Managing trade-offs in protein manufacturing

Citation for published version (APA):

Martagan, T. G., Krishnamurthy, A., & Leland, P. (2020). Managing trade-offs in protein manufacturing: how much to waste? *Manufacturing & Service Operations Management*, 22(2), 330-345.
<https://doi.org/10.1287/msom.2018.0740>

DOI:

[10.1287/msom.2018.0740](https://doi.org/10.1287/msom.2018.0740)

Document status and date:

Published: 01/03/2020

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Managing Trade-offs in Protein Manufacturing: How Much to Waste?

Tugce Martagan

School of Industrial Engineering, Eindhoven University of Technology, t.g.martagan@tue.nl,

Ananth Krishnamurthy

Department of Industrial Engineering, University of Wisconsin-Madison, ananth.krishnamurthy@wisc.edu,

Peter A. Leland

Aldevron, leland@aldevron.com,

We consider the challenges and trade-offs involved in the manufacturing of engineered proteins. Manufacturing these proteins involves high risk of financial losses due to the purity and yield trade-offs, uncertainty in the process outcomes, and high operating costs. In this setting, the biomanufacturer must determine how much protein to manufacture in the upstream fermentation operations, and then how much of it to waste in each subsequent purification operation because of the purity–yield trade-offs. We develop a Markov decision model to optimize three layers of interdependent decisions in protein manufacturing: the optimal amount of protein to be produced in upstream operations, the optimal choice of chromatography technique to be used in downstream operations, and the optimal choice of pooling windows during chromatography. The proposed stochastic model dynamically optimizes these three layers of interdependent decisions to maximize the expected profit. The structural analysis derives functional relationships between the purity–yield trade-offs and operating costs, and characterizes the optimal operating policies. The optimal policy also suggests when the biomanufacturer is better off failing early and cutting losses. We use a state aggregation scheme to reduce the computational efforts, and quantify the savings obtained from the use of the optimization model in industry practice at Aldevron.

Key words: Stochastic optimization, protein manufacturing, quality requirement, random yield.

1. Introduction

More than 8,000 therapeutics are in the global pharmaceutical research and development pipeline to treat cancer, diabetes, and many other diseases (Long and Works 2013). These therapeutics are the “next-generation” drugs produced by biomanufacturing technologies. Unlike traditional pharmaceuticals that are chemically synthesized, biopharmaceuticals are manufactured through live systems (e.g., bacteria, viruses, mammalian cells, etc.). The use of live systems introduces unique manufacturing challenges and trade-offs due to the uncertainty in production outcomes. In this paper, we focus on the biomanufacturing of

engineered proteins. These proteins are uniquely designed and manufactured as part of pharmaceutical research and development (R&D).

Protein manufacturing operations consist of two main steps: upstream fermentation and downstream purification processes. The upstream fermentation process results in a batch mixture containing the protein of interest along with several unwanted impurities. The downstream purification process consists of a series of chromatographic steps to eliminate the unwanted impurities. To satisfy the customer order, the biomanufacturer must meet the yield requirement (i.e., the desired amount of protein) and purity requirement (i.e., the minimum acceptable quality). However, meeting both of these requirements is challenging because of the purity–yield trade-offs. For example, the biomanufacturer needs to *waste* some amount of the protein at each chromatography step to improve the batch purity. If too much protein is wasted, then the yield requirement might not be achieved. In contrast, if too little protein is wasted, then the purity requirement might not be achieved. High purity requirements impose another layer of challenge. For example, the batch often needs to be *free* of impurities if the potential end users are humans. To achieve such high purity targets, biomanufacturers need to successfully manage the purity–yield trade-offs at multiple interdependent chromatography steps.

Although the purity–yield trade-off is inevitable, it is possible to control it by simultaneously optimizing the choice of *chromatography technique* and *pooling window* (i.e., the chromatography operating policy). In the presence of multiple impurities, simultaneously optimizing these two chromatography decisions is challenging because each impurity exhibits a distinct separation performance in different chromatography techniques. Therefore, it is important to choose the right chromatography technique and pooling window to simultaneously eliminate multiple impurities based on their purity–yield trade-offs.

An expensive but feasible way of addressing these challenges is to deliberately increase the amount of protein obtained from fermentation. Producing higher amounts of protein during fermentation might alleviate the failure risks and the purity–yield trade-offs involved in chromatography operations. However, it also leads to higher operating costs. In addition, each production order involves a predetermined yield requirement, and producing proteins in excess of this requirement does not guarantee additional revenue.

Therefore, a protein manufacturing project is required to successfully manage three layers of operating decisions: (1) optimizing the protein mass obtained from fermentation

(see Section 2.2), (2) optimizing the pooling window used at a chromatography technique (see Section 2.3), and (3) optimizing the choice of chromatography technique used in each purification step (see Section 2.4). In practice, these three layers of operating decisions are interdependent, and need to be optimized through a unifying framework. In this paper, we develop a stochastic optimization model that simultaneously addresses these three levels of interdependent decisions, and answer the following questions: When is it optimal to fail early and avoid future losses? What is the best protein amount to be manufactured during fermentation? What are the best choices of chromatography techniques and pooling windows? How much protein should be sacrificed at each chromatography step to improve the purity? How can biomanufacturers overcome challenges in conforming to high purity requirements? Answering these questions will shed light on the practical concern of biomanufacturers regarding *how much to waste*. The waste in this setting corresponds to (i) the amount of protein produced in excess of the yield requirement, and (ii) the amount of protein sacrificed at each chromatography step to improve the batch purity.

Main contributions and organization of the paper. Our contributions are summarized as follows: (i) We build a stochastic optimization model that simultaneously addresses the upstream protein mass decisions, the chromatography technique selection decisions, and the pooling window decisions in the presence of multiple impurities. (ii) We analyze the structural properties of the upstream protein mass problem, and characterize three critical protein threshold values. These thresholds provide formal guidelines to reduce the production costs of biomanufacturing operations. The structural results provide a rigorous framework to analyze the financial trade-offs in practice. (iii) We characterize functional relationships between the purity–yield trade-offs and operating costs, and determine the structural properties of the optimal policies. Such structural characteristics of the optimal policies have not been analyzed in the literature. (iv) The resulting model is complex to analyze analytically and computationally. We develop an aggregation scheme based on the structural results, and successfully simplify the problem by transforming its continuous state space into a binary state space for each impurity. The proposed approach provides a new perspective to handle industry-scale problems at a 100% purity requirement. (v) We quantify the potential impact of the optimization model with an industry case study. The case study shows that substantial improvement in the expected profit can be achieved using the optimization framework. This research has been conducted in collaboration with

Aldevron, a biomanufacturer specializing in protein manufacturing. The research outcomes have been shared and validated with a larger biotechnology community through industry working group sessions (BioWGS 2016).

To the best of our knowledge, Martagan et al. (2017) was the first to develop a Markov decision model for protein purification operations. However, they assume that the upstream protein mass problem and the downstream chromatography technique selection problem have been solved, and mainly focus on the pooling window selection problem. This aligns with the current practice where process improvements are limited to one of the decision layers in isolation. In contrast, this paper demonstrates the value obtained when these three decision layers are optimized simultaneously. The case study in Section 7 shows that the improvements achieved by the simultaneous optimization framework are almost double of that achieved by optimizing an individual decision layer alone.

The remainder of the paper is organized as follows. The problem setting is introduced in Section 2 and the relevant literature is discussed in Section 3. The optimization model is developed in Section 4, and the structural properties are analyzed in Section 5. Section 6 presents a state aggregation scheme. A case study from Aldevron is presented in Section 7, and concluding remarks are provided in Section 8.

2. Background on Protein Manufacturing

The first stage in biomanufacturing is the upstream fermentation operations, where the cells grow and produce the protein of interest. The batch obtained from fermentation consists of a mixture of the protein of interest along with several impurities. These impurities are unwanted byproducts, such as contaminants, dead cells, ammonia, lactate, and other metabolic residues. The batch proceeds with downstream purification operations, where these impurities are eliminated through a series of chromatography operations.

Chromatography operations are carried out in columns packed with special resins binding to either proteins or impurities. A chromatography technique uses the difference in the physical and chemical properties as a separation principle. For example, a chromatography technique could exploit the difference in molecular size, charge, hydrophobicity, etc. A typical purification process consists of 2 to 6 chromatography steps. In practice, the scientist first collects data on the performance of different chromatography techniques through scouting experiments, and then uses this data during the production runs.

2.1. Yield and Purity Requirements

Purity is a typical measure of batch quality, and it represents the ratio of the amount of the protein to the total amount of the protein and impurities available in the batch. Depending on the application, the final *purity requirement* specified by the client could be up to 99% or even 100%. For example, if the users are humans, then the batch often needs to be free of unwanted impurities to satisfy regulatory requirements. In this paper, we are specifically interested in proteins whose potential end users are humans, and hence the final batch should conform to high purity requirements (i.e., $\geq 99.9\%$ purity). *The yield requirement* represents the amount (mass) of the protein that needs to be produced. We use the terms “amount” and “mass” interchangeably hereinafter. The yield requirement is specified by the client along with the purity requirement. When the yield requirement is not achieved, the biomanufacturer incurs penalty costs per unit of protein in shortage.

