

RESILIENT CAPACITY PLANNING

A Dissertation
Presented to
The Academic Faculty

By

Junxuan Li

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
College of Engineering
H. Milton Stewart School of Industrial & Systems Engineering

Georgia Institute of Technology

May 2021

© Junxuan Li 2021

RESILIENT CAPACITY PLANNING

Thesis committee:

Co-chair: Dr. Chelsea C. White III
School of Industrial and System Engineering
Georgia Institute of Technology

Dr. Mathieu Dahan
School of Industrial and System Engineering
Georgia Institute of Technology

Co-chair: Dr. Ben Wang
School of Industrial and System Engineering,
Georgia Tech Manufacturing Institute
Georgia Institute of Technology

Dr. David Goldman
School of Industrial and System Engineering
Georgia Institute of Technology

Dr. Maria Papathanasiou
Department of Chemical Engineering
Imperial College London

Date approved: Jan 15, 2021

For my daughter Julia

ACKNOWLEDGMENTS

I would like to thank the members of my thesis committee for their helps during my Ph.D. years, without whom I would never be able to complete the research projects in this thesis. I would like to give special thanks to my advisors, Dr. Chelsea White and Dr. Ben Wang, who have inspired me of research ideas and advised me how to carry out research to transfer an idea into cutting edge technologies. I would also like to thank Drs. Mathieu Dahan, David Goldsman and Maria Papathanasiou who helped to shed lights on many of my research ideas and had priceless advice.

I'm also grateful to my wife, Yifei Ren, for her continuous supports to make me a better person. Additional thanks are due to my parents, Zhenti Li and Cuizhi Gu, and my mother-in-law, Hong Yang, who helped take care of my daughter in the first several months after she was born, and without whom I would not be able to have sufficient time to complete this work.

TABLE OF CONTENTS

Acknowledgments	iv
List of Tables	ix
List of Figures	x
Summary	xii
Chapter 1: Introduction and Background	1
1.1 Capacity Planning (CP) Problem	1
1.2 Relocatable Manufacturing/Service Module (REMO)	3
1.3 Centralized and Decentralized Networks	4
1.4 Reconfigurable Supply Chain	5
1.5 Supplier Disruption Risk	5
1.6 Dynamic Resilience	6
1.7 Domain Examples	7
1.7.1 Autologous cell therapy	7
1.7.2 Epidemic/pandemic control	10
1.8 Objectives and Organization	11
Chapter 2: Resilient Capacity Planning in a Centralized Network	14

2.1	Introduction	14
2.1.1	Literature review	16
2.1.2	Organization	19
2.2	Adjustable Bioreactor Quantity (ABQ) Model	19
2.2.1	Penalty cost approach	22
2.2.2	Chance constraint approach	25
2.2.3	Connections between ABQ-p and ABQ-c	27
2.3	Fixed Bioreactor Quantity (FBQ) Model	29
2.3.1	Penalty cost approach	30
2.3.2	Chance constraint approach	32
2.4	Case Study	37
2.5	Conclusions	41
	Chapter 3: Resilient Capacity Planning in a Centralized Network under Supplier Disruption Risks	42
3.1	Introduction	43
3.1.1	Literature review	44
3.1.2	Organization	47
3.2	Adjustable Bioreactor Quantity Model with Supplier Disruption Risks (ABQ-sdr)	48
3.2.1	Model formulation	48
3.2.2	A simplifying transformation	51
3.2.3	Attainability	53
3.2.4	Bounds	54

3.2.5	Solution algorithm	55
3.3	Fixed Bioreactor Quantity Model with Supplier Disruption Risks (FBQ-sdr)	56
3.3.1	Model formulation	56
3.3.2	Solution algorithm and bounds	57
3.4	Case studies	58
3.4.1	Case study 1: supplier disruption risk mitigation using ABQ-sdr and FBQ-sdr models	60
3.4.2	Case study 2: cost and benefit of resilience	62
3.5	Conclusion	66
Chapter 4:	Resilient Capacity Planning in a Decentralized Network under Sup- plier Disruption Risks	67
4.1	Introduction	68
4.1.1	Literature review	69
4.1.2	Organization	71
4.2	Problem Formulation	71
4.3	Analysis	78
4.4	Algorithms	80
4.4.1	Solving Co-Net-d	80
4.4.2	Solving Co-Net-o	81
4.5	Case Studies	84
4.5.1	Case study 1: the potential of resource sharing	84
4.5.2	Case study 2: the supplier correlations	89
4.6	Conclusions	92

Chapter 5: Conclusion and Future Research	94
Appendices	98
Appendix A: Complimentary Proofs	99
Appendix B: Complimentary Algorithms	111
References	113
Vita	121

LIST OF TABLES

2.1	System specs and cost parameters – Chapter 2 case study	38
3.1	System specs and cost parameters – Chapter 3 case study	59
3.2	Disruption parameters summary	60
3.3	System state statistics	62
4.1	System specs and cost parameters – Chapter 4 case study	85

LIST OF FIGURES

1.1	Demonstration of centralized and decentralized networks (Red dots represents demands and green dots represents manufacturing facilities	4
1.2	Centralized autologous cell therapy manufacturing network	7
1.3	Manufacturing process flow for autologous cell therapy [7]	8
1.4	WAVE bioreactor system 200	8
2.1	Bioreactor quantities versus t	38
2.2	Cost and bioreactor shortage probability versus bioreactor quantities	39
2.3	Reagent inventory levels for the first 5 simulations	40
2.4	Optimal bioreactor quantity and total cost v.s. α_B	40
3.1	Bioreactor quantities v.s. t	61
3.2	On-hand reagent v.s. t	63
3.3	Accounting cost increment and penalty cost reduction summary at different disruption profiles	65
3.4	Cost/benefit ratio of resilience	66
4.1	Case study 1: average expected total cost comparison	86
4.2	Case study 1: bioreactor quantity comparison	87
4.3	Case study 1: average specimen transshipment quantity comparison	88
4.4	Case study 1: average bioreactor relocation quantity comparison	88

4.5	Case study 1: average reagent transshipment quantity comparison	89
4.6	Case study 2: average expected total cost comparison	90
4.7	Case study 2: bioreactor quantity comparison	91
4.8	Case study 2: average specimen transshipment quantity comparison	92
4.9	Case study 2: average bioreactor relocation quantity comparison	92
4.10	Case study 2: average reagent transshipment quantity comparison	93

SUMMARY

We envision the next generation supply chain to be an efficient, resilient, and cybersecure system for the movement of goods, data, and money. An emerging design of the next generation supply chain is *reconfigurable supply chain*. A *reconfigurable supply chain* is a supply chain that permits material transshipment and adding, subtracting and relocating manufacturing/service modules. Relocatable manufacturing/service module is a trending concept in advanced manufacturing and service engineering.

In this Ph.D. thesis, we investigate a class of problems called *resilient capacity planning* in reconfigurable supply chains for build-to-order manufacturing systems and service-upon-request service systems. This emerging class of problems applies to a wide spectrum of novel application domains, such as personalized medicine, 3D-printing manufacturing, relocatable storage network, epidemic/pandemic control and etc. The objective of this thesis is to study capacity planning problems in reconfigurable supply chains with coordinating controls on consumable and reusable resources when customer demand is stochastic and suppliers are unreliable. We investigate several key fundamental research projects for both centralized and decentralized reconfigurable supply chains to provide resilient capacity planning decision supports. We introduce the approaches to model and quantify supplier disruption risks in both centralized and decentralized networks and reveal the potential of resilient capacity planning in enhancing system performances.

CHAPTER 1

INTRODUCTION AND BACKGROUND

The next generation supply chain is a trending topic in both academic and industrial world. A next-generation supply chain will be an efficient, resilient, and cybersecure system for the movement of goods, data, and money. Moreover, The next generation supply chain will extract value from (monetize) the significant increases in the volume, velocity, and variety of real-time data and computational capabilities for real-time supply chain control and competitive advantage. In this Ph.D. thesis, we aim to investigate a class of problems in the next generation supply chain, motivated by a variety of emerging application field, which is named as *Resilient Capacity Planning*.

1.1 Capacity Planning (CP) Problem

Capacity Planning (CP) [1, 2, 3] is a process of determining fulfillment capacity needed in a manufacturing/service system to meet demand of goods or services. Fulfillment capacity is the maximum amount of jobs that the system is capable to start given a system design of different kinds of resources. Two major categories of these resources often considered are *reusable* and *consumable*.

- A reusable resource is a resource that can be used by only one manufacturing/service process at a time and is not depleted by the completion of the process. The quantity of reusable resources are often limited. Examples of reusable resources include bioreactors in personalized medicine manufacturing, 3D printers, and mobile intensive care units and ventilators in epidemic/pandemic control.
- A consumable resource is a resource that is consumed (used up) by the completion of a manufacturing/service process. The quantity of consumable resources is often

named as "stock" which can be produced or replenished to supply future processes. Example of consumable resources include reagent in personalized medicine manufacturing, filament in 3D-printing manufacturing, and disposable medical supplies in epidemic/pandemic control.

A demand of good or service can be fulfilled only if both reusable and consumable resources are available. A combination of available reusable and consumable resources constitutes one unit of fulfillment capacity. A CP problem coordinates management of both reusable and consumable resources to ensure high fulfillment speed while minimizing the operational costs. A CP problem distinguishes from classic inventory control problem ([4] and others) where demand are directly fulfilled by on-hand stocks, as it considers joint effects of the reusables on the fulfillment capacity. A CP problem distinguishes from classic lot sizing problem ([4] and others) where goods can be build-to-stock to fulfill future demands, as it considers build-to-order and service-upon-request system where a fulfillment process is triggered by product order or service request.

A CP problem should consider a planning of resources over a time horizon to cope with demand variations, supplier uncertainties, economic climate changes and so on. Proper solution of a CP problem aim to ensure, with high likelihood, that fulfillment capacities are sufficient to permit a prompt manufacturing/service upon the arrival of a demand/request while maintain a low total cost of operating the facility network and its supply chain. There are two sub-classes of CP problems depending on the total number of the reusable resources in the system:

- *Adjustable*: the quantity of reusable resources at a manufacturing/service location can be adjusted at every decision epoch, for instance by leasing from a provider or by transshipping from other location in the network;
- *Fixed*: the quantity of reusable resources at a manufacturing/service location can only be adjusted periodically, e.g. yearly or semi-yearly, and is fixed between the

adjustable epochs.

The decisions of the two sub-classes are different. In an adjustable quantity model, the decision at each decision epoch includes reusable resource adjustment and consumable resource replenishment. While in the fixed quantity model, reusable resource adjustments are only allowed once a while, and consumable resource replenishment is determined at each decision epoch.

1.2 Relocatable Manufacturing/Service Module (REMO)

A standard technology that enable reconfigurability of a supply chain is the development of RElocatable manufacturing/service MOdule (REMO). REMO is a processing unit used for product manufacturing or request servicing. A REMO has the following key features [5, 6]:

- Each fulfillment occupies a REMO throughout its manufacturing or service. When increasing a system's fulfillment capacity with REMOs, it requires to "scale-out" capacities by increasing the number of REMOs, instead of "scale-up" capacities by increasing the amount that each REMO could provide;
- When necessary, the REMOs can be relocated geographically in the manufacturing/service network. In this research, we consider REMOs that can be swiftly relocated between different locations, and easily set up at new locations. The cost of relocating a REMO include the logistic cost and recalibration cost, which constrains the relocations unless necessary.

Examples of REMOs includes: bioreactors in personalized medicine manufacturing, 3D printers, smart lockers in relocatable storage network, mobile intensive care units (ICU) and ventilators in epidemic/pandemic control and etc.

A CP problem is complex as the system's fulfillment capacity process at each location is a modulated process affected by a REMO availability process. Breaking in a nutshell

of the REMO availability process, the number of ready-to-use REMOs is coupled with a number of REMO idle-busy processes. Challenges that REMOs bring to a CP problem are three-fold:

- (i) A system performance driver is the REMO availability process
- (ii) Determine a proper CP decision would require real time information of all REMOs in the system. The vast amount of streaming data should be properly saved, delivered and used efficiently and effectively
- (iii) The complexity of system dynamics leads to the complexity of finding an optimal or near-optimal solution of the CP.

1.3 Centralized and Decentralized Networks

A centralized network consists of a single fulfillment facility handling demands from multiple regions. While a decentralized network fulfill demands with multiple, geographically distributed fulfillment facilities. Figure 1.1 demonstrates examples of centralized and de-

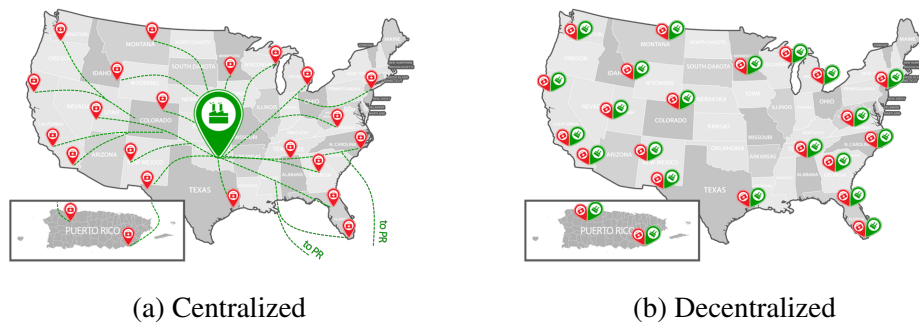


Figure 1.1: Demonstration of centralized and decentralized networks (Red dots represents demands and green dots represents manufacturing facilities)

centralized networks in the context of personalized regenerative medicine: (a) a centralized network fulfill demands of all regions with a single facility; (b) a decentralized network fulfill regional demands with multiple regional facilities.

A centralized network, comparing to a decentralized network, has less regulatory requirements, greater potential for economies of scale, greater consistency in operations, less total capital expenditure, but longer fulfillment time [7]. However, decentralized networks are able to fulfill demand with shorter fulfillment time, fewer transportation and storage of finished goods and work-in-processes (WIPs), and less handling and packaging [7].

1.4 Reconfigurable Supply Chain

A *reconfigurable supply chain* is a supply chain that permits material transshipment and allows for the adjustment and relocation of REMOs [8]. In our proposed studies of both centralized or decentralized networks, REMOs can be added or subtracted from a fulfillment facility, and can be relocated between facilities.

In a centralized reconfigurable network, adjusting the quantity of REMOS increases operation flexibility with potential cost and fulfillment time reduction. In a decentralized reconfigurable network, the facilities operate coordinately by transshipping both reusable and consumable resources. A reconfigurable supply chain is resilient against system risks such as demand uncertainty and supplier disruption. The proposed research in this thesis will investigate reconfigurable supply chains and study how to blend the benefits of centralized networks with the benefits of decentralized networks while mitigating their disbenefits.

1.5 Supplier Disruption Risk

Application domains of CP with REMOs are mostly emerging industries, whose supplier base will often have a small number of suppliers, e.g. a single supplier for a key consumable resource. This fact leads to supplier disruption risks of key consumable resources [9]. For instance, in the pharmaceutical industry in the US, the Food and Drug Administration often regulates the number of suppliers to ensure the consistency and quality of the final products, e.g. drugs and therapies. The proposed research in this thesis investigate how to model and quantify supplier disruption risks in both centralized and decentralized networks.

In the context of personalized medicine, a supply chain disruption due to the unavailability of reagent has been a severe risk for biomanufacturers and has prompted considerable interest and concern in the industry. In recent years, this industry has witnessed saline shortages due to Hurricane Maria, a severe flu season, and a 2-month shutdown of a major regenerative medicine supplier due to sterility issues [7]. These disruptions, particularly the delay due to sterility issues and its impact on our industry partners, e.g. The Center for Cellular Immunotherapies, University of Pennsylvania Perelman School of Medicine, have motivated our interest in reagent disruption risk. During the most recent COVID-19 outbreak in the US, severe shortages of reagents are reported due to the production surge of testing kits [10].

If a manufacturing/service system is operated lean without consideration of possible supplier disruptions, the consumable resource shortages would significantly diminish the fulfillment capacity of the system (regionally or globally). Mitigating the supplier disruption risks are challenging as the suppliers' states are stochastic: a future supplier state is unknown a priori, however it may be probabilistic inferred from the current supplier state.

1.6 Dynamic Resilience

The definition of supply chain network resilience has been defined in multiple perspectives, including accurate detecting, tracking, gracefully degrading and rapidly responding to and recovering from stochastic events in the supply chain. In this research, we consider enhancing network resilience against stochastic demand and supply disruptions. In a reconfigurable supply chain a dynamic form of resilience is possible. The design of a centralized or decentralized reconfigurable supply chain can be dynamically adjusted according to real-time system states. A dynamic resilient network reduces total resilience capital expenses and allows facilities to run lean when possible and resilient when necessary.

We remark that lean/low-cost supply chains may be unable to respond effectively to unexpected changes in demand or supply. However, such supply chains perform poorly

during a disruption, recover slowly after a disruption, and display a variety of previously hidden operational risks. Dynamic resilience is an emerging and incompletely understood feature of reconfigurable supply chains.

1.7 Domain Examples

1.7.1 Autologous cell therapy

We introduce an application domain that motivated our research on resilient capacity planning in reconfigurable supply chains. Autologous cell therapy is an emerging therapeutic method where a patient’s own cells are used as medical treatment [11, 12]. One of the biggest breakthroughs in cancer treatment in decades is chimeric antigen receptor T cell (CAR-T) therapy, which has demonstrated complete response rates of 69% - 90% in pediatric patients with relapsed or refractory acute lymphoblastic leukemia [13]. The Food & Drug Administration has approved three CAR-T products to treat certain blood cancers: Kymriah, made by Novartis, Yescarta, made by Kite Pharma, and Tecartus, made by UPMC Hillman Cancer Center. Building on these and other successes, large pharmaceutical companies (e.g. Pfizer, Roche) are investing heavily in gene therapy. The material and equipment planning problem can be modeled as a CP problem.

The production of autologous cell therapy is a build-to-order process. The simplified flow of materials and demand fulfillment in a centralized manufacturing network for an autologous cell therapy is presented in Figure 1.2, with solid lines indicating transportation of physical entities and dashed line indicating information flow. The manufacturing facility




Figure 1.2: Centralized autologous cell therapy manufacturing network

acquires a specimen of patient cells from a clinic of a specific tissue, following guidelines provided by the manufacturer. Upon the arrival of the cell specimen, an acceptance check is conducted to ensure that this specimen is able to yield a quality final product. Each patient's therapy is produced independently following a carefully designed process flow shown in 1.3. Target cells are first isolated from the specimen. Then these cells are expanded and




Figure 1.3: Manufacturing process flow for autologous cell therapy [7]

harvested in a *bioreactor* (a device or apparatus in which living organisms synthesize useful substances; see Figure 1.4 for an example that is small enough to be easily transportable), consuming units of *reagents* (a chemical ingredient, i.e., a compound or mixture, typically of inorganic or small organic molecules, introduced to cause the desired transformation of an organic substance) supplied from FDA approved suppliers. We remark that a bioreactor in autologous cell therapy production is used to manufacture a therapy for a specific, individual patient, and is portable, hence it can be considered as a REMO. Quality tests are conducted during the cell culturing processes. The cell batch is then washed, purified, and formulated into the customized final therapy product. A final quality test is conducted to ensure viability of the therapy, and the cell therapy is finally cryopreserved and delivered to the patient's clinic for infusion. The system is unique in the structure of a closed loop where



Figure 1.4: WAVE bioreactor system 200

the patient is both consumer and supplier of his/her autologous cell therapy. Operations in such a system is challenging and critical not only because of a lack of empirical knowledge of producing the unique therapy, but also because of the involvement of regulators and strict regulatory requirements.

There are over 900 companies developing cell therapies worldwide, with a total of 1028 clinical trials as of the end of 2018 [7]. There is considerable interest in transforming clinical trial empirical knowledge to large scale commercial production. Foundational research projects proposed in this proposal are motivated by both evolving trends in the autologous cell therapy industry and the need to transition cell therapy manufacturing to the commercial arena. The recent commercial approvals of some autologous cell therapies, e.g. Novartis' novel CAR-T cell therapy Kymriah and Kite Pharma's Yescarta in both the U.S. and in Europe, is raising real world manufacturing and supply chain issues, particularly due to the high cost of these therapies.

Currently, cost is a major barrier to the broad accessibility of these immunotherapies [14]. Kymriah has a list price of \$373,000 or \$475,000, depending on the type of cancer, and Yescarta has a list price of \$373,000. These list prices are only a portion of the total cost of treatment, which include the CAR-T infusion, doctors' services and hospital stays, but do not include time away from work and living expenses away from home if the patient has to travel to the infusion site, can easily exceed \$750,000. As a result, there is considerable interest in cost reduction with the help of good capacity planning to minimize capital investments and inventory costs. We note that cost reduction is a fundamental, but not exclusive, goal of the proposed research. Cost reduction must take into consideration various types of risk. Autologous cell therapy is an emerging industry with an emerging supplier base. Supply chain disruption due to the unavailability of a reagent has been a severe risk for biomanufacturers and has prompted considerable interest and concern in the industry.

The networks of interests are a centralized manufacturing network, a decentralized manufacturing network with independently operated facilities and a decentralized man-

ufacturing network allowing specimen/reagent transshipment and bioreactor relocation. Solving CP problem in these network will (1) provide direct insights to plant managers on determination of bioreactor adjustment and reagent replenishment decisions, (2) provide the pharmaceutical company about cost and resilience of each network designs, (3) refine the manufacturing network design with consideration of major supplier disruption risks, (4) cope with regulatory policies with respect to patient accessibility.

1.7.2 Epidemic/pandemic control

The recent rapid spread of COVID-19 virus has huge impacts on the healthcare network nationally and globally. The healthcare service network as well as its supply chain of medical countermeasures is ripped due to critical medical supplies shortages. The shortages of personal protection equipment (PPE) have positioned doctors and nurses on the front line of this invisible war in great danger; and the shortages of ventilators, treatment rooms and ICUs have led to an extensive number of patient mortalities in Wuhan and New York. The medical supplies and equipment planning problem can be modeled as a CP problem.

The treatment of a hospitalized patient can be modeled as a request-to-service process. When a patient develops acute respiratory distress syndrome, the patient requires occupation of life supporting equipment, e.g. a ventilator and other ICU equipment, for an extensive period of time; and the treatment of the patient consumes a bundle of drugs and medical supplies. The treatment of a critical condition patient can only be fulfilled when a life supporting equipment, drugs and other medical supplies are available. We highlight that a life supporting equipment can serve one patient at a time, and is sterilized before its use on another patient, and thus is considered reusable. Most of these equipment are portable. Other relocatable reusable resources in this domain include mobile ICUs, makeshift hospitals and portacabins. Hence these resources can be considered as REMOs, and epidemic/pandemic control network and its supply chain can be considered reconfigurable.

