.e., clean room charges, equipment costs, process analytics, etc). The difference between the upstream costs in Figure 5.7 (a) and (b) is as follows: In Figure 5.7 (a), the total cost of fermentation operation is roughly equivalent to the cost of one purification step when $p_d \leq p_1 \leq \hat{p}_1$;

while the fixed and variable cost considered in Figure 5.7 (b) are twice expensive than the ones in Figure 5.7 (a). Upstream operating costs in Figure 5.7 (b) are roughly equivalent to the total cost of downstream purification operations when $p_d \leq p_1 \leq \hat{p}_1$. We note that the value function of the downstream purification problem is $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{9,1}) = \$14$ for all $p_1 \geq \hat{p}_1 = 23$. When the upstream costs are taken into consideration, the maximum profit expected from the order reduces to $\mathcal{V}_0^* = \$10.57$ in Figure 5.7 (a) and $\mathcal{V}_0^* = \$7.27$ in Figure 5.7 (b). We note that the optimal protein mass is $p_1^* = 23$ milligrams in both Figure 5.7 (a) and (b), which is also identical to $\hat{p}_1$ that maximizes the expected profit of the downstream purification problem in Figure 5.6.

Figure 5.7 (c) and Figure 5.7 (d) consider linear operating costs with step-wise increments every $p_1 = 10$ milligrams. This setting represents the case where batch or fed-batch processes are operated either in parallel or series. For example, upstream operating costs in Figure 5.7 (a) and Figure 5.7 (c) are identical when $0 \leq p_1 \leq 10$ milligrams. However, producing $10 < p_1 \leq 20$ milligrams of protein in Figure 5.7 (c) requires two bioreactor runs, each producing $0 \leq p_1 \leq 10$ milligrams. Therefore, the step increments in Figure 5.7 (c) are mainly associated with the additional bioreactor runs. At high protein amount p_1, we note that the cost structure in Figure 5.7 (a) benefits from the economies of scale when compared to Figure 5.7 (c). Similarly, the upstream operating costs in Figure 5.7 (b) and Figure 5.7 (d) are identical when $0 \leq p_1 \leq 10$ milligrams. However, the manufacturing setting in Figure 5.7 (d) requires an additional bioreactor run every $p_1 = 10$ milligrams while the setting in Figure 5.7 (b) benefits from the economies of scale. Although the cost settings in Figure 5.7 (a) and Figure 5.7 (b) are more favorable than their counterparts in Figure 5.7 (c) and Figure 5.7 (d), they are not always feasible because of the bioreactor capacity constraints. As Figure 5.7 (c) and (d) illustrate, the total expected profit

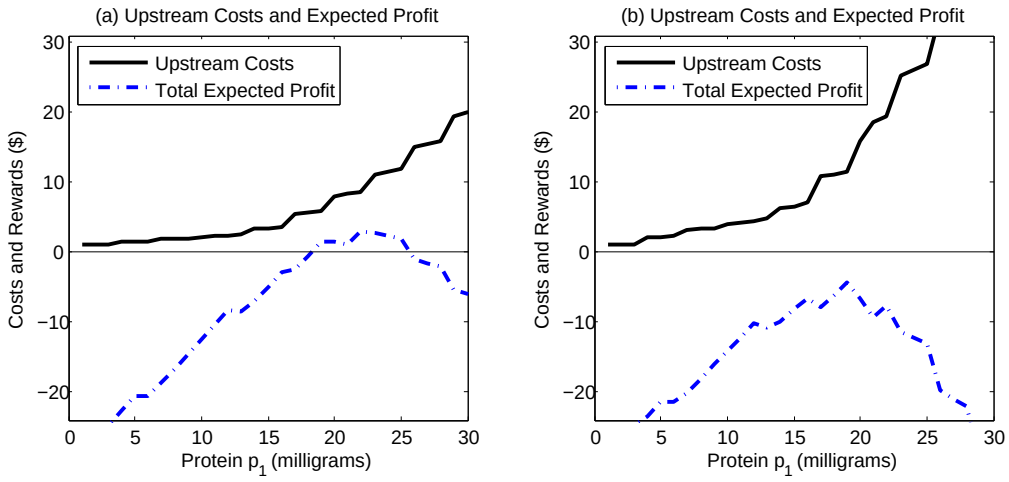

Figure 5.8: Total expected profit when higher protein mass is expensive

considering both upstream and downstream operations does not necessarily have a monotonic structure, although upstream operating costs are nondecreasing in p_1. The maximum expected profit is $\mathcal{V}_0^* = \$8.57$ in Figure 5.7 (c) and $\mathcal{V}_0^* = \$3.27$ in Figure 5.7 (d) with the optimal protein mass of $p_1^* = 23$ in both cases. Note that the abrupt change in the total profit at $p_1 = 21$ in Figure 5.7 (d) is associated with an additional bioreactor run required to manufacture more than 20 milligrams of protein. It is interesting to observe that all upstream costs considered in Figure 5.7 have led to the optimal protein mass $p_1^* = 23$ milligrams which is identical to the critical protein threshold value $\hat{p}_1$ of the downstream purification problem. However, this observation does not necessarily hold for all upstream cost settings, as discussed in Theorem 5.3.3.

Figure 5.8 illustrates two examples of upstream cost settings where $p_1^* < \hat{p}_1$. Upstream operating costs in Figure 5.8 are increasing exponentially in p_1. Compared to their counterparts in Figure 5.7 (a) and (b) where increasing the batch volume simply increases the protein mass of a bioreactor run, the manufacturing settings in Figure 5.8 (a) and (b) represent the cases where additional labor and material, special media and cells are required to boost the protein mass obtained from a single bioreactor run. At high protein mass where the bioreactor capacity is a constraint,

it is possible to encounter such costs structures for engineer-to-order drugs that are in research and development. Although the downstream purification operations have the value function $\mathcal{V}_1(p_1, 1, \ldots, 1) = \$14$ for $p_1 \geq 23$, the total expected profit considering both upstream and downstream operations drops significantly to $\mathcal{V}_0^* = \$3.07$ with $p_1^* = 22$ milligrams in Figure 5.8 (a), and $\mathcal{V}_0 = \$ - 4.33$ with $p_1^* = 19$ milligrams in Figure 5.8 (b). We note that the financial losses expected in Figure 5.8 (b) indicate that the price r charged per unit of protein delivered to the customer does not compensate the efforts in upstream operations. In such cases, Figure 5.8 (b) signals that some managerial actions are needed to turn financial losses into opportunities for profit. For example, these actions could include but not limited to outsourcing the starting material to subscontractors that charge $c_u(p_1)$ similar to those seen in Figure 5.7, re-evaluating the unit price r charged to the customer, identifying opportunities to reduce upstream operating costs, etc. Based on feedback from our partners, we note that the managerial insights obtained from Figures 5.6-5.8 could provide substantial basis to facilitate the way in which biomanufacturers communicate their manufacturing challenges with their customers.

5.6 Conclusions

Protein manufacturing typically involves upstream fermentation operations where the cell culture grows and produces the protein of interest, and the downstream purification operations where the batch of protein is purified by eliminating unwanted impurities (i.e., contaminants, metabolic residues, dead cells, etc). We focus on engineer-to-order proteins that are in the research and development phase. These research and development efforts are typically conducted by a large pharmaceutical company, but the manufacturing operations could be often performed by contract biomanufacturers because of the high failure risks, need for specialized labor and

equipment, and other manufacturing challenges involved in re-engineering proteins. These manufacturing challenges are typically associated with several factors, i.e., limitations in the purification capabilities of available chromatography techniques, randomness in the process outcomes and yield losses, stringent quality requirements for the final batch, expensive operating and penalty costs, interlinked nature of the manufacturing steps, and the strong interaction between the upstream protein mass and downstream purification decisions. There are excellent studies in the literature that contribute to the knowledge behind the biology and chemistry of these operations, but there is a room for improvement for a unified framework that combines the underlying biology and chemistry with the business implications of biomanufacturing decisions (i.e., process economics and financial trade-offs, manufacturing capabilities vs. production requirements, etc.). Because of these manufacturing challenges, current practices typically rely on historical experience and personal expertise to deal with the business implications of biological decisions. In this study, we build a stochastic optimization framework to optimize the profitability of engineer-to-order proteins considering the manufacturing system-level challenges and business implications of biomanufacturing decisions.