2.2. The Upstream Protein Mass Problem

The scientist operating the fermentation process does not have the ability to prevent or limit the formation of impurities, as they are natural metabolites. Several studies have documented evidence related to the impurities obtained during the fermentation process based on the physicochemical conditions (Tsao et al. 2004, Xing et al. 2010). The scientist has the ability to selectively increase the amount of the protein obtained from the fermentation using several controls, e.g., adjusting the harvesting times or increasing the productivity of the cell lines. The objective of this study is not to develop fermentation control policies to increase the protein amount. Instead, we focus on identifying the optimal amount of protein that should be obtained at the end of the fermentation process.

Increasing the protein amount obtained from a batch is an important research topic, and it is referred to as the *problem of increasing the titer*. Although increasing the protein mass obtained from a batch is possible, it may significantly boost the upstream operating costs, because these costs are nondecreasing in the amount of protein. Increasing the protein mass to more than is required could hurt the profitability of a batch, as the customers may not purchase the proteins manufactured in excess of their yield requirement. Identifying the best amount of protein to be obtained from the upstream batch is a challenging problem because of the random yield losses and uncertain purity outcomes in the downstream purification operations. Although higher protein mass increases the upstream costs, it alleviates the downstream purification risks. Due to these trade-offs, a formal decision-making framework is required in practice.

Figure 1 An example of chromatography outcome with multiple impurities (Industry data from Aldevron)

2.3. The Pooling Window Selection Problem

Figure 1 (a) shows an example of the outcome obtained from a chromatography operation. Each column in Figure 1 (a) is called a *lane*, and represents the volume flowing through the chromatography equipment during a specified time interval (e.g., one minute). The y-axis in Figure 1 (a) describes the composition of the volume collected as a function of the molecular size of the protein and impurity constituents. Each lane is comprised of some fraction of the total amount of the protein of interest, and some fraction of different types of impurities. For example, consider lane 8 in Figure 1 (a). In this lane, there are 8 different types of molecules; one of them is the protein of interest, and the remaining ones are different types of impurities (as shown by the arrows). We translate the chromatography outcome in Figure 1 (a) into the diagram in Figure 1(b). Figure 1(b) plots the fraction of total protein mass per lane and two representative impurities.

The scientist decides which *chromatography technique* and *pooling window* to use simultaneously. First, we introduce the pooling window decision for a given chromatography technique. For simplicity, we plot only two different impurity types in Figure 1(b), namely Impurity A and B, and use this example to discuss the purity–yield trade-off in pooling decisions. If the scientist pools lanes 3–15 in Figure 1(b), then she collects all protein along with Impurity A and B. On the other hand, if she pools lanes 9–15, she compromises on the yield (i.e., collects a smaller fraction of the protein) but completely eliminates Impurity A and some fraction of Impurity B. Alternatively, the scientist might decide to significantly compromise on the yield (i.e., lose $\approx 55\%$ of the protein) by pooling lanes

Figure 2 Difference in the separation capability of two chromatography techniques (Industry data from Aldevron)

12–15, but this decision helps to achieve 100% purity because all unwanted impurities are eliminated. This example illustrates the purity–yield trade-off typically encountered in chromatography operations.

Note that Figure 1 illustrates only one sample outcome of a chromatography step. However, in practice, the amount of protein and impurities associated with each pooling window is random. This randomness adds another layer of challenge in decision-making. In addition, pooling window decisions can become more complex when a chromatography technique demonstrates a distinct affinity toward different types of impurities. In practice, depending on the outcome of a chromatography step, the scientist makes decisions regarding the chromatography technique *and* the pooling window for the subsequent step. This leads to the related problem of chromatography technique selection.

2.4. The Chromatography Technique Selection Problem

Figure 2 shows the output of two different chromatography techniques based on industry data. Both outputs shown in Figure 2 use the same starting material but different techniques to separate the protein from unwanted impurities. The starting material contains the protein along with Impurity A and B. The x-axis in Figure 2 represents the lanes, and the y-axis denotes the fraction of the molecules corresponding to each lane. The solid curve in Figure 2 represents the protein, and the other two curves represent the impurities.

Note that the choice of chromatography technique influences the relative positions of the molecules and their corresponding amount in each lane. For example, Impurity A is located

on the left of the protein under the first chromatography technique in Figure 2. However, the same impurity overlaps completely with the protein under the second chromatography technique in Figure 2. It is clear that the second technique is not capable of separating Impurity A from the protein of interest. However, it provides a better separation outcome for Impurity B than first chromatography technique. The relative locations of the protein and impurities and their corresponding amounts in each lane are complex functions of the physical and chemical properties, and vary for each chromatography technique. Note that the choice of chromatography technique directly influences the underlying purity–yield trade-offs, as the relative locations of protein and impurities are different for each chromatography technique.

2.5. Interdependency between Chromatography Technique and Pooling Window

The pooling window decision is strongly influenced by the choice of chromatography technique. For example, if the scientist needs to completely eliminate Impurity B in Figure 2, then she could pool lanes 12–15 on the first chromatography technique (wasting $\approx 55\%$ of the protein of interest), or she could pool lanes 2–11 on the second chromatography technique without compromising the protein. This illustrates the complex interdependency between the choice of chromatography technique, pooling windows, and purity–yield trade-offs. In practice, a starting material could contain up to 100 different impurities to be separated using 5 to 10 available chromatography techniques. The combinatorial complexity arising from multiple impurities, chromatography techniques, pooling windows, and purity–yield trade-offs makes the purification decisions complex in practice.

2.6. Decision Making in Practice

We organized a series of working group sessions, and interviewed several small- and medium-sized biomanufacturers to understand the industry-wide practice. We learned that operating decisions are often made based on a combination of the industry guidelines and scientists’ experience. For example, a standard purification protocol suggests a three-step strategy with (1) a capture step to remove critical contaminants, (2) intermediate step(s) to remove bulk impurities, and (3) a polishing step to eliminate remaining impurities. The handbook of GE Healthcare provides an excellent summary of the common guidelines for purification processes (GE Healthcare 2010). These protocols often match an impurity type with a chromatography technique, e.g., if the impurity is affinity-tagged, then use the anion

exchange chromatography (GE Healthcare 2010). Popular heuristics also use the highest resolution criteria (i.e., the chromatography technique that has the least overlap between the protein and impurities) or the highest selectivity criteria (i.e., the chromatography technique that has the highest distance between the protein and impurity peaks) as a heuristic for decision-making. On a given chromatography technique, the scientist identifies the specific type of impurities to be eliminated, and selects the pooling window accordingly. Note that these guidelines and heuristics are defined based on the biological and chemical dynamics, but do not incorporate the financial trade-offs into decision-making.

For the upstream protein mass decisions, companies often plan for excess production to buffer against uncertainties in subsequent operations. Company-specific guidelines are often used to constrain a minimum and maximum limit on the production quantities, e.g., “do not produce more than threefold of the yield requirement,” “only buffer against the expected yield losses in downstream,” or “always run the bioreactor at maximum capacity.” Based on the working group sessions, we observed that Operations Research (OR) applications in the biomanufacturing industry are still in infancy. However, biomanufacturers across the world are realizing that they need to undergo a data-driven OR-based transformation to fully realize the societal benefits of bioscience research.

3. Literature Review

The relevant literature can be classified into two categories: fermentation optimization and chromatography optimization. In the field of fermentation optimization, several studies focus on developing models for cell growth and product formation (Patel et al. 2000, Tsao et al. 2004, Xing et al. 2010). These cell growth models are also incorporated into optimization models to determine the optimal control policies that increase the titer (Saucedo and Karim 1997, Yang et al. 2000, Peroni et al. 2005, Gnoth et al. 2007, Martagan et al. 2016). However, the literature typically focuses on the underlying biology of the fermentation process, and often does not account for the uncertainties, financial trade-offs, and interactions between the fermentation and chromatography decisions. There are limited number of simulation studies that model the manufacturing uncertainties (Saraph 2003, Petrides and Siletti 2004), but these simulation models do not yield insights on the structural characteristics of the optimal policies. To address this gap, we develop a stochastic optimization model and provide analytic results on the optimal policies and financial trade-offs.