Supplier disruptions lead to critical shortages of PPE during the coronavirus season.

The world's major PPE suppliers are located overseas, e.g. China and Vietnam. With the catastrophic outbreak of COVID-19 in China, the supply channels of PPE in the US are disrupted, causing life threatening PPE shortages in the disease control network. On the other hand, regional shortages of ventilators and hospital beds contribute to bottlenecks of regional (state-level) fulfillment capacities. The national and state-level stockpiles MCMs (e.g. PPE, ventilators and antibiotics) are insufficient compared to the boosting of confirmed cases of COVID-19. A resilient epidemic/pandemic control network with the ability to mitigate risks of supplier disruptions and regional demand surge is required.

The epidemic/pandemic control network can be considered as a decentralized service network. The state-level healthcare facilities can be simplified as a service node providing treatment for regional patients. On the federal level, national stockpiles of MCMs as well as other resources (e.g. military makeshift hospitals, mobile ICUs, mobile hospitals and etc.) can be risk-pooled to provide relocatable capacity across the national network. The networks of interests are a decentralized service network with independently operated regions and a decentralized service network allowing transshipment of consumable resources and relocation of reusable resources. Solving CP problem in these networks will (1) guide the capacity planning on both federal and state level to increase readiness for disease outbreaks and other natural disasters, (2) reveal the importance of federal level MCMs stockpile for cost-efficiency, (3) refine the service network design with consideration of major supplier disruption risks, (4) propose a inter-state coordination in combating the epidemic/pandemic.

1.8 Objectives and Organization

Our research objectives include but are not limited to

- (1) investigation on how to conduct capacity planning in centralized and decentralized networks with simultaneous control of both equipment quantity and inventory level of critical material;

- (2) Modeling and quantification of supplier disruption risks in both centralized and decentralized networks;
- (3) Revealing the costs and benefits of dynamic resilience in both centralized and decentralized networks.

The research findings and knowledge can be extended to a wide spectrum of application domains with managerial tasks in capacity planning of both reusable and consumable resources, in order to 1) reduce total manufacturing/service and logistic costs, 2) ensure high service level, and 3) ensure system resilience against supplier disruption risks. All proposed models and algorithms can be easily translated to applicable domains not limited to personalized medicine, 3D-printing manufacturing, relocatable storage network, and epidemic/pandemic control.

This thesis is organized as follows. In Chapter 2, we investigate approaches for determining the number of reusable resources and a replenishment policy for consumable resources in a centralized network, e.g. an autologous build-to-order cell therapy manufacturing system. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. For each model, we examine two variants, one where there are penalties for an insufficient number of bioreactors and units of reagents and one where there are chance constraints on an insufficient number of bioreactors and units of reagents. In a case study, we determine an optimal or near-optimal number of bioreactors and a reagent replenishment policy based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania.

In Chapter 3, We investigate approaches to determine the number of reusable resources and a replenishment policy for consumable resources in a centralized network, e.g. an autologous build-to-order cell therapy manufacturing system, when confronted with supplier disruption risks. We model a supplier disruption process as an exogenous Markov process. Incidents that cause reagent shortages include natural disasters, human errors and disease

outbreaks. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. Both models are formulated as Markov decision processes (MDP). We propose algorithms to solve the proposed models. In the case studies, given different supplier disruption scenarios, we determine bioreactors adjustment quantities and reagent replenishment policies based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania.

In Chapter 4, we investigate a capacity planning problem in a decentralized network when confronted with supplier disruption risks. We consider two kinds of decentralized networks. The first network has isolated regional facilities (Iso-Net), while the second network has coordinated regional facilities (Co-Net). In an Iso-Net, each facility fulfills its regional demand with its own capacity; while in a Co-Net, all facilities coordinate by sharing demand and transshipping resources (bioreactor and reagent). The second supply chain is reconfigurable as production capacities in different facilities are dynamically adjustable. We show that a coordinating and reconfigurable decentralized network exhibits greater resilience with lower costs in the face of supplier disruption risks. In the case studies, we construct a hypothetical decentralized network based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. We compare decentralized models and different heuristic policies based on the constructed decentralized network. Testing results suggest that instead of increasing resource redundancy in all facilities, it is beneficial to only restore limited level of redundancy and adaptively reconfigure the network.

In Chapter 5, we conclude this thesis and list some open research questions identified from Chapters 2-4.

CHAPTER 2

RESILIENT CAPACITY PLANNING IN A CENTRALIZED NETWORK

In this chapter, we introduce a Capacity Planning (CP) problem in a centralized network. We demonstrate problem settings and notations in the context of autologous cell therapy. Translating all models, algorithms and results to other domains is straightforward.

We investigate approaches for determining the number of reusable resources and a replenishment policy for consumable resources for an autologous build-to-order cell therapy manufacturing system. We assume therapy manufacturing for a patient can only begin if there is a patient specimen, an idle bioreactor (assumed to be reusable), and a unit of reagent (a consumable) at the manufacturing facility. Patients who qualify for such cell therapies have serious and even life-threatening diseases, and delays in therapy manufacturing can increase patient mortality rate. Hence, it is important for an idle bioreactor and a unit of reagent to be available when a patient's specimen arrives at the facility. However, cell therapy manufacturing is very expensive. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. For each model we examine two variants, one where there are penalties for an insufficient number of bioreactors and units of reagents and one where there are chance constraints on an insufficient number of bioreactors and units of reagents. In a case study, we determine an optimal or near-optimal number of bioreactors and a reagent replenishment policy based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania.

2.1 Introduction

Autologous cell therapy is an emerging therapeutic method that has proven to be dramatically effective for treating serious and even life-threatening diseases, such as various can-

cers, cardiovascular disease, liver disease, diabetes, neuro-degenerative disorders, bone repair, and spinal cord injuries. Autologous cell therapy is a form of "personalized" medicine, manufactured for a specific patient that harnesses the patient's own immune system in the production of the therapy. As the emerging cell therapy manufacturing industry moves from clinical trials to commercialization, the need for capacity planning becomes of critical importance. The intent of this thesis is to address a key capacity planning challenge for the industry – determining an optimal level of manufacturing capacity and an optimal replenishment policy for an important consumable required during the cell therapy manufacturing process.

The autologous cell therapy manufacturing supply chain begins by drawing a specimen from the patient at the patient's point-of-care (POC; e.g., clinic, hospital) and ends with infusing the resulting therapy in the same patient at the patient's POC [15]. Once the specimen is drawn from the patient, it is then transported to the cell therapy manufacturing facility, where initially an acceptance check is conducted to ensure that this specimen is able to yield a quality final product; e.g., the specimen must contain a sufficient amount of primitive multipotent stem or progenitor cells from the patient's bone marrow, adipose tissue, peripheral blood and other specific tissues to support therapy production. If the acceptance check is successful, target cells are first isolated from the specimen. Then these cells are expanded and harvested in a *bioreactor* (a device or apparatus, which we assume is reusable, in which living organisms synthesize useful substances), consuming units of *reagents* (a chemical ingredient, i.e., a compound or mixture, typically of inorganic or small organic molecules, which we consider a consumable, introduced to cause the desired transformation of an organic substance) supplied from FDA approved suppliers. Quality tests are conducted during the cell culturing processes. The cell batch is then washed, purified, and formulated into the customized final therapy product. A final quality test is conducted to ensure viability of the therapy. The cell therapy is finally cryopreserved and delivered to the patient's POC for infusion. Each patient's therapy is produced, separate

from the production of other patients' therapies, and is not interchangeable. A patient's therapy manufacturing can only begin after the patient's specimen is received at the manufacturing facility and once there is an available idle bioreactor and an available unit of reagent.

This process is build-to-order and complex due to the stochastic and dynamic natures of the biologic product and patient status. Operations in such a system are challenging and critical not only because of a lack of empirical knowledge of producing the unique therapy, but also because of the involvement of regulators and strict regulatory requirements. Additional detail regarding the manufacturing process and the rapidly emerging cell manufacturing industry can be found in [7] and [16].

The objective of this thesis is to develop and analyze a tractable (and hence necessarily simplified) stochastic optimization model of the cell therapy manufacturing process that can provide insights directly to the manager of the manufacturing facility for answering two key capacity planning questions, given cost information and information about the specimen arrival process: *What is the number of bioreactors needed for therapy production? What is a good reagent replenishment policy?* Further, we will show that the insights derived from the analytical model can be useful to reduce simulation computation time and the number of simulations for a simulation of the cell therapy manufacturing process. The value of such a simulation is that it can support a far more detailed and hence more realistic description of this process than the stochastic optimization model can and can produce detailed information valuable at an operational level.

2.1.1 Literature review

Autologous cell therapy manufacturing is an example of build-to-order (BTO), a manufacturing approach where product manufacturing may only begin after the order for the product is received. BTO is a particularly appropriate manufacturing approach for highly customized and/or low volume products such as autologous cell therapies. BTO has proved

to be a successful manufacturing approach for many companies in a variety of industries, e.g., consumer electronics, automotive ([1, 2, 3]).

Capacity planning (CP) is the process of determining the production capacity needed to meet product demand. Production capacity is variable for the domain considered in this thesis; at decision epoch t , production capacity equals the minimum of (i) the number of idle bioreactors at t and (ii) the number of units of reagents in inventory at t . Although workforce development is a major challenge for this industry, workforce capacity is not considered in this thesis for model simplicity and enhanced model tractability. A capacity planning that considers multiple product lines and facilities and assumes production capacity is pre-determined over the planning horizon has been studied by [17], [18], [19], and [20]. [17] modelled demand as a finite set of demand scenarios, [18] use point-estimate forecasts of demand, and [19] considered cumulative distribution functions to describe demand. [19] also assumed that the production capacity of a facility is dependent on workforce capacity, modelled the capacity planning problem as a multi-stage mixed integer program (MIP), and developed a two-stage heuristic for determining a near-optimal solution to the MIP. [21] extended the model presented in [19] to a multi-stage dynamic program (DP). [20] considered a multi-stage CP problem with service constraints to determine material planning in a multi-echelon production supply chain.

Capacity planning for a single production site has been studied by [22], [23], and [24]. [22] proposed a deterministic MIP to determine workforce plans so that build-to-order demands can be satisfied with a minimum number of workers. [23] developed a DP approach to manage inventory and production capacity adjustments based on sales forecasts, workforce constraints, and manufacturing lead time and proposed a rule-based heuristic to find a feasible solution to the computationally challenging DP. [24] minimized the total cost of capacity changes and operations over a finite planning horizon for a CP problem. We remark that none of these studies considered material management. However, each assumed additional staffing could reduce manufacturing lead time, which is currently not a realistic

assumption for autologous cell manufacturing.

As autologous cell therapy transitions from clinical trials to commercialization, capacity planning and material management have become of greater interest to industry. Challenges in production capacity planning and material management for autologous cell therapy manufacturing are discussed in [7], [16] and [25]. [7] proposed an agent-based simulation framework to construct a simulation ‘digital twin’ model of an autologous cell therapy manufacturing facility network and used the simulation model to examine several issues of importance to the emerging industry, e.g., supply chain cost reduction, supplier disruption impacts on network productivity, the impact on patient benefit of different queuing disciplines for patient specimens arriving at a manufacturing facility, the impact on network resilience and productivity on manufacturing capacity and reagent inventory policy, the number of bioreactors and units of reagent a cell manufacturing facility should have available. [16] presented an overview of the current state of the art and challenges in autologous cell manufacturing. They emphasized the importance of developing operational level analytical models to support supply chain capacity planning decisions using mathematical programming approaches. [25] discussed capacity allocation at a strategic level in an autologous cell therapy manufacturing network. None of this research examines analytical approaches for determining an optimal or near-optimal level of manufacturing capacity (e.g., the number of bioreactors) or an optimal or near-optimal reagent inventory replenishment policy.

Our interest in chance-constrained dynamic programming (CCDP) is due to we believe and based on our experience that chance constraints are easier for assessing key system parameters from healthcare providers than assessing (a different set of) key system parameters for a standard DP model of a cell manufacturing network. The CCDP was initially motivated by water management problems [26, 27]. Solution methods for constrained Markov decision processes due to [28] have been adopted to solve the CCDP if the chance constraints can be formulated so that they share the same additive structure of

the objective function [27, 29, 30]. Several reformulation approaches have been proposed by [31]. [30] considered a CCDP with a joint chance-constraint that violated the additive structure requirement using a heuristic developed by a conservative approximation of the joint probability using an approach based on Boole's inequality.

2.1.2 Organization

Our thesis is outlined as follows. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch (Section 2.2) and one where this number can only be adjusted periodically (Section 2.3). For each model we examine two variants, one where there are penalties for an insufficient number of bioreactors and units of reagents (Section 2.2.1 and Section 2.3.1) and one where there are chance constraints on an insufficient number of bioreactors and units of reagents (Section 2.2.2 and Section 2.3.2). For the case where the number of idle bioreactors can be adjusted at each epoch and for both variants of this case, we will show that this number and the reagent replenishment policy are both myopic and base stock. For the case where the number of bioreactors can only be adjusted periodically, we will present near-optimal heuristics for determining the number of bioreactors and the reagent replenishment policy. In a case study in Section 2.4, we determine an optimal or near-optimal number of bioreactors and a reagent replenishment policy based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania.

2.2 Adjustable Bioreactor Quantity (ABQ) Model

We now present an infinite horizon total discounted cost Markov Decision Process (MDP) model of the reagent replenishment and bioreactor quantity selection processes, assuming that at each decision epoch we can (i) replenish the reagent inventory instantaneously and (ii) adjust the number of idle bioreactors held at the facility instantaneously. Instantaneous replenishment of reagents is an unrealistic assumption; in reality, there are time lags (often

stochastic) between order placement and order receipt. The assumption that the number of idle bioreactors (and hence the total number of bioreactors at the facility) can be adjusted at each decision epoch instantaneously may be realistic if the facility manufactures several cell therapy lines, the bioreactors needed for each cell therapy line are interchangeable, and there is a substantial number of 'buffer' bioreactors in the facility. However, for a facility that manufactures only a single therapy, the total number of bioreactors at the facility would be more realistically adjusted only periodically. Nevertheless, the assumptions (i) and (ii) above ensure tractability and we will find the solution to this problem useful in addressing both of the questions: *What is the number of bioreactors needed for therapy production? What is a good reagent replenishment policy?* More specifically, we will find the solution to the adjustable bioreactor case useful in limiting the simulation search for an optimal number of bioreactors can only be adjusted periodically.

We now present the adjustable bioreactor quantity (ABQ) model. Assume at decision epoch t that

- $s_t \geq 0$ is the number of specimens in the arrival queue, waiting for therapy manufacturing to begin;
- $b_t^\tau \geq 0$ is the number of bioreactors that are τ epochs from completing the manufacturing process. Let $\mathbf{b}_t = (b_t^0, \dots, b_t^{T-1})$, where T epochs are required to complete the manufacturing process and b_t^0 is the number of idle bioreactors;
- $r_t \geq 0$ is the number of units of the reagent in inventory. In reality, multiple reagents (e.g., media, recombinant proteins, viral vectors, scaffolds, and other critical materials) are required for therapy manufacturing. We consider a unit of reagent to be a collection of the multiple reagents needed to manufacture a single therapy.

Let $x_t = (s_t, \mathbf{b}_t, r_t)$ be the system state at epoch t .

Further, assume

- $0 \leq m_t \leq \min\{s_t, b_t^0, r_t\}$, a decision variable, is the number of therapies starting at time t . Thus, a patient's therapy production can only begin if the patient's specimen is in the arrival queue and there is an idle bioreactor and one unit of each type of reagent at the manufacturing facility that can be assigned to the specimen
- d_t is the number of specimens that arrive between epoch t and epoch $t + 1$. We assume that the stochastic process $\{d_t, t \geq 0\}$ is a sequence of independent, identically distributed random variables with known distribution, where the realization of d_t becomes available to the decision maker just before epoch $t + 1$
- q_t , a decision variable, is the number of idle reactors either added to the current number of idle reactors (in which case q_t is non-negative) or subtracted from the current number of idle reactors (in which case q_t is non-positive) at epoch t and is received or removed before epoch $t + 1$
- $a_t \geq 0$, a decision variable, is the number of units of the reagent ordered at epoch t and received before epoch $t + 1$.

We assume m_t , q_t and a_t are determined based on knowledge of x_t . The resulting system dynamics are:

$$\begin{aligned}
 s_{t+1} &= s_t - m_t + d_t, \\
 b_{t+1}^\tau &= \begin{cases} b_t^0 - m_t + b_t^1 + q_t & \tau = 0, \\ b_t^{\tau+1} & 1 \leq \tau \leq T - 2, \\ m_t & \tau = T - 1, \end{cases} \\
 r_{t+1} &= r_t - m_t + a_t.
 \end{aligned}$$

from which we obtain an updated system state x_{t+1} . Since the number of idle bioreactors is non-negative, $q_t \geq -(b_t^0 - m_t + b_t^1)$, where $b_t^0 - m_t + b_t^1$ is the maximum number of

bioreactors that can be removed from the system.

In the remainder of this section, we introduce two modeling approaches, a classic penalty cost approach in Section 2.2.1 and a chance constraint approach in Section 2.2.2, to determine optimal capacity planning policies. The connection between the two approaches is presented in Section 2.2.3.

2.2.1 Penalty cost approach

Assume the single period costs accrued between epoch t and $t + 1$ are as follows:

- $c_R a_t$ is the reagent replenishment cost for the reagent, where c_R is the cost per unit of reagent
- $c_B q_t$ is the bioreactor adjustment cost or savings, where c_B is the cost or savings per bioreactor, depending on whether q_t is positive or negative
- $h_R(r_{t+1} - s_{t+1})^+$ is the random variable holding cost for the reagent, charged only if there are more units of reagent than specimens at epoch $t + 1$, where h_R is the holding cost per unit of reagent
- $h_B(b_{t+1}^0 - s_{t+1})^+$ is the random variable bioreactor holding cost, charged only if there are more bioreactors than specimens at epoch $t + 1$, where h_B is the bioreactor holding cost per bioreactor
- $p_R(s_{t+1} - r_{t+1})^+$ is the random variable penalty for having an insufficient number of units of reagent in order to start therapy manufacturing for all specimens that have arrived at the facility and are waiting for therapy production to begin
- $p_B(s_{t+1} - b_{t+1}^0)^+$ is the random variable penalty for having an insufficient number of idle bioreactors in order to start therapy manufacturing for all specimens that have arrived at the facility and are waiting for therapy production to begin.

We now present the optimality equation for the ABQ model with penalty cost (ABQ-p), assuming an infinite horizon total discounted cost MDP model with discount factor $\beta \in [0, 1)$. Let

$$[Hv](x_t) = \min_{\substack{0 \leq m_t \leq \min\{s_t, b_t^0, r_t\} \\ q_t \geq -(b_t^0 - m_t + b_t^1), a_t \geq 0}} \mathbb{E}\{C(q_t, a_t, m_t, d_t|x_t) + \beta v(x_{t+1})\} \quad (\text{ABQ-p})$$

where

$$\begin{aligned} C(q_t, a_t, m_t, d_t|x_t) = & c_R a_t + h_R (r_{t+1} - s_{t+1})^+ + p_R (s_{t+1} - r_{t+1})^+ \\ & + c_B q_t + h_B (b_{t+1}^0 - s_{t+1})^+ + p_B (s_{t+1} - b_{t+1}^0)^+ \end{aligned}$$

is the single period cost function. Results in [32] guarantee the existence of a unique fixed point v^* such that $\lim_{n \rightarrow \infty} \|v_n - v^*\| = 0$, where $\{v_n\}$ is such that $v_0 = 0$, $v_{n+1} = Hv_n$ for all n , and $\|\cdot\|$ is the sup-norm. Further, a policy that achieves the minimum in Hv^* is an optimal policy. Thus, by determining the fixed point of the operator H , the capacity planner obtains an optimal reagent replenishment policy and an optimal idle bioreactor adjustment policy. Our next result presents a simple, easily implemented and computed policy that is optimal for the ABQ-p, assuming Assumption (A1) holds, which we assume holds throughout the remainder of the thesis.

Assumption (A1). $r_1 \leq s_1 + F_d^{-1} \left(\frac{p_R - (1-\beta)c_R}{p_R + h_R} \right)$.

An interpretation of A1 is that the initial amount of reagent in inventory is insufficient to support cell therapy manufacturing for the initial number of specimens in the arrival queue added to a number that represents the number of specimens expected to arrive in the next period plus an amount of buffer reagent inventory determined by the cost structure and the cumulative density function (CDF) of demand. We now present an optimal policy for the ABQ-p.

Proposition 2.2.1. *If A1 holds, the following policy is optimal for the ABQ-p*

$$a^p(x_t) = s_t + F_d^{-1} \left(\frac{p_R - (1 - \beta)c_R}{p_R + h_R} \right) - r_t \quad (2.1a)$$

$$q^p(x_t) = s_t + F_d^{-1} \left(\frac{p_B - (1 - \beta)c_B}{p_B + h_B} \right) - (b_t^0 + b_t^1) \quad (2.1b)$$

$$m^p(x_t) = \min\{s_t, b_t^0, r_t\}. \quad (2.1c)$$

A proof of Proposition 2.2.1 is presented in Appendix A. We remark it is straightforward to show that if A1 holds for epoch t and if the policy presented in Proposition 2.2.1 (2.1a) is used, then A1 will hold for epoch $t+1$. Further, we note that policy (2.1) is myopic and easy to implement.

An alternative model of the penalty cost $p_R(s_{t+1} - r_{t+1})^+ + p_B(s_{t+1} - b_{t+1}^0)^+$ is $p(s_{t+1} - m_{t+1})^+$. We assume that $C(q_t, a_t, m_t, d_t|x_t)$ is our single period cost function throughout the remainder of this thesis due to the existence of the easily implemented and computed optimal policy presented in Proposition 2.2.1. However, for the moment, let

$$\begin{aligned} \hat{C}(q_t, a_t, m_t, d_t|x_t) &= C(q_t, a_t, m_t, d_t|x_t) - (p_R(s_{t+1} - r_{t+1})^+ + p_B(s_{t+1} - b_{t+1}^0)^+) \\ &\quad + p(s_{t+1} - m_{t+1})^+ \end{aligned}$$

and consider the resulting optimal equation

$$v(x_t) = \min_{\substack{0 \leq m_t \leq \min\{s_t, b_t^0, r_t\} \\ q_t \geq -(b_t^0 - m_t + b_t^1), a_t \geq 0}} \mathbb{E}\{\hat{C}(q_t, a_t, m_t, d_t|x_t) + \beta v(x_{t+1})\}. \quad (2.2)$$

A relationship between ABQ-p and (2.2) is presented in the following result, the proof of which is found in Appendix A.