The optimization problem is decomposed in two sub-problems: *the upstream protein mass problem* where the optimal amount of protein obtained from the fermentation is identified, and *the downstream purification problem* where the best choices of the chromatography techniques and polling windows are determined. First, we analyze the structural properties of the downstream purification problem. We characterize three critical threshold values, $\check{p}_t, \bar{p}_t, \bar{p}_t$ for the protein amount p_t involved in the starting material at a purification step t. These threshold values provide substantial basis for the biomanufacturing firm to evaluate and quantify the profit (or losses) expected from a specific order. We analyze the structural

properties of the optimal polling windows for a given chromatography technique, and identify the conditions under which the optimal policy is to abandon the purification operations. We then use the structural insight of the downstream purification problem to evaluate the characteristics of the upstream protein mass problem and assess the performance of popular protein mass policies used in practice. The cost of upstream fermentation operations could be a complex function of several design parameters, such as, the type and number of bioreactors used, cell harvesting and feeding strategies, etc. Therefore, we consider several different upstream fermentation configurations typically encountered in practice, and evaluate their optimal value function and their corresponding protein mass decisions.

This study has been conducted in close collaboration with industry through a series of working group sessions (BioWGS, 2014), and the research outcomes have been shared with a broader biotechnology community (BioForward, 2014; Engel, 2014). Since biomanufacturing operations are subject to strict regulations, the community response has been of cautious enthusiasm. However, as more biomanufacturing companies embrace operations research tools and techniques, we believe that the regulatory authorities might mandate the application of such tools to improve biomanufacturing practices. Future work could evaluate supply chain contract designs and pricing decisions for engineer-to-order proteins based on the insights obtained from this optimization framework.

5.7 Appendix: Proofs

Proof of Lemma 5.3.1. We evaluate several cases to assess the conditions under which Assumption 5.3.2 holds. We note that all other cases that are not enumerated in this proof are infeasible scenarios due to the characteristics of the pooling windows described in Assumption 5.3.1.

Case 1:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$.

Based on the stopping rewards considered in this case, the condition in Assumption 5.3.2 indicates

$$\begin{aligned} r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\ - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\ = \; r(p_d) - r(p_d) - r(p_d) + r(p_d) \\ = \; 0. \end{aligned}$$

Hence, Assumption (5.3.2) is satisfied in this case.

Case 2:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and

- $(\theta p^+|c, w_c^n) \geq p_d$ and $\frac{(\theta p^+|c,w_c^n)}{(\theta p^+ + \sum_k \psi_k i_k|c,w_c^n)} \geq \gamma_d$, and
- $\frac{(\theta p^-|c,w_c^n)}{(\theta p^- + \sum_k \psi_k i_k|c,w_c^n)} < \gamma_d$.

In this case, the stopping rewards indicate

$$\begin{aligned} & r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\ & \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^n) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^n) \\ = \;& r(p_d) - r(p_d) - r(p_d) - c_f \\ < \;& 0. \end{aligned}$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 3:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c,w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c,w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^-|c,w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c,w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^+|c,w_c^n)}{(\theta p^+ + \sum_k \psi_k i_k|c,w_c^n)} < \gamma_d$, and
- $\frac{(\theta p^-|c,w_c^n)}{(\theta p^- + \sum_k \psi_k i_k|c,w_c^n)} < \gamma_d$.

The stopping rewards yield to

$$\begin{aligned} & r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\ & \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^n) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^n) \\ = \;& r(p_d) - r(p_d) + c_f - c_f \\ = \;& 0. \end{aligned}$$

Hence, Assumption (5.3.2) holds in this case.

Case 4:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$.

$$
\begin{aligned}
r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
- r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad & r(p_d) - r(\theta p^-|c, w_c^{n+1}) + r_\ell(p_d - \theta p^-|c, w_c^{n+1})^+ - r(p_d) + r(p_d) \\
\geq \quad & 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 5:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \;& r(p_d) - r(\theta p^- | c, w_c^{n+1}) + r_\ell(p_d - \theta p^- | c, w_c^{n+1})^+ \\
& - r(p_d) + r(\theta p^- | c, w_c^{n}) - r_\ell(p_d - \theta p^- | c, w_c^{n}) \\
\geq \;& 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 6:

- $(\theta p^+ | c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+ | c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^- | c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^- | c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+ | c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+ | c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^- | c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \;& r(p_d) - r(\theta p^- | c, w_c^{n+1}) + r_\ell(p_d - \theta p^- | c, w_c^{n+1})^+ - r(p_d) - c_f \\
\leq \;& 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 7:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \;& r(p_d) - r(\theta p^-|c, w_c^{n+1}) + r_\ell(p_d - \theta p^-|c, w_c^{n+1})^+ + c_f - c_f \\
\geq \;& 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 8:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} < \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \;& r(p_d) + c_f - r(p_d) - c_f \\
= \;& 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 9:

- $(\theta p^+|c, w_c^{n+1}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} < \gamma_d$, and
- $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
- r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad & r(p_d) + c_f + c_f - c_f \\
\geq \quad & 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 10:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \; & r(\theta p^+ | c, w_c^{n+1}) - r_\ell(p_d - \theta p^+ | c, w_c^{n+1}) - r(\theta p^- | c, w_c^{n+1}) \\
& + r_\ell(p_d - \theta p^- | c, w_c^{n+1}) - r(p_d) + r(p_d) \\
\geq \; & 0.
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied in this case.

Case 11:

- $(\theta p^+ | c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+ | c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^- | c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^- | c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+ | c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+ | c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^- | c, w_c^{n}) < p_d$ and $\frac{(\theta p^- | c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n})} \geq \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \; & r(\theta p^+ | c, w_c^{n+1}) - r_\ell(p_d - \theta p^+ | c, w_c^{n+1}) - r(\theta p^- | c, w_c^{n+1}) \\
& + r_\ell(p_d - \theta p^- | c, w_c^{n+1}) - r(p_d) + r(\theta p^- | c, w_c^{n}) - r_\ell(p_d - \theta p^- | c, w_c^{n})
\end{aligned}
$$

In this case, Assumption (5.3.2) holds under the conditions specified in Lemma 5.3.1.

Case 12:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
\qquad -r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) - r(\theta p^-|c, w_c^{n+1}) \\
\qquad +r_\ell(p_d - \theta p^-|c, w_c^{n+1}) - r(p_d) - c_f \\
\leq 0
\end{aligned}
$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 13:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \quad & r(\theta p^+ | c, w_c^{n+1}) - r_\ell(p_d - \theta p^+ | c, w_c^{n+1}) - r(\theta p^- | c, w_c^{n+1}) + r_\ell(p_d - \theta p^- | c, w_c^{n+1}) \\
& - r(\theta p^+ | c, w_c^{n}) + r_\ell(p_d - \theta p^+ | c, w_c^{n}) + r(\theta p^- | c, w_c^{n}) - r_\ell(p_d - \theta p^- | c, w_c^{n}) \\
\leq 0 &
\end{aligned}
$$

Hence, Assumption (5.3.2) holds under the conditions specified in Lemma 5.3.1.