The literature on chromatography operations includes several studies on the chromatography technique selection problem. For example, Vasquez-Alvarez and Pinto (2004) presented two mixed-integer linear programming (MILP) models to optimize the choice of chromatography techniques. One of the MILP models minimizes the number of steps to achieve the desired purity, whereas the other MILP model maximizes the purity. The authors of Lienqueo et al. (2009) developed a mixed integer nonlinear programming (MINLP) model to optimize the choice of polypeptide tag and chromatography steps to maximize profit. A flowsheet optimization mechanism was developed by Nfor et al. (2013) to eliminate multiple impurities with the minimum number of steps. However, the aforementioned studies do not account for the interaction between the chromatography technique and pooling window decisions. There are only a few studies that consider the interdependency between these two operating decisions. For example, Polykarpou et al. (2011) developed an MINLP model that determines the best chromatography technique and pooling window to minimize the number of steps. Then, Polykarpou et al. (2012) extended this work by developing an MILP model that overcomes the drawbacks of the nonlinear model. The proposed MILP model helps to achieve a predetermined purity level with the minimum number of steps. However, existing studies often focus on deterministic models, and aim to minimize the number of steps. In contrast, we develop a stochastic optimization model that captures the process uncertainties and financial trade-offs involved in both the fermentation and purification operations.

In the context of stochastic optimization, Martagan et al. (2017) was the first to build a Markov decision model to optimize the pooling windows. However, the authors assume that the upstream production quantity and the choice of chromatography technique are predetermined, and only focus on optimizing the pooling windows. The authors characterize new performance guarantees for meeting the production requirements, and propose a zone-based decision-making approach. In contrast, this paper captures a broader scope, as it addresses the upstream protein mass problem, the chromatography technique selection problem, and the pooling window selection problem simultaneously.

4. Model Formulation

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

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

State Space: First, we focus on the state space for the downstream purification problem. Let $p_t \in \mathcal{P}$ represent the amount (mass) of the protein available in the batch at the beginning of the chromatography step $t \in \mathcal{T}$. Note that $p_t \in [0, p_1]$ for all chromatography steps $t \in \mathcal{T}$, because the amount of protein p_t available at the beginning of each step t is bounded by the starting material p_1 . Next, we define the different types of impurities and their corresponding amounts available in the batch. Let $\mathcal{K} = \{k : 1, 2, \dots, K\}$ be the set of K distinct types of impurities, where $K < \infty$. Let $i_{k,t} \in \mathcal{I}_k$ denote the amount of impurity type $k \in \mathcal{K}$ available in the batch at the beginning of chromatography step $t \in \mathcal{T}$. The impurity state $(i_{1,t}, \dots, i_{K,t}) \in \mathcal{I}_1 \times \dots \times \mathcal{I}_K$ represents the set of all distinct types of impurities $k \in \mathcal{K}$ and their corresponding amounts $i_{k,t}$ available in the batch at the beginning of the chromatography step $t \in \mathcal{T}$. Note that $i_{k,t} \in [0, i_{k,1}]$ for each impurity type $k \in \mathcal{K}$ at chromatography step $t \in \mathcal{T}$. The amount of protein p_1 and the amount of impurities $i_{k,1}$ for each impurity type $k \in \mathcal{K}$ at the beginning of the first chromatography step is determined by the upstream fermentation process. We define the state Δ as the stopping state, i.e., an absorbing state with zero cost representing the end of the optimization problem, where the batch is either shipped or scrapped. Therefore, $(p_t, i_{1,t}, \dots, i_{K,t}) \cup \Delta$ is the state of the downstream purification problem for all chromatography steps $t \in \mathcal{T}$.

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

process. Therefore, the starting state of the upstream protein mass problem is $p_0 = 0$ and $i_{k,0} = 0$ for all impurity types $k \in \mathcal{K}$.

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

In the upstream operations, the scientist has the ability to selectively increase the protein mass p_1 in the range $[0, p_{max}]$ through a set of controls. In practice, these controls correspond to manufacturing protocols related to the choice of buffers, adjustment of physical and biological parameters, feeding strategies, etc. Determining these controls itself is a process design problem (see Saucedo and Karim 1997, Yang et al. 2000, Xing et al. 2010, Martagan et al. 2016), and is outside our scope. Instead, we are interested in determining the protein amount p_1 to be produced in the upstream fermentation operation. Therefore, $p_1 \in \mathcal{P}$ is the decision variable in the upstream operation, and this production amount serves as starting material for the downstream purification operations.

Transitions: In the downstream purification problem, the state transitions define the amount of the protein of interest p_t and the amount $i_{k,t}$ of each impurity type $k \in \mathcal{K}$ that remains in the batch after the chromatography technique $c \in \mathcal{C}$ is performed using the

pooling window $w_c \in \mathcal{W}_c$ at purification step $t \in \mathcal{T}$. First, we model the changes in $i_{k,t}$ after the completion of the chromatography step $t \in \mathcal{T}$. For a given impurity type $k \in \mathcal{K}$, we let the random variable $(\Psi_k|c, w_c)$ denote the fraction of impurity amount that remains in the batch after using the chromatography technique c and the pooling window w_c . Therefore,

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

where $(\psi_k|c, w_c)$ is the realization of $(\Psi_k|c, w_c)$. The random variable $(\Psi_k|c, w_c)$ has the density function $g_k(\cdot|c, w_c)$ with support $[0, 1]$ for $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ and $k \in \mathcal{K}$. This density function is characterized based on the physicochemical properties of the impurity $k \in \mathcal{K}$ and its response to the chromatography technique $c \in \mathcal{C}$ (See Vasquez-Alvarez et al. (2001) and Polykarpou et al. (2011) for details). Because a chromatography technique exploits the physicochemical properties of the protein and impurities to separate one from the other, the density function $g_k(\cdot|c, w_c)$ is unique for each impurity type $k \in \mathcal{K}$, and is independent of time $t \in \mathcal{T}$ and impurity amount $i_{k,t}$ for all $k \in \mathcal{K}$. The density function $g_k(\cdot|c, w_c)$ is defined by the choice of chromatography technique c , pooling window w_c , and physicochemical characteristics of impurity type $k \in \mathcal{K}$.

Similarly, we let the random variable $(\Theta|c, w_c)$ denote the fraction of protein p_t that remains in the batch at the beginning of chromatography step $t + 1$, given that there are p_t units of protein at the beginning of the chromatography step t , and the action (c, w_c) is performed. This implies that the remaining amount of the protein is wasted during the chromatography step t . Therefore,

$$p_{t+1} = (\theta|c, w_c)p_t, \quad (2)$$

where $(\theta|c, w_c)$ is the realization of $(\Theta|c, w_c)$. The random variable $(\Theta|c, w_c)$ has the density function $f(\cdot|c, w_c)$ with support $[0, 1]$ for all $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ and $k \in \mathcal{K}$. Note that $f(\cdot|c, w_c)$ is a function of the chromatography technique $c \in \mathcal{C}$ and pooling window $w_c \in \mathcal{W}_c$, but is independent of step $t \in \mathcal{T}$ and impurities $i_{k,t}$ for all $k \in \mathcal{K}$. The state transitions as a result of the action $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$ are therefore captured by

$$(p_{t+1}, i_{1,t+1}, \dots, i_{K,t+1}) = (\theta p_t, \psi_1 i_{1,t}, \dots, \psi_K i_{K,t} | c, w_c). \quad (3)$$

The realizations $(\theta, \psi_1, \dots, \psi_K|c, w_c)$ in Equations (1)–(3) are independent of protein mass p_t , impurity mass i_t and chromatography step t , but are dependent on the chromatography

technique c , pooling window w_c and impurity type k . This is because a chromatography technique separates a specific protein from an impurity by exploiting the physicochemical properties, such as molecular weight, electric charge, and hydrophobicity. These physicochemical characteristics are specific to each molecule, and independent of other impurities (Vasquez-Alvarez et al. 2001, Polykarpou et al. 2011). The system transitions to the stopping state Δ when $a_t(p_t, i_{1,t}, \dots, i_{K,t}) = S$ at $t \in \mathcal{T} \cup \{0\}$. At the final time T , the only available action is to stop, i.e., $a_T(p_T, i_{1,T}, \dots, i_{K,T}) = S$ for all $p_T \in \mathcal{P}$, $i_{k,T} \in \mathcal{I}_k$, $k \in \mathcal{K}$.

In the upstream protein mass problem, let $(i_{1,1}, \dots, i_{K,1})$ denote the impurities that result from the fermentation process. Thus, if the scientist decides to produce p_1 units through a set of predetermined manufacturing protocols, then the system transitions from the initial state $(p_0, i_{1,0}, \dots, i_{K,0})$ to the state $(p_1, i_{1,1}, \dots, i_{K,1})$.

