Performance guarantees and optimal purification decisions for engineered proteins

Performance Guarantees and Optimal Purification Decisions for Engineered Proteins

Tugce Martagan

School of Industrial Engineering, Eindhoven University of Technology, Eindhoven, 5612 AZ, The Netherlands,
t.g.martagan@tue.nl

Ananth Krishnamurthy

Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706,
ananth.krishnamurthy@wisc.edu

Peter A. Leland

Aldevron, 5602 Research Park Blvd., Madison, WI 53719, leland@aldevron.com

Christos T. Maravelias

Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, WI 53706,
christos.maravelias@wisc.edu

We investigate protein purification operations conducted by biomanufacturers and pharmaceutical companies as part of their research and development efforts. Purification of these proteins involve unique challenges, such as, balancing the yield and purity trade-offs, dealing with uncertainty in the starting material, and estimating the impact of several interlinked decisions. We develop a Markov decision model and partition the state space into *decision zones* that provide managerial insights to optimize purification operations. We develop practical guidelines to quantify financial risks, and characterize the optimal operating decisions based on specific production requirements. The optimization framework has been implemented at Aldevron, a contract biomanufacturer specializing in proteins, and has resulted in 25% reduction in the total lead times and 20% reduction in the costs of protein purification operations on average.

Key words: Stochastic control, reachability of target set, guaranteed performance, biomanufacturing

1. Introduction

Recent advances in biomanufacturing have led to novel proteins used in the treatment of cardiovascular diseases, autoimmune disorders and cancer. In this paper, we focus on protein manufacturing operations in the pharmaceutical research and development. 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 biomanufacturing firm as part of its research and development efforts. Manufacturing of this protein at the biomanufacturing firm would then involve specialized fermentation operations followed by several purification operations. Our scope in this paper is the protein purification operations. In practice, purification of engineered proteins could be challenging for several reasons. For example, individual proteins have unique chemical and physical properties, and their end use sets constraints on the production methods needed to satisfy

Table 1   Current and proposed workflow for purification development

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

rigorous approval processes. Further, a purification order often has an associated yield requirement (i.e., the desired amount of the protein of interest) and a 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 accept yield shortages at a penalty cost as long as the purity requirement is satisfied.

Table 1 presents a typical workflow to purify an engineered protein. Upon the receipt of an order, the scientist at the biomanufacturing firm starts performing scouting runs at small scale. Scouting runs represent a set of experiments where the scientist collects data about the purification attributes of this protein on several alternative *chromatography techniques*. Once the performance of available chromatography techniques have been identified with respect to the protein of interest, the scientist performs validation runs. The role of the validation runs is to mitigate risks and quantify the yield and purity 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 development. 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 develop an optimization model that uses the information obtained from scouting runs and identifies the optimal purification polices for the production runs, thereby reducing costs and lead times. As shown in Table 1, reducing the time spent in the validation runs 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 in practice, 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. Randomness 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 impurity, then the specific requirements on the final purity and yield might never be satisfied, even though the biomanufacturer 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 paper, we provide an optimization framework that quantifies the risks and costs in protein purification operations and answers the following questions: (*i*) For a given starting material, can the biomanufacturer use the information from scouting runs to determine whether the purity and yield requirements specified by the customer are achievable at all? Can we provide performance guarantees for achieving these specific production requirements? Can we develop guidelines on the starting material to predict the batch failures? *(ii)* How easy or complex is the purification process likely to be, based on the starting material and purification capabilities of the chromatography techniques? How can the total profit be maximized for each purification order? By answering these questions using an optimization framework, we believe that biomanufacturing firms can significantly improve their profitability and reduce their lead times in protein purification operations.

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 chromatography step. 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 optimization model provides 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.
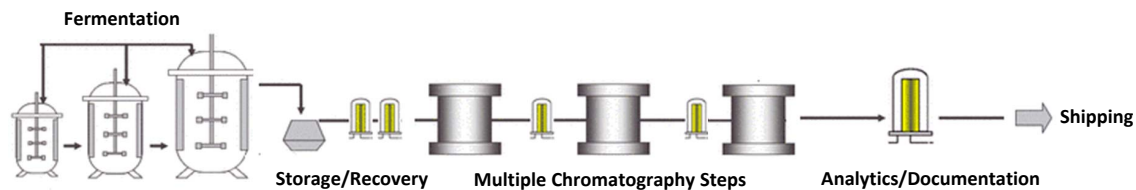
Figure 1   **Typical manufacturing stages in biomanufacturing**

This research is an outcome of an ongoing multi-year collaboration with Aldevron (2013-2016). 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 in protein purification, as discussed in Sections 7-8. Our research outcomes have also been shared and validated with a larger biomanufacturing community (BioWGS 2014, BioForward 2016). Through industry implementation, we observe that the optimization framework has the potential for significantly reducing 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 combines the knowledge from chemical engineering and stochastic modeling to derive guidelines that improve industry practices.

The remainder of the paper is organized as follows. Section 2 provides a background on purification operations and introduces the trade-offs and challenges. We develop a mathematical model in Section 3, and analyze its structural properties in Section 4 and Section 5. We present a state aggregation and action elimination scheme in Section 6. We discuss the implementation of the optimization model in Section 7 and Section 8, and provide concluding remarks in Section 9.

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 1). The primary output of fermentation is a batch mixture that includes the protein of interest and significant amount of unwanted impurity derived from the host cells or fermentation medium. After fermentation, this batch must be purified using multiple chromatography steps (typically, 2 to 6 steps) based on specific production requirements. The objective of each chromatography operation is to separate the protein of interest from unwanted impurity to achieve the desired purity level. In this paper, 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 trade-offs and operational challenges in practice.
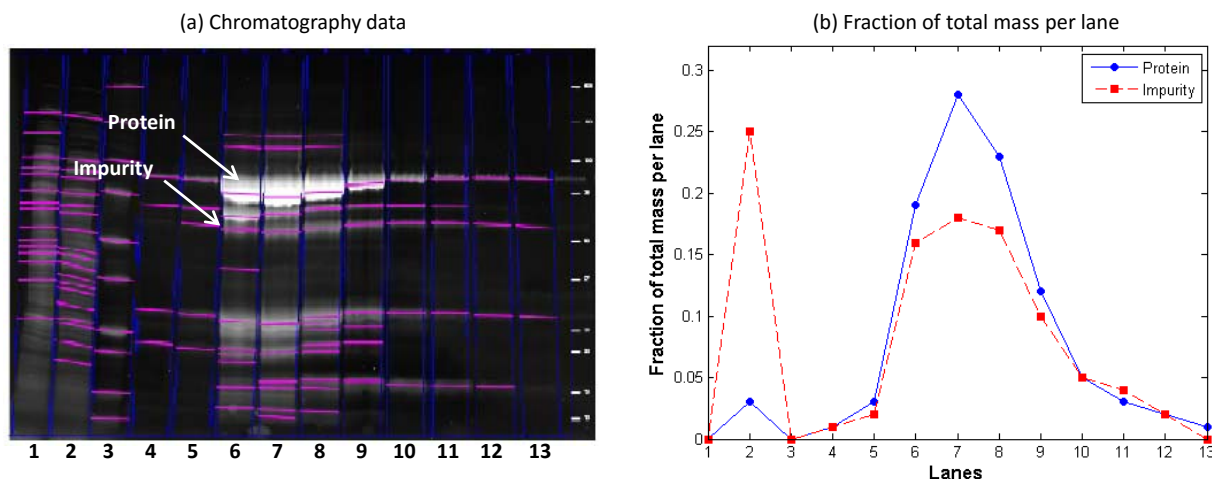
Figure 2   An example of chromatography output

2.1.   Chromatography Operations

Chromatography is one of the most common but also most challenging operations in biomanufacturing (Polykarpou et al. 2011, Liu et al. 2014). The objective of chromatography operations is to separate the protein of interest from unwanted impurities to meet a predetermined *purity requirement* specified by the end use or application. *Purity* represents the ratio of the total amount of protein of interest to the total amount of both protein and impurity contained in a batch. Purity requirement is defined by the end use or application of the purified protein. For example, a protein used in the treatment of a disease must be highly pure (i.e., 99.9% purity), whereas a protein used for a feed study in biomanufacturing 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. Most purification projects in research and development involve 2 to 3 chromatography steps, but in some cases they could require up to 6 chromatography steps. Each chromatography step often takes 6 to 8 hours or more, depending on the physico-chemical characteristics of the starting material, production requirements, and the process conditions.

Figure 2 (a) presents an example of chromatography data. This example uses the differential affinity of proteins to divalent metal ions as the separation principal. The y-axis in Figure 2 (a) denotes the molecular size. Each column on the x-axis is called a *lane*, and can be thought as

equivalent to a discrete time interval (i.e., close to 1 minute in practice). Each lane consists of some amount (mass) of the protein of interest and impurity. The size of the white pixels in Figure 2 (a) is correlated with the amount of the protein of interest and impurity contained in each lane.

The chromatography data in Figure 2 (a) is often translated into Figure 2 (b). The y-axis in Figure 2 (b) represents the expected fraction of the total mass of protein and the expected fraction of the total mass of impurity at each lane. For example, the second lane in Figure 2 (b) is expected to contain 3% of the total protein mass and 25% of the total impurity mass. Note that the sum of expected fractions at a specific lane does not necessarily need to be equal to 1. However, the sum of expected fractions over all lanes need to be equal to 1 for the protein of interest, and similarly for the impurity. The first and third lane in Figure 2 represent the load and the marker lane, respectively. These lanes serve as a reference point for data analytics, and do not represent the actual outcome of a chromatographic separation. Therefore, both of the protein and impurity fractions in these lanes are plotted as zero in Figure 2 (b).

2.1.1. Yield and Purity Trade-offs The scientist performing the chromatography step must decide which lane to *'pool'*. In this example, the scientist can choose to pool any consecutive lanes between lanes 4 and 13. For instance, lanes 5-10, lanes 6-9, and lanes 7-8 are examples of candidate *pooling windows.* In practice, the scientist often confronts with a challenging trade-off between the yield and purity at each chromatography step (Ngiam et al. 2003, Muller-Spath et al. 2013, Subramanian 2014). For example, consider lanes 6-9 and lanes 7-8 in Figure 2. Lanes 6-9 are expected to yield 16.4 milligrams (mg) of protein and 18.3 mg of impurity, leading to 47% purity. On the other hand, lanes 7-8 are expected to result in 10.2 mg of protein and 10.5 mg of impurity, leading to 49% purity. If the scientist pools the lanes 6-9, she collects larger fraction of protein along with larger fraction of impurity. However, if she pools the lanes 7-8, she collects smaller fraction of impurity at the expense of smaller fraction of protein. This illustrates one of the main trade-offs related with yield and purity of a chromatography step. Depending on the outcome of a chromatography step, the scientist could make decisions regarding the chromatography technique and the pooling window for each chromatography steps. In fact, identifying the sequence of chromatography techniques itself is a separate optimization problem. However, we consider purification settings where this sequence is predetermined based on scouting runs, and focus on the problem of selecting the best pooling window at each chromatography step.

2.1.2. Challenges in Practice Main challenges in chromatography operations can be summarized as follows: (1) *Yield and purity trade-offs.* Each order is associated with predetermined yield and purity requirements. However, the scientist often needs to compromise on the protein yield to achieve the desired purity level. (2) *Engineered proteins.* Each order is unique such that

the scientist re-engineers and manufactures each order for the first time. This requires to evaluate each order independently, unlike mass production. (3) *Uncertainty.* The amount of protein and impurity obtained at each chromatography step involves uncertainty due to the underlying biology and chemistry of the purification process, (4) *Interlinked decisions.* Purification involves multiple chromatography steps in series. The output of each step affects the possibility of successfully attaining the yield and purity requirements. (5) *Starting batch.* The starting material is 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 live cells introduces variability in the amount of protein and impurity obtained from the fermentation operation. These in turn affect the subsequent chromatography 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 and the action space increases exponentially in the number of chromatography steps.

2.2. Prior Work

Relevant prior work belongs to two categories: the literature on dynamic programming and the literature on chromatography operations. Bertsekas and Rhodes (1971), Puterman (1994), Bertsekas (2012) provide excellent overview of the dynamic programming approach. However, applications of the stochastic optimization methodologies in the context of protein purification are limited in the existing operations research literature. Therefore, the rest of our literature review mainly focuses on relevant optimization models from the chemical and biological engineering literature.