Corollary 2.2.2. *If $p_B + p_R = p$ and $\frac{p_R - (1 - \beta)c_R}{p_R + h_R} = \frac{p_B - (1 - \beta)c_B}{p_B + h_B}$, then ABQ-p and (2.2) have a same optimal policy (q^p, a^p, m^p) .*

We remark that $c_R, c_B, h_R,$ and h_B are cost values that are straightforward to determine, relative to the penalty parameters p_R and p_B , which require assessment from an individual or individuals involved in the decision-making process. The penalty parameters represent penalties, in monetary units, of delaying the beginning of therapy manufacturing and hence delaying the beginning of treatment one period for a single patient. A delay in beginning treatment for patients with extremely serious health conditions such as those treated by autologous cell therapies can increase the mortality risk of these patients. Based on our experience in assessing such parameters from human subjects [33], we anticipate that it may be difficult to directly assess the penalty parameters. Further, in this domain we have found human decision makers easily provide chance constraints useful for indirectly assessing the penalty parameters. For these reasons, we next consider an alternative chance constraint-based problem formulation that is straightforward to solve and serves as a path to indirectly determining the penalty parameters.

2.2.2 Chance constraint approach

Consider the following modification of the problem presented in Section 2.2.1. The approach suggests to remove the penalty terms from the single period costs and replace them with two chance constraints

$$\begin{aligned} Pr[s_{t+1} > b_{t+1}^0 | x_t] &\leq \alpha_B, & \forall t; \\ Pr[s_{t+1} > r_{t+1} | x_t] &\leq \alpha_R, & \forall t. \end{aligned}$$

where α_B and α_R are presumably small probabilities that must be assessed from the individual or individuals involved in the decision making process. We have found that clinicians provide these probabilities with little effort or concerns. The probability $(1 - \alpha_R)$ represents the probability that a patient whose specimen arrives at the facility between epochs t and $t + 1$ will have a unit of reagent available at epoch $t + 1$ so that therapy produc-

tion can begin at $t + 1$. Similarly, the probability $(1 - \alpha_B)$ represents the probability that a patient whose specimen arrives at the facility between epochs t and $t + 1$ will have a bioreactor available at epoch $t + 1$ so that therapy production can begin at $t + 1$. Hence the revised model becomes an infinite horizon chance constrained dynamic programming problem with discount factor $\beta \in [0, 1)$

$$\begin{aligned}
v(x_t) = & \min_{\substack{0 \leq m_t \leq \min\{s_t, b_t^0, r_t\} \\ q_t \geq -(b_t^0 - m_t + b_t^1), a_t \geq 0}} \mathbb{E}\{\tilde{C}(q_t, a_t, m_t, d_t|x_t) + \beta v(x_{t+1})\} & \text{(ABQ-c)} \\
\text{s.t. } & Pr[s_{t+1} > b_{t+1}^0|x_t] \leq \alpha_B, \\
& Pr[s_{t+1} > r_{t+1}|x_t] \leq \alpha_R,
\end{aligned}$$

where

$$\tilde{C}(q_t, a_t, m_t, d_t|x_t) = c_R a_t + h_R (r_{t+1} - s_{t+1})^+ + c_B q_t + h_B (b_{t+1}^0 - s_{t+1})^+.$$

Let Assumption A2 be defined as follows:

Assumption (A2). $r_1 \leq s_1 + F_d^{-1}(1 - \alpha_R)$.

Noting the close resemblance of A1 and A2, we now present an optimal policy for ABQ-c.

Proposition 2.2.3. *An optimal reagent replenishment policy and an optimal bioreactor adjustment policy for ABQ-c is to select*

$$a^*(x_t) = (s_t + F_d^{-1}(1 - \alpha_R) - r_t)^+ \quad (2.3a)$$

$$q^*(x_t) = s_t + F_d^{-1}(1 - \alpha_B) - (b_t^0 + b_t^1) \quad (2.3b)$$

$$m^*(x_t) = \min\{s_t, b_t^0, r_t\}, \quad (2.3c)$$

and if A2 holds, $a^*(x_t) = s_t + F_d^{-1}(1 - \alpha_R) - r_t$.

Recall that if A1 holds for epoch t and the policy (2.1a) in Proposition 2.2.1 is followed, then A1 holds for epoch $t + 1$. Similar, it is straightforward to show that if A2 holds for epoch t and policy (2.3a) in Proposition 2.2.3 is followed, then A2 holds for epoch $t + 1$. Similar to (2.1), we remark that policy (2.3) is myopic and easy to implement.

2.2.3 Connections between ABQ-p and ABQ-c

Our next result follows by comparing policies (2.1) and (2.3).

Corollary 2.2.4. *If $\alpha_B = \frac{h_B + (1-\beta)c_B}{p_B + h_B}$ and $\alpha_R = \frac{h_R + (1-\beta)c_R}{p_R + h_R}$, then ABQ-p and ABQ-c have identical optimal bioreactor adjustment and reagent replenishment policies.*

Corollary 2.2.4 suggests that probabilities α_B and α_R are equivalent to penalty costs

$$p_B(\alpha_B) = \frac{(1 - \alpha_B)h_B + (1 - \beta)c_B}{\alpha_B}, \quad (2.4a)$$

$$p_R(\alpha_R) = \frac{(1 - \alpha_R)h_R + (1 - \beta)c_R}{\alpha_R}. \quad (2.4b)$$

We now further present connections between ABQ-p and ABQ-c through a primal-dual method. Since an optimal production decision $m(x_t)$ is shown always to equal $\min\{s_t, b_t^0, r_t\}$, we substitute $m_t = \min\{s_t, b_t^0, r_t\}$ in ABQ-p and ABQ-c through the remainder this section. The chance constraints in ABQ-c can be rewritten as $q_t \geq Z_B + s_t - (b_t^0 + b_t^1)$ and $a_t \geq Z_R + s_t - r_t$, where $Z_B = F_d^{-1}(1 - \alpha_B)$ and $Z_R = F_d^{-1}(1 - \alpha_R)$. Therefore, ABQ-c can be reformulated as

$$\begin{aligned} \min_{q,a} \quad & \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\{\tilde{C}(q_t, a_t, m_t, d_t|x_t)\} \\ \text{s.t.} \quad & q_t \geq Z_B + s_t - (b_t^0 + b_t^1), \quad \forall t; \\ & a_t \geq Z_R + s_t - r_t, \quad \forall t; \\ & q_t \geq -(b_t^0 - m_t + b_t^1), \quad \forall t; \\ & a_t \geq 0, \quad \forall t. \end{aligned} \quad (2.5)$$

We relax the first and second set of constraints with Lagrangian multipliers $\mu_{B,t} \geq 0$ and $\mu_{R,t} \geq 0$, and let $\mu_{B,t} = \beta^{t-1}\lambda_{B,t}$ and $\mu_{R,t} = \beta^{t-1}\lambda_{R,t}$. We construct the Lagrangian dual of (2.5),

$$\begin{aligned} \max_{\{\lambda_{B,t}, \lambda_{R,t} : t \geq 0\}} \min_{q,a} \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\{L(q_t, a_t, m_t, d_t | x_t, \lambda_{B,t}, \lambda_{R,t})\} \quad (2.6) \\ \text{s.t. } q_t \geq -(b_t^0 - m_t + b_t^1), \quad \forall t; \\ a_t \geq 0, \quad \forall t, \end{aligned}$$

where

$$\begin{aligned} L(q_t, a_t, m_t, d_t | x_t, \lambda_{B,t}, \lambda_{R,t}) = \tilde{C}(q_t, a_t, d_t | x_t) + \lambda_{B,t}(Z_B + s_t - (q_t + b_t^0 + b_t^1)) \\ + \lambda_{R,t}(Z_R + s_t - r_t - a_t), \end{aligned}$$

The optimal objective value of the Lagrangian relaxation, i.e. the inner minimization problem

$$\begin{aligned} \min_{q,a} \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\{L(q_t, a_t, m_t, d_t | x_t, \lambda_{B,t}, \lambda_{R,t})\} \\ \text{s.t. } q_t \geq -(b_t^0 - m_t + b_t^1), \quad \forall t; \\ a_t \geq 0, \quad \forall t \end{aligned}$$

is jointly convex in $\lambda_{B,t}$ and $\lambda_{R,t}$ for all t , and hence there exists an optimal dual solution $\{\lambda_{B,t}^*, \lambda_{R,t}^* : t \geq 0\}$ (see [34]). By strong duality, (q_t^*, a_t^*) defined in (2.3) is optimal for

$$v(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1), a_t \geq 0} \mathbb{E}\{L_t(q_t, a_t, m_t, d_t | x_t, \lambda_{B,t}^*, \lambda_{R,t}^*) + \beta v(x_{t+1})\}. \quad (2.7)$$

Then by Corollary 2.2.4 and matching up the terms in (q_t^*, a_t^*) and (q_t^p, a_t^p) , there exists a

time-invariant set of optimal dual variables, such that $\lambda_{B,t}^* = \lambda_B^*$ and $\lambda_{R,t}^* = \lambda_R^*$, and

$$\lambda_B^* = \frac{(1 - \alpha_B)h_B + (1 - \beta)c_B}{\alpha_B} = p_B(\alpha_B) \quad (2.8a)$$

$$\lambda_R^* = \frac{(1 - \alpha_R)h_R + (1 - \beta)c_R}{\alpha_R} = p_R(\alpha_R) \quad (2.8b)$$

which are marginal costs of constraints $q_t \geq Z_B + s_t - (b_t^0 + b_t^1)$ and $a_t \geq Z_R + s_t - r_t$ for all t . We remark that λ_B^* and λ_R^* are functions of α_B and α_R respectively. We note that λ_B^* is monotonically decreasing in α_B , implying that the marginal cost of violating $Pr[s_{t+1} > b_{t+1}^0] \leq \alpha_B$ is monotonically increases in α_B . Similarly, since $p_B(\alpha_B)$ is monotonically decreasing in α_B , a reduction in α_B implies an increase in $p_B(\alpha_B)$. Similar results hold for the chance constraints on reagent outages and the penalty cost on an insufficient number of units of reagent. We remark that by establishing capacity outage probability thresholds, the chance constraint model implies underlying costs to ensure that the chance constraints are satisfied through their marginal costs (λ_B^* and λ_R^*). Since the marginal cost at (q_t^*, a_t^*) is directly related to penalty costs ($p_B(\alpha_B)$ and $p_R(\alpha_R)$) in the penalty cost model, the primal-dual results explicitly show how the two models are interrelated.

2.3 Fixed Bioreactor Quantity (FBQ) Model

We now consider the case where the number of bioreactors can only be readjusted periodically, i.e., every \mathcal{T} decision epochs, rather than at each epoch (the $\mathcal{T} = 1$ case). We expect \mathcal{T} to be significantly larger than T . For example, assume the time between decision epochs is one week, three weeks are required to manufacture a therapy, and the number of bioreactors can be adjusted quarterly. Then, $T = 3$ and $\mathcal{T} = 13$. We refer to this problem as the fixed bioreactor quantity (FBQ) model.

2.3.1 Penalty cost approach

Let the operator \bar{H} be defined as

$$[\bar{H}v](x_t) = \min_{a_t \geq 0} \mathbb{E}\{C(0, a_t, d_t|x_t) + \beta v(x_{t+1})\}.$$

Define the sequence $\{v_{\tau,n}\}$ as:

$$\begin{aligned} v_{\tau,n}(x_t) &= [\bar{H}v_{\tau+1,n}](x_t), & 2 \leq \tau \leq \mathcal{T}, \\ v_{1,n}(x_t) &= [Hv_{2,n}](x_t), \end{aligned} \tag{FBQ-p}$$

$v_{\mathcal{T}+1,n+1} = v_{1,n}$, and assume $v_{\mathcal{T}+1,0} = 0$. Results in [32] imply that there exists a set $\{v_\tau, \tau = 1, \dots, \mathcal{T} + 1\}$ such that this set is the unique solution to $v_\tau^* = \bar{H}v_{\tau+1}^*$ for $\tau = 2, \dots, \mathcal{T}$, $v_1^* = Hv_2^*$, and $v_1^* = v_{\mathcal{T}+1}^*$, and $\lim_{n \rightarrow \infty} \|v_{\tau,n} - v_\tau^*\| = 0$ for $\tau = 1, \dots, \mathcal{T} + 1$, where $\|\cdot\|$ is the sup-norm. Further, a policy π_τ^* that achieves the minimum in $\bar{H}v_{\tau+1}$ for $\tau = 2, \dots, \mathcal{T}$ and the minimum in Hv_2 for $\tau = 1$ is an optimal policy.

Unfortunately, the optimal policy structure in Proposition 2.2.1 (2.1b) does not extend to the $\mathcal{T} > 1$ case. This loss of structure and the resulting computational burden of determining an optimal policy for FBQ-sdr motivates interest in examining heuristic procedures.

Consider the following procedure for determining a near-optimal policy for FBQ-sdr. Let x_1 be such that $r_1 - s_1 \leq F_d^{-1}\left(\frac{p_R - (1-\beta)c_R}{p_R + h_R}\right)$. Fix the reagent replenishment decision as $\hat{a}_t(x_t) = s_t + F_d^{-1}\left(\frac{p_R - (1-\beta)c_R}{p_R + h_R}\right) - r_t$ for all $t \geq 0$, which is guaranteed to be nonnegative since $r_t - s_t \leq F_d^{-1}\left(\frac{p_R - (1-\beta)c_R}{p_R + h_R}\right)$ for all t . Given this policy, for $2 \leq \tau \leq \mathcal{T}$, let

$$\hat{v}_{\tau,n}(x_t) = \mathbb{E}\{C(0, \hat{a}_t, d_t|x_t) + \beta \hat{v}_{\tau+1,n}(x_{t+1})\},$$

where $r_{t+1} = r_t - m_t + \hat{a}_t = s_t - m_t + F_d^{-1}\left(\frac{p_R - (1-\beta)c_R}{p_R + h_R}\right)$ for all t . Then we determine the

bioreactor adjustment quantity by solving

$$\hat{v}_{1,n}(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1)} \mathbb{E}\{C(q_t, \hat{a}_t, d_t | x_t) + \beta \hat{v}_{2,n}(x_t)\}.$$

Thus, there exists a unique set $\{\hat{v}_\tau^*, \tau = 1, \dots, \mathcal{T} + 1\}$ such that $\hat{v}_1^* = \hat{v}_{\mathcal{T}+1}^*$,

$$\hat{v}_\tau^*(x_t) = \mathbb{E}\{C(0, \hat{a}_t, d_t | x_t) + \beta \hat{v}_{\tau+1}^*(x_{t+1})\},$$

for $1 \leq \tau \leq \mathcal{T} + 1$, and

$$\hat{v}_1^*(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1)} \mathbb{E}\{C(q_t, \hat{a}_t, d_t | x_t) + \beta \hat{v}_{\tau+1}^*(x_{t+1})\}.$$

Let the heuristic policy for adjusting the number of bioreactors achieve the above minimum.

Upper and lower bounds are

$$v^*(x_t) \leq v_\tau^*(x_t) \leq \hat{v}_\tau^*(x_t), \quad 1 \leq \tau \leq \mathcal{T} + 1,$$

where we recall that v^* is the fixed point of H and the optimal value function for ABQ-p. We now present a preliminary numerical analysis of this heuristic, based on CAR-T manufacturing facility specifications given in [7].

Example 1: Assume the demand distribution is Poisson with a mean of 250 patients per year, production duration is $T = 3$ weeks and the planning horizon $\mathcal{T} = 52$ weeks. Based on data presented in [14], let $c_R = \$42,174$, $h_R = \$113.50$, $c_B = \$25,000$ and $h_B = \$14.40$. By Corollary 2.2.4, if the total number of bioreactors is adjustable and $\beta = 0.9$, then it is equivalent to charge $p_R = \$50,273.60$ and $p_B = \$86,504.50$. Let $x_1 = (0, \mathbf{0}, 0)$, which assumes the facility starts with an empty queue and no reagent inventory. It follows that $v^*(x_1) = \$2,704,336.72$ for $v_1^*(x_1)$, and $\hat{v}_1^*(x_1) = \$2,799,168.49$ (with 21

bioreactors), The gap between the lower and upper bound is

$$Gap(x_1) = \frac{|v^*(x_1) - \hat{v}_1^*(x_1)|}{v^*(x_1)} = 0.035.$$

Thus, the expected cost accrued by the heuristic is no more than 3.5% higher than the optimal expected cost of FBQ-sdr.

2.3.2 Chance constraint approach

We now examine the FBQ model with chance constraints. Identical to Section 2.2.2, we assume $Pr[s_{t+1} > r_{t+1}|x_t] \leq \alpha_R$ for all t . Assuming t is a bioreactor adjustment decision epoch, we modify the chance constraint associated with the bioreactors as follows:

$$Pr[s_{t+\tau} > b_{t+\tau}^0|x_t] \leq \alpha_B$$

for $\tau = 1, \dots, \mathcal{T}$. Thus the choice of q_t must ensure these chance constraints hold until the next epoch when bioreactor readjustment is permitted.

We now consider a heuristic that considers only a single bioreactor adjustment period. Without loss of generality, assume bioreactor readjustment occurs at epoch $t = 1$. Thus, at epoch 1, the design phase, the fixed total number of bioreactors is determined over the finite planning horizon and the first reagent replenishment decision is selected. Then for epochs $t \geq 2$, the operations phase, we determine reagent replenishment decisions at each epoch. Let $x_1 = (s_1, \mathbf{b}_1, r_1)$ be a starting state at epoch 1. Then, the optimality equation for the FBQ design phase is

$$\begin{aligned} v_1(x_1) = & \min_{\substack{0 \leq m_1 \leq \min\{s_1, b_1^0, r_1\} \\ q_1 \geq -(b_1^0 - m_1 + b_1^1), a_1 \geq 0}} \mathbb{E}\{\tilde{C}(q_1, a_1, m_1, d_1|x_1) + \beta v_2(x_2)\} & \text{(FBQ-c-d)} \\ \text{s.t. } & Pr[s_{t+1} > b_{t+1}^0|x_1] \leq \alpha_B, \quad \forall t = 1, \dots, \mathcal{T}, \\ & Pr[s_2 > r_2|x_1] \leq \alpha_R. \end{aligned}$$

For $t = 2, \dots, \mathcal{T}$, the optimality equation for the FBQ operation phase is

$$v_t(x_t) = \min_{0 \leq m_t \leq \min\{s_t, b_t^0, r_t\}, a_t \geq 0} \mathbb{E}\{\tilde{C}(0, a_t, m_t, d_t|x_t) + \beta v_{t+1}(x_{t+1})\} \quad (\text{FBQ-c-o})$$

$$s.t. \quad Pr[s_{t+1} > r_{t+1}|x_t] \leq \alpha_R,$$

and $v_{\mathcal{T}+1}(x_{\mathcal{T}+1}) = 0$ is the boundary condition. The FBQ-c-d and FBQ-c-o problems constitute a chance constrained FBQ model FBQ-c. Let $B_1 = \sum_{\tau=0}^{\mathcal{T}-1} b_1^\tau$ be the initial total bioreactor quantities (either busy or idle). Then $B = B_1 + q_1$ is the total number of bioreactors after adjustment q_1 , a number that will remain fixed over the remainder of the planning horizon. Similar to the ABQ model, it is straight forward to show $m_t^* = \min\{s_t, b_t^0, r_t\}$. Therefore, we let $m_t = \min\{s_t, b_t^0, r_t\}$ in the FBQ-c models, which reduces the number of decision variables from three to two.

The FBQ-c problem can be solved using a two-step method:

Step 1. Given a bioreactor adjustment decision q_1 , and for the moment ignoring the chance constraints $Pr[s_{t+1} > b_{t+1}^0|x_1] \leq \alpha_B$, we solve FBQ-c-d and FBQ-c-o and obtain optimal reagent replenishment policies a^* . The minimized total cost for a given design q_1 is

$$u(x_1, q_1) = \min_{a_1 \geq 0} \mathbb{E}\{\tilde{C}(q_1, a_1, m_1, d_1|x_1) + \beta v_2(x_2)\}$$

$$s.t. \quad Pr[s_2 > r_2|x_1] \leq \alpha_R.$$

Step 2. We then seek an optimal q_1 that satisfies $Pr[s_{t+1} > b_{t+1}^0|x_1, q_1, a^*] \leq \alpha_B$ for $t \in [\mathcal{T}]$, and produces the minimum total cost $u(x_1, q_1)$.

In Step 1, an optimal reagent replenishment policy is separable from the bioreactor adjustment decision and has a simple analytical form.

Proposition 2.3.1. *Given q_1 and $q_t = 0$ for all $t \geq 2$, an optimal reagent replenishment*

policy for FBQ-c is to replenish

$$a^*(x_t) = (s_t + F_d^{-1}(1 - \alpha_R) - r_t)^+,$$

for $t \in [\mathcal{T}]$. In addition, if A2 holds, an optimal reagent replenishment policy is to replenish

$$a^*(x_t) = s_t + F_d^{-1}(1 - \alpha_R) - r_t.$$

We assume Assumption A2 holds for the remainder of this section. The FBQ-c model can be rewritten as

$$\begin{aligned} v_1(x_1) = & \min_{q_1 \geq -(b_1^0 - m_1 + b_1^1)} \mathbb{E}\{\tilde{C}(q_1, a^*(x_1), d_1|x_1) + \beta v_2(x_2)\} \\ \text{s.t. } & Pr[s_{t+1} > b_{t+1}^0|x_1] \leq \alpha_B, \quad \forall t = 0, 1, \dots, \mathcal{T}, \end{aligned} \quad (2.9)$$

where for $t = 2, \dots, \mathcal{T}$,

$$v_t(x_t) = \mathbb{E}\{\tilde{C}(0, a^*(x_t), d_t|x_t) + \beta v_{t+1}(x_{t+1})\}$$

and boundary condition such that $v_{\mathcal{T}+1}(x_{\mathcal{T}+1}) = 0$. Let q_1^* the minimum in (2.9). Then

$$v_1(x_1) = \mathbb{E}\{\tilde{C}(q_1^*, a^*(x_1), d_1|x_1) + \beta v_2(x_2)\},$$

and chance constraints $Pr[s_{t+1} > b_{t+1}^0|x_1, q_0^*] \leq \alpha_B$ hold for $t \in [\mathcal{T}]$.