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

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

$$r_s(p_t, i_{1,t}, \dots, i_{K,t}) = \begin{cases} r(p_d) & \text{if } p_t \geq p_d \text{ and } \gamma_t \geq \gamma_d, \\ r(p_t) - c_\ell(p_d - p_t) & \text{if } p_t < p_d \text{ and } \gamma_t \geq \gamma_d, \\ -c_f & \text{otherwise.} \end{cases} \quad (4)$$

If the batch meets the purity and yield requirements (i.e., $p_t \geq p_d$ and $\gamma_t \geq \gamma_d$ in Equation 4) then the biomanufacturer obtains revenue $r(p_d)$. In this case, the customer only purchases the amount ordered p_d , and does not pay for proteins produced in excess. On the other hand, if the batch satisfies the purity requirement but not the yield requirement (i.e., $\gamma_t \geq \gamma_d$ and $p_t < p_d$ in Equation 4), then the biomanufacturer obtains revenue $r(p_t)$ which is a function of the protein amount p_t produced. However, the biomanufacturer also incurs a yield penalty cost $c_\ell(p_d - p_t)$, which is a function of the yield shortage $(p_d - p_t)^+$. We assume that $r(p_t)$ is increasing in p_t when $p_d \geq p_t \geq 0$, and $c_\ell(p_d - p_t)$ is decreasing in p_t when $p_d \geq p_t \geq 0$. If the batch does not conform to the purity requirement (i.e., $\gamma_t < \gamma_d$ in Equation 4) then no revenue is obtained, and the biomanufacturer incurs a large penalty cost of failure c_f . Note that $c_f > r(p_d)$, and the yield penalty cost could be very large depending on the shortage amount $(p_d - p_t)^+$. Note that $c_f > c_\ell(p_d - p_t)$ for all $p_d > p_d - p_t \geq 0$ and $c_f \geq c_\ell(p_d)$.

Upstream Costs: The cost of upstream fermentation operations is captured by $c_u(p_1)$, and represents the operating costs (e.g., labor, equipment, and raw materials such as buffers and cell lines) required to obtain $p_1 \in \mathcal{P}$ units of protein at the end of the fermentation given that $p_0 = 0$ and $i_{k,0} = 0$ for all $k \in \mathcal{K}$. We assume that $c_u(p_1)$ is nondecreasing in p_1 , as additional resources are required to increase the protein mass during fermentation.

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

$$\mathcal{V}_T(p_T, i_{1,T}, \dots, i_{K,T}) = r_s(p_T, i_{1,T}, \dots, i_{K,T}). \quad (5)$$

For all $t \geq 1$, the value function of the downstream purification problem is

$$\mathcal{V}_t(p_t, i_{1,t}, \dots, i_{K,t}) = \max_{(c, w_c) \in \mathcal{C} \times \mathcal{W}} \left\{ -r_c + \mathbb{E}[\mathcal{V}_{t+1}(\Theta p_t, \Psi_1 i_{1,t}, \dots, \Psi_K i_{K,t}) | c, w_c], \right. \\ \left. r_s(p_t, i_{1,t}, \dots, i_{K,t}) \right\}, \quad (6)$$

where the expected value function at $t + 1$ is

$$\mathbb{E}[\mathcal{V}_{t+1}(\Theta p_t, \Psi_1 i_{1,t}, \dots, \Psi_K i_{K,t}) | c, w_c] \\ = \int_{\theta} \int_{\psi_1} \dots \int_{\psi_K} \mathcal{V}_{t+1}(\theta p_t, \psi_1 i_{1,t}, \dots, \psi_K i_{K,t}) \times \\ f(\theta | c, w_c) \times g_1(\psi_1 | c, w_c) \times \dots \times g_K(\psi_K | c, w_c) d\psi_K \dots d\psi_1 d\theta. \quad (7)$$

We combine the upstream protein mass problem at $t = 0$ with the downstream purification problem at $t = 1$, and obtain the objective function $\mathcal{V}_0(p_1)$ of the biomanufacturer:

$$\mathcal{V}_0(p_1) = -c_u(p_1) + \mathcal{V}_1(p_1, i_{1,1}, \dots, i_{K,1}). \quad (8)$$

We let $\mathcal{V}_0^*$ denote the optimal expected profit of the optimization problem in Equation (8), i.e., $\mathcal{V}_0^* = \max_{p_1 \in \mathcal{P}} \mathcal{V}_0(p_1)$. The model is a non-discounted finite-horizon model because it represents a short-term planning horizon (each chromatography step takes less than 6 hours) compared to the overall protein R&D timeline (often more than 3 months). Therefore, discounting the value function could lead to a bias in this setting.

5. Structural Analysis

We investigate the structural properties of the downstream purification problem at $t \geq 1$ (Section 5.1), and then use these insights to analyze the upstream protein mass problem at $t = 0$ (Section 5.2). In the remainder of the paper, we use a discretization scheme to analyze the structural properties of the optimal value function and policies. All proofs are available in the electronic companion. For notational convenience, we suppress the subscript t when possible in the remainder of the paper.

5.1. Analysis of the Downstream Purification Problem

First, we present the modeling assumptions on the downstream purification operations.

ASSUMPTION 1. Let $w_c^n \in \mathcal{W}_c$ be the n^{th} pooling window on a chromatography technique $c \in \mathcal{C}$ at any step $t \in \mathcal{T}$. For each technique c , the pooling windows $w_c^n \in \mathcal{W}_c$ can be ordered such that $\int_j^1 f(\theta | c, w_c^n) d\theta \leq \int_j^1 f(\theta | c, w_c^{n+1}) d\theta$ and $\int_m^1 g_k(\psi_k | c, w_c^n) d\psi_k \leq \int_m^1 g_k(\psi_k | c, w_c^{n+1}) d\psi_k$ for all $k \in \mathcal{K}$ on a given technique $c \in \mathcal{C}$, $0 \leq j \leq 1$, $0 \leq m \leq 1$.

ASSUMPTION 2. Let $p^+, p \in \mathcal{P}$ with $p^+ > p$, and $i_k^+, i_k \in \mathcal{I}$ with $i_k^+ > i_k$ for $k \in \mathcal{K}$. Then, $r_s(p, i_1, \dots, i_K) - r_s(p, i_1, \dots, i_k^+, \dots, i_K) \geq r_s(p^+, i_1, \dots, i_k, \dots, i_K) - r_s(p^+, i_1, \dots, i_k^+, \dots, i_K)$.

Using the index n , Assumption 1 presents an ordering scheme for the pooling windows $w_c^n \in \mathcal{W}_c$ for a given chromatography technique $c \in \mathcal{C}$. This assumption is a broader version of the increasing failure rate property used in Markovian deterioration models. In practice, Assumption 1 reflects the purity–yield trade-off discussed in Section 2, and makes the downstream purification problem practically relevant and challenging. This assumption is validated with industry data, and holds in practice because of the principles of chromatographic separation described in Section 2. Note that Assumption 1 does not guarantee a stochastic ordering in chromatography techniques. In practice, different chromatography techniques have different affinities to each impurity type, and therefore a stochastic ordering across chromatography techniques does not often exist. Assumption 2 is validated with industry data, and indicates that the penalty associated with an incremental increase in impurity is higher at lower levels of protein than higher levels of protein.

PROPOSITION 1. The value function $\mathcal{V}_t(p_t, i_{1,t}, \dots, i_{K,t})$ is nondecreasing in p_t for all $i_{k,t}$, $k \in \mathcal{K}$, and nonincreasing in $i_{k,t}$ for all $p_t \in \mathcal{P}$, $k \in \mathcal{K}$ at the chromatography step $t \geq 1$.

Proposition 1 indicates that the value function associated with the downstream purification problem never decreases as the protein amount p_t increases, and never increases as the impurity amount $i_{k,t}$ increases for all impurity types $k \in \mathcal{K}$ at chromatography step $t \in \mathcal{T}$. The monotonicity of the value function is used in Theorems 1–3 to characterize the structural properties of the optimal purification policies.

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

Theorem 1 shows that there exist three protein threshold values, $\check{p}_t$, $\bar{p}_t$, and $\hat{p}_t$, for a given impurity level $(i_{1,t}, \dots, i_{K,t})$ at chromatography step $t \in \mathcal{T}$. The threshold values associated with the first chromatography step have important managerial implications. For example, $p_1 \leq \check{p}_1$ corresponds to a deficient starting material that eventually leads to failure. If the upstream protein mass is $\check{p}_1 < p_1 < \bar{p}_1$, then it means that the amount of protein obtained

from the fermentation is *too little*, and the biomanufacturing company is expected to incur a financial loss rather than profit. The threshold value $\bar{p}_1$ represents the break-even point for the downstream operations. Finally, the threshold value $\hat{p}_1$ helps to prevent *overproduction*, as manufacturing more than $\hat{p}_1$ units of protein does not improve the expected profit. Using these insights, Theorem 2 identifies the batch states $(\check{p}_t, i'_{1,t}, \dots, i'_{K,t})$ that would lead to failure at $t \in \mathcal{T}$.

THEOREM 2. *There exists $(\check{p}_t, i'_{1,t}, \dots, i'_{K,t})$ such that: (i) $\frac{\check{p}_t}{\check{p}_t + \sum_k i'_{k,t}} < \gamma_d$, and (ii) for any state $(p_t, i_{1,t}, \dots, i_{K,t})$ where $p_t \leq \check{p}_t$ and $i_{k,t} \geq i'_{k,t}$ for $k = 1, 2, \dots, K$, it is optimal to stop the purification, i.e., $a_t^*(p_t, i_{1,t}, \dots, i_{K,t}) = S$.*

Theorem 2 shows that there exists a set of states, $\Sigma_t = \{p_t \leq \check{p}_t \text{ and } i_{k,t} \geq i'_{k,t}, \forall k \in \mathcal{K}\}$ at $t \in \mathcal{T}$, such that the biomanufacturing company has no financial incentive for conducting the purification operations when the starting batch state is in Σ_t . From a practical perspective, if failure is inevitable in the downstream, it is better to *fail earlier rather than later* to avoid operating costs.