Case 14:

- $(\theta p^+ | c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+ | c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^- | c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^- | c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^+ | c, w_c^{n}) < p_d$ and $\frac{(\theta p^+ | c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^- | c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \quad & r(\theta p^+ | c, w_c^{n+1}) - r_\ell(p_d - \theta p^+ | c, w_c^{n+1}) - r(\theta p^- | c, w_c^{n+1}) + r_\ell(p_d - \theta p^- | c, w_c^{n+1}) \\
& - r(\theta p^+ | c, w_c^{n}) + r_\ell(p_d - \theta p^+ | c, w_c^{n}) - c_f \\
\leq 0 &
\end{aligned}
$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 15:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $(\theta p^-|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad & r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) - r(\theta p^-|c, w_c^{n+1}) + r_\ell(p_d - \theta p^-|c, w_c^{n+1}) \\
& + c_f - c_f \\
\geq\ & 0
\end{aligned}
$$

Hence, Assumption (5.3.2) is satisfied.

Case 16:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} < \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) \geq p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
& \qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad & r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) + c_f - r(p_d) - c_f \\
\leq\ & 0
\end{aligned}
$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 17:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} < \gamma_d$, and
- $(\theta p^+|c, w_c^{n}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n+1}) \\
& \qquad -r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K|c, w_c^{n}) \\
= \quad & r(\theta p^+|c, w_c^{n+1}) - r_\ell(p_d - \theta p^+|c, w_c^{n+1}) + c_f - r(\theta p^+|c, w_c^{n}) \\
& \qquad +r_\ell(p_d - \theta p^+|c, w_c^{n}) - c_f \\
\leq & \, 0
\end{aligned}
$$

Hence, Assumption (5.3.2) does not hold in this case.

Case 18:

- $(\theta p^+|c, w_c^{n+1}) < p_d$ and $\frac{(\theta p^+|c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n+1})} \geq \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n+1})} < \gamma_d$, and
- $\frac{(\theta p^+|c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$, and
- $\frac{(\theta p^-|c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k|c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \quad & r(\theta p^+ | c, w_c^{n+1}) - r_\ell(p_d - \theta p^+ | c, w_c^{n+1}) + c_f + c_f - c_f \\
\geq 0 &
\end{aligned}
$$

Hence, Assumption (5.3.2) holds in this case.

Case 19:

- $\frac{(\theta p^+ | c, w_c^{n+1})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n+1})} < \gamma_d$, and
- $\frac{(\theta p^- | c, w_c^{n+1})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n+1})} < \gamma_d$, and
- $\frac{(\theta p^+ | c, w_c^{n})}{(\theta p^+ + \sum_k \psi_k i_k | c, w_c^{n})} < \gamma_d$, and
- $\frac{(\theta p^- | c, w_c^{n})}{(\theta p^- + \sum_k \psi_k i_k | c, w_c^{n})} < \gamma_d$.

$$
\begin{aligned}
& r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) - r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n+1}) \\
& \qquad\qquad - r_s(\theta p^+, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) + r_s(\theta p^-, \psi_1 i_1, \ldots, \psi_K i_K | c, w_c^{n}) \\
= \quad & 0
\end{aligned}
$$

Hence, Assumption (5.3.2) holds in this case. □

Proof of Proposition 5.3.1. (a) We first show that $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$ is nondecreasing in p_t for a given impurity level $i_{k,t}$ for all $k \in \mathcal{K}$ at time $t \in \mathcal{T}$. The proof is done by induction. At the end of the planning horizon $t = T$, the only available action is to stop, $a_T = S$. Clearly, the stopping rewards $r_s(p_T, i_{1,T}, \ldots, i_{K,T})$ given in Equation 5.5 are nondecreasing in p_T. Next, assume by induction hypothesis that

$\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$ is nondecreasing in p_t at time $t \in \mathcal{T}$. Let $p_t^- \in \mathcal{P}$ and $p_t \in \mathcal{P}$, such that, $p_t^- < p_t$ at $t \in \mathcal{T}$. Then,

$$
\begin{aligned}
&\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) \\
&= \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ -r_c + \int_\theta \int_{\psi_1} \ldots \int_{\psi_K} f(\theta|c,w_c) g_1(\psi_1|c,w_c), \ldots, g_K(\psi_K|c,w_c) \times \\
&\quad \mathcal{V}_t\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|c,w_c\big) \mathrm{d}\psi_K \ldots \mathrm{d}\psi_1 \mathrm{d}\theta, \quad r_s(p_t, i_{1,t}, \ldots, i_{K,t}) \Big\} \qquad (5.10) \\
&\geq \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ -r_c + \int_\theta \int_{\psi_1} \ldots \int_{\psi_K} f(\theta|c,w_c) g_1(\psi_1|c,w_c), \ldots, g_K(\psi_K|c,w_c) \times \\
&\quad \mathcal{V}_t\big(\theta p_t^-, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|c,w_c\big) \mathrm{d}\psi_K \ldots \mathrm{d}\psi_1 \mathrm{d}\theta, \quad r_s(p_t^-, i_{1,t} \ldots, i_{K,t}) \Big\} \qquad (5.11) \\
&= \mathcal{V}_t(p_t^-, i_{1,t}, \ldots, i_{K,t})
\end{aligned}
$$

Note that Equation (5.11) follows from Equation (5.2), Equation (5.7) and the induction hypothesis. The proof of the monotonicity of $\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$ in $i_{k,t}$ is entirely analogous and hence omitted.

(b) Note that the bounds on the value function are $-c_f \leq \mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) \leq r(p_d)$ for $p_t \in \mathcal{P}$ and $i_{k,t} \in \mathcal{I}_k$ for all $k \in \mathcal{K}$ at $t \in \mathcal{T}$. Clearly, these bounds follow from the stopping costs in Equation (5.5). Therefore, the existence of the threshold value $\bar{p}_t$, such that, $\mathcal{V}_t(\bar{p}_t, i_{1,t}, \ldots, i_{K,t}) = 0$, is a direct consequence of the monotonicity of the value function in p_t and its bounds. Next, we use induction to investigate the existence of the threshold value $\hat{p}_t$ at $t \in \mathcal{T}$. At time T, the only available action at state $(p_t, i_{1,t}, \ldots, i_{K,t})$ is to stop with rewards $r_s(p_T, i_{1_T}, \ldots, i_{K,T})$. The structure of stopping costs in Equation (5.5) indicates that $r_s(p_T, i_{1_T}, \ldots, i_{K,T})$ is nondecreasing in $p_T \in \mathcal{P}$ for a given impurity level $(i_{1,T}, \ldots, i_{K,T})$, and constant for all $p_T \geq p_d$ that satisfy the purity requirement. Next, at any time $t \in \mathcal{T}$, we observe that the value function is bounded by $r(p_d)$, i.e.,

$$\mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t})$$

$$\begin{aligned}
&= \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ -r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\big), r_s(p, i_{1,t} \ldots, i_{K,t}) + \mathcal{V}(\Delta) \Big\} \\
&\leq \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ -r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | c, w_c\big), r(p_d) \Big\} \qquad (5.12) \\
&= r(p_d).
\end{aligned}$$

Note that Equation (5.12) is a direct consequence of the monotonicity of the value function in p_t and the stopping cost structure in Equation (5.5), i.e., $r_s(p_t, i_{1,t} \ldots, i_{K,t}) = r(p_d)$ for all $p_t \geq p_d$ when the purity requirement is met at $t \in \mathcal{T}$. Hence, based on the monotonicity in part (a), the value function is nondecreasing in $p_t < \hat{p}_t$ and then constant for some $p_t \geq \hat{p}_t$. Next, we refer the reader to the proof of Theorem 5.3.1 for the existence of the threshold value $\check{p}$. Note that the ordering $\hat{p} \geq \bar{p} \geq \check{p}$ follows from the monotonicity of the value function in p. □