Several studies quantify the trade-off between purity and yield in chromatographic separation using physicochemical data of protein mixtures (Ngiam et al. 2001, 2003, Salisbury et al. 2006, Kraattli et al. 2013, Muller-Spath et al. 2013). Such process trade-offs obtained from chromatography data are often used as input for optimization models. For example, Vasquez-Alvarez et al. (2001) develop two mixed integer linear programming (MILP) models to determine the optimal synthesis of multi-step chromatography operations. One of the proposed MILP models focuses on minimizing the number of chromatography steps to achieve a desired purity level, whereas the other MILP model maximizes the final purity of a batch. Vasquez-Alvarez and Pinto (2003) extend this work by incorporating the yield and purity trade-off in a MILP model that identifies the optimal choice of chromatography techniques to achieve specific purity and yield requirements. Polykarpou et al. (2011) consider the problem of identifying the optimal pooling window, and develop a MILP model that minimizes the number of chromatography steps through optimal starting and finishing cut points. The proposed optimization model is then extended in Polykarpou et al. (2012) using

approximation techniques to overcome computational challenges. Note that aforementioned studies aim to minimize the number of chromatography steps but do not account for costs related to shortage (lost sale) and failures. Furthermore, MILP models assume that the outcome of a chromatography step is certain but our industry collaborators indicate that they all have a degree of uncertainty which challenges the decisions in practice.

In addition to the MILP models, there are studies that investigate cost reduction strategies in chromatography operations. For example, Simaria et al. (2012) propose a multi-level metaheuristic procedure that minimizes production costs by optimal column sizing decisions. Similarly, a mixed integer nonlinear programming model is developed by Liu et al. (2014) to minimize production costs through optimal chromatography sizing and sequencing decisions (i.e., the optimal choice of resins, column diameter, etc.). However, we observe that these studies focus on deterministic settings, and develop strategies that reduce only production costs through optimal process design (i.e., column sizing) or facility design decisions (Papageorgiou et al. 2001, Lakhdar et al. 2005).

Stochastic models for chromatography operations typically involve simulation of the biological and chemical dynamics to predict the yield and purity outcomes (Zhou et al. 2005, Chhatre et al. 2007, Nfor et al. 2009). There are only a few studies that capture the risks and uncertainties in biomanufacturing operations to determine cost reduction strategies. For example, Farid et al. (2007) develop a hierarchical framework for modeling biomanufacturing operations using a simulation software. The proposed framework is used to evaluate different alternatives on facility design, process design and capacity allocation decisions based on several performance parameters, such as, operating cost, lead time, and resource utilization. Similarly, Chhatre et al. (2006) develop a simulation model to assess the sensitivity of product yields and process times to several chromatography parameters such as affinity flow rate and matrix volume. Lim et al. (2006) and Martagan et al. (2016) develop models to evaluate the risks and production economics of fermentation systems. In the context of stochastic models, we observe that existing studies largely focus on simulation models to evaluate the risks and costs in biomanufacturing. Such simulation studies are not equipped to answer critical research questions identified in Section 1, namely those related to providing performance guarantees and determining optimal policies.

In this paper, we formulate a Markov decision model to optimize pooling windows and stopping decisions in chromatography operations. We investigate the structural characteristics of the model, and establish novel guidelines for practitioners. These guidelines provide a formal procedure to assess the starting material based on uncertainties and costs involved in chromatography operations. To our knowledge, such guidelines and performance guarantees have not been studied in the literature. We demonstrate the application of the model through implementation at Aldevron.

3. The Model

In this section, we formulate a finite horizon Markov decision model for purification decisions.

Decision epochs: $\mathcal{T} = \{t : 1, \ldots, T-1\}$ denotes the set of decision epochs. Each decision epoch $t \in \mathcal{T}$ represents the beginning of a chromatography step. Note that there are finite number of chromatography steps, and its sequence is predetermined based on scouting runs. We let T be the terminal step that corresponds to the end of the planning horizon. At step T, no chromatography operations are performed, and the batch is either shipped to the customer or scrapped.

States: The state space is defined as $\mathcal{X} = \mathcal{P} \times \mathcal{I} \cup \Delta$. The state $p_t \in \mathcal{P}$ denotes the amount of protein of interest available in the batch at the beginning of t^{th} chromatography step. Similarly, state $i_t \in \mathcal{I}$ represents the amount of impurity at the beginning of t^{th} chromatography step. Note that the starting material of the purification project is $(p_1, i_1) \in \mathcal{P} \times \mathcal{I}$ and corresponds to the protein and impurity amounts obtained from fermentation operations. A batch has the maximum possible amount of protein and impurity at the beginning of the first chromatography step, and hence $0 \leq p_t \leq p_1$, $0 \leq i_t \leq i_1$ at $t \in \mathcal{T} \cup T$. In practice, the amount of protein and impurity often ranges between milligrams and grams depending on the end use or application. The state Δ is defined as the stopping state for the project, and represents a batch which is either ready to be shipped to the customer or scrapped. The state Δ is an absorbing state with no rewards.

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 state (p_t, i_t) at the beginning of chromatography step $t \in \mathcal{T}$. 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 $\mathcal{W}_t = \{(i, \ldots, j) \subseteq \mathcal{L}_t : j = i + k, i = \{1, \ldots, L_t\}, k = \{0, 1, \ldots, L_t - i\}\}$. 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 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 step 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. 2011). We adopt these models to identify the amount of protein and impurity 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 is carried over the next chromatography step $t + 1$ when

the pooling window w_t is selected, implying that the remaining amount of impurity was eliminated by the chromatography step t. 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 $w_t \in \mathcal{W}_t$, $t \in \mathcal{T}$. If the scientist chooses the pooling window w_t at the chromatography step t and the realization of the random fraction is $\psi_t|w_t$, then the impurity state at the beginning of the chromatography step $t+1$ is

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

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 chromatography step $t+1$ when the pooling window w_t is selected, implying that the remaining amount of the protein was eliminated during that chromatography step. 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 $w_t \in \mathcal{W}_t$, $t \in \mathcal{T}$. Therefore, if the scientist pools the window w_t at the chromatography step t and the realization of the random fraction associated with the protein is $\theta_t|w_t$, then the protein state at the beginning of the chromatography step $t+1$ is

$$p_{t+1} = (\theta_t|w_t)p_t. \tag{2}$$

The probability density functions $f_t(\cdot|w_t)$ and $g_t(\cdot|w_t)$ and their finite support can be different for each chromatography step $t \in \mathcal{T}$, depending on physico-chemical characteristics of the proteins and impurities, and specific chromatography technique used at each step. 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. 2011). Chromatography techniques mainly differ in terms of how they exploit these unique characteristics to separate proteins from impurities. In practice, the probability density functions $f_t(\cdot|w_t)$ and $g_t(\cdot|w_t)$ and their finite support can be determined from scouting data collected as per recommended guidelines (Ellison and Willams 2012, ISO21748 2010).

One of the key performance measures for a chromatography technique is its *purification capability* under a pooling window. The purification capability is determined based on the fractions of protein and impurity that remain in the batch after performing a chromatography step. For example, $(\theta_t^u, \psi_t^l|w_t)$ represents the best possible purification capability of the chromatography step t under the pooling window w_t. Whereas, $(\theta_t^l, \psi_t^u|w_t)$ denotes the worst possible purification capability of the chromatography step t under the pooling window w_t. We define $(\bar{\theta}_t, \bar{\psi}_t|w_t)$ as the mean purification capability of chromatography step t under pooling window w_t. The purification capabilities are used to generate performance guarantees in Section 4 and Section 5. Note that the system transitions from state $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ to the stopping state Δ when the purification project is terminated at

chromatography step $t \in \mathcal{T}$ or at the terminal step T. This ensures that the decision making process is finalized since the state Δ is an absorbing state with no rewards.

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 part of the production requirement specified by the end use or application. Customers would not purchase the batch if it does not satisfy the purity requirement (i.e., $\gamma_t < \gamma_d$). Therefore, only batches that meet the purity requirement (i.e., $\gamma_t \geq \gamma_d$) are shipped to the customers. The purity requirement could range from 85% to 99.9% based on specific characteristics of each order. Biomanufacturing firms often do not receive additional rewards for attaining purity levels higher than the minimum requirement γ_d.

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 chromatography step t based on the type of resin, buffer, column, and other specifications of chromatography techniques used at each step (Farid 2007, 2009). If the batch does not meet the minimum purity requirement after the completion of a purification project, a penalty cost of failure c_f is incurred. The failure cost c_f could vary from company to company, and represents penalties associated with lost sales, loss of reputation and its impact on future orders.

Yield Requirement and Stopping Costs: In addition to the purity requirement γ_d, each order has a predetermined yield requirement p_d specified by the end use or application. At the completion of a purification project, the final reward obtained from a batch depends on its purity and yield. Let $r(p_t)$ be a function that represents the revenue obtained from p_t units of protein, and $c_\ell(p_d - p_t)$ be a function that denotes the yield penalty cost in case $p_t < p_d$. Then, the final reward $r_S(p_t, i_t)$ obtained from stopping the purification process at state $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ is

$$r_S(p_t, i_t) = \begin{cases} -c_f & \text{if } \gamma_t < \gamma_d, \\ r(p_d) & \text{if } \gamma_t \geq \gamma_d \text{ and } p_t \geq p_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } \gamma_t \geq \gamma_d \text{ and } p_t < p_d, \end{cases} \tag{3}$$

for $t \in \mathcal{T}$ when $a_t(p_t, i_t) = S$, and for $t = T$.

Equation (3) indicates that if the purity requirement is not achieved (i.e., $\gamma_t < \gamma_d$), the biomanufacturer incurs a penalty cost c_f. If the final batch satisfies the purity requirement and contains more protein than the yield requirement (i.e., $\gamma_t \geq \gamma_d$ and $p_t \geq p_d$), then the biomanufacturer obtains a fixed revenue, $r(p_d)$, regardless of the protein amount manufactured in excess. This means that the client does not pay for proteins produced in excess of the yield requirement. However, if the batch meets the purity requirement but fails to achieve the yield requirement (i.e., $\gamma_t \geq \gamma_d$

and $p_t < p_d$), then the biomanufacturer obtains a revenue $r(p_t)$ which is a function of the protein amount produced, and incurs a yield penalty cost $c_\ell(p_d - p_t)$ for the amount of protein in short. The function $r(p_t)$ is non-decreasing in p_t, and $c_\ell(p_d - p_t)$ is non-increasing in p_t for $p_t < p_d$ and $\gamma_t \geq \gamma_d$ at chromatography step $t \in \mathcal{T} \cup T$. The term $r(p_t) - c_\ell(p_d - p_t)$ can be negative depending on the amount of protein in short when $\gamma_t \geq \gamma_d$ and $p_t < p_d$ at $t \in \mathcal{T} \cup T$. Note that $r(p_d) < c_f$ and $0 < r(p_d) - \sum_{t=1}^{T-1} c_t$. The stopping state Δ is an absorbing state with no rewards, $r_S(\Delta) = 0$. This implies that if the purification project is terminated at chromatography step $t \in \mathcal{T}$ or at the end of the planning horizon T, then the stopping costs described in Equation (3) are incurred, and the system transitions to the stopping state Δ where the decision making process terminates.

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 $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$:

$$\mathcal{V}_t(p_t, i_t) = \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \underset{\theta_t, \psi_t | w_t}{\mathbb{E}} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\}, \text{ for } t = \{1, \dots, T-1\}, \tag{4}$$

$$\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T), \tag{5}$$

where the expectation is based on the probability distribution $f(\cdot|w_t)$ and $g(\cdot|w_t)$, i.e.,

$$\underset{\theta_t, \psi_t | w_t}{\mathbb{E}} \mathcal{V}_{t+1}(p_t \theta_t, \psi_t i_t) = \int_{\psi_t^l | w_t}^{\psi_t^u | w_t} \int_{\theta_t^l | w_t}^{\theta_t^u | w_t} f_t(\theta_t | w_t) g_t(\psi_t | w_t) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi. \tag{6}$$

Note that $\mathcal{V}_t(\Delta) = 0$ for $t \in \mathcal{T} \cup T$. Let π_t^* denote the optimal purification policy from step $t \in \mathcal{T}$ until the end of planning horizon T. If w_t^* maximizes the right hand side of Equation (4) for each (p_t, i_t) and t, the policy $\pi_1^* = \{w_1^*, \dots, w_T^*\}$ is optimal (Puterman 1994).

The purity and yield requirements are not modeled as explicit constraints in the mathematical model. Instead, they are captured through the stopping cost structure in Equation (3), which leads to a more realistic and flexible approach than imposing constraints on the final state. In practice, customers often understand the challenges involved in biomanufacturing operations, and they would be willing to compromise on the yield requirement at a certain penalty cost. In alignment with practice, the model allows shortages at the cost of penalty $c_l(\cdot)$, and implicitly captures the stringent purity requirement via the failure cost c_f. The model allows flexibility in stopping the purification project without meeting the customer requirements, which aligns with the notion of 'failing earlier than later'. We do not consider the discount factor in our model formulation because purification operations represent a short-term planning horizon compared to the overall protein manufacturing lead time. In this setting, discounting the value function could lead to a bias in decision making. Further, a finite horizon optimization model for each batch is reasonable since the motivating industry setting involves contract biomanufacturers where each batch represents an engineered protein uniquely made for a customer order.

4.   Structural Analysis of the State Space: Decision Zones

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 decision zones (namely *failure zone* in Section 4.1, *target zone* in Section 4.2 and *risk zone* in Section 4.3) and establish performance guarantees based on these zones. To do so, we first establish some important structural properties of the value function in Proposition 1.