THEOREM 3. *If Assumption 1 and Assumption 2 hold and the technique $c \in \mathcal{C}$ is used at $t \in \mathcal{T}$, then the optimal pooling policy $a_t^*(p_t, i_{1,t}, \dots, i_{K,t} | c) = w_c^{n*}$ is nondecreasing in p_t for $p_t > \check{p}_t$ at a given impurity level $(i_{1,t}, \dots, i_{K,t})$ for all $i_{k,t} < i'_{k,t}$, $k \in \mathcal{K}$.*

Theorem 3 indicates that the optimal pooling policy w_c^{n*} using a particular chromatography technique $c \in \mathcal{C}$ has a threshold-type structure when the pooling windows are ordered based on the ordering scheme described in Assumption 1. More specifically, Theorem 3 suggests that the optimal pooling policy at each step preserves at least some predetermined fraction θ^* of the protein, and the scientist tends to be less concerned in terms of yield losses and impurities that are carried along with the protein, as p_t increases. These optimal pooling threshold values provide guidelines that are easy to implement in practice, and have not been previously characterized in the literature.

5.2. Analysis of the Upstream Protein Mass Problem

Next, we provide insights related to the upstream protein mass decisions at $t = 0$. Theorem 4 presents the structural characteristics of the objective function $\mathcal{V}_0(p_1)$.

THEOREM 4. *Let $c_u(p_1) = c_{u1} \times p_1 + c_{u2}$, where $c_{u1} > 0$ and c_{u2} is a constant. Under the stopping cost structure in Equation (4), $\mathcal{V}_0(p_1)$ is decreasing in p_1 for $p_1 \geq \hat{p}_1$. For $p_1 < \hat{p}_1$, $\mathcal{V}_0(p_1)$ is not necessarily monotonic in p_1 and its rate of decrease is bounded by c_{u1} .*

Theorem 4 indicates that producing higher amounts of protein does not necessarily improve the expected profit until the production amount p_1 reaches a threshold value $\hat{p}_1$, after which point the expected profit decreases in p_1 . Theorem 4 also characterizes the bounds on its rate of decrease based on the upstream operating costs. This behavior is due to the complex interaction between the fermentation and chromatography operations, and demonstrates the challenges in optimizing the protein mass decisions in practice.

Next, Proposition 2 compares the performance of popular policies for the upstream protein mass decisions, and identifies the conditions under which a particular policy dominates its alternatives. The operating policies considered in Proposition 2 are identified based on the industry feedback (BioWGS 2016). To generate managerial insights, we use a discretization scheme δ in Proposition 2, where $p_\epsilon = \lceil p_d / \delta \rceil$ represents the specific yield requirement based on this discretization scheme. In practice, the discretization unit δ corresponds to the least count measured (often one milligram, depending on the application).

PROPOSITION 2. *Let Π_1 and Π_2 be two different upstream operating policies with the corresponding value functions $\mathcal{V}_0^{\Pi_1}$ and $\mathcal{V}_0^{\Pi_2}$, respectively.*

- (1) *Let Π_1 be the upstream policy $a_0 = p_1$, such that $p_1 > \hat{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the upstream policy $a_0 = \hat{p}_1$. Then, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$.*
- (2.a) *Let Π_1 be the upstream policy $a_0 = p_1$, such that $\bar{p}_1 \leq p_1 < \hat{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the policy $a_0 = \hat{p}_1$. Then, $\mathcal{V}_0^{\Pi_1} > \mathcal{V}_0^{\Pi_2}$ if the following condition holds: $c_u(\hat{p}_1) - c_u(p_1) > r(p_d)$.*
- (2.b) *Let Π_1 and Π_2 be the upstream policies identical to part (2.a). Then, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$ if the following condition holds: $c_u(\hat{p}_1) - c_u(p_1) < r(p_d) - r((p_\epsilon - 1)\delta) + c_\ell(\delta)$.*
- (3) *Let Π_1 be the upstream policy $a_0 = p_1$, such that $\check{p}_1 \leq p_1 < \bar{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the upstream policy $a_0 = \bar{p}_1$. Then, $\mathcal{V}_0^{\Pi_1} > \mathcal{V}_0^{\Pi_2}$ if $c_u(\bar{p}_1) - c_u(p_1) > c_f$.*
- (4) *Let Π_1 be the upstream policy $a_0 = p_1$, such that $p_1 \leq \check{p}_1$, $p_1 \in \mathcal{P}$. Let Π_2 be the upstream policy $a_0 = S$. Then, $\mathcal{V}_0^{\Pi_2} > \mathcal{V}_0^{\Pi_1}$.*

Proposition 2 assumes that the optimal purification actions are taken. Then, part (1) in Proposition 2 compares two upstream operating policies: Π_1 produces more than $\hat{p}_1$ units, and Π_2 produces only $\hat{p}_1$ units. Part (1) shows that Π_2 is always better than Π_1 . Part (2) compares the policy Π_1 , which produces $\bar{p}_1 \leq p_1 < \hat{p}_1$ units, against Π_2 , which creates $\hat{p}_1$ units. Part (2.a) states that Π_1 is better than Π_2 if the cost of increasing the upstream protein mass is expensive. For example, increasing the protein mass might require excessive

re-engineering in practice. The condition in Part (2.b) compares the incremental change in the revenue and yield penalty cost associated with protein mass $\bar{p}_1 \leq p_1 < \hat{p}_1$ against the incremental increase in the upstream operating costs associated with the protein mass $\hat{p}_1$. In a practical context, part (3) represents a setting where the failure cost is less critical than the additional cost needed to improve the protein mass. Part (4) shows that it is always better to abandon the project if the starting material has less than $\check{p}_1$ units of protein.

6. The Reduced-Dimension MDP Model

A typical industry setting could contain hundreds of different impurities with 5-10 candidate chromatography techniques, each having 50-100 candidate pooling windows. Although the action space is manageable, the size of the state space increases exponentially in the number of impurities. The state space could easily explode in most industry settings, making the optimization problem challenging to solve. To address this issue, we use an aggregation scheme to revise the state space, transitions and the value function of the optimization model, and refer to this revised version as the reduced-dimension model. Then, we use the structural insights obtained in Section 5 to identify the conditions under which the reduced-dimension model is exact. The state space, transitions, rewards and the value function of the reduced-dimension model are as follows.

State Space: Each impurity state $i_{k,t} \in \mathcal{I}_k$ is a binary variable $i_{k,t} \in \{0, 1\}$ for all $k \in \mathcal{K}$ at time $t \in \mathcal{T} \cup \{T\}$, such that, the state $i_{k,t}$ indicates whether the specific impurity type $k \in \mathcal{K}$ is present in the batch ($i_{k,t} = 1$) or has been completely eliminated ($i_{k,t} = 0$) by the time $t \in \mathcal{T} \cup \{T\}$. Therefore, the impurity state $(i_{1,t}, \dots, i_{K,t})$ is a 2^K dimensional vector representing which impurities are present in the batch and which ones have been completely removed. The state $p_t \in \mathcal{P}$ representing the amount of the protein available in the batch at time $t \in \mathcal{T} \cup \{T\}$ remains the same as Section 4.

State Transitions: Transitions in the protein state $p_t \in \mathcal{P}$ remain the same as the ones described in Section 3. However, the transitions associated with the impurity states $(i_{1,t}, \dots, i_{K,t})$ are simplified using the probability distribution function $P_k(i_{k,t+1}|i_{k,t}, c, w_c)$ for each impurity type $k \in \mathcal{K}$ at time $t \in \mathcal{T} \cup \{T\}$. $P_k(i_{k,t+1}|i_{k,t}, c, w_c)$ represents the probability of achieving the impurity state $i_{k,t+1} \in \{0, 1\}$ as a result of the chromatography technique c and the pooling window w_c , given that the state of the impurity type $k \in \mathcal{K}$ before

that chromatography step is $i_{k,t} \in \{0,1\}$. The transition probabilities for the upstream protein mass problem at time $t=0$ remain the same as the ones described on Section 3.

Since a chromatography technique exploits the difference in the physicochemical properties of each impurity type k as a separation principle, the probability distributions $P_k(i_{k,t+1}|i_{k,t}, c, w_c)$ are independently distributed for each impurity type k under the action $(c, w_c) \in \mathcal{C} \times \mathcal{W}_c$. Therefore, the state transitions associated with all impurities $(i_{1,t+1}, \dots, i_{K,t+1}) \in \mathcal{I}_1 \times \dots \times \mathcal{I}_K$ are captured by a joint probability distribution that can be written as product of marginals, i.e., $P_k(i_{1,t+1}|i_{1,t}, c, w_c) \times \dots \times P_K(i_{K,t+1}|i_{K,t}, c, w_c)$.