Proposition 2.3.2. *The policy (q_1^*, a^*) is an optimal policy for FBQ-c.*

Proof of Proposition 2.3.2 is presented in Appendix A.

Structural properties

We now present two monotonicity results and a lower bound on q_1^* that will lead to an algorithm for determining q_1^* (Step 2).

Proposition 2.3.3. For all x_1 ,

(i) $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1]$ is non-increasing in q_1 for all $t = 1, \dots, \mathcal{T}$.

(ii) $\mathbb{E}\{\tilde{C}(q_1, a^*(x_1), d_1 | x_1) + \beta v_2(x_2)\}$ is non-decreasing in q_1 .

Proof of Proposition 2.3.3 is given in Appendix A. Assuming a lower bound q_L on q_1^* , Proposition 2.3.3 implies that if we increase q_1 incrementally from q_L until FBQ-c becomes feasible, then the first value of q_1 that achieves this feasibility is optimal.

We now determine a lower bound $q_L = q_L(x_1, \alpha_B)$ on q_1^* . Our lower bound will be such that

for any $q_1 < q_L(x_1, \alpha_B)$, there exists at least one probability among $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1]$ for $t = 1, \dots, \mathcal{T}$ that is larger than α_B . Let $F_{(k)}$ be the CDF of demands in k consecutive periods, i.e. $F_{(k)}^{-1}(\theta) = \inf\{z : F_{(k)} \geq \theta\}$. Define $q_L(x_1, \alpha_B) := \max\{q_1^{(k)}(x_1, \alpha_B) : k = 1, \dots, T\}$, where

$$q_1^{(k)}(x_1, \alpha_B) = \begin{cases} s_1 + F_{(k)}^{-1}(1 - \alpha_B) - \sum_{\tau=0}^k b_1^\tau & 1 \leq k \leq T - 1, \\ s_1 + F_{(T)}^{-1}(1 - \alpha_B) - m_1 - B_1 & k = T. \end{cases}$$

Proposition 2.3.4. The function $q_L(x_1, \alpha_B)$ is a lower bound of q_1^* .

Proof of Proposition 2.3.4 is presented in Appendix A. The next result follows directly from Proposition 2.3.4.

Corollary 2.3.5. Assuming $x_1 = (0, \mathbf{0}, 0)$, $q_1^{(k)}(x_1, \alpha_B) = F_{(k)}^{-1}(1 - \alpha_B)$, and hence $q_L(x_1, \alpha_B) = F_{(T)}^{-1}(1 - \alpha_B)$.

Algorithm design

We now present Algorithm 1, which searches for q_1^* by incrementally increasing q_1 from $q_L(x_1, \alpha_B)$ until feasibility is achieved. We now present an approach for determining fea-

Algorithm 1: FBQ Solver

```

1 Initialization: set  $q_1 \leftarrow q_L(x_1, \alpha_B)$ ;
2 Calculate  $Z_R$  and derive reagent replenishment policy  $a^*$ ;
3 while  $\max_{t \in [\mathcal{T}]} Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] > \alpha_B$  do
4   |   Set  $q_1 \leftarrow q_1 + 1$ .
5 end
6 Output: bioreactor adjustment decision  $q_1$ .

```

sibility in Algorithm 1 (line 3) and briefly mention search for q_1^* can be accelerated by revising line 4 in Algorithm 1.

We use the following a Monte Carlo method to validate $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1]$ for a given bioreactor adjustment decision q_1 . N instances of \mathcal{T} -period simulation paths are sampled, given x_1 , q_1 and a^* . For the i th instance, define $\delta^{(i)}[s_{t+1} > b_{t+1}^0 | x_1, q_1] = 1$ if $s_{t+1} > b_{t+1}^0$ is true and zero-valued otherwise for $i \in [N]$ and $t \in [\mathcal{T}]$. Then the Monte Carlo average $\frac{\sum_{i \in [N]} \delta^{(i)}[s_{t+1} > b_{t+1}^0 | x_1, q_1]}{N}$ is an unbiased estimator of $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1]$. We consider two validation methods: a Monte Carlo average (MCA) method and a proportional statistical test (PST) method:

(i) **MCA:** Mark $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] \leq \alpha_B$ as valid if

$$\frac{\sum_{i \in [N]} \delta^{(i)}[s_{t+1} > b_{t+1}^0 | x_1, q_1]}{N} \leq \alpha_B.$$

(ii) **PST:** Define the null hypothesis $\mathcal{H}_0 : Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] \leq \alpha_B$ and its alternative hypothesis $\mathcal{H}_1 : Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] > \alpha_B$. We construct a test statistic

$$z_{t+1} = \frac{N\alpha_B - \sum_{i \in [N]} \delta^{(i)}[s_{t+1} > b_{t+1}^0 | x_1, q_1]}{N\sqrt{\alpha_B(1 - \alpha_B)/n}},$$

which can be approximated by a standard normal distribution ([35]; [36]). Let $\Psi^{-1}(\theta)$ be the θ quantile of the standard normal distribution $N(0, 1)$, where θ is a confidence level. Therefore, we can mark $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] \leq \alpha_B$ as invalid if $z_{t+1} > \Psi^{-1}(\theta)$. We note that the normal approximation in PST would require choos-

ing N such that $\min\{N\alpha_B, N(1 - \alpha_B)\} \geq 5$ ([36]). For instance, if $\alpha_B = 0.05$, $N = 100$ sample paths are required to perform an adequate PST.

Detailed feasibility checking procedures may refer to Algorithm 3, Appendix B.

There are a variety of techniques for accelerating the search routine in line 4 of Algorithm 1. Given an upper bound of q_1^* , i.e. $q_U(x_1, \alpha_B)$, we can replace the incremental search with more efficient search algorithms, such as bisection search and golden search. Since q_1^* is the smallest feasible q_1 , any feasible q_1 is essentially an upper bound of q_1^* . A practical upper bound can be identified utilizing information from an ABQ model. Since an ABQ model has the adjustability of idle bioreactor quantities, we can run a number of ABQ simulations to see how many bioreactors would require in the majority of decision epochs. This range would inform FBQ an upper bound candidate, and by checking its feasibility, we can easily obtain an upper bound of q_1^* . We have also identified an upper bound that appears always to be feasible in our numerical studies. The construction of this upper bound and Algorithm details are included in Appendix B.

2.4 Case Study

We now apply our capacity planning models to a supply chain described in [7] and comprised of a single CAR-T manufacturing facility, its healthcare network in charge of specimen collections and therapy transplants, and a single reagent supplier. Our CAR-T manufacturing facility model was based on the Clinical Cell and Vaccine Production Facility (CVPF) at the University of Pennsylvania. Our analysis uses the capacity planning research presented in this chapter and the AuCT-Sim simulation software described in [7]. We adopt system specifications (e.g. demand distribution, production duration, etc.) from [7], and adopt cost assessments from [14]. System specifications and cost parameters are summarized in Table 2.1.

The results in this section are based on 400 simulation scenarios. Five ABQ instances of the number of bioreactors at the facility over the 52-week planning horizon are pre-

Table 2.1: System specs and cost parameters – Chapter 2 case study

Parameter	F_d	T	\mathcal{T}	c_R	h_R	α_R	c_B	h_B	α_B
Type	distribution	scalar	scalar	scalar	scalar	scalar	scalar	scalar	scalar
Value	$Poisson(4.81)$	3 weeks	52 weeks	\$42,174	\$113.5	0.05	\$25,000	\$14.4	0.05

sented in Figure 2.1 (illustrated by the dotted lines). For the FBQ case, the lower bound using Proposition 2.3.4 is 21 bioreactors (the blue line in Figure 2.1). Use of Algorithm 1 determined that an optimal number of bioreactors equals 22 (the orange line in Figure 2.1).




Figure 2.1: Bioreactor quantities versus t

Figure 2.2 illustrates graphically how an optimal number of bioreactors for the FBQ case can be determined. The x-axis depicts the total number of fixed bioreactors over the problem horizon, which is varied from 16 to 25. The blue solid line is the discounted total expected cost. The blue dash-dotted line represents the worst-case bioreactor shortage probability $\max_{t \in [\mathcal{T}]} Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1]$. The red dash-dotted line is the feasibility threshold $\alpha_B = 0.05$. The black solid line is the discounted total expected cost for the ABQ case. We note from Figure 2 that the smallest integer on the x-axis where the blue dash-dotted line is below the red dash-dotted line is 22.

Comparing reagent inventory levels for the ABQ and FBQ cases, we note that the inven-

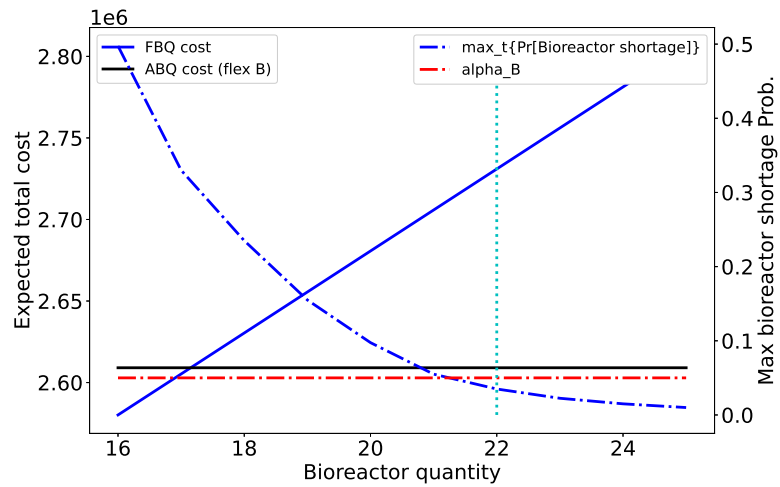


Figure 2.2: Cost and bioreactor shortage probability versus bioreactor quantities

tory replenishment policies are identical: order s_t minus r_t plus a constant that is identical for both cases. Further, the number of idle bioreactors for the ABQ case does not have an upper bound on the number of bioreactors that can be added. If, however, the number of specimens in the queue is large, exceeding the number of idle bioreactors for the FBQ case, the number of specimens waiting in queue for the FBQ case will exceed the number of specimens waiting in queue for the ABQ case. Hence, the number of specimens in queue for the FBQ case will tend to be larger than the number of specimens in queue for the ABQ case. As a result, reagent replenishment orders for the FBQ case will tend to be larger than reagent replenishment orders for the ABQ case, as depicted in Figure 2.3. This difference in part explains why for this case study the cost of ABQ is \$2,639,841.96 while the cost of FBQ is \$2,758,403.21, which is 4.3% higher than the ABQ cost.

We now examine the relationship between α_B , optimal bioreactor quantity, and total expected discounted cost. We vary α_B between 0.01 and 0.14 in increments of 0.01. As expected, as α_B increases, the optimal number of bioreactors and total expected discounted cost decrease, as shown in Figure 2.4.




Figure 2.3: Reagent inventory levels for the first 5 simulations

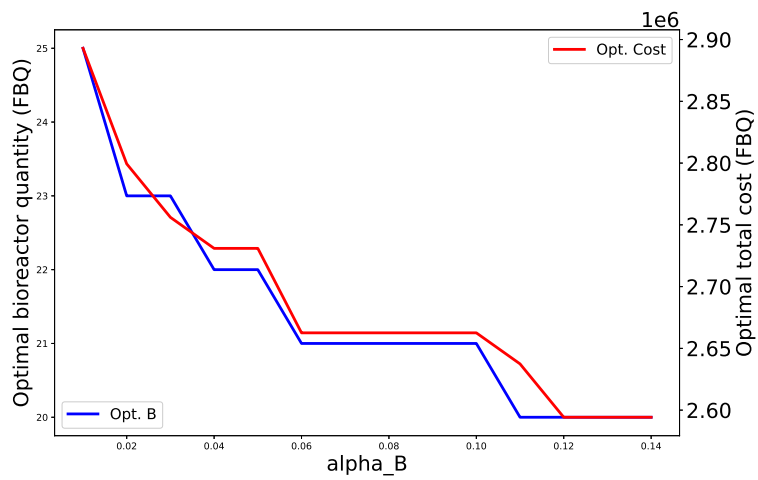


Figure 2.4: Optimal bioreactor quantity and total cost v.s. α_B

2.5 Conclusions

We have modeled and analyzed the capacity planning problem of determining the best number of bioreactors and the best reagent replenishment policy for a cell manufacturing facility. For the case where the number of idle bioreactors can be adjusted at each epoch, and for each variant of this case, the penalty cost approach and the chance constraint approach, we have shown that this number and the reagent replenishment policy are both myopic and base stock. For the case where the number of bioreactors can only be adjusted periodically, we have developed near-optimal heuristics for determining the number of bioreactors and the reagent replenishment policy. A case study based on a CAR-T manufacturing facility at the University of Pennsylvania was used to demonstrate these results.

We extend our model setting and consider supplier disruption risks in a centralized manufacturing network in the next chapter.

CHAPTER 3

RESILIENT CAPACITY PLANNING IN A CENTRALIZED NETWORK UNDER SUPPLIER DISRUPTION RISKS

In this chapter, we introduce a Capacity Planning (CP) problem in a centralized network under supplier disruption risks. We define the risk of supplier disruption as the risk of not getting the full amount of consumables ordered by the fulfillment facility to fulfill build-to-order/request-to-service demands. We demonstrate problem settings and notations in the context of autologous cell therapy. Translating all models, algorithms and results to other domains are straightforward.

We investigate approaches for determining the number of reusable resources and a replenishment policy for consumable resources for an autologous build-to-order cell therapy manufacturing system when confronted with supplier disruption risks. We model a supplier disruption process as an exogenous Markov process. Incidents that cause reagent shortages include nature disasters, human errors and disease outbreaks. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. Both models are formulated as Markov decision process (MDP). We propose algorithms to solve the proposed models. In the case studies, given different supplier disruption scenarios, we determine the number of bioreactors and a reagent replenishment policy based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania.

The solution of the proposed models introduce a concept of resilient manufacturing. We compare resilient designs with lean designs in the events of supplier disruptions to demonstrate costs and benefits of building extra resilience in a centralized network.

3.1 Introduction

Supplier disruptions in the cell manufacturing industry are of concern for two reasons. First, patients who qualify for cell therapy often have short term dire prognoses, a supplier disruption can cause a delay in the completion of a cell therapy, and hence a supplier disruption can significantly increase patient mortality risk. Second, supplier disruptions in this industry are not uncommon. The cell manufacturing industry is an emerging industry with an emerging supplier base for certain key reagents, and hence there may be only a single supplier for a key reagent. Further, the industry is heavily regulated in the U.S. by the Federal Drug Administration (FDA), and FDA requirements for reagent uniformity and quality are a challenge to satisfy for a single supplier. Satisfying these requirements for multiple suppliers would be much more of a challenge. Thus, a typical strategy for mitigating supplier disruption risk common in many other industries – sourcing from multiple suppliers for the same raw material or work in progress – is not common in the cell manufacturing industry. Further, the cell manufacturing industry has witnessed several major supplier disruptions. In 2017 alone, the cell manufacturing industry witnessed reagent supply disruptions due to Hurricane Maria, a severe flu season [37, 38], and a shutdown of a major cell therapy supplier due to sterility issues [39]. More recently, significant reagent supplier disruption has been caused by the COVID-19 pandemic [40] as demand for several key reagents [41] needed for COVID-19 antigen testing kits surged, leading to shortages of those same key reagents for autologous cell therapy manufacturing, shortages that are forecasted to persist into mid-2021 [42].

The objective of this chapter is to determine good bioreactor adjustment and reagent replenishment policies for a single cell manufacturing facility subject to supplier disruption risk. We consider two cases of bioreactor adjustment; adjustment can occur either (i) at each decision epoch or (ii) periodically. The analytic basis of our research is the infinite horizon expected total discounted cost Markov decision process (MDP). With regard

to supplier uncertainty, we assume that the maximum number of units of reagent that the supplier can deliver is a stochastic process described by an exogenous Markov chain, the realization of the current maximum value becomes available just before the current epoch, and all replenishment orders made are fulfilled just before the next epoch. We present exact and heuristic algorithms to solve the proposed models. Given different supplier disruption scenarios, we illustrate our results by determining optimal (or near-optimal) bioreactor adjustment and reagent replenishment policies based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. If reagent availability is completely certain in the future, we determine lean bioreactor adjustment and reagent replenishment policies. This special case of our model serves as a benchmark for determining the additional cost of (necessarily resilient) bioreactor adjustment and reagent replenishment policies when future supply availability is uncertain.

The usefulness of these results is two-fold. First, the results can provide useful managerial insights to the manager of a cell manufacturing facility. Second, the results can help improve the usefulness of much more detailed and realistic simulation models (see [7] for example) by significantly reducing simulation run time and computational effort.

3.1.1 Literature review

A review of how supply chain disruptions can be mitigated for a single production facility is presented in [9], where three distinct risk mitigation strategies are considered: inventory redundancy, the use of multiple suppliers, and product substitution. We focus on inventory redundancy since the autologous cell manufacturing industry invariably relies on single suppliers for key reagents for reasons given above and product substitution is not a possibility. We assume that inventory is periodically reviewed, a common practice in the industry. Continuous review models are surveyed in [9].

Periodic review inventory redundancy models fall into two categories, which we describe following several definitions. Assume A_t is the maximum number of units of reagent

that the supplier can deliver before epoch $t + 1$, is the order size at epoch t , and α_t is the order received just before epoch $t + 1$.

- Category 1: assumes that the process $\{A_t, t \geq 1\}$ is described by an exogenous Markov chain governed by the probability $Pr(A_{t+1}|A_t)$, the realization of A_t becomes available to the decision maker just before epoch t , $a_t \leq A_t$, and $\alpha_t = a_t$ for all t . Thus, the decision maker knows the maximum amount of reagent the supplier can deliver before placing a replenishment order, and assuming the order placed satisfies this constraint, the order is guaranteed to be fulfilled before the next epoch. This first category is assumed throughout this paper.
- Category 2: assumes A_t is infinite for all t and that $Pr(\alpha_t|a_t)$ is given. Thus, the decision maker can place a replenishment order of any size but is not guaranteed that the order placed will be the order received.

Several special cases of Category 1 have appeared in the literature; we present three general cases. First, we consider suppliers that are either ‘on’ with infinite capacity or ‘off’ (totally disrupted) with zero capacity, which we call ‘on-off’ suppliers. A finite-horizon inventory control problem with stationary random demand, zero lead time, backlogging demand, and geometric disruption duration (two-state Markov chain of supplier availability) is considered in [43], which assumes that the current supplier state is unknown but can be inferred from the previous supplier state. An (s, S) type policy is proven to be optimal when a fixed cost is assessed when placing an order, where s depends on the supplier state of the previous period; however, S is assumed to be independent of the previous supplier state. This model is generalized in [44] to the case where the unit ordering cost is stochastic and also extended to a perishable inventory control case with infinite horizon, deterministic demand, and backlogging. In [45] an ‘on-off’ disruption models random transportation delays with an exogenous Markovian process modeling the pipeline of in-transit replenishment. The authors consider stationary random demand and backlogging demand and prove

that a base-stock policy is optimal if there is no fixed cost and an (s, S) policy is optimal when a fixed cost exists. This result is applied to a port-of-entry supply disruption in [46]. Supplier disruption in [47] is modeled as a state dependent Bernoulli process. Given the state of supplier disruption, the availability of the supplier in the next period is a Bernoulli process whose parameter is state dependent, e.g. age of the current supplier disruption. Stationary random demand is assumed with lost-sales, and bounds are presented for myopic and look-ahead heuristics.

As a second example of Category 1, inventory control with fixed supplier capacity has been well studied. Finite and infinite horizon inventory control problems with stationary and backlogging demand, linear ordering cost, and finite positive supplier capacity are considered in [48, 49]. Results in [48, 49] show that an optimal inventory control policy is a modified base-stock policy: when the initial stock is below the base-stock level, the inventory is stocked up to the base-stock level, or as close to it as possible (given the limited capacity); otherwise, order nothing. The finite horizon model in [49] is extended in [50] to consider fixed ordering costs. It is shown that (s, S) policies may not be optimal and that there exists an optimal policy that is an $X - Y$ band type of policy: order the capacity when the initial stock is below X , and order nothing when the initial stock is above Y . If the initial stock is between X and Y , the optimal ordering quantity is problem-dependent, which requires an iterative method to solve the dynamic program. However, the $X - Y$ band structure significantly reduces the computational demands of the dynamic program. An extension of this model to an infinite horizon is presented in [51]. Key results of [51] are: (i) an $X - Y$ band type policy is optimal; (ii) the difference $Y - X \geq 0$ is no more than one unit of the supplier capacity. Similar results are also presented in [52].

As a third example of Category 1, disruption risks can also be modeled as a supplier with dynamic limited supply capacity. In [53] a model is presented that considers situations where supplier capacities at decision epochs are exogenous independent identically distributed random variables (IID). Stationary random backlogging demand is considered

in both finite and infinite period problems. It is shown that the cost-to-go functions (including the single period cost) may not be convex but are quasi-convex (convex on the left-hand side of the minimum, and nondecreasing on the right-hand side of the minimum). Thus, a base-stock policy is optimal. In the infinite horizon setting, it is shown that the structure of an optimal order policy is an extended myopic policy, requiring the consideration of review periods of uncertain length. Similar results are presented in [54]. The model of disruption risk presented in this paper extends the model considered in [53] to the case where the supplier capacities are described as an exogenous Markov chain.

With respect to Category 2, a comprehensive review of random yield is presented in [55]. Models using a yield rate are considered in [56, 57, 58, 59]. Many yield rate models assume $\alpha_t = a_t U$, where U is a random variable with given cumulative distribution function that is independent of a_t . For finite and infinite horizon problems with convex holding/penalty cost functions, a base-stock policy can be shown to be optimal; however, the base stock level is problem dependent. For stationary random and backlogging demand, the calculation of optimal base-stock levels is discussed in [58, 59]. Binomial yield is analyzed in [60] and considers stationary or cyclic demand with demand backlogging. A generic probabilistic modeling of random yield is presented in [61], where supplier uncertainty is modeled by $P(\alpha_t|a_t)$ and demand is assumed to be deterministic and can be backlogged. Conditions for the existence of an optimal policy that has a ‘staircase’ structure are presented, and this structure is shown to be computationally useful.

We are unaware if combinations of Categories 1 and 2 have been addressed in the literature, several of which would represent interesting directions for future research.