Proof of Theorem 5.3.1. Note that Theorem 5.3.1 defines the states $(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t})$, such that, $\gamma_d > \frac{\check{p}_t}{\check{p}_t + \sum_k i'_{k,t}}$ and $p_t \leq \check{p}_t$ for all impurity types $k \in \mathcal{K}$. At time T, the only available action is to stop with rewards $\mathcal{V}_T(p_T, i_{1,T}, \ldots, i_{K,T}) = -c_f$ for all $p_T \leq \check{p}_T$ and $i_k \geq i'_k$, $k \in \mathcal{K}$. Next, it is sufficient to show that if $a_t^*(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t}) = S$ then $a_t^*(p_t, i_{1,t}, \ldots, i_{K,t}) = S$ for $p_t \leq \check{p}_t$ and $i_{k,t} \geq i'_{k,t}$ for all $k \in \mathcal{K}$ at time $t \in \mathcal{T}$. Assume by contradiction hypothesis that $a_t^*(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t}) = S$ but $a_t^*(p_t, i_{1,t}, \ldots, i_{K,t}) = (c, w_c)$ for a given $p_t \leq \check{p}_t$ and $i_{k,t} \geq i'_{k,t}$, where $(c, w_c) \in \mathcal{C} \times \mathcal{W}$. Then,

$$r_S(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t}) > -r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta \check{p}_t, \psi_1 i'_{1,t}, \ldots, \psi_K i'_{K,t} | (c, w_c)\big)$$

and

$$-r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | (c, w_c)\big) > r_S(p_t, i_{1,t}, \ldots, i_{K,t})$$

which together imply

$$0 > \mathbb{E}\mathcal{V}_{t+1}\big(\theta\check{p}_t, \psi_1 i'_{1,t}, \ldots, \psi_K i'_{K,t}|(c, w_c)\big) - \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|(c, w_c)\big). \tag{5.13}$$

Since Theorem 5.3.1 defines the states $(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t})$ such that $\gamma_d > \frac{\check{p}_t}{\check{p}_t + \sum_k i'_{k,t}}$, note that the left hand side of the Equation (5.13) is a direct result of the stopping costs, i.e., $r_S(\check{p}_t, i'_{1,t}, \ldots, i'_{K,t}) = r_S(p_t, i_{1,t}, \ldots, i_{K,t}) = -c_f$. However, we observe that

$$\begin{aligned}
&\mathbb{E}\mathcal{V}_{t+1}\big(\theta\check{p}_t, \psi_1 i'_{1,t}, \ldots, \psi_K i'_{K,t}|(c, w_c)\big) - \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|(c, w_c)\big) \\
&\geq \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|(c, w_c)\big) \\
&\quad -\mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}|(c, w_c)\big) \qquad (5.14) \\
&= 0
\end{aligned}$$

which contradicts Equation (5.13), and hence the proof follows. Note that Equation (5.14) is obtained using the monotonicity of the value function in Proposition 5.3.1. □

Proof of Theorem 5.3.2. Consider two protein amounts, $p_t^+ \in \mathcal{P}$ and $p_t^- \in \mathcal{P}$, such that, $p_t^+ \geq p_t^- > \check{p}_t$. We use backward induction to study the structural properties of the optimal pooling window w_c^* on chromatography technique c. At time T, there is no decision on the pooling windows since the only available action is to stop with rewards described in Equation 5.5. At time $T-1$, assume that $a^*_{T-1}(p^-_{T-1}, i_{1,T-1}, \ldots, i_{K,T-1}|c) = w_c^{n+1}$ but $a^*_{T-1}(p^+_{T-1}, i_{1,T-1}, \ldots, i_{K,T-1}|c) = w_c^n$ at a given impurity level $(i_{1,T-1}, \ldots, i_{K,T-1})$ and chromatography technique $c \in \mathcal{C}$. This implies that,

$$
\begin{aligned}
0 \;>\; & \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^n))g_1(\psi_1|(c,w_c^n)),\dots,g_K(\psi_K|(c,w_c^n)) \\
& \qquad\qquad r_s\big(\theta p_T^-,\psi_1 i_{1,T},\dots,\psi_K i_{K,T}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& + \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& \qquad\qquad r_s\big(\theta p_T^+,\psi_1 i_{1,T},\dots,\psi_K i_{K,T}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& - \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n}))g_1(\psi_1|(c,w_c^{n})),\dots,g_K(\psi_K|(c,w_c^{n})) \\
& \qquad\qquad r_s\big(\theta p_T^+,\psi_{1,T} i_{1,T},\dots,\psi_K i_{K,T}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& - \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& \qquad\qquad r_s\big(\theta p_T^-,\psi_1 i_{1,T},\dots,\psi_K i_{K,T}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \qquad (5.15)
\end{aligned}
$$

We note that, if Assumption 5.3.1 and Assumption 5.3.2 hold, then the term on the right hand side of (5.15) is always positive, which contradicts the inequality in (5.15). Note that if the conditions stated in Lemma 5.3.1 are satisfied such that Assumption 5.3.2 holds, then Equation (5.15) leads to a contradiction.

At time $t \in \mathcal{T}$, assume that $a_t^*(p_t^-, i_{1,t},\dots,i_{K,t}|c) = w_c^{n+1}$ but $a_t^*(p_t^+, i_{1,t},\dots,i_{K,t}|c) = w_c^n$ at a given impurity level $(i_{1,t},\dots,i_{K,t})$ and chromatography technique $c \in \mathcal{C}$. This implies that,

$$
\begin{aligned}
0 \;>\; & \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^n))g_1(\psi_1|(c,w_c^n)),\dots,g_K(\psi_K|(c,w_c^n)) \\
& \qquad\qquad \mathcal{V}_{t+1}\big(\theta p_t^-,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& + \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& \qquad\qquad \mathcal{V}_{t+1}\big(\theta p_t^+,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& - \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n}))g_1(\psi_1|(c,w_c^{n})),\dots,g_K(\psi_K|(c,w_c^{n})) \\
& \qquad\qquad \mathcal{V}_{t+1}\big(\theta p_t^+,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& - \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& \qquad\qquad \mathcal{V}_{t+1}\big(\theta p_t^-,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
= \; & \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} \Big\{ f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& - f(\theta|(c,w_c^{n}))g_1(\psi_1|(c,w_c^{n})),\dots,g_K(\psi_K|(c,w_c^{n}))\Big\} \times \\
& \mathcal{V}_{t+1}\big(\theta p_t^+,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big)\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta \\
& - \int_{\theta}\int_{\psi_1}\dots\int_{\psi_K} \Big\{ f(\theta|(c,w_c^{n+1}))g_1(\psi_1|(c,w_c^{n+1})),\dots,g_K(\psi_K|(c,w_c^{n+1})) \\
& - f(\theta|(c,w_c^{n}))g_1(\psi_1|(c,w_c^{n})),\dots,g_K(\psi_K|(c,w_c^{n}))\Big\} \times \\
& \mathcal{V}_{t+1}\big(\theta p_t^-,\psi_1 i_{1,t},\dots,\psi_K i_{K,t})\mathrm{d}\psi_K\dots\mathrm{d}\psi_1\mathrm{d}\theta
\end{aligned}
\tag{5.16}
$$

Next, we evaluate the term on the right hand side of inequality (5.16) under three cases:

Case 1: $p_{t+1}^+ \geq p_{t+1}^- \geq \bar{p}_{t+1}$

In this case, $\mathcal{V}_{t+1}\big(\theta p_t^-,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big) \geq 0$ and $\mathcal{V}_{t+1}\big(\theta p_t^+,\psi_1 i_{1,t},\dots,\psi_K i_{K,t}\big) \geq 0$. Hence, monotonicity of the value function in Proposition 5.3.1 along with Assumptions 5.3.1-5.3.2 imply that the term on the right hand side of inequality (5.16) is positive, which contradicts (5.16).

Case 2: $\bar{p}_{t+1} > p^+_{t+1} \geq p^-_{t+1}$

In this case, $\mathcal{V}_{t+1}\big(\theta p^-_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}\big) < 0$ and $\mathcal{V}_{t+1}\big(\theta p^+_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}\big) < 0$. Monotonicity of the value function along with Assumptions 5.3.1-5.3.2 lead to a contradiction with inequality (5.16).