PROPOSITION 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 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 1 holds for any probability density functions $f_t(\cdot)$ and $g_t(\cdot)$ 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.

4.1.   Failure Zones

We analyze the minimum purity and yield required at the beginning of chromatography step $t \in \mathcal{T}$, such that, the biomanufacturer has no financial incentives to perform the purification if the batch does not meet these minimum requirements.

THEOREM 1.   *[Failure Zone] The optimal policy has the property that for some $(p'_t, i'_t) \in \mathcal{P} \times \mathcal{I}$ where $\frac{p'_t}{(p'_t + i'_t)} < \gamma_d$, the optimal action is $a^*_t(p_t, i_t) = S$ for all $p_t \leq p'_t$ and $i_t \geq i'_t$ at chromatography step $t \in \mathcal{T}$.*

Proof See Appendix.

Theorem 1 indicates that the biomanufacturer should stop the purification and scrap the batch, if the starting material does not satisfy some requirements on the amount of protein and impurity needed prior to running the chromatography step. More specifically, Theorem 1 shows that there exists some threshold values (p'_t, i'_t) at chromatography step $t \in \mathcal{T}$, such that, it is optimal to stop the purification if the state (p_t, i_t) of the starting material is $p_t \leq p'_t$ and $i_t \geq i'_t$. Note that Theorem 1 does not require any specific knowledge of the probability density functions $f_t(\cdot)$ and $g_t(\cdot)$; and only uses the monotonic behavior that follows from Equations (1)-(2), i.e., p_{t+1} is non-decreasing in p_t for a given $\theta_t|w_t$ at chromatography step $t \in \mathcal{T}$. Based on the insights from Theorem 1, we define the *failure zone* $\mathbb{F}_t$ of chromatography step $t \in \mathcal{T}$ as a set of states (p'_t, i'_t) where the optimal action is to stop the purification for all $p_t \leq p'_t$ and $i_t \geq i'_t$ despite $\frac{p'_t}{p'_t + i'_t} < \gamma_d$. Therefore,
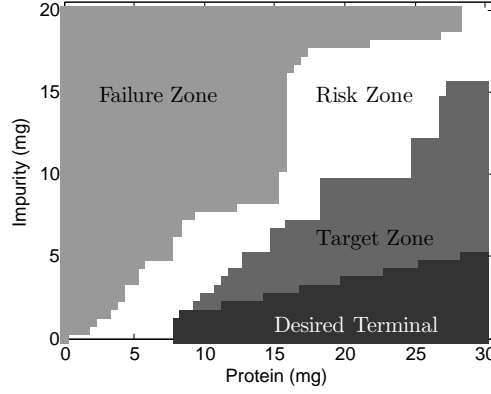
Figure 3   An example of the zones for a chromatography step using industry data

$\mathbb{F}_t = \{(p'_t, i'_t) \in \mathcal{P} \times \mathcal{I} : a^*_t(p_t, i_t) = S \text{ for all } p_t \leq p'_t \text{ and } i_t \geq i'_t; \frac{p'_t}{(p'_t + i'_t)} < \gamma_d\}$ is the failure zone at chromatography step $t \in \mathcal{T}$. Figure 3 illustrates an example of the failure zone using industry data.

Next, Proposition 2 characterizes the failure zone $\mathbb{F}_t$ at chromatography step $t \in \mathcal{T}$ in terms of the costs and the best purification capabilities of chromatography steps $t, t+1, \ldots, T-1$.

PROPOSITION 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 chromatography step* $t \in \mathcal{T}$, *if either of the following conditions hold:*

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

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

$$(ii) \quad r\big(p_t \prod_{w_j} (\theta^u_j | w_j)\big) - c_l\big(p_d - p_t \prod_{w_j} (\theta^u_j | w_j)\big) < c_t - c_f \quad \text{and} \quad i_t \leq p_t \frac{1 - \gamma_d}{\gamma_d} \prod_{w_j} \frac{(\theta^u_j | w_j)}{(\psi^l_j | w_j)}$$

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

Proof See Appendix.

Condition (i) in Proposition 2 represents the case where the purity requirement lies outside the purification capability of all possible pooling windows $w_j \in \mathcal{W}_j$ available in the subsequent chromatography steps $j = t, \ldots, T-1$. Condition (ii) corresponds to the case where none of the purification strategies w_j available in the subsequent steps $j = t, \ldots, T-1$ provide adequate financial incentives for continuing the purification process. Note that Proposition 2 provides a performance guarantee using the best possible realizations of the purification outcomes $(\theta^u_t, \psi^l_t | w_t)$ across all pooling windows w_t at all chromatography steps $t \in \mathcal{T}$. This analysis yields a conservative classification of states in $\mathbb{F}_t$, i.e., it identifies the states where failure or abandoning the purification is the best action. In practice, the value of the failure zone is to acknowledge the failure prior to

committing resources. Identifying the failure zone can help the biomanufacturer convince the client that failure is unavoidable, and might help redefine the expectations from the project.

4.2. Target Zones

We characterize a particular subset of the state space called the *target zone* $\mathbb{T}_t$ at chromatography step $t \in \mathcal{T}$. The target zone $\mathbb{T}_t$ represents a set of states $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$, such that, if $(p_t, i_t) \in \mathbb{T}_t$ at the beginning of chromatography step $t \in \mathcal{T}$, then both of the yield and purity requirements can be achieved with certainty by the end of the planning horizon T given that the optimal pooling windows are chosen at each chromatography step $t, \ldots, T-1$ (See Section 5 for a discussion on the optimal pooling windows). Such guaranteed performance measures are critical in most biomanufacturing applications to justify customer expectations and guard against manufacturing inefficiencies. To characterize the target zone $\mathbb{T}_t$ at chromatography step t, we use recursion based on the worst possible outcomes corresponding to each pooling window $w_t \in \mathcal{W}_t$ at chromatography steps $t, t+1, \ldots, T-1$. First, we define the terminal zone $\mathbb{S}$ of the purification project in Definition 1.

DEFINITION 1. The terminal zone $\mathbb{S}$ corresponds to the set of protein and impurity states that meet both of the yield and purity requirements specified by the end use or application, i.e.,

$$\mathbb{S} = \left\{ (p_t, i_t) \in \mathcal{P} \times \mathcal{I} : p_t \geq p_d, \; \frac{1 - \gamma_d}{\gamma_d} p_t \geq i_t \right\} \text{ at } t \in \mathcal{T} \cup T. \tag{7}$$

It follows that, if the batch is in the terminal zone at the beginning of chromatography step $t \in \mathcal{T}$, i.e., $(p_t, i_t) \in \mathbb{S}$, then the purification can be stopped at the chromatography step t, and the batch can be shipped to the customer since it satisfies both of the yield and purity requirements. Clearly, based on the terminal zone $\mathbb{S}$ in Definition 1, the target zone $\mathbb{T}_T$ at the end of the planning horizon T is $\mathbb{T}_T = \left\{ (p_T, i_T) \in X : p_T \geq p_d, \; \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T \right\}$. Next, we let $X = [0, p_1] \times [0, i_1]$ and characterize the target zone $\mathbb{T}_t$ at each chromatography step $t \in \mathcal{T}$ in Proposition 3.

PROPOSITION 3. *The target zone $\mathbb{T}_t$ at chromatography step $t \in \mathcal{T}$ is defined as*

$$\mathbb{T}_T = \left\{ (p_T, i_T) \in X : p_T \geq p_d, \; \frac{1 - \gamma_d}{\gamma_d} p_T \geq i_T \right\}, \tag{8}$$

$$J_{t,w} = \left\{ (p_t, i_t) \in X : p_t = \frac{p_{t+1}}{\theta_t^l | w}, i_t = \frac{i_{t+1}}{\psi_t^u | w}, (p_{t+1}, i_{t+1}) \in \mathbb{T}_{t+1} \right\} \text{ for } w \in \mathcal{W}_t, \tag{9}$$

$$\mathbb{T}_t = \bigcup_{w \in \mathcal{W}_t} J_{t,w} \text{ for } t = 1, \ldots, T-1. \tag{10}$$

Proof See Appendix.

The target zone $\mathbb{T}_t$ in Proposition 3 is obtained recursively using the worst-case outcomes $(\theta_t^l, \psi_t^u | w_t)$ for all pooling windows $w_t \in \mathcal{W}_t$ available in chromatography steps $t, t+1, \ldots, T-1$. This enables to establish performance guarantees. Therefore, Proposition 3 implies that if a batch

(p_t, i_t) belongs to the target zone $\mathbb{T}_t$ at the beginning of chromatography step t, then there *exists* a sequence of actions that will guarantee that both the yield and purity requirements can be achieved by the end of the planning horizon T. Figure 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 3, and provide important managerial insights (Bertsekas and Rhodes 1971):

(i) At the beginning of t^{th} chromatography step, if the starting material (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 end of the planning horizon T.

(ii) The target zone provides some threshold values $(\hat{p}_t, \hat{i}_t)$ on the starting material (p_t, i_t) at chromatography step t, such that, if $(\hat{p}_t, \hat{i}_t) \in \mathbb{T}_t$ then $(p_t, i_t) \in \mathbb{T}_t$ for all $p_t \geq \hat{p}_t$ and $i_t \leq \hat{i}_t$ at $t \in \mathcal{T}$.

The characteristics listed above have practical implications for managing chromatography operations. For example, item (i) indicates that target zones provide performance guarantees in terms of achieving both of the purity and yield requirements. Item (ii) indicates that the target zone has a threshold-type structure, and hence can be easily interpreted and implemented in practice. Due to limitations and inherent uncertainties of chromatography operations, such performance guarantees are valuable for both the biomanufacturer and its client. For example, the potential for eventual success provides visibility in the production pipeline and ensures customer satisfaction. Most customers recognize the challenges involved in biomanufacturing operations and highly value the analysis of such performance guarantees.

4.3. Risk Zones and Bounds on the Value Function

As a direct consequence of the target zone $\mathbb{T}_t$ and failure zone $\mathbb{F}_t$ at chromatography step $t \in \mathcal{T}$, we define the risk zone, $\mathbb{R}_t = \{(p_t, i_t) \in \mathcal{P} \times \mathcal{I} : (p_t, i_t) \notin \mathbb{F}_t \text{ and } (p_t, i_t) \notin \mathbb{T}_t\}$ at $t \in \mathcal{T}$. The risk zone includes all states $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ that are neither in the target zone $\mathbb{T}_t$ nor in the failure zone $\mathbb{F}_t$ at the beginning of chromatography step $t \in \mathcal{T}$. Next, we characterize the bounds on the optimal value function $\mathcal{V}_t^*(p_t, i_t)$ based on the zones at each chromatography step $t \in \mathcal{T}$ as follows:

$$\mathcal{V}_t^*(p_t, i_t) = -c_f \text{ for all } (p_t, i_t) \in \mathbb{F}_t, t \in \mathcal{T}. \tag{11}$$

$$\sum_{j=t}^{T-1} -c_j + r(p_d) \leq \mathcal{V}_t^*(p_t, i_t) \leq r(p_d) \text{ for all } (p_t, i_t) \in \mathbb{T}_t, t \in \mathcal{T}. \tag{12}$$

$$-c_f \leq \mathcal{V}_t^*(p_t, i_t) \leq r(p_d) \text{ for all } (p_t, i_t) \in \mathbb{R}_t, t \in \mathcal{T}. \tag{13}$$

Note that Equation (11) is a direct consequence of Theorem 1. Similarly, the cost bounds on the target zone in inequality (12) follow from Proposition 3 and the stopping cost structure in Equation (3). The cost bounds on the risk zone in inequality (13) follow from the monotonicity of the value function in Proposition 1 and the definition of the failure and target zones. These

bounds provide managerial insights to quantify the risks and costs of states within each zone. For example, a stating material which an element of the failure zone will result in large penalty cost $-c_f$; whereas a starting material in the target zone can lead to a large reward up to $r(p_d)$. On the other hand, a batch that is in the risk zone $\mathbb{R}_t$ at chromatography step t could either achieve the purity and yield requirements or fail to do so leading to large penalties associated with shortage costs or quality failures. Insights from the bounds are used in the structural analysis of the optimal purification policies in Section 5, and also provide basis for a state aggregation scheme in Section 6.

5. Structural Analysis of the Optimal Policy

In this section, we identify the structural properties of the optimal policies by exploiting the structural properties of the state space discussed in Section 4.

5.1. Optimal Policies in the Failure Zone and Risk Zone

Recall that, if the starting material is $(p'_t, i'_t) \in \mathbb{F}_t$ at chromatography step $t \in \mathcal{T}$, Theorem 1 indicates that the optimal policy is $a^*_t(p_t, i_t) = S$ for all $p_t \leq p'_t$ and $i_t \geq i'_t$. In this section, we analyze the structural properties of the optimal policy when the starting material is in the risk zone $\mathbb{R}_t$ at chromatography step $t \in \mathcal{T}$. To do so, we first define the *effective purity set* $\mathbb{P}_t$ at chromatography step $t \in \mathcal{T}$ that corresponds to all protein-impurity pairs which can lead to the purity requirement by the end of the planning horizon T. Let $X = [0, p_1] \times [0, i_1]$, then the effective purity set $\mathbb{P}_t$ at chromatography step t is defined in Proposition 4.