3.1.2 Organization

This chapter is outlined as follows. Section 3.2 presents the adjustable bioreactor quantity (ABQ-sdr) model. We show that an optimal bioreactor adjustment and a reagent inventory control policies are base-stock type. There exists a myopic base-stock policy to be opti-

mal for bioreactor adjustment. We discuss the computation of optimal and near-optimal reagent inventory policies. In Section 3.3, we present the fixed bioreactor quantity (FBQ-sdr) model. The convexity of total cost in the number of bioreactors guarantee the existence of an efficient search algorithm to determine an optimal bioreactor quantity. Case studies of the CAR-T cell therapy facility in the University of Pennsylvania is revisited in Section 3.4 given variant supplier disruption scenarios. Section 3.5 concludes this chapter.

3.2 Adjustable Bioreactor Quantity Model with Supplier Disruption Risks (ABQ-sdr)

3.2.1 Model formulation

We now present an infinite horizon total discounted cost Markov Decision Process (MDP) model of the reagent replenishment and bioreactor quantity selection processes. We assume that reagent replenishment and idle bioreactor adjustment decisions made at epoch t can be fulfilled before epoch $t + 1$, an assumption made for model simplicity and tractability. The level of realism of this assumption is situation dependent. Typically, periodic review in this industry occurs weekly and the time between order placement and order receipt for reagent replenishment may take longer than one week. With respect to idle bioreactor adjustment, this assumption may be realistic if the facility manufactures several cell therapies, the bioreactors needed for each cell therapy are interchangeable, and there is a substantial number of 'buffer' bioreactors in the facility. However, for a facility that manufactures only a single therapy, the total number of bioreactors at the facility would likely be a capital expenditure decision more realistically determined over a longer than one week period, e.g., quarterly, every 6 months, annually, a situation that we consider later in this chapter. In any case, the solution to this problem will be useful in addressing a key capacity planning question: *How many bioreactors are needed and what is a good reagent replenishment policy for therapy production in light of supplier disruption risks?*

We now present the adjustable bioreactor quantity model with supplier disruption risks

(ABQ-sdr). Assume at decision epoch t that

- $s_t \geq 0$ is the number of specimens in the arrival queue, waiting for therapy manufacturing to begin
- $b_t^\tau \geq 0$ is the number of bioreactors that are τ epochs from completing the manufacturing process. Let $\mathbf{b}_t = (b_t^0, \dots, b_t^{T-1})$, where T is the production time required and b_t^0 is the number of idle bioreactors
- $r_t \geq 0$ is the number of units of reagent in stock. We consider a unit of reagent to be a collection of the multiple reagents needed to complete the production of a therapy
- $0 \leq m_t = \min\{s_t, b_t^0, r_t\}$ is the number of therapies starting at time t . A patient's therapy production can only begin if the patient's specimen is in the arrival queue and there is an idle bioreactor and one unit of each type of reagent at the manufacturing facility that can be assigned to the specimen.

Further, assume

- d_t is the number of specimens that arrive between epoch t and epoch $t + 1$. We assume that the stochastic process $\{d_t, t \geq 0\}$ is a sequence of independent, identically distributed distributed random variables with known cumulative distribution function F , where the realization of d_t becomes available to the decision maker just before epoch $t + 1$. We assume $\mathbb{E}[d_t] < \infty$;
- q_t , a decision variable, is the idle bioreactor adjustment, which is determined at epoch t , and added or subtracted from the system before epoch $t + 1$;
- $a_t \geq 0$, a decision variable, is the number of units of reagent ordered at epoch t and received before epoch $t + 1$;
- $A_t \geq 0$ is the maximum units of reagent that can be supplied to the manufacturing facility at epoch t . The dynamic of the process $\{A_t, t \geq 0\}$ is described by the epoch and action invariant Markov chain $Pr(A_{t+1}|A_t)$.

Let $x_t = (s_t, \mathbf{b}_t, r_t, A_t)$ be the system state at epoch t . We assume q_t and a_t are determined based on knowledge of x_t . The resulting system dynamics are:

$$\begin{aligned}
 s_{t+1} &= s_t - m_t + d_t, \\
 b_{t+1}^\tau &= \begin{cases} b_t^0 - m_t + b_t^1 + q_t & \tau = 0, \\ b_t^{\tau+1} & 1 \leq \tau \leq T - 2, \\ m_t & \tau = T - 1, \end{cases} \\
 r_{t+1} &= r_t - m_t + a_t, \\
 Pr(A_{t+1}|A_t),
 \end{aligned}$$

from which we obtain an updated system state x_{t+1} . Since the number of idle bioreactors is non-negative, $q_t \geq -(b_t^0 - m_t + b_t^1)$, where $b_t^0 - m_t + b_t^1$ is the maximum number of bioreactors that can be removed from the system. In addition, $0 \leq a_t \leq A_t$.

Assume the single period cost occurred between epoch t and $t + 1$ are as follows:

- $c_R a_t$ is the reagent replenishment cost for the reagent, where c_R is the cost per unit of reagent
- $c_B q_t$ is the bioreactor adjustment cost, where c_B is the cost or savings per bioreactor, depending on whether q_t is positive or negative
- $h_R (r_{t+1} - s_{t+1})^+$ is the random variable holding cost for the reagent, charged only if there are more units of reagent than specimens at epoch $t + 1$, where h_R is the holding cost per unit of reagent
- $h_B (b_{t+1}^0 - s_{t+1})^+$ is the random variable bioreactor holding cost, charged only if there are more bioreactors than specimens at epoch $t + 1$, where h_B is the bioreactor holding cost per bioreactor
- $p_R (s_{t+1} - r_{t+1})^+$ is the random variable penalty for having an insufficient number of

units of reagent in order to start therapy manufacturing for all specimens that have arrived at the facility and are waiting for therapy production to begin

- $p_B(s_{t+1} - b_{t+1}^0)^+$ is the random variable penalty for having an insufficient number of idle bioreactors in order to start therapy manufacturing for all specimens that have arrived at the facility and are waiting for therapy production to begin.

We now present the optimality equation for ABQ-sdr, assuming an infinite horizon total discounted cost MDP model with discount factor $\beta \in [0, 1]$. Let

$$[Hv](x_t) = \min_{\substack{q_t \geq -(b_t^0 - m_t + b_t^1), \\ 0 \leq a_t \leq A_t}} \mathbb{E}\{C(q_t, a_t, d_t|x_t) + \beta v(x_{t+1})\} \quad (\text{ABQ-sdr})$$

where

$$\begin{aligned} C(q_t, a_t, d_t|x_t) = & c_R a_t + h_R(r_{t+1} - s_{t+1})^+ + p_R(s_{t+1} - r_{t+1})^+ \\ & + c_B q_t + h_B(b_{t+1}^0 - s_{t+1})^+ + p_B(s_{t+1} - b_{t+1}^0)^+ \end{aligned}$$

is the single period cost function. Results in [32] guarantee the existence of a unique fixed point v^* such that $\lim_{n \rightarrow \infty} \|v_n - v^*\| = 0$, where $\{v_n\}$ is such that $v_0 = 0$, $v_{n+1} = Hv_n$ for all n , and $\|\cdot\|$ is the sup-norm. Further, a policy that achieves the minimum in Hv^* is an optimal policy. Our objective is to determine an optimal or near-optimal policy.

3.2.2 A simplifying transformation

Let $z_t = q_t + b_t^0 + b_t^1$ and $y_t = a_t + r_t$, and define an operator \bar{H} as follows

$$[\bar{H}w](x_t) = \min_{\substack{z_t \geq m_t, \\ r_t \leq y_t \leq r_t + A_t}} \mathbb{E}\{G(z_t, y_t, d_t|x_t) + \beta w(x_{t+1})\}, \quad (3.1)$$

where

$$G(z_t, y_t, d_t | x_t) = (1 - \beta)c_B z_t + h_B(z_t - s_t - d_t)^+ + p_B(s_t + d_t - z_t)^+ \\ + (1 - \beta)c_R y_t + h_R(y_t - s_t - d_t)^+ + p_R(s_t + d_t - y_t)^+.$$

We show that (ABQ-sdr) and (3.1) are optimal policy invariant by constructing a shaping function [62]

$$\phi(x_t) = c_R(s_t - r_t) + c_B(s_t - b_t^0) - \sum_{n=0}^{\infty} \beta^n \{c_B m_{t+n-(T-1)} + (c_R + c_B) \mathbb{E}\{d_t\}\},$$

so that $G(z_t, y_t, d_t | x_t) = C(z_t - b_t^0 - b_t^1, y_t - r_t, d_t | x_t) + \phi(x_t) - \beta\phi(x_{t+1})$, and $w(x_t) = v(x_t) - \phi(x_t)$. It follows that $w^* = v^* + \phi$ is the unique fixed point of \bar{H} and a policy that is optimal for (3.1) is optimal for (ABQ-sdr) and conversely. Thus, in seeking an optimal policy, it is sufficient to analyze the operator \bar{H} . Let

$$g(z_t, y_t | x_t) = \mathbb{E}\{G(z_t, y_t, d_t | x_t) + \beta w(x_{t+1})\}.$$

Lemma 3.2.1. $g(z_t, y_t | x_t)$ is convex in z_t and y_t .

A proof of Lemma 3.2.1 is presented in Appendix A. The existence of a base-stock type policy is a direct result of Lemma 3.2.1.

Proposition 3.2.2. *There exists an optimal policy for ABQ-sdr that is base-stock, and hence there exist two state dependent critical values, $z^*(x_t)$ and $y^*(x_t)$, such that optimal actions in epoch t are*

$$q^*(x_t) = z^*(x_t) - b_t^0 - b_t^1, \quad a^*(x_t) = \min\{(y^*(x_t) - r_t)^+, A_t\}.$$

3.2.3 Attainability

We now present conditions that guarantee the existence of an optimal policy that is base-stock and myopic. Let

$$G(z_t, y_t, d_t | x_t) = G_z(z_t, d_t | x_t) + G_y(y_t, d_t | x_t),$$

where

$$G_z(z_t, d_t | x_t) = (1 - \beta)c_B z_t + h_B(z_t - s_t - d_t)^+ + p_B(s_t + d_t - z_t)^+,$$

$$G_y(y_t, d_t | x_t) = (1 - \beta)c_R y_t + h_R(y_t - s_t - d_t)^+ + p_R(s_t + d_t - y_t)^+.$$

Then we can rewrite the optimality equation as a multistage stochastic programming problem that minimizes

$$\mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k G_z(z_{t+k}, d_{t+k} | x_{t+k}) \right] + \mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k G_y(y_{t+k}, d_{t+k} | x_{t+k}) \right].$$

Let $\bar{z}(s_t)$ minimize $\mathbb{E}[G_z(z_t, d_t | x_t)]$ with respect to z_t . Then, $\bar{z}(s_t) = F^{-1}(\bar{\rho}_B) + s_t$, where $\bar{\rho}_B = \frac{p_B - (1-\beta)c_B}{p_B + h_B}$ and

$$F^{-1}(\rho) = \inf\{x : F(x) \geq \rho\}.$$

Since $m_t \leq s_t$, the following lemma holds trivially.

Lemma 3.2.3. $\bar{z}(s_t) \geq m_t$ for all t .

The next result follows directly from Lemma 3.2.3, which illustrates the attainability for bioreactor adjustment.

Proposition 3.2.4. *The policy $\bar{z}(s_t) = F^{-1}(\bar{\rho}_B) + s_t$ minimizes $\mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k G_z(z_{t+k}, d_{t+k} | x_{t+k}) \right]$.*

We next show that the bioreactor control and reagent control can be decoupled. Proof of Proposition 3.2.5 is presented in Appendix A.

Proposition 3.2.5. *Suppose there exists a reagent replenishment policy, i.e. \tilde{y} , that minimizes $\mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k G_y(y_{t+k}, d_{t+k} | x_{t+k}) \right]$. Then, (\bar{z}, \tilde{y}) is an optimal policy for ABQ-sdr.*

Propositions 3.2.4 and 3.2.5 indicate the optimality the myopic bioreactor adjustment policy, i.e. $z^* = \bar{z}$. Let $\bar{y}(s_t)$ minimize $\mathbb{E}[G_y(y_t, d_t | x_t)]$ with respect to y_t . Then, $\bar{y}(s_t) = F^{-1}(\bar{\rho}_R) + s_t$, where $\bar{\rho}_R = \frac{p_R - (1-\beta)c_R}{p_R + h_R}$. The following result provides sufficient conditions for $r_t \leq \bar{y}(s_t) \leq r_t + A_t$ for all $t \geq 1$; see proof in Appendix A.

Lemma 3.2.6. *Assume $r_1 \leq \bar{y}(s_1) \leq r_1 + A_1$ and $Pr(d_t \leq A_t) = 1$ for all t . Then, $r_t \leq \bar{y}(s_t) \leq r_t + A_t$ for all t .*

We define Condition 1 as follows.

Condition 1. (C1) $r_1 \leq \bar{y}(s_1) \leq r_1 + A_1$ and $Pr(d_t \leq A_t) = 1$ for all t .

(C1) is the attainability condition for reagent replenishment, and the next result follows directly from Lemma 3.2.6.

Proposition 3.2.7. *Assume (C1) holds. Then, $\bar{y}(s_t) = F^{-1}(\bar{\rho}_R) + s_t$ minimizes $\mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k G_y(y_{t+k}, d_{t+k} | x_{t+k}) \right]$.*

Propositions 3.2.4, 3.2.5 and 3.2.7 indicate the following Corollary.

Corollary 3.2.8. *Assume $r_1 \leq \bar{y}(s_1) \leq r_1 + A_1$ and $Pr(d_t \leq A_t) = 1$ for all t . Then, (\bar{z}, \bar{y}) is an optimal policy for ABQ-sdr.*

3.2.4 Bounds

We obtain a myopic policy by solving a single period problem

$$\begin{aligned} \min_{\substack{z_t \geq m_t, \\ r_t \leq y_t \leq r_t + A_t}} \mathbb{E}\{G(z_t, y_t, d_t | x_t)\}. \end{aligned} \tag{3.2}$$

A myopic reagent replenishment policy is $\bar{y}(s_t) = F^{-1}(\bar{\rho}_R) + s_t$. The myopic base-stock level is affine in s_t and is independent of A_t . The myopic policy is a idealistic policy assuming there are no or negligible supplier disruption risks.

In contrast idealistic planner, a pessimistic planner at the epoch t makes the decisions as if the supplier would be totally disrupted in the following periods. Define a pessimistic single period function

$$J(z_t, y_t | x_t) = (1 - \beta)c_B z_t + h_B(z_t - s_t - d_t)^+ + p_B(s_t + d_t - z_t)^+ + (1 - \beta)c_R y_t + \sum_{k=0}^{\infty} \beta^k \left[h_R(y_t - s_t - \sum_{j=0}^k d_{t+j})^+ + p_R(s_t + \sum_{j=0}^k d_{t+j} - y_t)^+ \right].$$

and a pessimistic policy is obtained by solving an pessimistic simple period problem

$$\begin{aligned} \min \quad & \mathbb{E}\{J(z_t, y_t | x_t)\}. \\ & z_t \geq m_t, \\ & r_t \leq y_t \leq r_t + A_t \end{aligned} \tag{3.3}$$

We remark that J is convex in z_t and y_t , and therefore, there also exists an base-stock policy that minimizes (3.3). An optimal reagent replenishment $\hat{y}(s_t)$ is obtained by solving a root finding problem

$$\hat{y}(s_t) = \inf \left\{ y_t | (1 - \beta)c_R + \sum_{k=0}^{\infty} \beta^k [-p_R + (p_R + h_R)F_{(k+1)}(y_t - s_t)] \geq 0 \right\},$$

where $F_{(k)}$ is the convoluted distribution of i consecutive periods' demand.

We now consider (C1) does not hold, and present bounds of optimal reagent base-stock level y^* . Proof of Proposition 3.2.9 can be found in Appendix A.

Proposition 3.2.9. $\bar{y}_t(s_t) \leq y^*(x_t) \leq \hat{y}_t(s_t)$.

3.2.5 Solution algorithm

We now introduce an accelerated algorithm to solve (3.1). We consider two acceleration methods.

- Action elimination: the action space of y_t is reduced to

$$[\max\{\bar{y}_t, r_t\}, \min\{\hat{y}_t, r_t + A_t\}]$$

- Reduced-space reformulation: substituting $z_t^*(x_t) = \bar{z}(s_t) = F^{-1}(\bar{\rho}_B) + s_t$ in the state dynamic equation, we have $b_{t+1}^0 = s_t - m_t + F^{-1}(\bar{\rho}_B)$. Therefore, (s_t, b_t^0, r_t) is a sufficient statistic. Define

$$v^-(x_t^-) = \min_{\max\{\bar{y}_t, r_t\} \leq y_t \leq \min\{\hat{y}_t, r_t + A_t\}} \mathbb{E}\{G^-(y_t, d_t | x_t^-) + \beta v^-(x_{t+1}^-)\}, \quad (3.4)$$

where

$$\begin{aligned} G^-(y_t, d_t | x_t^-) = & (1 - \beta)c_B(F^{-1}(\bar{\rho}_B) + s_t) + h_B(F^{-1}(\bar{\rho}_B) - d_t)^+ + p_B(d_t - F^{-1}(\bar{\rho}_B))^+ \\ & + (1 - \beta)c_R y_t + h_R(y_t - s_t - d_t)^+ + p_R(s_t + d_t - y_t)^+. \end{aligned}$$

We note that (3.4) has a reduction of $T - 1$ dimension in its state space, which would accelerate the speed of iterative methods by reducing the number of cost-to-go function evaluations. *Value/policy iteration algorithms* can be applied to obtain an optimal base-stock level for any given reduced state x_t^- , which is also optimal to all x_t 's which can be reduced to x_t^- .

3.3 Fixed Bioreactor Quantity Model with Supplier Disruption Risks (FBQ-sdr)

3.3.1 Model formulation

We now consider the case where the number of bioreactors can only be readjusted periodically, i.e. every \mathcal{T} epochs, rather than at each epoch (the $\mathcal{T} = 1$ case). We expect \mathcal{T} to be significantly larger than T . For example, assume the time between decision epochs is one week, three weeks are required to manufacture a therapy, and the number of bioreactors

can be adjusted yearly. Then, $T = 3$ and $\mathcal{T} = 52$. We refer to this problem as the fixed bioreactor quantity model with supplier disruption risks (FBQ-sdr).

Let the operator \bar{H} be defined as

$$[\bar{H}v](x_t) = \min_{0 \leq a_t \leq A_t} \mathbb{E}\{C(0, a_t, d_t|x_t) + \beta v(x_{t+1})\}.$$

Define the sequence $\{v_{\tau,n}\}$ as:

$$\begin{aligned} v_{\tau,n}(x_t) &= [\bar{H}v_{\tau+1,n}](x_t), & 2 \leq \tau \leq \mathcal{T}, \\ v_{1,n}(x_t) &= [Hv_{2,n}](x_t), \end{aligned} \tag{FBQ-sdr}$$

$v_{\mathcal{T}+1,n+1} = v_{1,n}$, and assume $v_{\mathcal{T}+1,0} = 0$. Results in [32] imply that there exists a set $\{v_\tau, \tau = 1, \dots, \mathcal{T} + 1\}$ such that this set is the unique solution to $v_\tau^* = \bar{H}v_{\tau+1}^*$ for $\tau = 2, \dots, \mathcal{T}$, $v_1^* = Hv_2^*$, and $v_1^* = v_{\mathcal{T}+1}^*$, and $\lim_{n \rightarrow \infty} \|v_{\tau,n} - v_\tau^*\| = 0$ for $\tau = 1, \dots, \mathcal{T} + 1$, where $\|\cdot\|$ is the sup-norm. Further, a policy π_τ^* that achieves the minimum in $\bar{H}v_{\tau+1}$ for $\tau = 2, \dots, \mathcal{T}$ and the minimum in Hv_2 for $\tau = 1$ is an optimal policy.

3.3.2 Solution algorithm and bounds

The optimal policy structure in Proposition 3.2.5 does not extend to the $\mathcal{T} > 1$ case. This loss of structure and the resulting computational burden of determining an optimal policy for FBQ-sdr motivate interest in examining heuristic procedures. Consider the following easily implemented procedure for determining a near-optimal policy for FBQ-sdr. Given a state x_t , we adopt the reagent replenishment policy of ABQ-sdr, i.e. $y^*(x_t)$, and replenish $\tilde{a}(x_t) = (y^*(x_t) - s_t)^+$. Given this policy, for $2 \leq \tau \leq \mathcal{T}$, let

$$\tilde{v}_{\tau,n}(x_t) = \mathbb{E}\{C(0, \tilde{a}(x_t), d_t|x_t) + \beta \tilde{v}_{\tau+1,n}(x_{t+1})\},$$

while determine the bioreactor adjustment quantity by solving

$$\tilde{v}_{1,n}(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1)} \mathbb{E}\{C(q_t, \tilde{a}(x_t), d_t | x_t) + \beta \tilde{v}_{2,n}(x_t)\}.$$

Thus, there exists a unique set $\{\tilde{v}_\tau^*, \tau = 1, \dots, \mathcal{T} + 1\}$ such that $\tilde{v}_1^* = \tilde{v}_{\mathcal{T}+1}^*$,

$$\tilde{v}_\tau^*(x_t) = \mathbb{E}\{C(0, \tilde{a}(x_t), d_t | x_t) + \beta \tilde{v}_{\tau+1}^*(x_{t+1})\},$$

for $1 \leq \tau \leq \mathcal{T} + 1$, and

$$\tilde{v}_1^*(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1)} \mathbb{E}\{C(q_t, \tilde{a}(x_t), d_t | x_t) + \beta \tilde{v}_{\mathcal{T}+1}^*(x_{t+1})\}.$$

Let the heuristic policy for adjusting the number of bioreactors achieve the above minimum.

We now present a bound result.

Proposition 3.3.1. *Upper and lower bounds of optimal value functions*

$$v^*(x_t) \leq v_\tau^*(x_t) \leq \tilde{v}_\tau^*(x_t), \quad 1 \leq \tau \leq \mathcal{T} + 1,$$

where v^* is the fixed point of H and the optimal value function for ABQ-sdr.

$v^*(x_t)$ is a lower bound since ABQ-sdr is a relaxation of FBQ-sdr by allowing bioreactor adjustment at each decision epoch. $\tilde{v}_\tau^*(x_t)$ is an upper bound since $\tilde{a}(x_t)$ is a feasible policy.