Rewards and the Value Function: The costs and rewards remain the same as Section 4. We note that the revenue $r_s(p, i_1, \dots, i_K)$ obtained from stopping the batch at state $(p, i_1, \dots, i_K)$ is a special case of Equation (4) where $\gamma_d = 100\%$. Therefore, the value function of the downstream purification problem and the upstream protein mass problem remain the same as Equations (5)-(8). The expectation operation in this model is

$$\begin{aligned} & \mathbb{E}[\mathcal{V}_{t+1}(p_{t+1}, i_{1,t+1}, \dots, i_{K,t+1}) | c, w_c] \\ &= \int_{\theta} \sum_{i_{1,t+1}=0}^1 \dots \sum_{i_{K,t+1}=0}^1 f(\theta|c, w_c) \times P_1(i_{1,t+1}|i_{1,t}, c, w_c) \times \dots \times P_K(i_{K,t+1}|i_{K,t}, c, w_c) \\ & \quad \times \mathcal{V}_{t+1}(\theta p_t, i_{1,t+1}, \dots, i_{K,t+1}) d\theta. \end{aligned}$$

Proposition 3 summarizes the state aggregation scheme used in the reduced-dimension model, and identifies the conditions under which this aggregation scheme is exact.

PROPOSITION 3. *For each impurity type $k \in \mathcal{K}$, the values of the impurity state $i_{k,t} \in \mathcal{I}_k$ can be aggregated and viewed as a binary variable, $i_{k,t} \in \{0,1\}$, such that, $i_{k,t} = 0$ represents the case where the impurity type k has been removed from the batch by the time $t \in \mathcal{T} \cup \{T\}$, and $i_{k,t} = 1$ denotes the case where a positive amount $i_{k,t} \in (0, i_{k,1}]$ of impurity type k is present in the batch at time $t \in \mathcal{T} \cup \{T\}$. This aggregation scheme is exact when the purity requirement is equal to 100%.*

Proposition 3 indicates that the state aggregation scheme is exact for a special instance of the optimization problem where the final purity requirement is $\gamma_d = 100\%$. Proposition 3 uses the specific characteristics of the transition probabilities and stopping costs under the 100% purity requirement, and indicates that it is sufficient to keep track of which impurity types are present in the batch rather than capturing their corresponding amounts. The idea

with the aggregation scheme is intuitive. It relies on the fact that the transition probabilities depend on the type of impurities but not their amounts (i.e., chromatography performs the separation based on the physical and chemical properties) and that all impurity types need to be removed under the 100% purity requirement.

Proposition 3 provides a way to address the curse of dimensionality for a special instance of the optimization problem which is frequently encountered in practice. For example, if the potential end users are humans, these drugs should abide by very high purity standards (i.e., $\geq 99.9\%$). Such high standards are also required to conduct Phase I-III clinical trials. Because of high penalty costs, process uncertainties, and manufacturing trade-offs, dealing with the 100% purity requirement could be a significant challenge in practice.

All structural results discussed in Section 5 are valid for the reduced-dimension model. For example, the monotonicity in Proposition 1, the threshold values in Theorem 1, the structural characteristics in Theorem 3 and Theorem 4, and the operating policies in Proposition 2 are not impacted by the aggregation scheme. Theorem 2 can be interpreted such that it is optimal to stop the purification operations when $p_t \leq \check{p}_t$ for a given impurity state $(i_{1,t}, \dots, i_{k,t})$, $i_{k,t} = \{0, 1\}$ for all $k \in \mathcal{K}$ at $t \in \mathcal{T}$.

7. Insights from an Industry Case Study

7.1. Problem Setting

The purification data used in this case study is obtained from Aldevron, a biomanufacturer specializing in the production of proteins. The production requirement is 8 mg of protein at a 100% purity. The target protein is manufactured for in-vitro studies. Scouting runs indicate that the purification process involves 6 candidate chromatography techniques with an average of 170 possible pooling windows per chromatography technique. The starting material consists of a mixture of 6 impurities along with the protein of interest. Statistical analysis of the scouting data indicates that the protein and impurity fractions per lane are uniformly distributed, and the expected outcome of each technique is shown in Figure 3. The solid line in Figure 3 represents the expected fraction of the protein of interest corresponding to each lane using a specific chromatography technique. The dotted lines represent the different types of impurities contained in the starting material. Figure 3 explicitly plots one specific impurity type (Impurity A) on different chromatography techniques, and masks the information on remaining impurities to protect client confidentiality.

Figure 3 Expected separation outcomes of the candidate chromatography techniques

Figure 3 illustrates the complex interdependency between the choice of chromatography technique and pooling window. For example, the relative position of the protein and Impurity A is different at each chromatography technique. In this setting, simultaneously determining the right chromatography technique and pooling window could be challenging in practice due to the complexity arising from the differences in the separation outcomes.

The cost and revenue information used in this case study represent typical industry values obtained from a cross section of several local biomanufacturing companies (BioWGS 2016), and is normalized for confidentiality purposes. The normalized values of costs and

revenue are as follow: chromatography operating costs are $r_c = \$3$ for each chromatography technique, as they use similar types of resins and buffers in this case study, revenue $r = \$2$ per milligram of protein produced at 100% purity, yield penalty cost $c_\ell = \$2$ per milligram of protein short, and failure cost $c_f = c_\ell \times p_d = \16 . The upstream operating cost is $c_u(p) = 0.1 + p \times 0.05$. We refer to these cost and revenue values as the *base case*.

7.2. Overview of Analysis

We use the case study to demonstrate the potential impact of the optimization framework in practice, and to provide an understanding of the problem configurations that have high potential for improvement. We compare the performance of the following five strategies:

1. **Current Practice (CP)**: Upstream protein mass, downstream chromatography technique and pooling window selection decisions are made based on standard practice.
2. **Fermentation Optimization (FO)**: Downstream chromatography technique and pooling window selection decisions are made based on the current practice, and only the upstream protein mass decisions are optimized.
3. **Pooling Window Optimization (PO)**: Upstream protein mass and downstream chromatography technique selection decisions are made based on the current practice, and only the pooling window selection decisions are optimized.
4. **Fermentation and Pooling Window Optimization (FO+PO)**: Downstream chromatography technique selection decisions are made based on the current practice, and only the upstream protein mass and the pooling window selection decisions are optimized.
5. **Joint Optimization (JO)**: Upstream protein mass, downstream chromatography technique and pooling window selection decisions are optimized simultaneously.

Feedback received from industry working group sessions indicate that current process improvement activities only focus on a single layer, e.g., the pooling window optimization problem. Therefore, this section focuses on (1) helping operations managers understand how going beyond CP, FO, and PO makes a difference, and (2) quantifying the potential gains in the expected profit and operational efficiency through the use of the JO framework. For this purpose, a key performance metric used in this section is the percentage improvement (IMP%) that could have been achieved in practice by adopting one of the FO, PO, FO+PO, or JO strategies instead of CP.

Because companies like Aldevron incur strict penalties on yield shortages and quality failures, we focus on a sensitivity analysis to revenue and penalty costs. For this purpose, we evaluate an extensive set of configurations with $r \in \{1, 1.5, 2, 2.5, 3\}$ and $\frac{c_\ell}{r} \in$

$\{0, 0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2\}$. These configurations are identified based on the input received from our industry partners. For brevity, we only report the results for *low-revenue*, *low-penalty* projects ($r = 1.5$ and $\frac{c_\ell}{r} \in \{0, 0.5, 1\}$), *base-case variations* ($r = 2$ and $\frac{c_\ell}{r} \in \{0.5, 1, 1.5, 2\}$), and *high-revenue, high-penalty* projects ($r = 3$ and $\frac{c_\ell}{r} \in \{1, 1.5, 2\}$). Scenarios with $c_\ell = 0$ are also reported as a benchmark.

Decision Making in Practice. In this case study, the scientist adopts the three-step purification approach described in Section 2.6. The chromatography technique selection decisions are made based on the highest resolution criteria. Numerical experiments also confirm that the highest resolution criteria outperforms the highest selectivity criteria in this specific case study. Therefore, chromatography techniques (3), (5) and (1), shown in Figure 3, are used in the first, second, and third steps, respectively. In the first step, the scientist pools lanes 1-8. As a result, she aims to eliminate impurity types 2, 3, 5, and 6 with an expected yield loss of 35%. In the second step, she aims to eliminate the remaining impurities. Therefore, she pools lanes 1-10 with an expected yield loss of 35%. If trace amounts of impurities remain, then the scientist tends towards a yield-conservative strategy with 10% yield loss in the last step (lanes 7-18). In the upstream operations, 75% excess is typically planned as a buffer (BioWGS 2016). This leads to 28 mg of protein production (i.e., $8 \times 1.35 \times 1.35 \times 1.1 \times 1.75 \approx 28$) in the upstream process.