PROPOSITION 4. *The effective purity set at the beginning of the chromatography step t is*

$$\mathbb{P}_T = \left\{ (p_T, i_T) \in X : \frac{1 - \gamma_d}{\gamma_d} p_T \geq i_T \right\}, \tag{14}$$

$$K_{t,w} = \left\{ (p_t, i_t) \in X : p_t = \frac{p_{t+1}}{\theta^u_t | w}, i_t = \frac{i_{t+1}}{\psi^l_t | w}, (p_{t+1}, i_{t+1}) \in \mathbb{P}_{t+1} \right\} \text{ for } w \in \mathcal{W}_t, \tag{15}$$

$$\mathbb{P}_t = \bigcup_{w \in \mathcal{W}_t} K_{t,w} \text{ for } t = 1, \ldots, T-1. \tag{16}$$

Proof See Appendix.

Note that if the state (p_t, i_t) belongs to the effective purity set $\mathbb{P}_t$ at chromatography 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 level under the best purification capabilities. Note that Proposition 4 uses the best-case realizations $(\theta^u_t, \psi^l_t | w_t)$ to ensure that the set $\mathbb{P}_t$ includes all states (p_t, i_t) at chromatography step $t \in \mathcal{T}$ from which the final purity requirement can be achieved by the end of the planning horizon T. Proposition 4 is used to identify the characteristics of the optimal policy in Theorem 2.

THEOREM 2. *[Risk Zone] If $(p_t, i_t) \in \mathbb{R}_t$ at chromatography step $t \in \mathcal{T}$ and $\gamma_t < \gamma_d$, the optimal action has the property that $a^*_t(p_t, i_t) = \{ w^*_t \in \mathcal{W}_t : (p_{t+1}, i_{t+1} | p_t, i_t, w^*_t) \in \mathbb{P}_{t+1} \}$ for all $t \in \mathcal{T}$.*

Theorem 2 indicates that if the batch state is in the risk zone $\mathbb{R}_t$ at chromatography step t, then the optimal policy selects the pooling windows in such a way as to keep the batch state (p_{t+1}, i_{t+1}) within the effective purity set $\mathbb{P}_{t+1}$ of the next chromatography step $t+1 \in \mathcal{T}$. Theorem 2 provides guidelines to choose the best candidates for pooling windows in the risk zone. We note that the purification example in Section 8 illustrates the lack of threshold-type optimal policies for industry data. However, the guidelines obtained from Theorem 2 can help the scientists evaluate and understand which pooling windows are good or bad choices for a chromatography step.

5.2. Optimal Policies in the Target Zone

We explore the optimal policies when the starting state of the batch is in the target zone at the beginning of chromatography step $t \in \mathcal{T}$, i.e., $(p_t, i_t) \in \mathbb{T}_t$. We break this analysis into two cases: In Case 1, the biomanufacturer is committed to meeting both of the purity and yield requirements, and yield shortages are not allowed when $(p_t, i_t) \in \mathbb{T}_t$. In Case 2, yield shortages are permitted even though the batch state is in the target zone, i.e., the biomanufacturer might meet the purity requirement but not the yield requirement at the expense of incurring shortage penalties. 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 2. The target set $\mathbb{T}_T$ is said to be *reachable* at step T from the state (p_t, i_t) at chromatography step $t \in \mathcal{T}$, if there exists at least one sequence of pooling windows $\pi_t = (w_t, \cdots w_{T-1})$ such that the state (p_T, i_T) of the dynamic system $(p_{t+1}, i_{t+1}) = (\Theta_t p_t, \Psi_t i_t | w_t)$ is contained in $\mathbb{T}_T$ at step T for all possible purification outcomes at chromatography steps $t, t+1, \ldots, T-1$.

Definition 2 indicates that both of the yield and purity requirements are said to be reachable from state (p_t, i_t) and chromatography step $t \in \mathcal{T}$, only if there exists a pooling policy that attains these minimum requirements by step T, despite incurring the worst possible purification capabilities in all chromatography steps. As a direct consequence of Definition 2 and Proposition 3, we make the following observation (Bertsekas and Rhodes 1971):

OBSERVATION 1. *The target zone $\mathbb{T}_T$ is reachable at step T from all points of the target zone $\mathbb{T}_t$ defined in Proposition 3 for $t = \{1, \ldots, T-1\}$.*

Observation 1 indicates that the yield and purity requirements can be attained by step T as long as the batch state at the chromatography step $t \in \mathcal{T}$ is an element of the target zone $\mathbb{T}_t$ defined in Proposition 3. Consequently, the reachability problem from chromatography step t to step T can be reduced to the reachability problem from chromatography step t to 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 exists 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 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 of the yield and purity requirements at the end of the planning horizon T, if the starting state (p_t, i_t) at chromatography step $t \in \mathcal{T}$ is an element of the target zone $\mathbb{T}_t$. Then, the problem is equivalent to the problem of reachability of a target set described in Definition 2. Characteristics of the reachability problem are used to analyze the optimal pooling policies in Theorem 3.

THEOREM 3. *[Target Zone, Case 1] If $(p_t, i_t) \in \mathbb{T}_t$ at chromatography step $t \in \mathcal{T}$ and $\gamma_t < \gamma_d$, the necessary condition of the optimal policy is $a_t^*(p_t, i_t) = \left\{ w_t^* \in \mathcal{W}_t : (\theta_t^l p_t, \psi_t^u i_t | w_t^*) \in \mathbb{T}_{t+1} \right\}$ for $t \in \mathcal{T}$.*

Proof See Appendix.

Theorem 3 provide guidelines to select the optimal pooling window w_t at a chromatography step $t \in \mathcal{T}$. Theorem 3 indicates that the optimal action at chromatography step $t \in \mathcal{T}$ is to 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 batch state is in the target zone $\mathbb{T}_t$ at the beginning of chromatography step $t \in \mathcal{T}$. Recursive application of Theorem 3 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 $(p_{t+1}, i_{t+1}), \ldots, (p_{T-1}, i_{T-1})$ are in their respective target zones $\mathbb{T}_{t+1}, \ldots, \mathbb{T}_{T-1}$ in all subsequent 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 3 ensures that there exists at least one optimal policy that satisfies Theorem 3. Also note that the optimal policy is to stop if the state (p_t, i_t) is in the terminal zone $\mathbb{S}$ at chromatography step $t \in \mathcal{T}$.

Optimal Policies for Case 2 (Yield shortage allowed): We define a new reachability problem by allowing yield shortages (i.e., $p_T \leq p_d$) despite the batch state (p_t, i_t) being in the target zone $\mathbb{T}_t$ at chromatography step $t \in \mathcal{T}$. Compromising on yield might not be ideal, especially when it is know that the batch state (p_t, i_t) is in the target zone $\mathbb{T}_t$. However, compromising on yield could help reduce the number of purification steps in practice due to the purity-yield trade-off described in Section 2. To analyze the optimal policies in Case 2, we relax the yield requirement from Case 1. Then, the structural analysis becomes similar to Section 5.1, except that, we establish guaranteed performance for achieving the purity requirement in Theorem 4. Let $\hat{\mathbb{P}}_T = \left\{ (p_T, i_T) \in X : \frac{1-\gamma_d}{\gamma_d} p_T \geq i_T \right\}$ and $G_{t,w} = \left\{ (p_t, i_t) \in X : p_t = \frac{p_{t+1}}{\theta_t^l | w}, i_t = \frac{i_{t+1}}{\psi_t^u | w}, (p_{t+1}, i_{t+1}) \in \hat{\mathbb{P}}_{t+1} \right\}$ for $w \in \mathcal{W}_t$, and hence $\hat{\mathbb{P}}_t = \bigcup_{w \in \mathcal{W}_t} G_{t,w}$ for $t = 1, \ldots, T-1$. Then, Theorem 4 provides the necessary condition of the optimal policy in the target zone (Case 2).

THEOREM 4. *[Target Zone, Case 2] If $(p_t, i_t) \in \mathbb{T}_t$ at chromatography step $t \in \mathcal{T}$ and $\gamma_t < \gamma_d$, the necessary condition of the optimal policy is $a_t^*(p_t, i_t) = \left\{ w_t^* \in \mathcal{W}_t : (\theta_t^l p_t, \psi_t^u i_t | w_t^*) \in \hat{\mathbb{P}}_{t+1} \right\}$ for $t \in \mathcal{T}$.*

Proof See Appendix.

Note that Theorem 4 provides guaranteed performance for achieving the final purity requirement since it takes into consideration the worst-case realizations of the purification capabilities, i.e., $(\theta_t^l, \psi_t^u | w_t^*)$ for all $w_t \in \mathcal{T}$ and $t \in \mathcal{T}$. Note that the optimal policy is to stop when $\gamma_t \geq \gamma_d$ at chromatography step $t \in \mathcal{T}$. In Case 2, although the biomanufacturing firm has the capability of achieving both the purity and yield requirements by T, the optimal policy can choose to reduce the number of chromatography steps (and hence operating costs) at the expense of shortage penalties. In practice, in order to maintain good long-term relationships with the customers, the biomanufacturing firm might decide to meet both of the yield and purity requirements whenever they can – even if this decision might not be the best decision that increases the expected profit of a particular order. In such cases, the decision maker will proceed with the optimal policy suggested in Theorem 3, instead of Theorem 4. In practice, note that the optimal policies and the optimal value function can be determined for each protein and impurity pair contained on the state space by solving the MDP model; whereas the target zones are generated based on the worst-case analysis to establish performance guarantees.

6. State Aggregation, Action Elimination and Ordering Scheme

We use insights from the structural analysis of the state space to construct a state aggregation and action elimination procedure for the Markov decision model to improve computational efficiency. 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 in solving industry problems. Additionally, we define a stochastic ordering scheme for the pooling windows $w_t \in \mathcal{W}$ at a chromatography step $t \in \mathcal{T}$. This action ordering scheme provides a consistent method for labeling the pooling windows.

6.1. State Aggregation

The state aggregation scheme groups certain subset of the original system states into a single aggregate state. We first define an aggregate state called 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 5.

PROPOSITION 5. *All batch states $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ that are an element of the failure zone $\mathbb{F}_t$ at chromatography step $t \in \mathcal{T}$ can be grouped and viewed as a single state called the failure state d_t with reward $r(d_t) = -c_f$.*

Proof See Appendix.

Proposition 5 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 4.3 indicate that the optimal value $\mathcal{V}_t^*(p_t, i_t)$ is constant over the (sub)set $\left\{ \mathbb{1}_{p_t \leq p_t', i_t \geq i_t'} | (p_t', i_t') \in \mathbb{F}_t \right\}$ of the original state space $\mathcal{P} \times \mathcal{I}$ at each $t \in \mathcal{T} \cup T$, where $\mathbb{1}$ is the indicator function. More specifically, we have $\mathcal{V}_t^*(p_t, i_t) = -c_f$ for all $(p_t, i_t) \in \left\{ \mathbb{1}_{p_t \leq p_t', i_t \geq i_t'} | (p_t', i_t') \in \mathbb{F}_t \right\}$. Since all protein and impurity pair that satisfy Proposition 5 are already an element of the failure zone $\mathbb{F}_t$, the aggregation scheme in Proposition 5 is exact, in the sense that the aggregate state d_t encompasses subsets of the original system states that have equal costs and transitions (Bertsekas 2012).

6.2. Action Elimination and Ordering

Next, we discuss an action elimination procedure in Proposition 6, which is then used to develop a stochastic ordering scheme for pooling windows $w_t \in \mathcal{W}_t$ at chromatography step $t \in \mathcal{T}$.

PROPOSITION 6. *Let w_t^i and w_t^j be two distinct pooling windows at 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 all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ at $t \in \mathcal{T}$.

(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 all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ at $t \in \mathcal{T}$.

Proof See Appendix.

Conditions in Proposition 6 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 $t \in \mathcal{T}$. Proposition 6 indicates that the pooling window w_t^i is expected to result in strictly lower profits than the pooling window w_t^j for all $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ at the chromatography step $t \in \mathcal{T}$, i.e., $\mathcal{V}_t(p_t\theta_t, \psi_t i_t|w_t^i) < \mathcal{V}_t(p_t\theta_t, \psi_t i_t|w_t^j)$. 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., $\hat{\mathcal{W}}_t \subseteq \mathcal{W}_t$ at chromatography step $t \in \mathcal{T}$.