3.4 Case studies

We now apply our capacity planning models to a supply chain described in [7] and comprised of a single CAR-T manufacturing facility, its healthcare network in charge of specimen collections and therapy transplants, and a single reagent supplier. Our CAR-T manufacturing facility model is based on the Clinical Cell and Vaccine Production Facility

(CVPF) at the University of Pennsylvania. We adopt system specification (e.g. demand distribution, production duration, etc.) from [7], and adopt cost assessments from [14]. The penalty costs are carefully evaluated using system simulations (see [7]) to ensure high

Table 3.1: System specs and cost parameters – Chapter 3 case study

Parameter	F_d	T	\mathcal{T}	c_R	h_R	p_R	c_B	h_B	p_B
Type	distribution	scalar	scalar	scalar	scalar	scalar	scalar	scalar	scalar
Value	$Poisson(4.81)$	3 weeks	52 weeks	\$42,174	\$113.5	\$121,106.3	\$25,000	\$14.4	\$70,383.04

service levels. Two case studies are included in this section. The first study compares ABQ-sdr and FBQ-sdr to demonstrate the use of the models. The second study compares FBQ-sdr and a risk-free FBQ model to demonstrate cost and benefit of building extra resilience into the manufacturing network.

In both case studies, we investigate a 52-week planning horizon, and consider a simplified supplier disruption process with two states – undisrupted ($A_t = \infty$) and disrupted ($A_t = 0$) with supplier state transition matrix shown below.

$$\begin{array}{c}
 A_t = \infty \\
 A_t = 0
 \end{array}
 \left|
 \begin{array}{cc}
 A_t = \infty & A_t = 0 \\
 1 - \epsilon_d & \epsilon_d \\
 \epsilon_u & 1 - \epsilon_u
 \end{array}
 \right|$$

ϵ_d is the probability that a supplier transits from an undisrupted state to a disrupted state, and ϵ_u is the probability that a supplier transits from a disrupted state to an undisrupted state. The supply disruption process is a two-state Markov process whose parameters, i.e. ϵ_d and ϵ_u , are generally referred to as the *disruption* and *recovery probabilities* (see [9]). We adopt this supply process modeling as [9] and references therein argue that "this simple supply model is flexible enough to accommodate a spectrum of disruption profiles". Indeed, frequency of occurrence of the supplier disruption is characterized by ϵ_d ; a larger ϵ_d indicates a more frequent disruption. On the other hand, duration of the supplier disruption is correlated with ϵ_u ; a smaller ϵ_u suggests longer supply outage.

3.4.1 Case study 1: supplier disruption risk mitigation using ABQ-sdr and FBQ-sdr models

We now demonstrate the use of ABQ-sdr and FBQ-sdr models under four supplier disruption profiles:


- I. Infrequent and long disruptions: small ϵ_d and ϵ_u
- II. Infrequent and short disruptions: small ϵ_d and large ϵ_u
- III. frequent and long disruptions: large ϵ_d and small ϵ_u
- IV. frequent and short disruptions: large ϵ_d and ϵ_u .

Parameter values are summarized in Table 3.2.


Table 3.2: Disruption parameters summary

Instance	I	II	III	IV
ϵ_d	0.1	0.1	0.9	0.9
ϵ_u	0.1	0.9	0.1	0.9


Test results in this section are based on 400 simulation scenarios. Five ABQ-sdr instances of the number of bioreactors at the facility over the 52-week planning horizon are presented in Figure 3.1 (illustrated by the dotted lines). For the FBQ-sdr models, the optimal number of bioreactors are illustrated by the solid line in Figure 3.1. When bioreactor is adjustable, idle bioreactors are stocked up when long disruptions occurred, and are outlet after the disruption. Otherwise, the total bioreactor quantities stays with in a range of $[15, 25]$. When bioreactor adjustment is fixed, an excessive number of bioreactors are kept to handle the specimen queue accumulated during the supplier disruptions. Figures 3.1a and 3.1c suggest 3 times more bioreactor when disruptions are expected to be long (10 weeks in expectation) comparing to Figures 3.1b and 3.1d when supplier disruptions are short (1.1 weeks in expectation). Moreover, FBQ-sdr suggests more bioreactors when supplier disruptions are more frequent.




(a) Infrequent and long disruptions



(b) Infrequent and short disruptions



(c) Frequent and long disruptions



(d) Frequent and short disruptions

Figure 3.1: Bioreactor quantities v.s. t

Comparing reagent inventory levels for ABQ-sdr and FBQ-sdr cases, we observe that the more frequent and longer supplier disruptions are expected, the more units of reagent are ordered from the supplier. Further, since the number of idle bioreactors for the ABQ-sdr model does not have an upper bound on the number of bioreactors that can be added. If the specimen queue is long, exceeding the number of idle bioreactors for the FBQ-sdr case, the number of specimens waiting in queue for the FBQ-sdr case will tend to be larger than the number of specimens in queue for the ABQ-sdr case. As a result, reagent replenishment orders for the FBQ-sdr case will tend to be larger than reagent replenishment orders for the ABQ-sdr case, as depicted in Figure 3.2. Statistics of system state are summarized in Table 3.3.

Table 3.3: System state statistics


ϵ_d	ϵ_u	Bioreactor Quantity		Avg. Reagent Inv. Lvl.	Avg. Specimen Queue Length
		Avg.	Max.		
ABQ-sdr					
0.1	0.1	28.14	264	37.60	15.18
0.1	0.9	18.11	40	10.40	4.86
0.9	0.1	26.20	279	55.79	16.01
0.9	0.9	19.06	42	14.15	4.98
FBQ-sdr					
0.1	0.1		69	38.44	16.02
0.1	0.9		22	10.43	4.88
0.9	0.1		71	56.53	17.76
0.9	0.9		22	14.16	5.02

3.4.2 Case study 2: cost and benefit of resilience


In this case study, we first introduce a definition of cost/benefit ratio, \mathcal{R} . Then report a heatmap of \mathcal{R} to assist capacity planner in evaluating the cost and benefit of building extra resilience into the manufacturing network.

Given actions $\{q_t, a_t\}_{t \in \mathcal{T}}$, let


$$\Omega_a = \sum_{t \in [\mathcal{T}]} \beta^{t-1} [c_R a_t + h_R (r_{t+1} - s_{t+1})^+ + c_B q_t + h_B (b_{t+1}^0 - s_{t+1})^+]$$




(a) Infrequent and long disruptions



(b) Infrequent and short disruptions



(c) Frequent and long disruptions



(d) Frequent and short disruptions

Figure 3.2: On-hand reagent v.s. t

be an discounted total accounting cost, which includes purchasing and holding costs of both bioreactor and reagent. Let

$$\Omega_p = \sum_{t \in [\mathcal{T}]} \beta^{t-1} [p_R(s_{t+1} - r_{t+1})^+ + p_B(s_{t+1} - b_{t+1}^0)^+]$$


be an discounted total penalty cost. Using the penalty cost FBQ model in Chapter 2 as a baseline, a risk-ignored capacity planner determine actions $\{q_t, a_t\}_{t \in \mathcal{T}}$ with the myopic replenishment policy \bar{z} and \bar{y} . Let Ω_a^{FBQ} and Ω_p^{FBQ} be accounting cost and penalty cost of an FBQ model, and $\Omega_a^{FBQ-sdr}$ and $\Omega_p^{FBQ-sdr}$ be accounting cost and penalty cost of an FBQ-sdr model, we define a cost/benefit ratio as

$$\mathcal{R} = \frac{\Omega_a^{FBQ-sdr} - \Omega_a^{FBQ}}{\Omega_p^{FBQ} - \Omega_p^{FBQ-sdr}}$$


which measures the accounting cost increment per dollar value saved for the penalty cost.

We variate ϵ_d and ϵ_t from 0.05 to 0.95 with a step size of 0.05. Test results are calculated based on 400 simulations for each supplier disruption profile. Account cost increment and penalty reduction are illustrated in Figure 3.3. FBQ-sdr model reports larger accounting cost increment if supplier disruptions are liekly to be more frequent and longer, see Figure 3.3a. On the other hand, the penalty cost is more significant if supplier disruptions are liekly to be more frequent and longer, see Figure 3.3b.

A heatmap of \mathcal{R} is depicted in Figure 3.4. When supplier disruption is infrequent (i.e. ϵ_d close to 0) and short (i.e. ϵ_u close to 1), building up extra resilience is less beneficial. However, if disruption is more frequent (i.e. larger ϵ_d) and longer (i.e. smaller ϵ_u), building up extra resilience is preferred.



(a) Accounting cost increment



(b) Penalty cost reduction

Figure 3.3: Accounting cost increment and penalty cost reduction summary at different disruption profiles




Figure 3.4: Cost/benefit ratio of resilience

3.5 Conclusion

We have modeled and analyzed the capacity planning problem under supplier disruption risk to determine the best number of bioreactors and the best reagent replenishment policy for a cell manufacturing facility. For the case where the number of idle bioreactors can be adjusted at each epoch, we show that there exist a myopic bioreactor adjustment policy to be optimal, while the optimal policy of reagent replenishment will not be myopic. We have developed optimal and near-optimal algorithms to solve this problem. For the case where the number of bioreactors can only be adjusted once, we have developed near-optimal heuristics to determining the number of bioreactors and the reagent replenishment policy. Case studies based in part on data collected form a CAR-T cell therapy manufacturing facility at the University of Pennsylvania are studied demonstrate these results.

CHAPTER 4

RESILIENT CAPACITY PLANNING IN A DECENTRALIZED NETWORK UNDER SUPPLIER DISRUPTION RISKS

In this chapter, we investigate a capacity planning problem in a decentralized network under supplier disruption risks. Our research is motivated by the emerging need to decentralize a cell therapy manufacturing facility to increase customer access and reduce therapy fulfillment leadtime [7].

We consider two kinds of decentralized networks. The first network has isolated regional facilities (Iso-Net), while the second network has coordinated regional facilities (Co-Net). In an Iso-Net, each facility fulfills its regional demand with its own capacity; while in a Co-Net, all facilities coordinate by sharing demand and transshipping resources (bioreactor and reagent). The second supply chain is reconfigurable as production capacities in different facilities are dynamically adjustable.

A key concept we propose in this chapter is *dynamic resilience*. We show that a coordinating and reconfigurable decentralized network exhibit greater resilience with lower costs when confronted with supplier disruption risks. In this chapter, we model the two decentralized networks, Iso-Net and Co-Net, and introduce computational challenges of solving these models, which motivate our research in design of efficient and effective heuristic algorithms. In the case studies, we construct a hypothetical decentralized network based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. We compare decentralized models and different heuristic policies based on the constructed decentralized network. Testing results suggest that instead of increasing resource redundancy in all facilities, it is beneficial to only restore limited level of redundancy and adaptively reconfigure the network.

4.1 Introduction

Since the approvals of several CAR-T therapies, the public education and awareness of these advanced medicines boosted the demands of autologous cell therapies. A centralized manufacturing network has less demanding regulatory requirements, greater potential for economies of scale, greater consistency in operations, but less patient access and longer fulfillment time. A decentralized network is able to manufacture cell therapies closer to the demand locations, resulting in higher patient access, shorter fulfillment time, fewer transportation and storage of work-in-processes, and less handling and packaging. The emerging need for decentralized networks is motivated by the necessity to improve patient service quality.

Supplier disruption risks, introduced in Chapter 3, is also a major challenge for any decentralized manufacturing network. Supply chain resilience can be strengthened by, for example, increasing the inventory levels of raw material, work-in-progress, and the final product and/or adding manufacturing and/or storage capacity. However, such procedures for strengthening supply chain resilience and agility also can be expensive. Competitive advantage is enhanced if a firm's supply chain resilience and agility are at least as good as the competition's but at lower cost. Different from handling reagent supplier disruption in a centralized facility (see Chapter 3), a decentralized network can rebalance specimen and reagent inventories among facilities. Supply chain resilience can be further enhanced by also relocate manufacturing modules (e.g. bioreactors) based on real-time data-driven demand analysis. The reconfigurability of a decentralized network could blend the advantages of a distributed supply chain system of easy customer access and fast fulfillment, and a centralized supply chain system to enable economies of scale, resource risk "pooling" and cost reduction. Therefore, we propose a dynamic resilient decentralized manufacturing network for autologous cell therapy (and other build-to-order products) which allows the supply chain system to respond to and recover from supplier disruptions by sharing

demand, transshipping inventory and relocating manufacturing capacities where needed in the network. Such a reconfigurable supply chain will be dynamically resilient and either lean or agile, depending on need.

The objective of this chapter is to investigate ways to reduce the cost of supply chain resilience in a decentralized autologous cell therapy manufacturing network under supplier disruption risks. We propose analytical and numerical methods for determining (i) the bioreactor quantities and reagent replenishment policies for the reagent at each manufacturing facility and (ii) specimen and reagent transshipment and bioreactor relocation policies. The solution of the proposed models gain insights for capacity planners on

- How many bioreactors should be invested at each facility?
- How many units of reagent should each facility order?
- When and where to transship specimen or reagent and relocate bioreactors to ensure a desired level of resilience with lower cost?

4.1.1 Literature review

The problem considered in this chapter involves capacity planning, resource sharing and management under supplier disruption risks. Resource sharing operations considered in this chapter includes demand sharing (transshipping specimens), inventory transshipping (transshipping reagent) and equipment relocation (relocating bioreactors). Related literature on capacity planning is presented in Chapter 2. Therefore, we focus on reviewing multi-location resource sharing and mitigating supplier disruption risks in a decentralized network.

Resource sharing. Multi-period inventory replenishing and transshipping in a supply chain network is proposed in [63] and [64]. [65] demonstrates the potentials of transshipping strategies outperforming a centralized system using a two facility example. [66] and [67] investigate how inventory transshipping is able to increase supply chain leagility

corresponding to the number of fulfillment nodes in a supply chain. [68] discuss inventory sharing in a two-location network for make-to-stock product. Different structure of supply chain allowing inventory transshipping is compared in [69]. A two-echelon single-warehouse multiple-retail supply chain with transshipping is considered in [70] and [71]. Equipment (usable resource) relocation problem is similar to a dynamic facility location problem. Single facility (batched capacity) relocation problem is studied in [72] and [73], and multiple facility relocation problem is studied in [74, 75, 76]. A production capacity relocation problem and inventory control problem is studied in [5], and a storage capacity relocation problem is studied in [6]. Our research distinguish from existing literature: (i) We consider demand sharing as our product is build-to-order and personalized, the transshipping of patient specimen is unique; (ii) The total production capacity in our system is not fixed or predetermined. The availability process of bioreactors is a dynamic process, and the total number of production capacity is a decision variable in our problem; (iii) We consider mitigating supplier disruption risk with resource sharing and extra buffer inventory.

Mitigating supplier disruption risks in a decentralized network. Modeling supplier disruptions and inventory control problems under supplier disruption risks in a centralized network are reviewed in Chapter 3. Due to regulatory requirement multi-sourcing [77, 78] is an invalid strategy to mitigate supplier disruption risks. [79] shows that a decentralized network structure, even without transshipping inventory, reduces cost variance through the risk diversification effect. Given inventory replenishment plans, inventory transshipping in a two-retailer network with supplier disruptions is studied in [80]. [81] and [82] extend the two-location problem to n -retailer ($n > 2$) cases. [83] extends [80] with joint inventory replenishing and transshipping decisions. [84] studies inventory transshipping with demand distribution updates. Mitigating supplier disruption with demand sharing and production capacity relocation is rarely studied, which distinguish our research from existing literatures.

4.1.2 Organization

We model capacity planning problems in decentralized networks in Section 4.2. Analysis of the proposed models are presented in Section 4.3. Solution algorithms are presented in Section 4.4. In Section 4.5, we compare different decentralized models and different heuristic policies based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. Testing results suggest that instead of increasing resource redundancy in all facilities, the Co-Net model only restore limited level of redundancy and adaptively reconfigure the network with lower investment and operational costs. We conclude this chapter in Section 4.6.

4.2 Problem Formulation

Consider a network of geographically distributed manufacturing facilities, each of which uses reusable resources (bioreactors) and consumable resources (reagent) during manufacturing. There are multiple ways to manage such a network; we consider two. First, we consider the case where each facility is managed independently of the other facilities, is focused on serving its regional demand, and hence is essentially isolated from the other facilities. We refer to this case as the Iso-Net case. Second, we consider the case where the facilities coordinate and collaborate with each other, when appropriate. We refer to this case as the Co-Net case.

In an Iso-Net, the decisions we consider include:

- (i) The capital investment plan (determining the total amount of bioreactors in the network) and the assignment plan (assigning bioreactors to different facilities)
- (ii) For each facility, the consumable replenishment plan (determining the reagent replenishment policy for each facility).

We assume the capital investment and assignment plans are design decisions made initially that hold over the entire planning horizon and that the consumable replenishment decisions

are operational decisions made at each decision epoch over the planning horizon. Decisions that we consider for Co-Net include all the decisions in the Iso-Net and additionally :

- (iii) Demand sharing plans (determining policies for transshipping specimens among different regions)
- (iv) Manufacturing capacity relocation plans (determining policies for bioreactor relocation)
- (v) Material transshipping plans (determining policies for reagent transshipment),

where the additional plans, all operational, make decisions at each decision epoch.

We assume the network has L facilities, the finite problem horizon contains \mathcal{T} decision epochs, and T epochs are required for therapy manufacturing. We assume that T is considerably smaller than \mathcal{T} . For example, assume epochs occur weekly, three weeks are required to manufacture a therapy ($T = 3$), and the problem horizon contains 52 epochs ($\mathcal{T} = 52$). Let $[\cdot]$ be the set operator such that $[z] = \{1, \dots, z\}$ for any integer $z \geq 1$. At epoch t let

- $\mathbf{s}_t = \{s_t^l, l \in [L]\}$, where $s_t^l \geq 0$ is the number of specimens waiting for therapy manufacturing to begin at facility l at epoch t
- $\mathbf{b}_t = \{\mathbf{b}_t^l, l \in [L]\}$, where $\mathbf{b}_t^l = (b_t^{l,0}, \dots, b_t^{l,T-1}), b_t^{l,\tau} \geq 0$ is the number of bioreactors at facility l at epoch t that are τ epochs from completing therapy manufacturing, and where $b_t^{l,0}$ is the number of idle bioreactors at facility l at epoch t
- $\mathbf{r}_t = \{r_t^l, l \in [L]\}$, where $r_t^l \geq 0$ is the number of units of reagent at facility l at epoch t
- $\mathbf{A}_t = \{A_t^l, l \in [L]\}$, where A_t^l is the maximum number of units of reagent that can be supplied to facility l at epoch t . The dynamic of the process $\{\mathbf{A}_t, t \geq 0\}$ is described by the epoch and action invariant Markov chain $Pr(\mathbf{A}_{t+1}|\mathbf{A}_t)$.

We assume there are five sets of action to select at each epoch $t \in [\mathcal{T}]$.

- $m_t^l \in \{0\} \cup [\min\{s_t^l, b_t^{l,0}, r_t^l\}]$ is the number of therapies that are selected to begin manufacturing at epoch $t \forall l \in [L]$
- $w_t^l \in \mathbb{Z} = \{\dots, -1, 0, 1, \dots\}$ is the number of specimens transshipped into location l after m_t^l is decided and before epoch $t + 1$. A negative valued w_t^l indicates specimens are transshipped out from facility l . We assume the number of specimens in the network is identical before and after specimen transshipment; hence, $\sum_l w_t^l = 0$
- $a_t^l \in \{0\} \cup [A_t^l]$ is the number of units of reagent ordered at epoch t and received before epoch $t + 1, \forall l \in [L]$
- $e_t^l \in \mathbb{Z}$ is the number of units of reagent transshipped after a_t is received and before epoch $t + 1$. We assume the number of units of reagent in the network is identical before and after reagent transshipment; hence, $\sum_l e_t^l = 0$
- $q_t^l \in \mathbb{Z}$ is the number of bioreactors relocated after $b_t^{l,1}$ bioreactors become to idle and before epoch $t + 1$. We assume the number of bioreactors in the network is identical before and after bioreactor relocation; hence, $\sum_l q_t^l = 0$.

Note that we assume transshipped specimens and reagent and relocated bioreactors at epoch t arrive before epoch $t + 1$. For example, assume epochs occur weekly and transshipping/relocation are a few days. However, if epochs occur daily or twice a day, it require to model transshipping/relocation lead time. Let $\mathbf{d}_t = (d_t^1, \dots, d_t^L)$ be stochastic demands occurred during period t . We assume regional demands are independent, and leave correlated regional demand as a topic of future research. The system dynamics are

$$s_{t+1}^l = s_t^l - m_t^l + d_t^l + w_t^l, \quad r_{t+1}^l = r_t^l - m_t^l + a_t^l, \quad Pr(\mathbf{A}_{t+1}|\mathbf{A}_t),$$

$$b_{t+1}^{l,\tau} = \begin{cases} b_t^{l,0} - m_t^l + b_t^{l,1} + q_t^l & \tau = 0, \\ b_t^{l,\tau+1} & 1 \leq \tau \leq T - 2, \\ m_t^l & \tau = T - 1. \end{cases}$$

The state of the network then becomes $\mathbf{x}_t = (\mathbf{s}_t, \mathbf{b}_t, \mathbf{r}_t, \mathbf{A}_t)$. We use the $\{\mathbf{A}_t, t \in [\mathcal{T}]\}$ process to model a variety of ways that supply can be restricted, uncertain, or correlated across facilities. This model can consider situations that include: (i) the l th supplier can only allocate up to A_t^l units of reagent to the firm because of supplier manufacturing capacity constraints and/or supply commitment to other firms, (ii) the supplier cannot totally predict how many units of reagent it can provide at the next decision epoch due to an unpredictable production rate and/or unpredictable orders from other customers with higher priority, (iii) a quality evaluation failure at the supplier, (iv) correlations across facilities. As an example of the fourth situation, if two facilities use the same supplier, then a reduction in production rate at the supplier will reduce the availability of reagent for both facilities, a positive correlation. We remark that there is a close relationship between models of relocatable manufacturing capacity for manufacturing supply chains and models of mobile storage capacity for urban package express networks [6]. For package express networks, pickup and delivery rates and hence storage demand for residential areas and business areas are often negatively correlated.

We assume that transshipment/relocation decisions are made at the beginning of each decision epoch before the realization of patient demands. Facility l is able to transship at most $s_t^l - m_t^l$ specimens or receive at most $\sum_{i \neq l} (s_t^i - m_t^i)$ specimens in period t , so that

$$-(s_t^l - m_t^l) \leq w_t^l \leq \sum_{i \neq l} (s_t^i - m_t^i).$$

Given the fact that $\sum_l w_t^l = 0$, the right-hand side inequality is redundant and thus the action space of specimen transshipment is

$$\{\mathbf{w}_t \in \mathbb{Z}^L : w_t^l \geq -(s_t^l - m_t^l), \sum_l w_t^l = 0\}.$$

Similarly, the action space of reagent transshipment and bioreactor relocation are,

$$\{\mathbf{e}_t \in \mathbb{Z}^L : e_t \geq -(r_t^l - m_t^l + a_t^l), \sum_l e_t^l = 0\}$$

and

$$\{\mathbf{q}_t \in \mathbb{Z}^L : q_t^l \geq -(b_t^{l,0} - m_t^l + b_t^{l,1}), \sum_l q_t^l = 0\}$$

respectively.