Optimal Policies. Operating policies corresponding to other strategies are as follows. Table 1 shows the optimal upstream production quantities under JO. These production quantities are further analyzed in Section 7.3. Under JO, chromatography technique 5 with 30% yield loss is expected to be used (lanes 1-11), followed by chromatography technique 2 with 20% yield loss (lanes 1-12). The optimal chromatography techniques recommended by JO are robust to the scenarios. This is mainly because all techniques have the same operating cost. Optimal pooling windows under PO are state dependent, but chromatography technique 3 with 30% yield loss is expected to be used (lanes 1-9), followed by chromatography technique 5 with 30% yield loss (lanes 1-11). If trace amounts of impurities remain, then it is expected that chromatography technique 1 with 15% yield loss will be used (lanes 8-18). PO produces 27 mg protein in the upstream (i.e., $8 \times 1.3 \times 1.3 \times 1.15 \times 1.75 \approx 27$). The optimal upstream production quantities p_1^* under FO are reported in Table 1.

Table 1 Potential benefits of the joint optimization framework in practice

r	c_ℓ/r	CP	FO			PO		FO+PO			JO				
		$\mathcal{V}_0(p_1 = 28)$	p_1^*	$\mathcal{V}_0(p_1^*)$	IMP%	$\mathcal{V}_0(p_1^* = 27)$	IMP%	p_1^*	$\mathcal{V}_0(p_1^*)$	IMP%	p_1^*	$\hat{p}_1$	$\bar{p}_1$	$\mathcal{V}_0(p_1^*)$	IMP%
1.5	0	2.3	30	2.4	3%	2.6	15%	25	2.7	17%	21	25	10	3.0	33%
1.5	0.5	1.6	30	1.8	9%	2.1	29%	25	2.1	31%	23	26	12	2.7	68%
1.5	1	1.2	32	1.3	8%	1.6	33%	28	1.7	40%	25	27	14	2.5	105%
2	0	5.9	30	6.0	2%	6.3	7%	25	6.3	8%	21	26	6	6.8	17%
2	0.5	5.0	32	5.3	6%	5.6	12%	28	5.7	14%	25	27	10	6.5	30%
2	1	4.1	35	4.7	14%	4.9	19%	30	5.1	24%	25	30	12	6.1	48%
2	1.5	3.3	37	4.1	27%	4.2	30%	31	4.5	40%	25	30	13	5.7	76%
2	2	2.4	38	3.5	49%	3.6	49%	33	4.0	68%	25	30	14	5.3	124%
3	0	13.0	32	13.3	2%	13.6	5%	28	13.7	5%	25	27	4	14.5	11%
3	1	10.4	38	11.5	11%	11.6	11%	33	12.0	16%	25	30	11	13.3	28%
3	1.5	9.1	40	10.7	18%	10.5	16%	34	11.2	23%	25	31	12	12.8	41%
3	2	7.8	40	9.8	26%	9.5	22%	34	10.4	34%	27	31	14	12.3	58%

7.3. Potential Impact in Practice

Table 1 presents the upstream protein production amount p_1 and the expected profit $\mathcal{V}_0(p_1)$ under different settings. The columns labelled “IMP%” in Table 1 quantify the percentage improvement in the expected profit that could be achieved in practice by using the FO, PO, FO+PO, or JO strategies instead of CP. The bold entries in Table 1 represent the base case.

In the base case, we observe that optimizing the pooling windows alone provides considerable savings compared to CP, with 19% improvement in the expected profit. This result aligns with Martagan et al. (2017). In addition, it is important to note that optimizing the fermentation operations alone yields notable savings of 14% improvement in practice. However, the percentage improvement obtained by the JO framework significantly outperforms other strategies. For example, the JO framework yields 48% improvement in the base case, which is the double the 24% improvement achieved by the combined strategy FO+PO. Even if $c_\ell = 0$ at $r = 2$, JO continues to outperform other strategies, with 17% improvement while other strategies achieve at most 8% improvement. This indicates that there is significant room for improvement in current practice through the use of JO. This is mainly because JO provides a comprehensive framework that incorporates the financial trade-offs, process uncertainties and the underlying chemical dynamics into decision-making.

We observe that FO does not provide significant savings for low-revenue, low-penalty projects (e.g., $\text{IMP}\% \leq 9\%$). In such projects, purification optimization mostly dominates the potential savings in Table 1. The intuition behind this behavior is that at lower levels of revenue, the cost of producing higher amounts of protein in the upstream does not necessarily outweigh the resulting benefits in the downstream. This is mainly why the production quantities and expected profit of FO and CP are very close for low-revenue, low-penalty projects in Table 1. In contrast, we observe that JO enables a large room for improvement on CP, and its magnitude increases as margins become lower. For example, at a fixed c_ℓ/r in Table 1, the potential savings through JO increase as revenue decreases. Even in case of no penalty costs at $r = 1.5$, JO achieves a 33% improvement over CP. The intuitive explanation is that it becomes more critical to meet client requirements with minimal waste as margins decrease. Therefore, projects with lower margins receive higher benefits from the use of JO.

For high-revenue, high-penalty projects, the improvements using FO and PO are comparable, but FO outperforms PO as c_ℓ increases. This indicates that, if penalties and revenues are high, it can be effective to deliberately increase the upstream production to alleviate downstream purity–yield trade-offs. Note that FO can reduce the expected yield shortages only by producing higher amounts of protein; whereas PO can reduce it only by managing the purity–yield trade-offs. Therefore, FO increases the production amount at high c_ℓ in order to reduce shortage penalties, and the corresponding upstream costs pay off due to high revenues and reduced penalties. On the other hand, JO outperforms all strategies in Table 1. Even compared to the combined strategy FO+PO, JO provides an addition of 6%–24% improvement at $r = 3$. Moreover, Table 1 indicates that the percentage improvements increase as c_ℓ increases at a fixed value of r . This is because it becomes critical to reduce the expected yield shortages as c_ℓ increases. At high penalty levels, Table 1 reports very high improvements on CP. This clearly indicates that CP can be very inefficient at high penalties because it does not consider the financial trade-offs of each decision layer.

Next, we compare the production quantities p_1^* under the five manufacturing strategies to obtain insights related to how much to waste in the upstream. In Table 1, we observe that FO produces higher amounts of protein compared to FO+PO and JO. This is mainly because FO tends to buffer against the possible inefficiencies in downstream operations. The results in Table 1 indicate that expanding the scope of the optimization model from

FO to JO helps to reduce the upstream production quantities. For example, in the base case, FO recommends the production of 35 mg of protein, whereas FO+PO recommends 30 mg, and JO enables savings of an additional 5 mg by producing 25 mg. In addition, Table 1 presents the break-even points $\bar{p}_1$ and the maximum threshold values $\hat{p}_1$ under JO. Note that the values of $\bar{p}_1$, p_1^* and $\hat{p}_1$ are nondecreasing in c_ℓ to buffer against yield shortages. The difference between $\hat{p}_1$ and $\bar{p}_1$ ranges between 13-27 mg in Table 1. This shows that there is a large amount of room for profit in this case study. The optimal production quantities are closer to $\hat{p}_1$ than $\bar{p}_1$ in Table 1. Results from additional experiments indicate that p_1^* approaches $\hat{p}_1$ as the upstream costs reduce. In Table 1, we observe that p_1^* and $\hat{p}_1$ values do not change significantly in r and c_ℓ under JO. We believe this is mainly due to the robustness of the optimal downstream policies at high protein levels (i.e., the optimal policy uses the same chromatography techniques in all scenarios, and the pooling window selections are not sensitive at high protein levels). Note that the optimal production amount of the base case ($p_1^* = 25$ mg) is almost three times greater than the demand requirement ($p_d = 8$ mg). This illustrates how complex and challenging it is to satisfy production requirements in the biomanufacturing industry.

Table 1 shows that the percentage improvement using JO can be more than the double the amount achieved by FO or PO. This is mainly because JO helps to better address the practical concern of “*how much to waste*” in the upstream and downstream operations. For example, JO enables the reduction of upstream production quantities by 3 mg compared to CP (in the base case, on average), and is still able to meet the demand. In the downstream operations, the total expected yield loss under JO is 6.35 mg less than that under CP (in the base case). Overall, JO produces less, wastes less, and achieves the highest expected profit compared to other strategies. This underscores the importance of JO.

7.4. Benefits to Downstream Purification Operations

Figure 4 considers four representative scenarios from Table 1, and plots the value function of JO, PO, and CP in the first chromatography step. Note that JO optimizes the chromatography technique and pooling window selection decisions at $t = 1$, and the behavior of $\mathcal{V}_1$ in p_1 impacts the upstream protein mass decisions. To quantify the potential benefits to the downstream operations, we analyze (1) the difference between the value functions of PO and CP at $t = 1$, and (2) the difference between the value functions of JO and PO

Figure 4 The value function of JO, PO, and CP in the first chromatography step.

at $t = 1$. The former quantifies the impact of pooling window optimization, and the latter represents the added value of optimizing the chromatography technique selection.