Figure 4 shows an example of a strictly dominated pooling window using industry data described in more detail in Section 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 $(\bar{\theta}, \bar{\psi}|w^i) = (0.71, 0.53)$
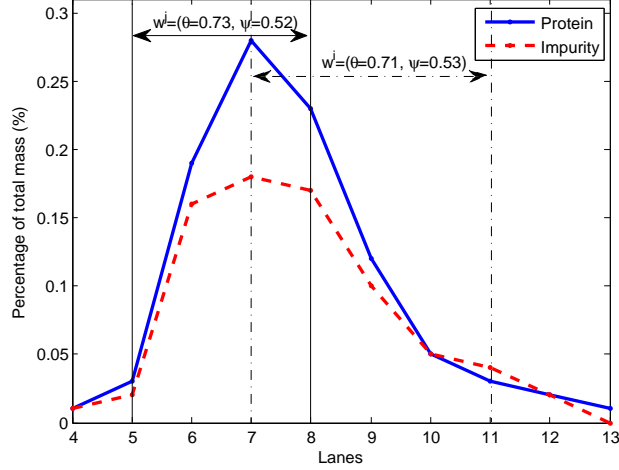
Figure 4   Example of a dominated pooling window (based on chromatography data from Aldevron)

with the bounds $(\theta^l, \psi^l | w^i) = (0.64, 0.47)$ and $(\theta^u, \psi^u | w^i) = (0.78, 0.58)$. The window w^j corresponds to the lanes 5 to 8 with the purification capability $(\bar{\theta}, \bar{\psi} | w^j) = (0.73, 0.52)$, and the bounds $(\theta^l, \psi^l | w^j) = (0.65, 0.46)$ and $(\theta^u, \psi^u | w^j) = (0.80, 0.51)$. 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 6 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 in this specific chromatography step.

Next, we provide a stochastic ordering scheme for ranking and labeling the pooling windows $w_t \in \hat{\mathcal{W}}_t$ at chromatography step $t \in \mathcal{T}$. Let $\hat{N}_t$ be the number of pooling windows at chromatography step $t \in \mathcal{T}$ after performing the action elimination procedure in Proposition 6. Let the action index n denote the position of the pooling window w_t^n in our ordering scheme, i.e., the pooling window w_t^n is the n^{th} pooling window among $\hat{N}_t$ windows that are stochastically ordered at chromatography step $t \in \mathcal{T}$. Property 1 and Assumption 1 provide necessary conditions for a stochastic ordering of pooling windows $w_t^n \in \hat{\mathcal{W}}_t$ at chromatography step $t \in \mathcal{T}$.

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 $w_t^{n-1}, w_t^n, w_t^{n+1} \in \hat{\mathcal{W}}_t$ at chromatography step $t \in \mathcal{T}$.

ASSUMPTION 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 at chromatography step $t \in \mathcal{T}$ increase in the action index n. When Property 1 and Assumption 1 hold together, it results in a stochastic ordering scheme where the pooling windows having higher

action index n at chromatography step $t \in \mathcal{T}$ lead to stochastically higher amount of protein and impurity at that chromatography step. In this ordering scheme, a pooling window $w_t^n \in \hat{\mathcal{W}}_t$ is said to be *larger* window as its action index n approaches to $\hat{N}_t$, and *smaller* window as its action index n approaches to 1 at chromatography step $t \in \mathcal{T}$. This ordering scheme is indeed in alignment with what is observed in practice at Aldevron. Property 1 and Assumption 1 formalize the well-known trade-off between purity and yield involved in chromatographic separation, and also align with the chemical engineering literature (Ngiam et al. 2001, Vasquez-Alvarez et al. 2001, Ngiam et al. 2003, Muller-Spath et al. 2013, Subramanian 2014). In practice, Property 1 and Assumption 1 have been validated using scouting data obtained from Aldevron.

Note that the structural analysis and insights in Section 4 and Section 5 hold regardless of the state aggregation, action elimination and ordering schemes. These schemes are developed to facilitate the numerical analysis in Section 8. For example, the application of the action ordering procedure enables a consistent mechanism for ranking and labeling the pooling windows in practice.

7. Implementation at Aldevron

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

7.1. Timeline

The optimization framework has been constructed, revised, validated, and implemented over a three-year period (2013-2016) 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.

7.2. Implementation Results

Three 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. *Reduction in the number of the validation runs.* The optimization model has allowed to reduce 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 indicated 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 chromatography run. This information is especially critical in communicating the manufacturing challenges with the customers. 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 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.

7.3. Feedback from Biomanufacturing Community and 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, BioForward 2016). 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, 2016). 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. 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 (WEDC 2014, BioForward 2016).

Understanding the theory of Markov decision processes and computing the zones and optimal polices could be challenging for most purification scientists. To facilitate the industry implementation, we developed a decision support tool using Java. The tool provides a user-friendly interface

for the purification scientists to easily enter the problem parameters (including the results of the scouting experiments). The tool takes this information as input, applies the theory described in the paper, and computes the decision zones. The tool also generates a file where the optimal pooling policies and the optimal value function are reported for each state. Using this tool does not require any specific knowledge on stochastic optimization, and hence the purification scientists were comfortable with using this tool to generate the zones and optimal policies in practice.

Other implementation challenges were related to the formatting of the scouting data required to run the optimization model. Initially, the data obtained from scouting runs were in the format of gel pictures as shown in Figure 2 (a). 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 used as input for the Java tool 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, we automated this process using the Java tool. 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.

8. A Case Study Illustrating Results for an Engineered Protein

Since each purification order is custom-engineered and unique, each order has its own operating policies and managerial insights. Therefore, we believe that it would not be useful to explain the optimal policies and insights for every single protein considered in the implementation process at Aldevron. Instead, we elaborate on one of the custom-engineered purification orders that involves two chromatography steps (Section 8.1), and explain the way how the optimization framework was implemented at Aldevron for that order. More specifically, we demonstrate the decision zones, identify the optimal policies, compare the optimal policies with current practice, and discuss the managerial insights (Sections 8.2-8.3). Furthermore, we briefly provide another example from Aldevron that involves three chromatography steps (Section 8.4), and conclude with quantifying computational savings due to action elimination and state aggregation (Section 8.5). To protect client confidentiality, actual data and cost information obtained from Aldevron are masked.

8.1. Problem Setting and Parameters

The protein of interest considered 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 5. The first step uses the binding affinities between
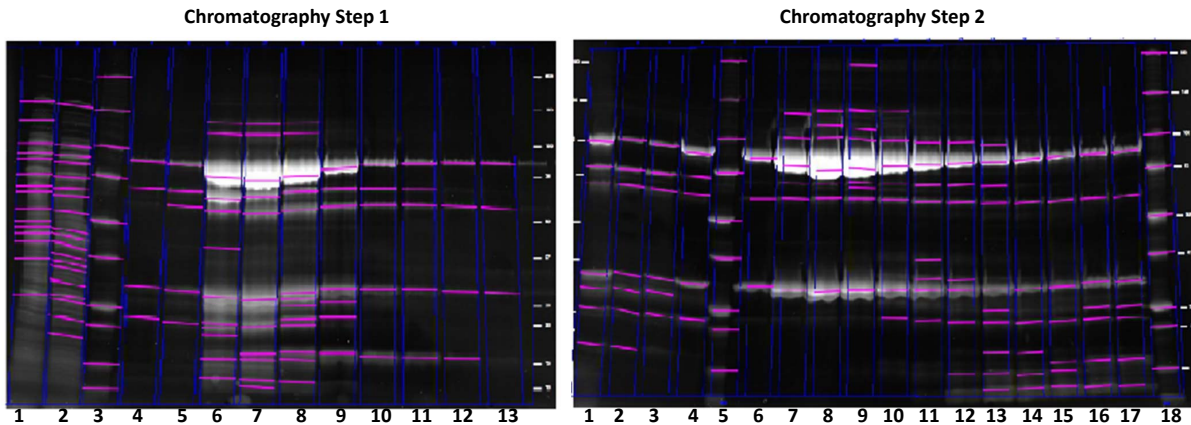
Figure 5   Two step chromatography outputs

proteins and metal ions as a separation principle, and the second step uses separation based on electric charge. Figure 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. In this case study, statistical analysis of the scouting data indicates that the purification capabilities are uniformly distributed within 10% of their mean $(\bar{\theta}_t, \bar{\psi}_t | w_t)$ for all $w_t \in \mathcal{W}_t$ at the chromatography step $t \in \{1, 2\}$. In the first chromatography step, the smallest pooling window w_1^1 has the mean purification capability $(\bar{\theta}_1, \bar{\psi}_1 | w_1^1) = (0.010, 0.003)$, and the largest pooling window w_1^{55} has $(\bar{\theta}_1, \bar{\psi}_1 | w_1^{55}) = (0.939, 0.745)$. In the second chromatography step, the smallest pooling window w_2^1 corresponds to $(\bar{\theta}_2, \bar{\psi}_2 | w_2^1) = (0.047, 0.003)$, and the largest pooling windows w_2^{78} has $(\bar{\theta}_2, \bar{\psi}_2 | w_2^{78}) = (0.928, 0.671)$. We note that the scouting data is collected and analyzed as per recommended guidelines (Ellison and Willams 2012, ISO21748 2010). All pooling windows satisfy the data characteristics and assumptions in Section 6.

The production 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, and representative values are used instead. The operating costs of a chromatography step is $c_t = \$15$ for $t \in \{1, 2\}$. These include costs associated with labor, materials, equipment, inspection and analytics. The revenue structure is $r(p_t) = \$5 \times p_t$ for $p_t < 8$ mg and $r(p_t) = \$40$ for $p_t \geq 8$ mg. Shortage cost is $c_l(p_d - p_t) = \$48 - \$6 \times p_t$ for $p_t < 8$ mg, $c_l(p_d - p_t) = 0$ otherwise. Penalty cost of failure is $c_f = \$48$, which is equivalent to the maximum possible shortage cost considered in our purification setting. The state space is discretized based on the least count measured in each project, and the model is solved using the backward induction algorithm. During industry implementation, sensitivity analysis is conducted to ensure that the decision zones, the optimal value function, and the optimal policies are robust to finer discretization levels. We note that Proposition 2 and Proposition 3 also allow to generate the zone-based performance guarantees without discretizing the state space.
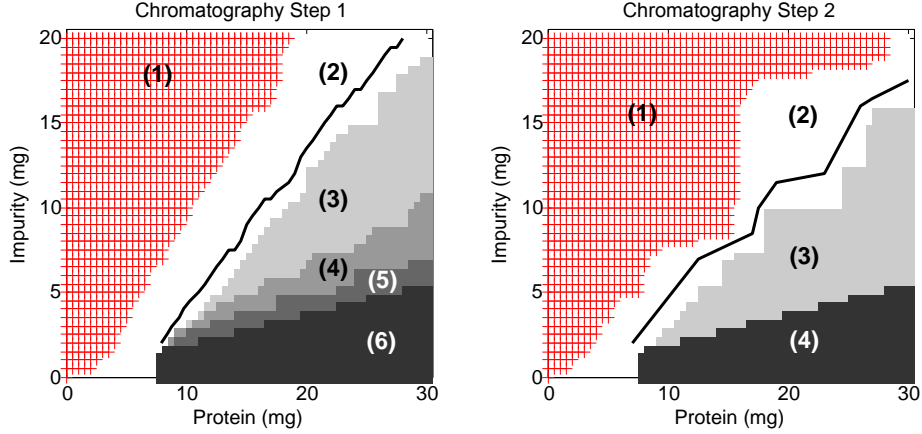
Figure 6 Optimal value function for the first step: (1) $\mathcal{V}_1^*(p_1, i_1) = -48$, (2) $-48 < \mathcal{V}_1^*(p_1, i_1) < 10$, (3-6) $10 \leq \mathcal{V}_1^*(p_1, i_1) \leq 40$, and the solid line for $\mathcal{V}_1^*(p_1, i_1) = 0$. Optimal value function for second step: (1) $\mathcal{V}_2^*(p_2, i_2) = -48$, (2) $-48 < \mathcal{V}_2^*(p_2, i_2) < 25$, (3-4) $25 \leq \mathcal{V}_2^*(p_2, i_2) \leq 40$, and the solid line for $\mathcal{V}_2^*(p_2, i_2) = 0$.

Table 2   Summary of the insights based on Figure 6

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

8.2. Decision Zones and Their Financial Implications

We investigate the financial implications of a starting material 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 6 presents the decision zones and the optimal value function at each chromatography step. All managerial insights discussed below are derived from Figure 6, and summarized in Table 2.

In Figure 6, the region (1) corresponds to the failure zone $\mathbb{F}_t$ of the chromatography step $t \in \{1, 2\}$. This region represents the protein and impurity pairs 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 6 represents the risk zone $\mathbb{R}_t$ of the chromatography step $t \in \{1, 2\}$. When the starting material is an element of the risk zone, the biomanufacturing firm can expect to incur financial losses due to combined impact of shortage costs and operating costs. For example, the solid line passing through the risk zone in Figure 6 corresponds to all protein-impurity pairs (p_t, i_t) having $\mathcal{V}_t^*(p_t, i_t) = 0$ 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. Due to the monotonicity of the value function (Proposition 1), the expected profit is nondecreasing in protein amount 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 state of the starting material is on the left hand side of the solid line, then the firm might prefer to scrap that starting material, rework in-house or request the provider to send a new starting material.