Let the single period cost accrued in period t be additive over facilities

$$C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) = \sum_l C^l(w_t^l, a_t^l, e_t^l, q_t^l, m_t^l, d_t^l | x_t^l),$$

and each single facility cost, i.e. $C^l(w_t^l, a_t^l, e_t^l, q_t^l, m_t^l, d_t^l | x_t^l)$, consist multiple cost components

- $c_R a_t^l$ is the reagent replenishment cost for the reagent, where c_R is the cost per unit of reagent
- $h_R (r_{t+1}^l - s_{t+1}^l)^+$ is the reagent overstock holding cost, charged if there are more units of reagent than specimens at epoch $t + 1$, where h_R is the holding cost per excess unit of reagent
- $h_B (b_{t+1}^{l,0} - s_{t+1}^l)^+$ is the bioreactor overstock holding cost, charged if there are more idle bioreactors than specimens at epoch $t + 1$, where h_B is the bioreactor holding cost per excess idle bioreactor
- $p_R (s_{t+1}^l - r_{t+1}^l)^+$ is the reagent understock penalty, charged if there are less units of reagent than specimens at epoch $t + 1$, where p_R is the penalty cost per insufficient unit of reagent
- $p_B (s_{t+1}^l - b_{t+1}^{l,0})^+$ is the bioreactor understock penalty, charged if there are less idle bioreactors than specimens at epoch $t + 1$, where p_B is the penalty cost per insufficient

idle bioreactor

- $K_S \frac{|w_t^l|}{2}$ is the cost of transshipping w_t^l specimens from/to facility l that is charged on facility l
- $K_R \frac{|e_t^l|}{2}$ is the cost of transshipping e_t^l units of reagent from/to facility l that is charged on facility l
- $K_B \frac{|q_t^l|}{2}$ is the cost of relocating q_t^l bioreactors from/to facility l that is charged on facility l .

Therefore,

$$C^l(w_t^l, a_t^l, e_t^l, q_t^l, m_t^l, d_t^l | x_t^l) = c_R a_t^l + K_R \frac{|e_t^l|}{2} + h_R (r_{t+1}^l - s_{t+1}^l)^+ + p_R (s_{t+1}^l - r_{t+1}^l)^+ \\ + K_B \frac{|q_t^l|}{2} + h_B (b_{t+1}^{l,0} - s_{t+1}^l)^+ + p_B (s_{t+1}^l - b_{t+1}^{l,0})^+ + K_S \frac{|w_t^l|}{2}$$

Then, a capacity planning operation phase problem of Co-Net can be modeled as a dynamic programming problem

$$v_t^*(\mathbf{x}_t) = \min_{\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t} \mathbb{E}[C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) + \beta v_{t+1}^*(\mathbf{x}_{t+1})] \quad (\text{Co-Net-o})$$

$$s.t. \quad 0 \leq \mathbf{a}_t \leq \mathbf{A}_t, \quad (4.1a)$$

$$\sum_l w_t^l = 0, \sum_l e_t^l = 0, \sum_l q_t^l = 0, \quad (4.1b)$$

$$w_t^l \geq -(s_t^l - m_t^l) \quad \forall l, \quad (4.1c)$$

$$e_t^l \geq -(r_t^l - m_t^l + a_t^l) \quad \forall l, \quad (4.1d)$$

$$q_t^l \geq -(b_t^{l,0} - m_t^l + b_t^{l,1}) \quad \forall l, \quad (4.1e)$$

$$0 \leq m_t^l \leq \min\{s_t^l, b_t^{l,0}, r_t^l\} \quad \forall l, \quad (4.1f)$$

where β is a discount factor and boundary condition $v_{T+1}(x) = 0$. A capacity planning

design phase problem of Co-Net model is

$$\min_{b_1^{1,0}, \dots, b_1^{L,0} \geq 0} c_B \sum_l b_1^{l,0} + v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0})). \quad (\text{Co-Net-d})$$

We remark that by setting $\mathbf{w}_t = \mathbf{e}_t = \mathbf{q}_t = 0$ for all t in (Co-Net-o), the Co-Net model becomes an Iso-Net model, with a operation phase problem

$$\begin{aligned} v_t^\dagger(\mathbf{x}_t) &= \min_{\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t} \mathbb{E}[C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) + \beta v_{t+1}^\dagger(\mathbf{x}_{t+1})] & (\text{Iso-Net-o}) \\ \text{s.t.} \quad & 0 \leq \mathbf{a}_t \leq \mathbf{A}_t, \\ & \mathbf{w}_t = \mathbf{e}_t = \mathbf{q}_t = 0, \end{aligned}$$

and a design phase problem

$$\min_{b_1^{1,0}, \dots, b_1^{L,0} \geq 0} c_B \sum_l b_1^{l,0} + v_1^\dagger(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0})). \quad (\text{Iso-Net-d})$$

Existences of optimal policy for (Co-Net-o) and (Iso-Net-o) are guaranteed by results in [32]. Similar with results in Chapters 2 and 3, it is straightforward to show $m_t^{l*}(x_t^l) = \min\{s_t^l, b_t^{l,0}, r_t^l\}$ by showing (i) if there exist some $m_t^l < \min\{s_t^l, b_t^{l,0}, r_t^l\}$, we can always find a larger feasible m_t^l that gives no higher single period cost; (ii) inductively the same fact hold for multi-period cost-to-go functions. Therefore, we reduce the action space by transforming decision variables m_t^l to deterministic variables. This suggest that in a Co-Net, each facility should process at most of its own pending productions and only consider transshipping slack materials and equipment. For an Iso-Net, $m_t^{l\dagger}(x_t^l) = \min\{s_t^l, b_t^{l,0}, r_t^l\}$ is guaranteed by the results in Chapter 3. Therefore, we set $m_t^l = \min\{s_t^l, b_t^{l,0}, r_t^l\}$ throughout the remainder of this chapter.

4.3 Analysis

We first present bound results on an optimal expected total cost of a Co-Net. Let $(b_1^{1,0*}, \dots, b_1^{L,0*})$ be an optimal design of a Co-Net and an optimal expected total cost is

$$u^*(b_1^{1,0*}, \dots, b_1^{L,0*}) = c_B \sum_{l \in [L]} b_1^{l,0*} + v_1^*(\mathbf{x}_1(b_1^{1,0*}, \dots, b_1^{L,0*})).$$

Let $(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger})$ be an optimal design of an Iso-Net with the same cost parameters; and an optimal expected total cost is

$$u^\dagger(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger}) = c_B \sum_{l \in [L]} b_1^{l,0\dagger} + v_1^\dagger(\mathbf{x}_1(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger})).$$

Since a Co-Net model is a relaxation of an Iso-Net model, the Iso-Net optimal cost upper bounds the optimal cost of the Co-Net model, i.e. $u^*(b_1^{1,0*}, \dots, b_1^{L,0*}) \leq u^\dagger(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger})$.

We then present a lower bound construction on the expected total cost of a Co-Net. The construction include two parts (i) set resource sharing costs are zero, $K_S = K_R = K_B = 0$ and (ii) assume specimens, bioreactors (at any production stage) and reagent can be immediately shared before therapy production \mathbf{m}_t for all t . The constructed model is a centralized single facility model. Let $a_t = \sum_{l \in [L]} a_t^l$, $s_t = \sum_{l \in [L]} s_t^l$, $b_t^\tau = \sum_{l \in [L]} b_t^{l,\tau}$ and $r_t = \sum_{l \in [L]} r_t^l$, $m_t = \sum_{l \in [L]} m_t^l$, $d_t = \sum_{l \in [L]} d_t^l$ for all t and τ , and a revised system state $x_t = (s_t, b_t, r_t, \mathbf{A}_t)$, where $b_t = \{b_t^\tau\}$. Construct a single period cost

$$\begin{aligned} C^\sharp(a_t, d_t | x_t) = & c_R a_t + h_R (r_t + a_t - s_t - d_t)^+ + p_R (s_t + d_t - r_t - a_t)^+ \\ & + h_B (b_t^0 + b_t^1 - s_t - d_t)^+ + p_B (s_t + d_t - b_t^0 - b_t^1)^+, \end{aligned}$$

then a constructed single facility model has an operation phase problem

$$\begin{aligned}
v_t^\#(x_t) &= \min_{a_t, m_t \in \mathbb{Z}} \mathbb{E}[C^\#(a_t, d_t|x_t) + \beta v_{t+1}^\#(x_{t+1})] & \text{(L1-Apx-o)} \\
s.t. \quad & 0 \leq a_t \leq \sum_{l \in [L]} A_t^l, \\
& 0 \leq m_t \leq \min \{s_t, b_t^0, r_t\},
\end{aligned}$$

and a design phase problem

$$\min_{b_1^0 \geq 0} c_B b_1^0 + v_1^\#(x_1(b_1^0)). \quad \text{(L1-Apx-d)}$$

We name the constructed problem Single Location Approximation (L1-Apx) problem. Let $b_1^{0\#}$ be an optimal solution of (L1-Apx-d) and an optimal expected total cost is

$$u^\#(b_1^{0\#}) = c_B b_1^{0\#} + v_1^\#(x_1(b_1^{0\#})).$$

The following theorem present upper and lower bounds of Co-Net total cost.

Theorem 4.3.1. *Given the same cost parameters, then*

$$u^\#(b_1^{0\#}) \leq u^*(b_1^{1,0*}, \dots, b_1^{L,0*}) \leq u^\dagger(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger}).$$

Proof of Theorem 4.3.1 is presented in Appendix A. The lower and upper bounds can be solved by solving $L = 1$ model(s): (i) L1-Apx is a single facility model with modified transition dynamic of $\{A_t\}$ and (ii) Iso-Net is separable by facility, which can be decomposed into L single facility models. We note that an $L = 1$ model is discussed extensively in Chapter 3 as an FBQ-sdr model.

Next we present a key result which advise the algorithm design of search algorithm to determine an optimal Co-Net system design (total number of bioreactors at each location).

We show the convexity of an expected total cost given system design $b_1^{1,0}, \dots, b_1^{L,0}$.

Theorem 4.3.2. $u^*(b_1^{1,0}, \dots, b_1^{L,0}) = c_B \sum_l b_1^{l,0} + v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0}))$ is convex in design variables, i.e. $b_1^{1,0}, \dots, b_1^{L,0}$.

The above theorem implies an efficient search algorithm, e.g. coordinate decent search, to find a global minimum of (Co-Net-d) (see [85]). Proof of Theorem 4.3.2 is included in Appendix A. The following result is straightforward.

Corollary 4.3.3. $u^\dagger(b_1^{1,0}, \dots, b_1^{L,0}) = c_B \sum_l b_1^{l,0} + v_1^\dagger(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0}))$ is convex in design variables, i.e. $(b_1^{1,0}, \dots, b_1^{L,0})$.

4.4 Algorithms

Solving an Iso-Net model is equivalent to solve L FBQ-sdr models (with $n = 1$) proposed in Chapter 3. Therefore, we focus on solution algorithms for Co-Net models, i.e. (Co-Net-d) and (Co-Net-o).

4.4.1 Solving Co-Net-d

Suppose we are able to evaluate an expected total operation cost v_1^* for any system design, i.e. $b_1^{1,0}, \dots, b_1^{L,0}$. Theorem 4.3.2 suggests that a global minimum can be iteratively searched using a coordinate descent algorithm. The optimality and convergence of the proposed coordinate descent algorithm is guaranteed by Theorem 4.3.2.

In iteration k , the coordinate descent algorithm sample an improving coordinate, choose a searching step size δ_k , and evaluate

$$v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{i_k,0} - \delta_k, \dots, b_1^{L,0})) \quad \text{and} \quad v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{i_k,0} + \delta_k, \dots, b_1^{L,0})).$$

We remark that the learning step size δ_k is a non-increasing integer such that $\lim_{k \rightarrow \infty} \delta_k = 1$. Update $b_1^{i_k,0}$ to the design candidate from $\{b_1^{i_k,0} - \delta_k, b_1^{i_k,0}, b_1^{i_k,0} + \delta_k\}$ that produces the

smallest total cost u^* . Keep iterating until a stopping criterion, e.g. maximum iteration or cost improvement threshold, is satisfied. Details of the coordinate descent search is summarized in Algorithm 2.

Algorithm 2: Coordinate descent search for (Co-Net-d)	
1	Initialization: choose an initial bioreactor investment and assignment plan $(b_1^{l,0}, \dots, b_1^{L,0})$; choose an initial searching step size δ_1 ; choose maximum iteration K , set $k \leftarrow 1$;
2	Evaluate $v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0}))$, and calculate $u^*(b_1^{1,0}, \dots, b_1^{L,0})$;
3	while $k \leq K$ do
4	Sample an improving coordinate $i_k \in [L]$;
5	Evaluate $v_1^*(\mathbf{x}_t(b_1^{1,0}, \dots, b_1^{i_k,0} - \delta_k, \dots, b_1^{L,0}))$ and $v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{i_k,0} + \delta_k, \dots, b_1^{L,0}))$;
6	Calculate $u^*(b_1^{1,0}, \dots, b_1^{i_k,0} - \delta_k, \dots, b_1^{L,0})$ and $u^*(b_1^{1,0}, \dots, b_1^{i_k,0} + \delta_k, \dots, b_1^{L,0})$;
7	if $u^*(b_1^{1,0}, \dots, b_1^{i_k,0} - \delta_k, \dots, b_1^{L,0}) < u^*(b_1^{1,0}, \dots, b_1^{i_k,0}, \dots, b_1^{L,0})$ then
8	$b_t^{i_k,0} \leftarrow b_t^{i_k,0} - \delta_k$;
9	else if $u^*(b_1^{1,0}, \dots, b_1^{i_k,0} + \delta_k, \dots, b_1^{L,0}) < u^*(b_1^{1,0}, \dots, b_1^{i_k,0}, \dots, b_1^{L,0})$ then
10	$b_t^{i_k,0} \leftarrow b_t^{i_k,0} + \delta_k$
11	end
12	Set $k \leftarrow k + 1$ and update searching step size δ_k ;
13	end
14	Output: Bioreactor investment and assignment plan $(b_1^{l,0}, \dots, b_1^{L,0})$ and expected total cost $u^*(b_1^{1,0}, \dots, b_1^{L,0})$.

4.4.2 Solving Co-Net-o

Solving a Co-Net problem is more complex than the L1-Apx and Iso-Net problems due to state space complexity and the stochasticity of demand process and supplier disruption process at all locations. Given a bioreactor system design, $(b_1^{1,0}, \dots, b_1^{L,0})$, it is often a challenge to solve an exact dynamic programming problem to obtain $v_1^*(\mathbf{x}_1(b_1^{1,0}, \dots, b_1^{L,0}))$, i.e. Co-Net-o. The output of solving a Co-Net-o problem is a policy which is a functional mapping a system state to a control action. The computational challenge motivate our research in design of near-optimal heuristics. We propose three heuristics based on Approximate Dynamic Programming (ADP) [86, 87, 88, 89], which solves surrogate models with ap-

proximated cost-to-go functions. Let $\pi = \{\pi_t\}_{t \in [\tau]}$ be a policy obtained from an ADP heuristic, and $v_1^\pi(\mathbf{x}_t(b_1^{1,0}, \dots, b_1^{L,0}))$ be an expected total operation phase cost when policy π is implemented. We then obtain a solution of Co-Net-d by substituting $v_1^*(\mathbf{x}_t(b_1^{1,0}, \dots, b_1^{L,0}))$ with $v_1^\pi(\mathbf{x}_t(b_1^{1,0}, \dots, b_1^{L,0}))$ in the coordinate descent search (Algorithm 2). The proposed heuristic algorithms are named *myopic*, *extended myopic* and *mean demand lookahead*.

Myopic heuristic (MYO)

Let $\Pi(\mathbf{x}_t)$ be an action space at epoch t , i.e.

$$\begin{aligned} \Pi(\mathbf{x}_t) = \{(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t) : 0 \leq \mathbf{a}_t \leq \mathbf{A}_t, \sum_l w_t^l = 0, \sum_l e_t^l = 0, \sum_l q_t^l = 0, \\ w_t^l \geq -(s_t^l - m_t^l + d_t^l), e_t^l \geq -(r_t^l - m_t^l + a_t^l), q_t^l \geq -(b_t^{l,0} - m_t^l + b_t^{l,1}), \forall l.\} \end{aligned}$$

The MYO heuristic minimizes a single period cost at epoch t

$$\min_{(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t) \in \Pi(\mathbf{x}_t)} \mathbb{E}[C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t)] \quad (\text{MYO})$$

This heuristic is myopic as the cost-to-go function is ignored. (MYO) can be reformulated as a mix integer programming (MIP) problem.

Extended myopic heuristic (E-MYO)

Due to the fact that a myopic policy would potentially understock reagent if a supplier disruption occurs, we extend the myopic heuristic by hybridizing

- (i) Reagent inventory replenishment decision by solving an computational efficient Iso-Net model, denoted as $\hat{\mathbf{a}}_t$
- (ii) Resource sharing decision by solving a single period cost minimization problem, given reagent inventory replenishment policy $\mathbf{a}_t = \hat{\mathbf{a}}_t$.

Let $\Pi(\mathbf{x}_t|\hat{\mathbf{a}}_t)$ be a hybrid action space at epoch t , i.e.

$$\begin{aligned} \Pi(\mathbf{x}_t|\hat{\mathbf{a}}_t) = \{ & (\mathbf{w}_t, \hat{\mathbf{a}}_t, \mathbf{e}_t, \mathbf{q}_t) : \sum_l w_t^l = 0, \sum_l e_t^l = 0, \sum_l q_t^l = 0, \\ & w_t^l \geq -(s_t^l - m_t^l + d_t^l), e_t^l \geq -(r_t^l - m_t^l + \hat{a}_t^l), q_t^l \geq -(b_t^{l,0} - m_t^l + b_t^{l,1}), \forall l. \} \end{aligned}$$

The E-MYO heuristic minimizes a single period cost over the hybrid action space at epoch t

$$\min_{(\mathbf{w}_t, \hat{\mathbf{a}}_t, \mathbf{e}_t, \mathbf{q}_t) \in \Pi(\mathbf{x}_t|\hat{\mathbf{a}}_t)} \mathbb{E}[C(\mathbf{w}_t, \hat{\mathbf{a}}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t|\mathbf{x}_t)]. \quad (\text{E-MYO})$$

(E-MYO) can also be reformulated as a MIP.

Mean demand lookahead heuristic (MDL)

An MDL heuristic extends a MYO heuristic by approximating the cost-to-go function as the total cost over a lookahead horizon. We calculate the lookahead period cost based on deterministic state transitions assuming that the realized demands in each lookahead period equals to the mean demand. For example, given \mathbf{s}_t , we cast deterministic state transition $\mathbf{s}_{t+1} = \mathbf{s}_t - \mathbf{m}_t + \bar{\mathbf{d}}$, where $\bar{\mathbf{d}}$ is the mean demands per period. However, we use distributional demand to calculate the expect cost in each lookahead period.

Let $\mathcal{A}(\mathcal{L}) = \{\mathbf{A}_{t+1}, \dots, \mathbf{A}_{t+\mathcal{L}}\}$ be a realization of supplier capacities in an \mathcal{L} -period lookahead horizon, and $\bar{x}_{t+j}(\mathcal{A}(j))$ be a hypothesis state at epoch $t+j$ assuming $\bar{\mathbf{d}}$ and $\{\mathbf{A}_{t+1}, \dots, \mathbf{A}_{t+j}\}$ are realized in the first j lookahead periods. An \mathcal{L} -lookahead MDL

(MDL- \mathcal{L}) heuristic solves a multi-period problem

$$\begin{aligned}
& \min_{(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t) \in \Pi(\mathbf{x}_t),} & \mathbb{E} \left[C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) \right. \\
& (\mathbf{w}_{t+1}(\mathcal{A}(1)), \mathbf{a}_{t+1}(\mathcal{A}(1)), \mathbf{e}_{t+1}(\mathcal{A}(1)), \mathbf{q}_{t+1}(\mathcal{A}(1))) \in \Pi(\bar{\mathbf{x}}_{t+1}(\mathcal{A}(1))), \dots, \\
& (\mathbf{w}_{t+\mathcal{L}}(\mathcal{A}(\mathcal{L})), \mathbf{a}_{t+\mathcal{L}}(\mathcal{A}(\mathcal{L})), \mathbf{e}_{t+\mathcal{L}}(\mathcal{A}(\mathcal{L})), \mathbf{q}_{t+\mathcal{L}}(\mathcal{A}(\mathcal{L}))) \in \Pi(\bar{\mathbf{x}}_{t+\mathcal{L}}(\mathcal{A}(\mathcal{L}))) \\
& \left. + \sum_{i=1}^{\mathcal{L}} \beta^i C(\mathbf{w}_{t+1}, \mathbf{a}_{t+1}, \mathbf{e}_{t+1}, \mathbf{q}_{t+1}, \mathbf{m}_{t+1}, \mathbf{d}_{t+1} | \bar{\mathbf{x}}_{t+1}) \right]. \tag{MDL- \mathcal{L} }
\end{aligned}$$

(MDL- \mathcal{L}) can be reformulated as a MIP, however, if the cardinality of \mathbf{A} is huge, the number of decision variables in (MDL- \mathcal{L}) grows exponentially. Therefore, we would rely on cut generation methods (i.e. Benders decomposition) [90, 91, 92] to accelerate the computational process of proper decisions $(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t)$.