Case 3: $p^+_{t+1} \geq \bar{p}_{t+1} > p^-_{t+1}$

In this case, $\mathcal{V}_{t+1}\big(\theta p^-_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}\big) < 0$ but $\mathcal{V}_{t+1}\big(\theta p^+_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t}\big) \geq 0$. Hence, the proof follows from monotonicity of the value function and Assumptions 5.3.1-5.3.2. □

Proof of Theorem 5.3.3. We evaluate the value function of the design problem $\mathcal{V}_0^{\Pi_1}$ and $\mathcal{V}_0^{\Pi_2}$ under the upstream operating policies Π_1 and Π_2. For this purpose, we use the structural properties of the downstream problem defined in Proposition 5.3.1.

(1) Proposition 5.3.1 indicates that the value function associated with the downstream problem's first step operation is $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) = r(p_d)$ when $p_1 > \hat{p}_1$. It is easy to observe that $\mathcal{V}_0^{\Pi_1} = -c_u(p_1) + \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) = -c_u(p_1) + r(p_d)$ where $p_1 > \hat{p}_1$. Hence $\mathcal{V}_0^{\Pi_1}$ is nondecreasing in p_1 when $p_1 > \hat{p}_1$. On the other hand, $\mathcal{V}_0^{\Pi_2} = -c_u(\hat{p}_1) + r(p_d)$. Therefore, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$ since $c_u(p_1) > c_u(\hat{p}_1)$ when $p_1 > \hat{p}_1$.

(2.a) For $\bar{p}_1 \leq p_1 < \hat{p}_1$, we have $\mathcal{V}_0^{\Pi_1} = -c_u(p_1) + \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1})$ and $\mathcal{V}_0^{\Pi_2} = -c_u(\hat{p}_1) + r(p_d)$ due to Proposition 5.3.1. Assume by contradiction that $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$. This implies

$$\begin{aligned} -c_u(\hat{p}_1) + r(p_d) &> -c_u(p_1) + \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) \\ &\geq -c_u(p_1) + \mathcal{V}_1(\bar{p}_1, i_{1,1}, \ldots, i_{K,1}) \\ &= -c_u(p_1) \end{aligned} \tag{5.17}$$

which contradicts the condition given in part (2.a) of the theorem.

(2.b) Following on the part (2.a), we now investigate the conditions under which $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$. Assume by contradiction hypothesis that $\mathcal{V}_0^{\Pi_1} > \mathcal{V}_0^{\Pi_2}$, then, this implies that

$$\begin{aligned} c_u(\hat{p}_1) - c_u(p_1) &> r(p_d) - \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) \\ &\geq r(p_d) - r(p_{\epsilon-1}) + c_\ell(p_{\epsilon-1}) \end{aligned} \tag{5.18}$$

which contradicts the condition specified in part (2.b) of the theorem. Note that $0 \leq \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) < r(p_d)$ in the range $\bar{p}_1 \leq p_1 < \hat{p}_1$, as shown in Proposition 5.3.1. Hence, we use a discretization scheme ϵ to derive managerial insights on the maximum value $\mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1})$ that could be achieved in the interval $\bar{p}_1 \leq p_1 < \hat{p}_1$.

(3) In this case, we have $\mathcal{V}_0^{\Pi_1} = -c_u(p_1) + \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1})$ where $\check{p}_1 \leq p_1 < \bar{p}_1$, $p_1 \in \mathcal{P}$, and $\mathcal{V}_0^{\Pi_2} = -c_u(\bar{p}_1) + \mathcal{V}_1(\bar{p}_1, i_{1,1}, \ldots, i_{K,1}) = -c_u(\bar{p}_1) + 0$ due to Proposition 5.3.1. Assume by contradiction that $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$. This implies that,

$$\begin{aligned} c_u(p_1) - c_u(\bar{p}_1) &> \mathcal{V}_1(p_1, i_{1,1}, \ldots, i_{K,1}) \\ &\geq \mathcal{V}_1(\check{p}_1, i_{1,1}, \ldots, i_{K,1}) \qquad (5.19) \\ &= -c_f \qquad (5.20) \end{aligned}$$

which contradicts the condition specified in Theorem 5.3.3, part (3). Note that $\check{p}_1 \leq p_1 < \bar{p}_1$, and hence Equation (5.19) follows from the monotonicity of the value function.

(4) For all $p_1 \leq \check{p}_1$, we have $\mathcal{V}_0^{\Pi_1} = -c_u(p_1) - c_f$ and $\mathcal{V}_0^{\Pi_2} = -c_f$ due to Proposition 5.3.1. Hence, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$. □

Proof of Proposition 5.4.1. Consider the original problem described in Section 5.2. Let $\mathbb{I}_x$ be a partition of the state space where $\mathbb{I}_x \equiv \{(p_t, i_{1,t}, \ldots, i_{K,t}) | p_t \in \mathcal{P}$ and $i_{k,t} > 0$ for at least one impurity type $k \in \mathcal{K}\}$ at time $t \in \mathcal{T}$. For the proposed aggregation

scheme to be exact, it is enough to show that the value function is piecewise constant over the partition $\mathbb{I}_x$ of the state space. At time $t = T$, the stopping costs under $\gamma_d = 100\%$ purity requirement are:

$r_s(p_T, i_{1,T}, \ldots, i_{K,T})$

$$= \begin{cases} r(p_d) & \text{if } p_T \geq p_d \text{ and } i_{k,T} = 0 \text{ for all } k \in \mathcal{K}, \\ r(p_T) - c_\ell(p_d - p_T) & \text{if } p_T < p_d \text{ and } i_{k,T} = 0 \text{ for all } k \in \mathcal{K}, \\ -c_f & \text{otherwise.} \end{cases} \tag{5.21}$$

The stopping cost structure in Equation (5.21) is constant over the partition $\mathbb{I}_x$ of the state space, i.e., the value function over the partition $\mathbb{I}_x$ is $\mathcal{V}_T(\mathbb{I}_x) = -c_f$.

At time $t = T - 1$, the value function is:

$\mathcal{V}_{T-1}(p_{T-1}, i_{1,T-1}, \ldots, i_{K,T-1})$

$$\begin{aligned} = \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ & r_s(p_{T-1}, i_{1,T-1} \ldots, i_{K,T-1}) + \mathcal{V}(\Delta), \\ & -r_c + \mathbb{E}\mathcal{V}_T\big(\theta p_{T-1}, \psi_1 i_{1,T-1}, \ldots, \psi_K i_{K,T-1} | (c, w_c)\big) \Big\} \end{aligned} \tag{5.22}$$

which is piecewise constant over the partition $\mathbb{I}_x$ since r_c is constant and independent of the state, and also the transition probabilities $g_k(\Psi_k = 0|c, w)$ are independent of the protein p_{T-1} and impurity amount $i_{k,T-1}$ for all $i_k \in I_x$. Hence, the proof follows from the structural characteristics of $\mathcal{V}_T(I_x)$ and Equation (5.21). Next, at any time $t \in \mathcal{T}$, we have

$$\begin{aligned} \mathcal{V}_t(p_t, i_{1,t}, \ldots, i_{K,t}) = \max_{(c,w_c)\in\mathcal{C}\times\mathcal{W}} \Big\{ & -r_c + \mathbb{E}\mathcal{V}_{t+1}\big(\theta p_t, \psi_1 i_{1,t}, \ldots, \psi_K i_{K,t} | (c, w_c)\big), \\ & r_s(p, i_{1,t} \ldots, i_{K,t}) + \mathcal{V}(\Delta) \Big\}, \end{aligned} \tag{5.23}$$

which is piecewise constant over the partition $\mathbb{I}_x$ due to the induction hypothesis, structural properties of the costs r_c and $r_s(p, i_1 \ldots, i_K)$, and the structural character-

istics of the transition probabilities $g_k(\Psi_k = 0|c, w)$. Hence, the proof follows from Bertsekas (2012)(Vol 2, pages 476-479). □

Chapter 6

Future Directions

6.1 Biomanufacturing Supply Chain Contracts

The optimization models developed in Chapters 3, 4 and 5 can be used to refine pricing decisions, reduce risks, and provide more value to contract biomanufacturers and their clients. As a future work, one can investigate how to incorporate the optimization models developed in Chapter 3, 4, and 5 into contract design. More specifically, stochastic games between the large pharmaceutical company (client) and the contract biomanufacturer (supplier) can be analyzed. The large pharmaceutical company outsources the manufacture of an engineered protein, but the manufacturing protocol is often not specified since the protein is uniquely engineered for clinical trials. If the contract biomanufacturer agrees to accept that order, then they often perform several initial test runs (named as scouting experiments) to determine if and how the protein of interest can be manufactured to meet customer specifications.