Regions $(3-6)$ in the first chromatography step and regions $(3-4)$ in the second chromatography step represent the target zone $\mathbb{T}_t$ where the firm is capable of meeting both the purity and yield requirements at $t \in \{1, 2\}$. For example, if the starting material (p_1, i_1) is in the region (3) of the first chromatography step, the firm can expect to achieve the final yield and purity requirements through two chromatography steps using the optimal policies, resulting in $\mathcal{V}_1^*(p_1, i_1) = 10$. However, the optimal policy in region (4) of the first step suggests that the firm might be better off compromising on yield to achieve the final purity requirement at the end of the first step, despite incurring some shortage penalties. In this case, the operating cost of the second step is greater than the expected shortage costs. Note that although both the yield and purity requirements could have been met in the region (4), it is financially better off to choose pooling windows that can achieve the purity requirement but also lead to yield shortages by the end of the first step, i.e., $10 < \mathcal{V}_1^*(p_1, i_1) < 25$. 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 3), with $\mathcal{V}_1^*(p_1, i_1) = 10$. In the region (5), the firm can expect to achieve the final purity and yield requirements at the end of the first chromatography step, $\mathcal{V}_1^*(p_1, i_1) = 25$. In this case, the second step chromatography is not required. In Figure 6, note that the size of the target zone expands while the failure zone shrinks in the first chromatography step compared to the zones in the second chromatography step. 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.

8.3. Optimal Policies and Comparison with Current Practice

We present the optimal policies for the batch states $p_1 \in [10, 30]$ mg and $i_1 \in [0, 20]$ mg. In this case study, the actual starting material processed at Aldevron contained $(p_1, i_1) = (27.5, 17.5)$. Figure 6 indicates that the starting material is in the risk zone of the first chromatography step. Therefore,

the firm can not provide any guarantees for achieving the production requirements in this specific example. We quantify the risks and costs associated with the starting material, and compare the optimal policy with the one used in practice.

State-dependent optimal policies: Let π^* denote the optimal policy, and $\mathcal{V}_t^*(p_t, i_t | \pi^*)$ represent the optimal value function at chromatography step t. Table 3 presents a snapshot of the optimal policy for the first chromatography step. Table 3 only displays the optimal policies at selected states (i.e., in the intervals of 2.5 mg) to improve readability. In Table 3, S represents the stopping action. For other actions in Table 3, we present the starting lane, ending lane, and the corresponding action index assigned based on the action ordering procedure in Section 6. For example, $L6\text{-}8\,(21)$ means pooling the lanes 6-8, and this action is the 21^{st} action out of 55 pooling windows at the first chromatography step. Cells colored in gray represent the target zone based on Figure 6, and the entries in bold correspond to the failure zone.

Table 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

We make the following observations regarding the optimal policy. First, we observe that the optimal action is to stop the purification process when the batch state is in the failure zone $\mathbb{F}_1$ (i.e., top left of Table 3) or in the terminal zone $\mathbb{S}$ (i.e, bottom right of Table 3). Second, we see that the optimal policies in the target zone $\mathbb{T}_1$ do not have any threshold-type structure, but they do satisfy the necessary conditions in Theorem 3-4. Third, in the risk zone $\mathbb{R}_1$, we can observe a non-decreasing trend in the action index as the impurity amount decreases for a given protein amount. For example, at $p_1 = 17.5$, the optimal policy chooses actions with higher index as the impurity amount decreases. However, this monotonic trend is not present for all protein-impurity pairs. For example, at $p_1 = 20$ mg and $i_1 = 20$ mg, the optimal policy suggests to pool lanes $4 - 13$ with the action index 36. When $i_1 \in \{15, 17.5\}$, 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. We also observe the lack of threshold-type policies as the protein amount increases for a fixed impurity level. Such deviations in optimal policies are also observed in the second step.

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

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 $\pi_1'' = \{$Lanes $6-9$, Lanes $7-10\}$ denote the 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'', 13.3 mg of protein and 2.2 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 $\theta_2 = 0.582$, $\psi_2 = 0.210$). Therefore, we evaluate the performance of the policy π_1'' even though yield and purity requirements were satisfied in our example production 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.

8.4. Three-Step Chromatography Example

It is most common to adopt two or three step protocol in practice, however, difficult proteins may require several additional steps (Healthcare 2010). Since the implementation of the model, purification orders received at Aldevron required either two or three chromatography steps. Therefore, we provide another example from Aldevron involving three chromatography steps. The protein of interest considered in this example is used for in vitro studies. Each chromatography step separates the protein of interest based on its charge, hydrophobicity, and size, respectively. Scouting experiments indicate that purification capabilities are uniformly distributed within 10% of their mean values $(\bar{\theta}_t, \bar{\psi}_t)$ for $t \in \{1, 2, 3\}$. Total number of available pooling windows in each chromatography step is $153, 120$ and 105, respectively. The production requirement is 7.5 mg of protein with at least 90% purity. Information about pooling windows and costs are masked for confidentiality, and representative values are used instead. Operating cost is $c_t = \$15$ for $t \in \{1, 2, 3\}$. The revenue structure is $r(p_t) = \$12 \times p_t$ for $p_t < 7.5$ mg and $r(p_t) = \$90$ for $p_t \geq 7.5$ mg, $t \in \{1, 2, 3\}$. Shortage cost is $c_l(p_d - p_t) = \$108 - \$14.4 \times p_t$ for $p_t < 7.5$ mg, $c_l(p_d - p_t) = 0$ otherwise. Penalty cost of failure is $c_f = \$108$.
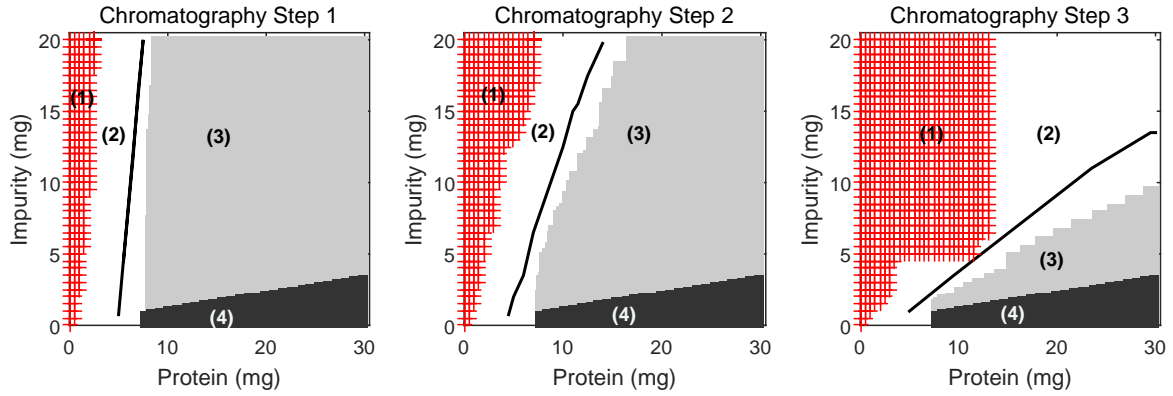
Figure 7 Decision zones for a purification project with three chromatography steps: (1) failure zone, (2) risk zone, (3) target zone, (4) terminal zone, and the solid line for $\mathcal{V}_t^*(p_t, i_t) = 0$.

Figure 7 represents the decision zones of each chromatography step. We observe that the size of the target zone expands while the failure zone shrinks as t approaches to the first chromatography step. The behavior of target and failure zones in Figure 7 is intuitive and aligns with the structural analysis since each additional step offers an opportunity to eliminate the remaining amount of impurity. In this case study, the decision zones associated with the first chromatography step indicate a promising business case for processing this order since the failure zone is relatively small compared to the target zone. In practice, several factors might affect the sizes of the decision zones (e.g., the number of chromatography steps, purification capabilities at each step, production requirements, etc.). If the sizes of both failure and target zones are small due to specific process parameters, then the starting material is more likely to belong to the risk zone, and this would still be an important insight for both the biomanufacturer and the client.

In this example, the starting material is in the target zone of the first step with 9.5 mg protein and 17 mg impurity. The scientist managed to achieve both purity and yield requirements using this starting material in practice. However, comparison of the optimal value function against the value function associated with current practice reveals 14% improvement in the expected profit. In addition, the biomanufacturer is capable of providing performance guarantees using our decision support tool since the starting material is in the target zone. Such promise on guaranteed success provides significant advantage to both the client and the biomanufacturer.

8.5. Impact of State Aggregation and Action Elimination

Using the state aggregation scheme in Proposition 5 and the action elimination procedure in Proposition 6, we obtain significant savings in the computational effort required to obtain solutions to industry size problems. For example, applying the state aggregation scheme to the purification project presented in Section 8.1 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 6, the total number of pooling windows reduced from 55 to 36 windows in the first step, and from 78 to 20 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. In the three-step example presented in Section 8.4, the total number of actions reduced from 378 to 197 pooling windows, and the combined impact of the state aggregation and action elimination resulted in 50% reduction in the CPU time. In general, the cardinality of the state space ranged between $40,000 - 80,000$ during implementation at Aldevron. All computations presented in this paper have been executed on a system with 2.9 GHz CPU and 16 GB of RAM. The CPU time is 21.1 seconds for the purification project in Section 8.1, and 56.2 seconds for the project in Section 8.4. Overall, the CPU time has been between 10 seconds to 60 seconds during implementations at Aldevron.

9. Conclusions

We focus on protein purification operations conducted by biomanufacturers and pharmaceutical companies. Each order for the engineered protein has specific purity and yield requirements defined by the end use or application, and the biomanufacturer incurs penalty costs when these specific requirements are not achieved. However, achieving both of the purity and yield requirement is challenging due to the purity and yield trade-off involved in chromatography operations. Furthermore, the starting material often involves variability in terms of the protein and impurity amounts, which affects subsequent purification decisions. Limitations in the available chromatography techniques further challenge the purification decisions. Due to high penalty costs and strict requirements on purity, biomanufacturers need help with two levels of decisions: *(i)* determine whether the purity and yield requirements specified by the customer are achievable at all, and if so, *(ii)* determine the optimal purification strategies that maximize the expected profit. This paper addresses both issues and provides guidelines for practitioners.

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. Our structural analysis partitions the state space into decision zones (i.e., failure zone, risk zone and target zone) that provide a strong basis to analyze the financial trade-offs and business risks associated with the starting material. The decision zones also establish the nature and type of performance guarantees that can be provided in practice. For each zone, we then provide practical guidelines for optimal purification decisions that maximize the expected profit. The proposed zone-based decision making approach is particularly easy to implement in practice.

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. Our optimization framework provides a rigorous analysis of the risks and financial trade-offs involved in chromatography operations. Applications of operations research techniques are mostly new to the biomanufacturing community. As more companies like Aldevron embrace operations research and integrate it into practice, regulatory authorities might mandate the use of such approaches to improve the biomanufacturing research and development.

Future research could explore the interaction between fermentation and purification decisions. For example, some R&D projects require the biomanufacturer to first produce the starting material through fermentation, and then perform the purification operations. In such projects, the scientist can benefit from a comprehensive framework that links the complex dynamics between fermentation and purification operations. Furthermore, the model assumes that the sequence of chromatography techniques is predetermined based on scouting runs, which could be a limitation for the decision maker. As another future research direction, one could explore the optimal number of chromatography steps and the optimal choice of chromatography technique in each step.

Acknowledgments

References

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.

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.

BioForward. 2016. BioForward speaker series: Aldevron collaborates with University of Wisconsin to develop decision support tools for biomanufacturing. URL http://www.bioforward.org/events/aldevron-collaborates-with-university-of-wisconsin-to-develop-decision-support-tools-for-biomanufacturing/.

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

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.

Chhatre, S., C. Jones, R. Francis, K. O'Donovan, N. Titchener-Hooker, A. Newcombe, E. Keshavarz-Moore. 2006. The integrated simulation and assessment of the impacts of process change in biotherapeutic antibody production. *Biotechnology Progress* **22**(6) 1612–1620.

Ellison, S.L.R., A. Willams. 2012. EURACHEM/CITAC Guide CG 4: Quantifying uncertainty in analytical measurement. URL http://www.citac.cc/QUAM2012_P1.pdf.

Farid, S. S. 2009. Process economic drivers in industrial monoclonial antibody manufacture. Uwe Gottschalk, ed., *Process Scale Purification of Antibodies*. John Wiley & Sons, 239–261.

Farid, S. S., J. Washbrook, Nigel J. Titchener-Hooker. 2007. Modelling biopharmaceutical manufacture: Design and implementation of SimBiopharma. *Computers & Chemical Engineering* **31**(9) 1141–1158.

Farid, S.S. 2007. Process economics of industrial monoclonal antibody manufacture. *Journal of Chromatography B* **848** 8–18.

Healthcare, GE. 2010. Strategies for protein purification: Handbook by GE Healthcare. Tech. rep.

ISO21748. 2010. Guidance for the use of repeatability, reproducibility and trueness estimates in measurement uncertainty estimation. URL http://www.iso.org/iso/catalogue_detail.htm?csnumber=46373.

Kraattli, M.n, T. Muller-Spath, N. Ulmer, G. Strohlein, M. Morbidelli. 2013. Separation of lanthanides by continuous chromatography. *Industrial & Engineering Chemistry Research* **52**(26) 8880–8886.