The value function of JO, PO, and CP converge to a constant value at higher levels of protein in Figure 4 (see Theorem 1). As an extreme case, consider $p_1 = 40$. In Figure 4, adopting PO instead of CP results in 2%–6% improvement in the value function at state $p_1 = 40$. However, using JO instead of PO yields another 7%–15% improvement. This implies that choosing the right chromatography technique drives the potential room for improvement at high levels of protein. Similar to Section 7.3, we see that projects with higher penalties or lower margins encounter the most benefit from JO. In addition, Figure 4 shows that the gap among the value functions of JO, PO, and CP is higher when the starting material is in the range $\bar{p}_1 \leq p_1 \leq p_1^*$. This corresponds to a risky business case where the biomanufacturer is likely to incur yield shortages. In such cases, Figure 4 shows that smartly managing the purity–yield trade-offs through the optimal selection of chromatography techniques and pooling windows plays a critical role in reducing penalties.

Figure 4 shows that JO reaches the break-even point $\bar{p}_1$ and the maximum profit threshold $\hat{p}_1$ by using less protein than PO. For example, in the base case (Figure 4c), JO uses 2

mg less protein than PO to break even, and needs 7 mg less than PO to reach the threshold $\hat{p}_1$. The average of all scenarios considered in the numerical analysis indicates that JO can break even by using 2 mg less protein than PO, and 4 mg less than CP. We observe that scenarios with higher penalties or lower margins achieve higher reduction in the break-even points. These differences are one of the key indicators of the improved operational efficiency under JO.

In the downstream operations, the expected profit depends on the starting amount p_1 . To quantify the potential impact, we focus on the average improvement in the range $21 \leq p_1 \leq 35$. This range includes the optimal production quantities of JO and FO+PO. In the base case (Figure 4c), optimizing the pooling windows enables an average of 22% improvement in practice. Optimizing the chromatography techniques together with the pooling windows provides another 18% improvement in the value function of PO, on average. In other scenarios, the average percentage improvement achieved by using the JO framework instead of PO are 29% (Figure 4a), 14% (Figure 4b), and 40% (Figure 4d). These improvements underscore that an OR-based decision-making approach would make a significant difference in the survival and competitiveness of SMEs.

8. Conclusions

We focus on biomanufacturing decisions for engineered proteins. In this setting, research and development are often conducted by a large pharmaceutical company, but the biomanufacturing operations are performed by a contract biomanufacturer due to high failure risks and specific domain knowledge. Manufacturing challenges for engineered proteins can be associated with several factors, such as limitations in chromatography operations, randomness in process outcomes, yield losses, stringent quality requirements, and failure risks. There are excellent studies that contribute to the knowledge behind the biology and chemistry of these operations, but there is a significant need for an optimization framework that links the underlying biology and chemistry with the financial trade-offs and manufacturing challenges encountered in practice.

In this study, we build a stochastic optimization model that addresses three layers of interdependent decisions: (1) the amount of protein to be produced in fermentation, (2) the chromatography techniques to be used in purification, and (3) the choice of pooling windows based on the purity–yield trade-offs. We characterize the structure of the optimal

policies, and derive functional relationships related to costs and purity–yield trade-offs. We use a state-aggregation mechanism to reduce the computational efforts to solve complex problems at 100% purity. Studies in the literature and feedback received from industry partners indicate that most process improvement activities focus on one of these decision layers in isolation. However, the case study example in Section 7 shows that the percentage improvement under the joint optimization framework is almost double that which can be achieved by optimizing a single decision level. This is mainly because the joint optimization model better addresses the practical concern of minimizing waste in upstream and downstream operations. The results obtained from the case study are encouraging to support the survival and competitiveness of SMEs. Nevertheless, the potential impact of this framework extends beyond SMEs. As Tom Foti, the vice president of Aldevron states “We are producing 50-liter cultures here, but our clients [large pharmaceutical companies] are dealing with 5,000-liter cultures. If we can build optimization models here, and demonstrate the feasibility of how it works, our clients could also do that. If they can reduce health care costs, that could be directly passed off to the patients” (Aldevron 2017).

Insights obtained from this framework can lead to several future research directions. First, one can extend the model to make decisions with limited data. Second, pricing decisions and renegotiation schemes can be analyzed under the joint optimization framework.

References

- Aldevron. 2017. A new tool for biomanufacturing. URL <https://www.youtube.com/watch?v=CWoQ5ofTZfY>.
- BioWGS. 2016. Biomanufacturing Working Group Sessions: Aldevron collaborates with University of Wisconsin to develop decision support tools for biomanufacturing. URL <http://www.bioforward.org/past-events-aldevron-collaboration/>.
- GE Healthcare. 2010. Strategies for protein purification: Handbook by GE Healthcare. Tech. rep.
- Gnoth, S., M. Jenzsch, R. Simutis, A. Lubbert. 2007. Process Analytical Technology: Batch-to-batch reproducibility of fermentation process by robust process operational design and control. *Journal of Biotechnology* **132** 180–186.
- Lienqueo, M.E., J.C. Salgado, O. Giaverini, J.A. Asenjo. 2009. Computer-aided design to select optimal polypeptide tags to assist the purification of recombinant proteins. *Separation and Purification Technology* **65**(1) 86–94.
- Long, G., J. Works. 2013. Innovation in the biopharmaceutical pipeline: A multidimensional view. Analysis Group Industry Report.

- Martagan, T., A. Krishnamurthy, P. A. Leland, C. T. Maravelias. 2017. Performance guarantees and optimal purification decisions for engineered proteins. *Operations Research* **0**(0). doi:10.1287/opre.2017.1661. URL <https://doi.org/10.1287/opre.2017.1661>.
- Martagan, T., A. Krishnamurthy, C. Maravelias. 2016. Optimal condition-based harvesting policies for biomanufacturing operations with failure risks. *IIE Transactions* (48:5) 440–461.
- Nfor, B. K., T. Ahamed, G. W. K. van Dedem, P. D. E. M. Verhaert, L. A. M. van der Wielen, M. H. M. Eppink, E. J.A.X. van de Sandt, M. Ottens. 2013. Model-based rational methodology for protein purification process synthesis. *Chemical Engineering Science* **89** 185–195.
- Patel, S. D., E. T. Papoutsakis, J. N. Winter, W. M. Miller. 2000. The lactate issue revisited: Novel feeding protocols to examine inhibition of cell proliferation and glucose metabolism in hematopoietic cell cultures. *Biotechnology Progress* **16** 885–892.
- Peroni, C. V., N. S. Kaisare, J. H. Lee. 2005. Optimal control of fed-batch bioreactor using simulation based approximate dynamic programming. *IEEE Transactions* **13** 786–790.
- Petrides, D., C. A. Siletti. 2004. The role of process simulation and scheduling tools in the development and manufacturing of biopharmaceuticals. *Proceedings of the Winter Simulation Conference*, 2046–2051.
- Polykarpou, E. M., P. A. Dalby, L. G. Papageorgiou. 2011. Optimal synthesis of chromatographic trains for downstream protein processing. *Biotechnology Progress* **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.
- Saraph, P. V. 2003. Shared resource capacity analysis in biotech manufacturing. *Proceedings of the Winter Simulation Conference*, 1247–1250.
- Saucedo, V. M., M. N. Karim. 1997. Experimental optimization of a real time fed-batch fermentation process using Markov decision process. *Biotechnology and Bioengineering* **55** 317–327.
- Tsao, Y. S., A. G. Cardoso, R. G. G. Condon, M. Voloch, P. Lio, J. C. Lagos, B. G. Kearns, Z. Liu. 2004. Monitoring chinese hamster ovary cell culture by the analysis of glucose and lactate metabolism. *Journal of Biotechnology* **118** 316–327.
- Vasquez-Alvarez, E., M. E. Lienqueo, J. M. Pinto. 2001. Optimal synthesis of protein purification processes. *Biotechnology Progress* **17** 685–696.
- Vasquez-Alvarez, E., J.M. Pinto. 2004. Efficient MILP formulations for the optimal synthesis of chromatographic protein purification processes. *Journal of Biotechnology* **110**(3) 295–311.
- Xing, Z., N. Bishop, K. Leister, Z. J. Li. 2010. Modeling kinetics of a large-scale fed-batch CHO cell culture by Markov chain monte carlo method. *Biotechnology Progress* **26** 208–219.
- Yang, J. D., Y. Angelillo, M. Chaudhry, C. Goldenberg. 2000. Achievement of high cell density and high antibody productivity by controlled-fed perfusion bioreactor process. *Biotechnology and Bioengineering* **69** 74–82.

Acknowledgments

The authors would like to thank the editor, associate editor, and two anonymous referees for the many helpful comments that greatly improved the paper. This research was conducted due to the generous support provided by the National Science Foundation under grant CMMI 1334933, and the Marie S. Curie Fellowship provided by the European Commission under Horizon 2020 program.