While negotiating with the clients, the contract biomanufacturer could benefit from the optimal harvesting time analysis in Chapter 3, the decision-zones (risk, target and failure zones) derived in Chapter 4, and the optimization framework

linking the upstream and downstream decisions in Chapter 5. Therefore, future research can incorporate these insights into designing/refining the terms of contracts. For example, it could be interesting to explore the following issues:

Refining pricing schemes for failure and risk zones: Information obtained from scouting runs would enable the contract biomanufacturer to predict batch failures (i.e., the failure zone in Chapter 4) before committing to the production runs. Knowing the failures before starting the production runs would significantly help reduce time delays and penalty costs. For example, if a starting material is identified to be in the failure zone, then the biomanufacturer and the client could take several corrective actions before stating production runs, and hence prevent costly delays and expensive penalties. These actions would include (but not limited to) re-negotiating the production requirements with the client, requesting better starting material from the client, re-manufacturing the starting material in-house, outsourcing the starting material, re-negotiating the price with the client, or terminating the contract. Feedback from our industry partners indicate that a formal and rigorous assessment of risks and batch failures based on scouting runs would significantly facilitate the way how biomanufacturers communicate their challenges with their clients.

Refining pricing schemes based on target zones: Negotiating the pricing scheme after scouting runs could be beneficial when the starting material is in the target zone. For example, if the starting material is in the target zone, the contract biomanufacturer would be able to provide performance guarantees to the client using the information from scouting runs. This would add significant value to the client since they will be able to confidently schedule future experiments for the subsequent phases of clinical trials, and they will also have significant visibility for the successful completion of their order. This would also add significant value to the contract

biomanufacturer since they would be able to provide performance guarantees despite process uncertainties and also enhance customer satisfaction. We believe that the contract biomanufacturer's capability of providing performance guarantees after performing scouting runs would increase their competitive advantage. Therefore, refining the terms of contracts based on these performance guarantees and decision-zones would significantly improve the communication with the clients, hedge against failure risks and expensive penalty costs, provide better visibility for the production pipeline and ensure customer satisfaction.

6.2 Performance Evaluation of the Decision Zones

In Chapter 4, we analyze the pooling window optimization problem related to downstream purification operations. We study the structural properties of the state space, and then partition the state space into distinct decision zones. Note that the failure zone represents all impurity and protein amounts in the starting batch, such that, the biomanufacturing company has no financial incentives for performing the purification operations. Furthermore, the target zone provides a guaranteed performance, such that, the biomanufacturing company can be confident about meeting the customer requirements on purity and yield when the starting material is in that zone. The risk zone, on the other hand, represents all protein and impurity amounts in the starting material such that the biomanufacturing company has high risk of incurring large yield penalty costs or quality failures. The analysis in Chapter 4 characterizes and derives closed form expressions on these zones, and also investigates the optimal purification policies corresponding to each zone.

As a future work, one can investigate the main factors that affect the size and structure of the decision zones in Chapter 4 and quantify their sensitivity to criti-

cal process parameters (i.e., the number of purification steps, yield and purity requirements, operating and penalty costs, and the separation capabilities of the chromatography techniques, etc). Specifically, the following research questions could be interesting to explore:

- What are the main factors that drive the size and structure of the failure zone, risk zone and target zone? How sensitive is each zone to critical process parameters, such as, the number of purification steps, yield and purity requirements, operating and penalty costs, etc?
- How does the size and structure of each zone change as a function of the critical process parameters? What is the impact of the chromatographic separation capabilities on these decision-zones?

6.3 Analysis of Biomanufacturing Systems Under Capacity Constraints

Chapter 3 develops stochastic models to identify the optimal harvesting time for upstream bioreactor operations, and Chapter 4 and Chapter 5 develop optimization models for downstream purification operations. Future work could build optimization models that capture manufacturing system-level constraints, such as, storage capacity constraints, production capacity constraints, scheduling issues, availability of equipment and labor, etc. Future work could analyze the following questions:

- Should the equipment availability in downstream operations be considered while making the harvesting decisions in upstream? For example, should a batch harvested earlier/later than its optimal time identified in Chapter 3 based on the equipment availability in downstream operations? Or should the harvesting

decision take advantage of the storage option and not consider the equipment availability in downstream operations?

- How can one develop production planning and scheduling strategies that consider the interaction between upstream and downstream operations and also take into account the randomness in yield, quality, processing times and costs?

To answer these research questions, future work could involve building simulation models or queuing networks to analyze the manufacturing system-level interactions between upstream and downstream operations. For example, simulation studies or performance evolution of the queuing networks could reveal several managerial insights related with the impact of equipment availability at downstream operation on the harvesting decision at the upstream operations. Furthermore, simulation studies or queuing network models could help evaluate several performance parameters, such as, the total lead time taking into consideration the complex interaction between upstream and downstream operations, throughput of the biomanufacturing system and the amount of the protein delivered to customers, etc. Such performance measures derived by the simulation models or queuing networks could provide important insights in evaluating several scheduling policies based on the equipment availability constraints and randomness in biomanufacturing operations.

Bibliography

Barlow, R. E., F. Proschan. 1965. *Mathematical Theory of Reliability*. John Wiley & Sons, New York.

Bertsekas, D. P. 2012. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, Massachusetts.

Bertsekas, D. P., I. B. Rhodes. 1971. On the minimax reachability of target sets and target tubes. *Automatica* **7**(2) 233–247.

Beuzekom, B., A. Arundel. 2009. Biotechnology statistics. Tech. rep., Organization for Economic Cooperation and Development.

BioForward. 2014. Decision support tools to reduce lead times and costs in biomanufacturing, Bioscience Vision Summit. URL `http://www.bioforward.org/general/custom.asp?page=Biomanufacturing`.

BioWGS. 2014. Biomanufacturing Working Group Sessions, Organized by the University of Wisconsin-Madison and Aldevron. URL `http://www.bioforward.org/event/id/392756/Challenges-and-Opportunities-in-Managing-Biomanufacturing-Operations.htm`.

Boland, Philip J, Frank Proschan. 1983. Optimum replacement of a system subject to shocks. *Operations Research* **31**(4) 697–704.

Brunet, R., G. Guillén-Gosálbez, J. R. Pérez-Correa, J. A. Caballero, L. Jiménez. 2012. Hybrid simulation-optimization based approach for the optimal design of single-product biotechnological processes. *Computers & Chemical Engineering* **37** 125–135.

CBI. 2010. *CBI's BioMAN Program Hosts Summit on Science to Manage Risk in Biomanufacturing Summary: Science to Manage Risk in Biomanufactur-*

ing. URL `http://cbi.mit.edu/bioman-news-events/manufacturing-summit/`. Center for Biomedical Innovation at MIT.

Chhatre, S., R. Francis, K. O'donovan, N. J. Titchener-Hooker, A. R. Newcombe, E. Keshavarz-Moore. 2007. A prototype software methodology for the rapid evaluation of biomanufacturing process options. *Biotechnology and applied biochemistry* **48**(2) 65–78.