Lakhdar, K., Y. Zhou, J. Savery, N. J. Titchener-Hooker, L. G. Papageorgiou. 2005. Medium term planning of biopharmaceutical manufacture using mathematical programming. *Biotechnology Progress* **21**(5) 1478–1489.

Lim, A. C., J. Washbrook, N. J. Titchener-Hooker, S. S. Farid. 2006. A computer-aided approach to compare the production economics of fed-batch and perfusion culture under uncertainty. *Biotechnology and Bioengineering* **93**(4) 687–697.

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* .

Martagan, T., A. Krishnamurthy, C. T. Maravelias. 2016. Optimal condition-based harvesting policies for biomanufacturing operations with failure risks. *IIE Transactions* **48**(5) 440–461.

Muller-Spath, T., N. Ulmer, L. Aumann, G. Strohlein, M. Bavand, L. J.A. Hendriks, J. de Kruif, M. Throsby, A.B.H. Bakker. 2013. Purifying common light-chain bispecific antibodies. *BioProcess International* **11**(5) 36–45.

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.

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.

Papageorgiou, L. G., G. E. Rotstein, N. Shah. 2001. Strategic supply chain optimization for the pharmaceutical industries. *Industrial & Engineering Chemistry Research* **40**(1) 275–286.

Polykarpou, E. M., Paul A. D., Lazaros G. P. 2011. Optimal synthesis of chromatographic trains for downstream protein processing. *Biotechnol. Prog.* **27 (6)** 1653–1660.

Polykarpou, E. M., P. A. Dalby, L. G. Papageorgiou. 2012. An MILP formulation for the synthesis of protein purification processes. *Chemical Engineering Research and Design* **90**(9) 1262–1270.

Puterman, M. L. 1994. *Markov Decision Processes Discrete Stochatic Dynamic Programming*. John Wiley & Sons, New York.

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.

Simaria, A. S., R. Turner, S. S. Farid. 2012. A multi-level meta-heuristic algorithm for the optimisation of antibody purification processes. *Biochemical Engineering Journal* **69** 144–154.

Subramanian, G. 2014. *Continuous Processing in Pharmaceutical Manufacturing*. John Wiley & Sons.

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. 2003. A mixed integer linear programming model for the optimal synthesis of protein purification processes with product loss. *Chemical and Biochemical Engineering Quarterly* **17**(1) 77–84.

WEDC. 2014. Wisconsin Economic Development Corporation (WEDC), company feature: Aldevron. URL http://inwisconsin.com/entrepreneurs-and-innovators/launch_blog/company-feature-aldevron/.

Zhou, Y., S. H. Ngiam, N. Titchener-Hooker. 2005. Modelling of chromatographic processes and its application in therapeutic protein. *Control Applications, 2005. CCA 2005. Proceedings of 2005 IEEE Conference on*. IEEE, 699–704.

Tugce Martagan is an Assistant Professor and Marie S. Curie Research Fellow in the School of Industrial Engineering at Eindhoven University of Technology. She received her Ph.D. in Industrial Engineering from the University of Wisconsin-Madison. Her research interests include stochastic modeling and optimization with applications in the pharmaceutical industry. She is the recipient of the first prize in the 2016 POMS Applied Research Challenge, the honorable mention in the 2016 George B. Dantzig dissertation award, and the finalist of 2015 MSOM student paper competition.

Ananth Krishnamurthy is a Professor in the Department of Industrial and Systems Engineering at the University of Wisconsin- Madison. His research targets the development and application of performance modeling techniques in the design and analysis of manufacturing systems and supply chains. Dr. Krishnamurthy also serves as the Director of the Center for Quick Response Manufacturing. He received is masters and Ph.D. from the University of Wisconsin-Madison.

Peter A. Leland is a Senior Manager of Protein Services at Aldevron LLC. He obtained his Ph.D. in Biochemistry from the University of Wisconsin-Madison in 2001. His expertise lies in the expression and purification of proteins utilizing automated and high throughput strategies, and the design and commercialization of protein refolding screens.

Christos T. Maravelias is a Vilas Distinguished Achievement Professor in the Department of Chemical and Biological Engineering at the University of WisconsinMadison. His research interests lie in the areas of chemical production scheduling, supply chain optimization, chemical process synthesis and analysis, and computational methods for novel material discovery. He is the recipient of the 2008 David Smith and the 2013 Outstanding Young Researcher Award from the CAST division of AIChE.

Appendix

Proof of Proposition 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 the step 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 (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 Equations (4)-(6), we have, for $i_t \in \mathcal{I}$ and $t \in \mathcal{T}$,

$$\mathcal{V}_t(p_t, i_t) = \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \mathbb{E}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\}$$

$$\geq \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t^-, i_t), -c_t + \mathbb{E}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\} \tag{17}$$

$$\geq \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t^-, i_t), -c_t + \mathbb{E}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t^-, \psi_t i_t) \right\} \tag{18}$$

$$= \mathcal{V}_t(p_t^-, i_t) \tag{19}$$

where, Equation (17) follows from the stopping cost structure in Equation (3), and Equation (18) is obtained from the induction hypothesis. Proof for monotonicity of the value function in $i_t \in \mathcal{I}$ for a given $p_t \in \mathcal{P}$ at $t \in \mathcal{T}$ is entirely analogous, and hence omitted. $\square$

Proof of Theorem 1 It is sufficient to show that if $a_t^*(p_t', i_t') = S$ then $a_t^*(p_t, i_t) = S$ for all $p_t \leq p_t'$ and $i_t \geq i_t'$ at $t \in \mathcal{T}$, $p_t \in \mathcal{P}, i_t \in \mathcal{I}$.

Note that at the end of the planning horizon T, the only available action is to stop with rewards $\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T)$. Next, assume by contradiction hypothesis that $a_t^*(p_t', i_t') = S$ but $a_t^*(p_t, i_t) = w$ for a given $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ where $p_t \leq p_t'$ and $i_t \geq i_t'$, $t \in \mathcal{T}$, $w \in \mathcal{W}_t$ and $w \neq S$. This implies that,

$$r_S(p_t', i_t') > -c_t + \int_{\psi_t^l | w}^{\psi_t^u | w} \int_{\theta_t^l | w}^{\theta_t^u | w} f_t(\theta_t | w) g_t(\psi_t | w) \mathcal{V}_{t+1}(\theta_t p_t', \psi_t i_t') \mathrm{d}\theta \mathrm{d}\psi \tag{20}$$

and

$$-c_t + \int_{\psi_t^l | w}^{\psi_t^u | w} \int_{\theta_t^l | w}^{\theta_t^u | w} f_t(\theta_t | w) g_t(\psi_t | w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi > r_S(p_t, i_t) \tag{21}$$

which together imply

$$r_S(p_t', i_t') - r_S(p_t, i_t)$$

$$> \int_{\psi_t^l | w}^{\psi_t^u | w} \int_{\theta_t^l | w}^{\theta_t^u | w} f_t(\theta_t | w) g_t(\psi_t | w) \mathcal{V}_{t+1}(\theta_t p_t', \psi_t i_t') \mathrm{d}\theta \mathrm{d}\psi$$

$$- \int_{\psi_t^l | w}^{\psi_t^u | w} \int_{\theta_t^l | w}^{\theta_t^u | w} f_t(\theta_t | w) g_t(\psi_t | w) \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \mathrm{d}\theta \mathrm{d}\psi. \tag{22}$$

Note that $r_S(p'_t, i'_t) - r_S(p_t, i_t) = 0$ due to stopping cost structure in Equation (3). Theorem 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 (22) indicates that the term on its right hand side is negative. However,

$$\int_{\psi^l_t|w}^{\psi^u_t|w} \int_{\theta^l_t|w}^{\theta^u_t|w} f_t(\theta_t|w)g_t(\psi_t|w)\mathcal{V}_{t+1}(\theta_t p'_t, \psi_t i'_t)\mathrm{d}\theta\mathrm{d}\psi$$

$$- \int_{\psi^l_t|w}^{\psi^u_t|w} \int_{\theta^l_t|w}^{\theta^u_t|w} f_t(\theta_t|w)g_t(\psi_t|w)\mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t)\mathrm{d}\theta\mathrm{d}\psi \tag{23}$$

$$\geq \int_{\psi^l_t|w}^{\psi^u_t|w} \int_{\theta^l_t|w}^{\theta^u_t|w} f_t(\theta_t|w)g_t(\psi_t|w)\mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t)\mathrm{d}\theta\mathrm{d}\psi$$

$$- \int_{\psi^l_t|w}^{\psi^u_t|w} \int_{\theta^l_t|w}^{\theta^u_t|w} f_t(\theta_t|w)g_t(\psi_t|w)\mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t)\mathrm{d}\theta\mathrm{d}\psi \tag{24}$$

$$= 0.$$

Therefore, the term on the right hand side of inequality (22) is non-negative, which contradicts the inequality (22), and hence the proof follows. Note that Equation (24) follows from the monotonicity of the value function in Proposition 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 $p_t \leq p'_t$ and $i_t \geq i'_t$. $\square$

Proof of Proposition 2 We prove Proposition 2 by induction. First, we 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 chromatography step $t \in \mathcal{T}$. Assume by induction hypothesis that (p_t, i_t) at $t \in \mathcal{T}$ satisfies the condition (i) for all $\pi_t = (w_t, w_{t+1}, \ldots, w_{T-1})$. Then, in the last chromatography step $T - 1$, we have,

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

$$= \max_{w_{T-1} \in \mathcal{W}_{T-1}} \left\{ r_S(p_{T-1}, i_{T-1}), -c_{T-1} + \mathbb{E}_{\theta_{T-1}, \psi_{T-1}|w_{T-1}} r_S(\theta_{T-1}p_{T-1}, \psi_{T-1}i_{T-1}) \right\}$$

$$= \max\left\{ -c_f, -c_{T-1} - c_f \right\} \tag{25}$$

$$= -c_f.$$

Note that Equation (25) follows from the induction hypothesis and the stopping costs structure defined in Equation (3).

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

$$\mathcal{V}_t(p_t, i_t) = \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \mathbb{E}_{\theta_t, \psi_t|w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\}$$

$$\leq \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \mathcal{V}_{t+1}(\theta^u_t p_t, \psi^l_t i_t|w_t) \right\} \tag{26}$$

$$\leq \max\left\{ -c_f, -c_t - c_f \right\} \tag{27}$$

$$= -c_f$$

where, Equation (26) follows from monotonicity of the value function and indicates the best-case analysis of purification capabilities, and Equation (27) follows from the induction hypothesis and stopping cost structure as condition (i) holds. Hence, abandoning the purification at state (p_t, i_t) and step $t \in \mathcal{T}$ leads to less financial losses than continuing the purification under condition (i).

Next, we investigate the condition (ii) in Proposition 2. Let $(p_t, i_t) \in \mathcal{P} \times \mathcal{I}$ with $\gamma_d > \frac{p_t}{p_t + i_t}$ at chromatography step $t \in \mathcal{T}$, and assume by the induction hypothesis that condition (ii) holds for all $\pi_t = (w_t, w_{t+1}, \ldots, w_{T-1})$. In the last chromatography step $T-1$, we have,

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

$$= \max_{w_{T-1} \in \mathcal{W}_{T-1}} \left\{ r_S(p_{T-1}, i_{T-1}), -c_{T-1} + \mathop{\mathbb{E}}_{\theta_{T-1}, \psi_{T-1} | w_{T-1}} r_S(\theta_{T-1} p_{T-1}, \psi_{T-1} i_{T-1}) \right\}$$

$$\leq \max_{w_{T-1} \in \mathcal{W}_{T-1}} \left\{ r_S(p_{T-1}, i_{T-1}), -c_{T-1} + r_S(\theta_{T-1}^u p_{T-1}, \psi_{T-1}^l i_{T-1} | w_{T-1}) \right\} \tag{28}$$

$$= \max_{w_{T-1} \in \mathcal{W}_{T-1}} \left\{ -c_f, -c_{T-1} + r(\theta_{T-1}^u p_{T-1} | w_{T-1}) - c_l(p_d - \theta_{T-1}^u p_{T-1} | w_{T-1}) \right\} \tag{29}$$

$$= -c_f. \tag{30}$$

Note that Equation (28) follows from the monotonicity of the value function and represents the best-case analysis of purification outcomes. Equation (29) and Equation (30) follow from the induction hypothesis and the stopping cost structure as condition (ii) holds.