Dove, A. 2002. Uncorking the biomanufacturing bottleneck. 777–779.

Elwany, A. H., N. Z. Gebraeel, L. M. Maillart. 2011. Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research* **59**(3) 684–695.

Engel, J. 2014. Madison biotech supplier goes lean to stave off foreign competitors. Xconomy. URL `http://www.xconomy.com/wisconsin/2014/02/19/madison-biotech-supplier-goes-lean-to-stave-off-foreign-competitors/`.

Engell, S., A. Toumi. 2005. Optimisation and control of chromatography. *Computers & Chemical Engineering* **29**(6) 1243–1252.

Farid, S. 2007. Process economics of industrial monoclonal antibody manufacture. *Journal of Chromatography B* **848** 8–18.

Farid, S. 2008. Economic drivers and trade-offs in antibody purification processes. *BioPharm International Supplements* .

Farid, S. 2009. Process economic drivers in industrial monoclonial antibody manufacture. Uwe Gottschalk, ed., *Process Scale Purification of Antibodies*. John Wiley & Sons, 239–261.

Gao, W., S. Engell. 2005. Iterative set-point optimization of batch chromatography. *Computers & Chemical Engineering* **29**(6) 1401–1409.

Gnoth, S., M. Jenzsch, R. Simutis, A. Lubbert. 2007. Process Analytical Technology: Batch-to-batch reproducibility of fermentation process by robust process operational design and control. *Journal of Biotechnology* **132** 180–186.

Gottlieb, Gary. 1982. Optimal replacement for shock models with general failure rate. *Operations Research* **30**(1) 82–92.

Herper, M. 2012. The truly staggering cost of inventing new drugs. URL `http://www.forbes.com/`.

Huang, W., R. G. Askin. 2003. Reliability analysis of electronic devices with multiple competing failure modes involving performance aging degradation. *Quality and Reliability Engineering International* **19**(3) 241–254.

Huuk, T. C., T. Hahn, A. Osberghaus, J. Hubbuch. 2014. Model-based integrated optimization and evaluation of a multi-step ion exchange chromatography. *Separation and Purification Technology* **136** 207–222.

Huynh, K. T., A. Barros, C. Berenguer, I. T. Castro. 2011. A periodic inspection and replacement policy for systems subject to competing failure modes due to degradation and traumatic events. *Reliability Engineering & System Safety* **96**(4) 497–508.

Jang, J. D., J. P. Barford. 2000. An unstructured kinetic model of macromolecular metabolism in batch and fed-batch cultures of hybridoma cells producing monoclonal antibody. *Biochemical Engineering Journal* **4** 153–168.

Jenzsch, M., R. Simutis, A. Luebbert. 2006. Generic model control of the specific growth rate in recombinant Escherichia Coli cultivations. *Journal of Biotechnology* **122** 483–493.

Kalyanpur, M. 2002. Downstream processing in the biotechnology industry. *Molecular Biotechnology* **22** 87–98.

Kawohl, M., T. Heine, R. King. 2007. Model based estimation and optimal control of fed-batch fermentation processes for the production of antibiotics. *Chemical Engineering and Processing* **46** 1223–1241.

Kelley, B. 2009. Industrialization of mab production technology. *mAbs* **1:5** 443–452.

Kharoufeh, J. P., S. M. Cox. 2005. Stochastic models for degradation-based reliability. *IIE Transactions* **37**(6) 533–542.

Klutke, G.A., Y. Yang. 2002. The availability of inspected systems subject to shocks and graceful degradation. *IEEE Transactions on Reliability* **51**(3) 371–374.

Kurt, M., J. P. Kharoufeh. 2010. Monotone optimal replacement policies for a Markovian deteriorating system in a controllable environment. *Operations Research Letters* **38**(4) 273–279.

Langer, E. 2008. Biotech facilities average a batch failure every 40.6 weeks. *Bioprocess International* **6(8)** 28–30.

Langer, E. 2009. The future of biomanufacturing: Towards more and better therapies. URL `http://www.pharmamanufacturing.com/articles/2009/176/`.

Li, W., H. Pham. 2005. Reliability modeling of multi-state degraded systems with multi-competing failures and random shocks. *IEEE Transactions on Reliability* **54**(2) 297–303.

Lienqueo, M.E., J.C. Salgado, O. Giaverini, J.A. Asenjo. 2009. Computer-aided design to select optimal polypeptide tags to assist the purification of recombinant proteins. *Separation and Purification Technology* **65**(1) 86–94.

Liu, J. Z., L. P. Weng, Q. L. Zhang, H. Xu, L. N. Ji. 2003. A mathematical model for gluconic acid fermentation by aspergillus niger. *Biochemical Engineering Journal* **14** 137–141.

Liu, S., A. S. Simaria, S. S. Farid, L. G. Papageorgiou. 2014. Optimising chromatography strategies of antibody purification processes by mixed integer fractional programming techniques. *Computers & Chemical Engineering* .

Liu, S., S. A. Simaria, S. S. Farid, L. G. Papageorgiou. 2013a. Designing cost-effective biopharmaceutical facilities using mixed-integer optimization. *Biotechnology progress* **29**(6) 1472–1483.

Liu, X., J. Li, N. Al-Khalifa, S. H. Hamouda, D. W. Coit, E. Elsayed. 2013b. Condition-based maintenance for continuously monitored degrading systems with multiple failure modes. *IIE Transactions* **45**(4) 422–435.

Ludemann, I., R. Portner, H. Markl. 1994. Effect of NH_3 on the cell growth of a hybridoma cell line. *Cytotechnology* **14** 11–30.

Luus, R. 1993a. Application of dynamic programming to differental-algebraic process systems. *Computers and Chemical Engineering* **17** 373–377.

Luus, R. 1993b. Optimization of fed-batch fermentations by iterative dynamic programming. *Biotechnology and Bioengineering* **41** 599–602.

Maiti, S. M., R. K. Srivastava, M. Bhushan, P. P. Wangikar. 2009. Real time phase detection based online monitoring of batch fermentation processes. *Process Biochemistry* **44** 799–811.

McNeil, B., L. M. Harvey. 2008. *Practical Fermentation Technology*. John Wiley & Sons, New York.

Montoro-Cazorla, D., R. Pérez-Ocón. 2010. System availability in a shock model under preventive repair and phase-type distributions. *Applied Stochastic Models in Business and Industry* **26**(6) 689–704.

Montoro-Cazorla, D., R. Pérez-Ocón. 2011. Two shock and wear systems under repair standing a finite number of shocks. *European Journal of Operational Research* **214**(2) 298–307.

Nakagawa, T. 2007. *Shocks and Damage Models in Reliability Theory*. Springer, London, UK.

Nakagawa, T., S. Osaki. 1975. The discrete Weibull distribution. *IEEE Transactions on Reliability* **24**(5) 300–301.

Narhi, M., K. Nordstrom. 2005. Manufacturing, regulatory and commercial challanges of biopharmaceuticals production: a Finnish perspective. *European Journal of Pharmaceutics and Biopharmaceutics* **59** 397–405.

Neuts, M. F., M. C. Bhattacharjee. 1981. Shock models with phase type survival and shock resistance. *Naval Research Logistics Quarterly* **28**(2) 213–219.

Newland, M., M. N. Kamal, P. F. Greenfield, L. K. Nielsen. 1994. Ammonia inhibition of hybridomas propagated in batch, fed batch and continuous culture. *Biotechnology and Bioengineering* **43**.

Nfor, B. K., P. D.E.M. Verhaert, Luuk A.M. Van der W., J. Hubbuch, M. Ottens. 2009. Rational and systematic protein purification process development: the next generation. *Trends in Biotechnology* **27**(12) 673–679.

Nfor, Beckley K, Tangir Ahamed, Gijs WK van Dedem, Peter DEM Verhaert, Luuk AM van der Wielen, Michel HM Eppink, Emile JAX van de Sandt, Marcel Ottens. 2013. Model-based rational methodology for protein purification process synthesis. *Chemical Engineering Science* **89** 185–195.