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

$$\mathcal{V}_t(p_t, i_t) = \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \mathop{\mathbb{E}}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\}$$

$$\leq \max_{w_t \in \mathcal{W}_t} \left\{ r_S(p_t, i_t), -c_t + \mathcal{V}_{t+1}(\theta_t^u p_t, \psi_t^l i_t | w_t) \right\} \tag{31}$$

$$\leq \max_{w_t \in \mathcal{W}_t} \left\{ -c_f, -c_t + r\left(p_t \prod_{j=t}^{T-1} (\theta_j^u | w_j) \right) - c_l\left(p_d - p_t \prod_{j=t}^{T-1} (\theta_j^u | w_j) \right) \right\} \tag{32}$$

$$= -c_f \tag{33}$$

where, Equation (31) follows from monotonicity of the value function and indicates the best-case analysis of purification capabilities, and Equation (32)-(33) follow from the induction hypothesis and the stopping cost structure as condition (ii) holds. Therefore, stopping the purification at state (p_t, i_t) and chromatography step $t \in \mathcal{T}$ leads to less financial losses than continuing the purification under condition (i) or (ii), and hence the proof follows from Theorem 1. $\square$

Proof of Proposition 3 We use backward induction. By definition, the target zone at the end of the planning horizon T is

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

At $T-1$, for the state (p_{T-1}, i_{T-1}) to be element of $\mathbb{T}_T$ by step T, we have

$$J_{T-1,w} = \left\{ (p_{T-1}, i_{T-1}) \in X : p_{T-1} = \frac{p_T}{\theta_{T-1}^l | w}, i_{T-1} = \frac{i_T}{\psi_{T-1}^u | w}, (p_T, i_T) \in \mathbb{T}_T \right\} \text{ for } w \in \mathcal{W}_{T-1}, \quad (34)$$

$$\mathbb{T}_{T-1} = \bigcup_{w \in \mathcal{W}_{T-1}} J_{T-1,w}. \quad (35)$$

Repeated application of the same procedure leads to

$$J_{t,w} = \left\{ (p_t, i_t) \in X : p_t = \frac{p_{t+1}}{\theta_t^l | w}, i_t = \frac{i_{t+1}}{\psi_t^u | w}, (p_{t+1}, i_{t+1}) \in \mathbb{T}_{t+1} \right\} \text{ for } w \in \mathcal{W}_t, \quad (36)$$

$$\mathbb{T}_t = \bigcup_{w \in \mathcal{W}_t} J_{t,w} \text{ for } t = 1, \dots, T-1. \quad (37)$$

which is equivalent to Equation (10) in Proposition 3. $\quad \square$

Proof of Proposition 4: We use backward induction to generate the effective purity set $\mathbb{P}_t$. By definition, the effective purity set at the end of the planning horizon T is

$$\mathbb{P}_T = \left\{ (p_T, i_T) \in X : \frac{1 - \gamma_d}{\gamma_d} p_T \geq i_T \right\}.$$

At $T-1$, for the state (p_{T-1}, i_{T-1}) to be an element of $\mathbb{P}_T$ by step T, it is sufficient that we have

$$K_{T-1,w} = \left\{ (p_{T-1}, i_{T-1}) \in X : p_{T-1} = \frac{p_T}{\theta_{T-1}^u | w}, i_{T-1} = \frac{i_T}{\psi_T^l | w}, (p_T, i_T) \in \mathbb{P}_T \right\} \text{ for } w \in \mathcal{W}_{T-1}, \quad (38)$$

$$\mathbb{P}_{T-1} = \bigcup_{w \in \mathcal{W}_{T-1}} K_{T-1,w}. \quad (39)$$

Using backward induction, repeated application of the same procedure leads to

$$K_{t,w} = \left\{ (p_t, i_t) \in X : p_t = \frac{p_{t+1}}{\theta_t^u | w}, i_t = \frac{i_{t+1}}{\psi_t^l | w}, (p_{t+1}, i_{t+1}) \in \mathbb{P}_{t+1} \right\} \text{ for } w \in \mathcal{W}_t, \quad (40)$$

$$\mathbb{P}_t = \bigcup_{w \in \mathcal{W}_t} K_{t,w} \text{ for } t = 1, \dots, T-1. \quad (41)$$

$\square$

Proof of Theorem 2: Theorem 2 identifies the characteristics of the optimal policies for states in the risk zone $\mathbb{R}_t$ at chromatography step $t \in \mathcal{T}$. We note that all protein and impurity pairs where $a_t^*(p_t, i_t) = S$ at $t \in \mathcal{T}$ are classified as $(p_t, i_t) \in \mathbb{F}_t$ by the definition of the failure zone.

First, we classify the pooling actions into two distinct sets: $\bar{W}_t = \{\bar{w}_t \in \mathcal{W}_t : (\theta_t^u p_t, \psi_t^l i_t | \bar{w}_t) \notin \mathbb{P}_{t+1}\}$, and $\check{W}_t = \{\check{w}_t \in \mathcal{W}_t : (\theta_t^u p_t, \psi_t^l i_t | \check{w}_t) \in \mathbb{P}_{t+1}\}$ for all $t \in \mathcal{T}$. An example of action type $\bar{w}_t$ could be a pooling window that leads from the risk zone to 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, the value function and the stopping cost structure of this revised problem can be rewritten as:

$$\mathcal{V}_t(p_t, i_t) = \max_{w_t \in \{\bar{W}_t \cup \check{W}_t\}} \left\{ r_S(p_t, i_t), -c_t + \mathbb{E}_{\theta_t, \psi_t | w_t} \mathcal{V}_{t+1}(\theta_t p_t, \psi_t i_t) \right\} \tag{42}$$

and

$$\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T) \tag{43}$$

where, at the end of the planning horizon T, we have

$$r_S(p_T, i_T) = \begin{cases} -c_f & \text{if } (p_T, i_T) \notin \mathbb{P}_T, \\ r(p_d) & \text{if } (p_T, i_T) \in \mathbb{P}_T \text{ and } p_t \geq p_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } (p_T, i_T) \in \mathbb{P}_T \text{ and } p_t < p_d. \end{cases} \tag{44}$$

As a result of Equations (42)-(44), we observe that the optimal pooling actions at step $T-1$ have the characteristic of keeping the system state in the effective purity set of the next period, i.e., $a^*_{T-1}(p_{T-1}, i_{T-1}) = \{\check{w}_{T-1} \in \check{\mathcal{W}}_{T-1} : (p_T, i_T | p_{T-1}, i_{T-1}, \check{w}_{T-1}) \in \mathbb{P}_T\}$ for all $(p_{T-1}, i_{T-1}) \in \mathbb{R}_{T-1}$.

At step $t \in \mathcal{T}$, by the definition of the desired purity set in Proposition 4, we observe that a batch state $\{(p_{t+1}, i_{t+1}) \notin \mathbb{P}_{t+1} | (p_t, i_t) \in \mathbb{R}_t\}$ has no chance of meeting the final purity requirement by step T, even under the best-case realizations of the purification capabilities. Hence, the cost structure in Equation (44) indicates that $\mathcal{V}_{t+1}(p_{t+1}, i_{t+1}) = r_S(p_{t+1}, i_{t+1}) = -c_f$ for all $(p_{t+1}, i_{t+1}) \notin \mathbb{P}_{t+1}$. As a result, the optimal action at step $t \in \mathcal{T}$ has the characteristic that $a^*_t(p_t, i_t) = \{\check{w}_t \in \check{\mathcal{W}}_t : (p_{t+1}, i_{t+1} | p_t, i_t, \check{w}_t) \in \mathbb{P}_{t+1}\}$ for all $(p_t, i_t) \in \mathbb{R}_t$ at chromatography step $t \in \mathcal{T}$. $\square$

Proof of Theorem 3 Theorem 3 analyzes the necessary condition of the optimal policy for $(p_t, i_t) \in \mathbb{T}_t$, $t \in \mathcal{T}$ in Case 1. We use backward induction. At the end of the planning horizon T, we have $\mathcal{V}_T(p_T, i_T) = r_S(p_T, i_T)$ where

$$r_S(p_T, i_T) = \begin{cases} -c_f & \text{if } \gamma_T < \gamma_d, \\ r(p_d) & \text{if } (p_T, i_T) \in \mathbb{T}_T \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } \gamma_T \geq \gamma_d \text{ and } p_t < p_d. \end{cases} \tag{45}$$

Hence, the optimal pooling action at $T-1$ is to perform the purification in such a way as $a^*_{T-1}(p_{T-1}, i_{T-1}) = \{w^*_{T-1} \in \mathcal{W}_{T-1} : (\theta^l_{T-1} p_{T-1}, \psi^u_{T-1} i_{T-1} | w^*_{T-1}) \in \mathbb{T}_T \mid (p_{T-1}, i_{T-1}) \in \mathbb{T}_{T-1}\}$ 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 3, Definition 2 and Observation 1 ensure that there exists at least one such policy. Analysis at step $t \in \mathcal{T}$ proceeds similarly. Note that the bounds on the value function in Section 4.3 indicate that $\mathcal{V}^*_t(p_t, i_t) = -c_f$ for all $(p_t, i_t) \in \mathbb{F}_t$ at $t \in \mathcal{T}$, and $\sum_{j=t}^{T-1} -c_j + r(p_d) \leq \mathcal{V}^*_t(p_t, i_t) \leq r(p_d)$ for all $(p_t, i_t) \in \mathbb{T}_t$ at $t \in \mathcal{T}$. Hence, based on the bounds of the value function, the necessary condition for the optimal pooling policy is $a^*_t(p_t, i_t) = \{w^*_t \in \mathcal{W}_t : (\theta^l_t p_t, \psi^u_t i_{t+1} | w^*_t) \in \mathbb{T}_{t+1} \mid (p_t, i_t) \in \mathbb{T}_t\}$ for all $(p_t, i_t) \in \mathbb{T}_t$ at $t \in \mathcal{T}$. $\square$

Proof of Theorem 4   Since Case 2 is relaxing the yield requirement from Case 1, The proof is analogous to that of Theorem 3, and hence omitted.   □

Proof of Proposition 5   Let $(p'_t, i'_t) \in \mathbb{F}_t$ and the (sub)set $\{I_{p_t \leq p'_t, i_t \geq i'_t}\}$ represent all protein-impurity pairs satisfying Proposition 5. Note that the conditions in Proposition 5 correspond to the failure zone. Therefore, based on Theorem 1, we have $\mathcal{V}^*_t(p_t, i_t) = -c_f$ for the states $\{I_{p_t \leq p'_t, i_t \geq i'_t}\}$ specified in Proposition 5. Hence, the aggregate failure state d_t 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 the same costs and transitions. The proof follows from Bertsekas (2012), Vol.1, page 321.   □

Proof of Proposition 6:   First, we fix any protein-impurity pair $(p_t, i_t) \in \mathcal{P} \times \mathcal{T}$ at chromatography step $t \in \mathcal{T}$. Let w^i_t and w^j_t be two distinct pooling windows at chromatography step $t \in \mathcal{T}$, such that, $F_t(\Theta|w^i_t) \geq_{st} F_t(\Theta|w^j_t)$, $G_t(\Psi|w^i_t) \leq_{st} G_t(\Psi|w^j_t)$, and $(\theta^l_t|w^i_t) < (\theta^l_t|w^j_t)$, $(\theta^u_t|w^i_t) < (\theta^u_t|w^j_t)$, and $(\psi^l_t|w^i_t) > (\psi^l_t|w^j_t)$, $(\psi^u_t|w^i_t) > (\psi^u_t|w^j_t)$, as specified in Proposition 6. Next, we evaluate the value function $\mathcal{V}_t(p_t, i_t|w^j_t)$ of state (p_t, i_t) under the pooling action w^j_t at chromatography step $t \in \mathcal{T}$:

$$
\begin{aligned}
\mathcal{V}_t(p_t, i_t|w^j_t) &= \max\Big\{r_s(p_t, i_t), -c_t + \int_{\psi^l_t|w^j_t}^{\psi^u_t|w^j_t} \int_{\theta^l_t|w^j_t}^{\theta^u_t|w^j_t} f_t(\theta_t|w^j_t)g_t(\psi_t|w^j_t)\mathcal{V}_{t+1}(p_t\theta_t, \psi_t i_t|w^j_t)\mathrm{d}\theta\mathrm{d}\psi\Big\} \\
&> \max\Big\{r_s(p_t, i_t), -c_t + \int_{\psi^l_t|w^i_t}^{\psi^u_t|w^i_t} \int_{\theta^l_t|w^i_t}^{\theta^u_t|w^i_t} f_t(\theta_t|w^i_t)g_t(\psi_t|w^i_t)\mathcal{V}_{t+1}(p_t\theta_t, \psi_t i_t|w^i_t)\mathrm{d}\theta\mathrm{d}\psi\Big\} (46) \\
&= \mathcal{V}_t(p_t, i_t|w^i_t).
\end{aligned}
$$

Note that Equation (46) follows from the conditions in Proposition 6 and the monotonicity of the value function in Proposition 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 $\mathcal{V}_t(p_t, i_t|w^j_t)$ under the pooling window w^j_t denotes strictly higher profit then the value function $\mathcal{V}_t(p_t, i_t|w^i_t)$ under the pooling window w^i_t. Hence, w^i_t is said to be strictly dominated by w^j_t at chromatography step $t \in \mathcal{T}$ since $\mathcal{V}_t(p_t, i_t|w^j_t) > \mathcal{V}_t(p_t, i_t|w^i_t)$, and thus $a^*_t(p_t, i_t) \neq w^i_t$.   □