Ngiam, S. H., D. G. Bracewell, Y. Zhou, N. J. T. Hooker. 2003. Quantifying process tradeoffs in the operation of chromatographic sequences. *Biotechnology Progress* **19** 1315–1322.

Ngiam, S.H., Y.H. Zhou, M.K. Turner, N.J. Titchener-Hooker. 2001. Graphical method for the calculation of chromatographic performance in representing the trade-off between purity and recovery. *Journal of Chromatography A* **937**(1) 1–11.

Omasa, T., K.I. Higashiyama, S. Shioya, K.I. Suga. 1992. Effects of lactate concentration on hybridoma cell culture in lactate controlled fed batch operation. *Biotechnology and Bioengineering* **39** 556–564.

Özekici, Süleyman. 1988. Optimal periodic replacement of multicomponent reliability systems. *Operations Research* **36**(4) 542–552.

Ozturk, S. S., M. R. Riley, B. O. Palsson. 1992. Effects of ammonia and lactate on hybridoma growth, metabolism, and antibody production. *Biotechnology and Bioengineering* **39**.

Patel, S. D., E. T. Papoutsakis, J. N. Winter, W. M. Miller. 2000. The lactate issue revisited: Novel feeding protocols to examine inhibition of cell proliferation and glucose metabolism in hematopoietic cell cultures. *Biotechnology Progress* **16** 885–892.

Peng, H., Q. Feng, D. W. Coit. 2010. Reliability and maintenance modeling for systems subject to multiple dependent competing failure processes. *IIE Transactions* **43**(1) 12–22.

Peroni, C. V., N. S. Kaisare, J. H. Lee. 2005. Optimal control of fed-batch bioreactor using simulation based approximate dynamic programming. *IEEE Transactions* **13** 786–790.

Petrides, D., C. A. Siletti. 2004. The role of process simulation and scheduling tools in the development and manufacturing of biopharmaceuticals. *Proceedings of the Winter Simulation Conference,* 2046–2051.

Polykarpou, E. M., P. A. Dalby, L. G. Papageorgiou. 2011a. Optimal synthesis of chromatographic trains for downstream protein processing. *Biotechnology Progress* **27**(6) 1653–1660.

Polykarpou, E. M., P. A. Dalby, L. G. Papageorgiou. 2012a. An MILP formulation for the synthesis of protein purification processes. *Chemical Engineering Research and Design* **90**(9) 1262–1270.

Polykarpou, E. M., P. A. Dalby, L. G. Papageorgiou. 2012b. A novel efficient optimisation system for purification process synthesis. *Biochemical Engineering Journal* **67** 186–193.

Polykarpou, Eleftheria M., Paul A. Dalby, Lazaros G. Papageorgiou. 2011b. Optimal synthesis of chromatographic trains for downstream protein processing. *Biotechnology Progress* **27 (6)** 1653–1660.

Puterman, M. L. 1994. *Markov Decision Processes Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York.

Radhakrishnan, T.K., S. Sundaram, M. Chidambaram. 1999. Non-linear control of continuous bioreactors. *Bioprocess Engineering* **20** 173–178.

Rafiee, K., Q. Feng, D. W. David. 2013. Reliability modeling for dependent competing failure processes with changing degradation rate. *IIE Transactions* (Forthcoming).

Rangan, A., D. Thyagarajan, et al. 2006. Optimal replacement of systems subject to shocks and random threshold failure. *International Journal of Quality & Reliability Management* **23**(9) 1176–1191.

Rathore, A. S., P. Latham, H. Levine, J. Curling, O. Kaltenbrunne. 2004. Costing issues in the production of biopharmaceuticals. *BioPharm International* **17(2)** 46–55.

Rathore, A. S., H. Winkle. 2009. Quality by design for biopharmaceuticals. *Nature Biotechnology* **27**.

Salisbury, R. S., D. G. Bracewell, N. J. Titchener-Hooker. 2006. A methodology for the graphical determination of operating conditions of chromatographic sequences incorporating the trade-offs between purity and yield. *Journal of Chemical Technology and Biotechnology* **81**(11) 1803–1813.

Saraph, P. V. 2001. Simulating biotech manufacturing operations: issues and complexities. *Winter Simulation Conference Proceedings* 530–534.

Saraph, P. V. 2003. Shared resource capacity analysis in biotech manufacturing. *Proceedings of the Winter Simulation Conference,* 1247–1250.

Saraph, P. V. 2004. Future of simulation in biotechnology industry. *Winter Simulation Conference Proceedings* 2052–2054.

Saucedo, V. M., M. N. Karim. 1997. Experimental optimization of a real time fed-batch fermentation process using Markov decision process. *Biotechnology and Bioengineering* **55** 317–327.

Saucedo, V. M., M. N. Karim. 1998. Real time optimal feeding in a fermentor using a Markov decision algorithm. *Computers and Chemical Engineering* **22**.

Sheui, S. H., W. S. Griffith. 2002. Extended block replacement policy with shock models and used items. *European Journal of Operational Research* **140**(1) 50–60.

Shivhare, M., G. McCreath. 2010. Practical considerations for DoE implementation in quality by design. *BioProcess International*. 22–30.

Tsao, Y. S., A. G. Cardoso, R. G. G. Condon, M. Voloch, P. Lio, J. C. Lagos, B. G. Kearns, Z. Liu. 2004. Monitoring chinese hamster overy cell culture by the analysis of glucose and lactate metabolism. *Journal of Biotechnology* **118**.

Ulukus, M. Y., J. P. Kharoufeh, L. M. Maillart, D. Bertsekas, S. Shreve, B. Çekyay, S. Özekici, J. Cochran, A. Cox, P Keskinocak, et al. 2012. Optimal replacement policies under environment-driven degradation. *Probability in the Engineering and Informational Sciences* **26**(3) 405–424.

Vasquez-Alvarez, E., M. E. Lienqueo, J. M. Pinto. 2001. Optimal synthesis of protein purification processes. *Biotechnology Progress* **17** 685–696.

Vasquez-Alvarez, E., J.M. Pinto. 2004. Efficient MILP formulations for the optimal synthesis of chromatographic protein purification processes. *Journal of Biotechnology* **110**(3) 295–311.

Wang, Z., H-Z. Huang, Y. Li, N-C. Xiao. 2011. An approach to reliability assessment under degradation and shock process. *IEEE Transactions on Reliability* **60**(4) 852–863.

WEDC. 2014. Company feature: Aldevron. URL `http://inwisconsin.com/entrepreneurs-and-innovators/launch_blog/company-feature-aldevron/`.

Werner, R. G. 2004. Economic aspects of commercial manufacture of biopharmaceuticals. *Journal of Biotechnology* **30** 171–182.

Xing, Z., N. Bishop, K. Leister, Z. J. Li. 2010. Modeling kinetics of a large-scale fed-batch CHO cell culture by Markov chain monte carlo method. *Biotechnology Progress* **26** 208–219.

Yang, J. D., Y. Angelillo, M. Chaudhry, C. Goldenberg. 2000. Achievement of high cell density and high antibody productivity by a controlled-fed perfusion bioreactor process. *Biotechnology and Bioengineering* **69**.

Ye, Z. S., L. C. Tang, H. Y. Xu. 2011. A distribution-based systems reliability model under extreme shocks and natural degradation. *IEEE Transactions on Reliability* **60**(1) 246–256.

Yu, M., Y. Tang, W. Wu, J. Zhou. 2014. Optimal order replacement policy for a phase type geometric process model with extreme shocks. *Applied Mathematical Modelling* .

Zhou, Y., S. H. Ngiam, N. Titchener-Hooker. 2005. Modelling of chromatographic processes and its application in therapeutic protein. *Proceedings of 2005 IEEE Conference on Control Applications*. IEEE, 699–704.