4.5 Case Studies

4.5.1 Case study 1: the potential of resource sharing

We investigate cost efficiency of a dynamic resilient network, i.e. a Co-Net, comparing to a traditional node resilient network, i.e. an Iso-Net, under the same resiliency attitude, i.e. identical p_B and p_R in the Co-Net and Iso-Net problems. We consider a two facility network ($L = 2$), and the specifications of the facilities are based on the Clinical Cell and Vaccine Production Facility (CVPF) at the University of Pennsylvania (see Table 4.1). We adopt system specification (e.g. demand distribution, production duration, etc.) from [7], and adopt cost assessments from [14]. Resource sharing costs are estimated based on consultations with researchers at the NSF Engineering Research Center for Cell Manufacturing Technology (CMaT). We consider a simply supplier disruption profile described by two independent Bernoulli processes, $Ber(p_1)$ and $Ber(p_2)$: at each decision epoch, the supplier of facility 1 has a probability of p_1 to be disrupted and be not able to supply any reagent; and the supplier of facility 1 has a probability $1 - p_1$ to be undisrupted and be capable to supply as much reagent as required; and similar for the supplier disruption process

Table 4.1: System specs and cost parameters – Chapter 4 case study

Parameter	Type	Value
F_d	distribution	$Poisson(4.81)$
T	scalar	3 weeks
\mathcal{T}	scalar	52 weeks
c_R	scalar	\$42,174
h_R	scalar	\$113.5
p_R	scalar	\$121,106.3
c_B	scalar	\$25,000
h_B	scalar	\$14.4
p_B	scalar	\$70,383.04
K_S	scalar	\$600
K_R	scalar	\$200
K_B	scalar	\$2,000

$Ber(p_2)$. Using Iso-Net as a baseline model, we compare performances of four heuristics: MYO, E-MYO, MDL-1 and MDL-2. MDL-1 is an MDL heuristic with one lookahead periods, and MDL-2 is an MDL heuristic with two lookahead periods. We consider four supplier disruption profiles, let $p = [p_1, p_2]$:

- (i) Mild: $p_1 = p_2 = 0.1$
- (ii) Moderate: $p_1 = p_2 = 0.3$
- (iii) Severe: $p_1 = p_2 = 0.6$
- (iv) Asymmetric: $p_1 = 0.3$ and $p_2 = 0.6$.

Testing results in this section are based on 400 simulation scenarios. Several performance measures reported include: average expected total cost, number of bioreactors, average specimen transshipment, average bioreactor relocation and average reagent transshipment. Figure 4.1 depicts the average expected total costs of the baseline Iso-Net and Co-Net heuristics under each supplier disruption profiles. The MYO heuristic has degraded performance comparing to the Iso-Net baseline. MYO heuristic fails to build sufficient reagent safety stock and hence results in a higher total cost. The E-MYO heuristic utilizes reagent

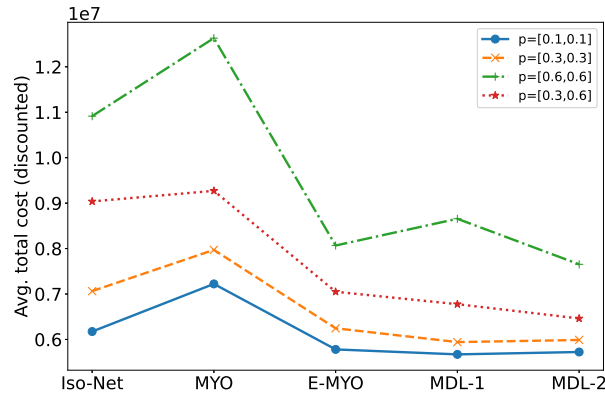



Figure 4.1: Case study 1: average expected total cost comparison


decision provided by the Iso-Net baseline to build safety stock which results in significant cost reductions comparing to MYO. On the other hand, the flexibility of resource sharing also results in significant cost reduction comparing to the Iso-Net baseline. In the cases of mild and moderate supplier disruption risk, MDL-1 produces the lowest cost; and in the cases of severe and asymmetric supplier disruption risk, MDL-2 outperforms the rest of the methods. This is because of the fact that an MDL heuristic is able to (i) build up sufficient safety stock, and (ii) utilize the flexibility of resource sharing during the lookahead periods.

Figure 4.2 compares bioreactor quantities suggested by different methods. All Co-Net heuristics suggest less total bioreactors, see Figure 4.2a. While producing much lower expected total cost, E-MYO, MDL-1 and MDL-2 reports 9.5% less bioreactor investment when supplier disruption risk is mild, and E-MYO reports 25% less bioreactor investment when supplier disruption risk is severe. Allowing bioreactor relocation results in better bioreactor utilization. In the first three cases, where supplier disruption profiles are symmetric, all methods suggest larger bioreactor quantity as the suppliers suffer worse disruption risks. When the supplier disruption profile is asymmetric, the Co-Net methods intend to keep buffer bioreactors at the location with more severe supplier risk while keep the other facility lean, see red dotted lines in Figures 4.2b and 4.2c.


Next, we present the comparison results on resource sharing. Figure 4.3 presents the



(a) Sum of two facilities



(b) Facility 1



(c) Facility 2

Figure 4.2: Case study 1: bioreactor quantity comparison

comparison of average specimen transshipment quantity when different methods are applied. We note that E-MYO suggests negligible specimen transshipment quantity, since E-MYO keeps higher redundant reagent stock at both facilities and transshipping reagent is much cheaper than transshipping patient specimen in our problem setting. This can be verified by the fact that E-MYO transships more reagent than other methods (see Figure 4.5). MYO transships the largest number of reagent among all investigated methods. Figure




Figure 4.3: Case study 1: average specimen transshipment quantity comparison

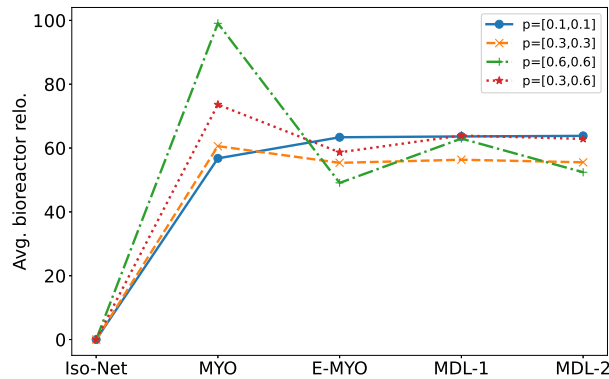


Figure 4.4: Case study 1: average bioreactor relocation quantity comparison

4.4 presents the comparison of average bioreactor transshipment quantities when different methods are applied. We observe the largest bioreactor transshipment quantities for MYO among all methods. The comparison of average reagent transshipment is presented in Fig-

ure 4.5. The E-MYO method has the largest reagent transshipment quantities in all testing scenarios.




Figure 4.5: Case study 1: average reagent transshipment quantity comparison

4.5.2 Case study 2: the supplier correlations

We investigate how a decentralized network could benefit from a dynamic resilient design (i.e. Co-Net) if the suppliers of facilities are independent, positively correlated or negatively correlated. We compare performances of four heuristics, MYO, E-MYO, MDL-1 and MDL-2, and use Iso-Net as a baseline model. We consider three supplier correlation scenarios:

- (i) Independent: $p_1 = p_2 = 0.3$, and for all t , $Pr(A_t^1 = 0|A_t^2) = 0.3$ for any A_t^2 , and $Pr(A_t^2 = 0|A_t^1) = 0.3$ for any A_t^1
- (ii) Positively correlated (identical): $p_1 = 0.3$, and for all t , $Pr(A_t^2 = 0|A_t^1 = 0) = 1$ and $Pr(A_t^2 = \infty|A_t^1 = \infty) = 1$
- (iii) Negatively correlated (flipped): $p_1 = 0.3$, and for all t , $Pr(A_t^2 = \infty|A_t^1 = 0) = 1$ and $Pr(A_t^2 = 0|A_t^1 = \infty) = 1$.

We note that the Iso-Net baseline solves two centralized network models for scenario (i) and (ii) with $p_1 = p_2 = 0.3$, and solves two centralized network models for scenario (iii)

with $p_1 = 0.3$ and $p_2 = 0.7$. Testing results in this section are based on 400 simulation scenarios. We compare performances of different methods based on average expected total cost, number of bioreactors, average specimen transshipment, average bioreactor relocation and average reagent transshipment.

Figure 4.6 depicts average expected total cost of the baseline Iso-Net and Co-Net heuristics in each supplier correlation scenario. When the suppliers are independent, the MYO





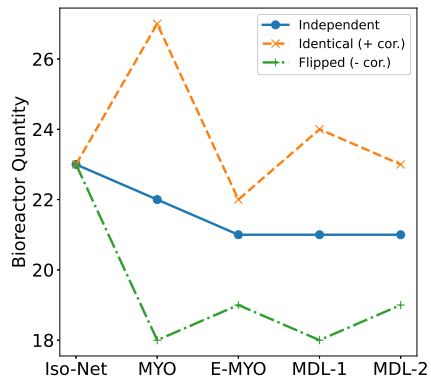
Figure 4.6: Case study 2: average expected total cost comparison

heuristic has degraded performance comparing to the baseline Iso-Net, while the other three Co-Net heuristic outperform the baseline with cost reduction as large as 14.3% (MDL-1). In the case of identical (strong positively correlated) suppliers case, MYO and MDL-1 both failed to reduce the total cost comparing to the baseline Iso-Net. E-MYO and MDL-2 have negligible cost reduction comparing to the baseline. When suppliers are strong negatively correlated, flipped supplier states in our case, resource sharing produce the largest cost reduction. All Co-Net methods outperform the baseline Iso-Net with the largest cost reduction of 38.2% (MDL-1).


Figure 4.7 compares bioreactor quantities in different supplier correlation scenarios. The largest bioreactor investment saving occurs when the suppliers are negatively correlated, i.e. the E-MYO achieves 33.3% less bioreactors comparing to the baseline Iso-Net while reducing the expected total cost at the same time. We observe less resource sharing



(a) Sum of two facilities



(b) Facility 1



(c) Facility 2

Figure 4.7: Case study 2: bioreactor quantity comparison

when the suppliers are positively correlated, and more resource sharing when suppliers are negatively correlated (see Figures 4.8-4.10 for comparison of average specimen transshipping, bioreactor relocation and reagent transshipping quantities).




Figure 4.8: Case study 2: average specimen transshipping quantity comparison




Figure 4.9: Case study 2: average bioreactor relocation quantity comparison

4.6 Conclusions

We have modeled and analyzed the capacity planning problem in a decentralized network under supplier disruption risk to determine the best number of bioreactors in each facility, the best reagent replenishment policy, and the best resource sharing plans. For the case




Figure 4.10: Case study 2: average reagent transshipment quantity comparison

where facilities are operated isolatedly (Iso-Net), we show it is equivalent to solve multiple centralized regional networks. For the case where facilities are coordinating by resources sharing (Co-Net), we analyze structural properties of Co-Net, discuss computational challenges and develop heuristic algorithm to solve the Co-Net models. In the case studies, we compare different decentralized models and different heuristic policies based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. Testing results suggest that instead of increasing resource redundancy in all facilities, the Co-Net model only restore limited level of redundancy and adaptively reconfigure the network with lower investment and operational costs.

CHAPTER 5

CONCLUSION AND FUTURE RESEARCH

In this Ph.D. thesis, we investigate a class of problems called resilient capacity planning in reconfigurable supply chains for build-to-order manufacturing system and service-upon-request service system. Several key fundamental research projects are investigated on decision supports to achieve dynamic resiliency in reconfigurable supply chains.

We introduce key concepts, such as capacity planning, reusable resource, consumable resource, relocatable manufacturing/service module, centralized/decentralized network, reconfigurability, supplier disruption and dynamic resilience. Using personalized medicine as motivating example, our research address modeling, analysis and algorithm design in both centralized and decentralized network with and without supplier disruption risks. In Chapter 2, we investigate approaches for determining the number of reusable resources and a replenishment policy for consumable resources in a centralized network, e.g. an autologous build-to-order cell therapy manufacturing system. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. For each model, we examine two variants, one where there are penalties for an insufficient number of bioreactors and units of reagents and one where there are chance constraints on an insufficient number of bioreactors and units of reagents. In a case study, we determine an optimal or near-optimal number of bioreactors and a reagent replenishment policy based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. In Chapter 3, We investigate approaches to determine the number of reusable resources and a replenishment policy for consumable resources in a centralized network, e.g. an autologous build-to-order cell therapy manufacturing system, in the face of supplier disruption risks. We model a supplier disruption process as an exogenous Markov process. Incidents that

cause reagent shortages include natural disasters, human errors and disease outbreaks. We consider two capacity planning models, one where the number of bioreactors is adjustable at each decision epoch and one where this number can only be adjusted periodically. Both models are formulated as Markov decision processes (MDP). We propose algorithms to solve the proposed models. In the case studies, given different supplier disruption scenarios, we determine bioreactors adjustment quantities and reagent replenishment policies based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. In Chapter 4, we investigate a capacity planning problem to determine the number of reusable resources, a replenishment for consumable resources and resource sharing plans in a decentralized network under supplier disruption risks. We consider two kinds of decentralized networks. The first network has isolated regional facilities (Iso-Net), while the second network has coordinated regional facilities (Co-Net). In an Iso-Net, each facility fulfills its regional demand with its own capacity; while in a Co-Net, all facilities coordinate by sharing demand and transshipping resources (bioreactor and reagent). The second supply chain is reconfigurable as production capacities in different facilities are dynamically adjustable. We show that a coordinating and reconfigurable decentralized network exhibit greater resilience with lower costs in the face of supplier disruption risks. We introduce computational challenges of solving Iso-Net and Co-Net models, which motivate our research in design of efficient and effective heuristic algorithms. In the case studies, we construct a hypothetical decentralized network based in part on data collected from a CAR-T cell therapy manufacturing facility at the University of Pennsylvania. We compare decentralized models and different heuristic policies based on the constructed decentralized network. Testing results suggest that instead of increasing resource redundancy in all facilities, it is beneficial to only restore limited level of redundancy and adaptively reconfigure the network.

We finalize this thesis by discussing open questions and proposing future research topics.

- (I) **Perishability:** Perishable consumable resource with fixed lifetimes has to be used before their expiration date or otherwise will become outdated and removed from inventory. Waste and financial loss due to outdating can be significant. For example, approximately 10% of the total inventory of certain important blood products outdated in 2017 (see [93]) and approximately 17% of unsellable products in the health and beauty, food, and beverage industries is due to outdating (see [94]). These facts support the claim that inventory control for products that outdate is an important topic for analysis, particularly when the consumable represents a significant portion of the cost of goods. An extension of works in this thesis is to consider reagent as a perishable inventory.
- (II) **Patient withdrawal:** Patients waiting for cell therapy transplant may withdraw their requests if the patient is not eligible for the treatment due to degradation of the patients' health states or patient mortalities (see [7]). Understanding and modeling a withdraw process of the specimen queue is a valuable extension. Penalty costs would be redesigned to relate patient withdrawal to some direct and indirect cost measures.
- (III) **Intentional supplier disruption:** Instead of modeling a supplier disruption process as an exogenous stochastic process, some supplier disruptions may be intentional. An adversary could disrupt a set of suppliers in the network. Then at each decision epoch, we can model a two agent game: a capacity planner chooses a reagent replenishment and a resource sharing plan, and an adversary chooses an attacking plan. Then the multi-period capacity planning problem can be extended to a dynamic game.
- (IV) **Demand variation and correlation:** Instead of stationary demand, total/regional demand may surge or reduce significantly. In addition, regional demands may be correlated. Modeling demand variation and correlation and capacity planning with nonstationary and correlated demand bring additional challenges in resilient capacity

planning.

- (V) **Real-world decentralized network case study:** A more complex real-world decentralized network data can be collected from industrial collaborators. Applying the resilient capacity methods to real-world case studies will generate more insights of the proposed models and benefit all investigated systems.

Appendices

APPENDIX A
COMPLIMENTARY PROOFS

A.1 Chapter 2 Proofs

A.1.1 Proof of Proposition 2.2.1

Rewrite the single period cost

$$C(q_t, a_t, m_t, d_t|x_t) = c_B q_t + c_R a_t + h_B(q_t + b_t^0 + b_t^1 - s_t - d_t)^+ + h_R(r_t + a_t - s_t - d_t)^+ \\ + p_B(s_t + d_t - q_t - b_t^0 - b_t^1)^+ + p_R(s_t + d_t - a_t - r_t)^+,$$

whose value does not depend on m_t . Since only the difference of s_{t+1} and b_{t+1}^0 and the difference of s_{t+1} and r_{t+1} affected the costs for all t , we can easily verify that the total cost does not depend on the choice of m_t . Thus, $m^p(x_t)$ is optimal. Let $y_t = r_t + a_t$ and $z_t = q_t + b_t^0 + b_t^1$ and

$$C'(z_t, y_t, d_t|x_t) = C(z_t - (b_t^0 + b_t^1), y_t - r_t, d_t|x_t).$$

Then ABQ-p becomes

$$v(x_t) = \min_{z_t \geq m_t, y_t \geq r_t} \mathbb{E}\{C'(z_t, y_t, d_t|x_t) + \beta v(x_{t+1})\}. \quad (\text{A.1})$$

Let

$$v^*(x_t) = \min_{z_t \geq m_t, y_t \geq r_t} \mathbb{E}\{G(z_t, y_t, d_t|x_t) + \beta v^*(x_{t+1})\}. \quad (\text{A.2})$$

where

$$G(z_t, y_t, d_t | x_t) = (h_R + (1 - \beta)c_R)(y_t - (s_t + d_t))^+ + (p_R - (1 - \beta)c_R)(s_t + d_t - y_t)^+ \\ + (h_B + (1 - \beta)c_B)(z_t - (s_t + d_t))^+ + (p_B - (1 - \beta)c_B)(s_t + d_t - z_t)^+.$$

Lemma A.1.1. *Assume $\mathbb{E}\{d_t\} < \infty$ and m_t is uniformly bounded for all t . Then, a policy optimal for (A.1) is optimal for (A.2), and conversely.*

Proof. (Lemma A.1.1) We construct a function as the following

$$\phi(x_t) = c_R(s_t - r_t) + c_B(s_t - b_t^0) - \sum_{n=0}^{\infty} \beta^n \{c_B m_{t+n-(T-1)} + (c_R + c_B) \mathbb{E}\{d_t\}\},$$

and in addition, we can examine that $G(z_t, y_t, d_t | x_t) = C'(z_t, y_t, d_t | x_t) + \beta\phi(x_{t+1}) - \phi(x_t)$.

Suppose there exists an optimal policy δ , that maps the state space to the action space at each epoch t , i.e. $\delta(x_t) = (z_t^\delta, y_t^\delta)$, then

$$v(x_t) = \mathbb{E}\{C'(z_t^\delta, y_t^\delta, d_t | x_t) + \beta v(x_{t+1})\} = \mathbb{E}\{G(z_t^\delta, y_t^\delta, d_t | x_t) - \beta\phi(x_{t+1}) + \phi(x_t) + \beta v(x_{t+1})\},$$

and hence let $\bar{v}(x_t) = v(x_t) - \phi(x_t)$, we have $\bar{v}(x_t) = \mathbb{E}\{G(z_t^\delta, y_t^\delta, d_t | x_t) + \beta\bar{v}(x_{t+1})\}$, which satisfied the optimality equation for (A.1). To show the converse direction result, simply construct $\bar{\phi}(x_t) = -\phi(x_t)$ and simply apply the same proof with the roles of (A.1) and (A.2) interchanged. \square

We remark that m_t is uniformly bounded if the number of bioreactors is finite. We note that $\mathbb{E}\{G(z_t, y_t, d_t | x_t)\}$ is convex in z_t and y_t . Let $z^*(x_t)$ and $y^*(x_t)$ be the smallest value of z_t and y_t that achieve the minimum of $\mathbb{E}\{G(z_t, y_t, d_t | x_t)\}$. It follows that $z^p(x_t) = s_t + F_d^{-1}\left(\frac{p_B - (1-\beta)c_B}{p_B + h_B}\right)$ and $y^p(x_t) = s_t + F_d^{-1}\left(\frac{p_R - (1-\beta)c_R}{p_R + h_R}\right)$. We remark that $m_t \leq s_t$ implies $m_t \leq z_t^p(x_t)$. Thus, in the single stage problem a myopic policy has a base stock policy structure: $q^p(x_t) = z^p(x_t) - (b_t^0 + b_t^1)$ and $a^p(x_t) = (y^p(x_t) - r_t)^+$. Therefore, it's

straight forward to verify that $m_t \leq z^p(x_t)$ and $r_t \leq y^p(x_t)$, and

$$m_{t+1} = \min\{s_{t+1}, b_{t+1}^0, r_{t+1}\} \leq s_{t+1} \leq s_t + F_d^{-1} \left(\frac{p_B - (1 - \beta)c_B}{p_B + h_B} \right) = z^p(x_{t+1}),$$

and

$$\begin{aligned} r_{t+1} &= y^p(x_t) - m_t = s_t + F_d^{-1} \left(\frac{p_R - (1 - \beta)c_R}{p_R + h_R} \right) - m_t \\ &\leq s_t + F_d^{-1} \left(\frac{p_R - (1 - \beta)c_R}{p_R + h_R} \right) - m_t + d_t = y^p(x_{t+1}). \end{aligned}$$

Therefore if (A1) holds, (A.2) can be decomposed into a series of single period newsvendor problem, where each newsvendor problem is solved by the myopic base stock policy. \square

A.1.2 Proof of Corollary 2.2.2

Let (q^*, a^*, m^*) and $(q^\dagger, a^\dagger, m^\dagger)$ be optimal policies of ABQ-p and (2.2), respectively. $m^* = m^\dagger = m^p$ holds trivially. Let $y_t = r_t + a_t$ and $z_t = q_t + b_t^0 + b_t^1$, then

$$\begin{aligned} \hat{C}(q_t, a_t, m_t, d_t | x_t) &= c_B q_t + c_R a_t + h_B (z_t - s_t - d_t)^+ + h_R (y_t - s_t - d_t)^+ \\ &\quad + p (s_t + d_t - \min\{z_t, y_t\})^+ \end{aligned}$$

and

$$\begin{aligned} C(q_t, a_t, m_t, d_t | x_t) &= c_B q_t + c_R a_t + h_B (z_t - s_t - d_t)^+ + h_R (y_t - s_t - d_t)^+ \\ &\quad + p_B (s_t + d_t - z_t)^+ + p_R (s_t + d_t - y_t)^+. \end{aligned}$$

If $\frac{h_R + (1 - \beta)c_R}{p_R + h_R} = \frac{h_B + (1 - \beta)c_B}{p_B + h_B}$, then by Proposition 2.2.1, we have $y_t^* = z_t^* = s_t + F_d^{-1} \left(\frac{h_R + (1 - \beta)c_R}{p_R + h_R} \right)$.

Therefore $q_t^* = z_t^* - (b_t^0 + b_t^1) = q^p$ and $a_t^* = y_t^* - r_t = a^p$. On the other hand, if $p_B + p_R = p$, it is easy to see that $\hat{C}(q_t, a_t, m_t, d_t | x_t) \geq C(q_t, a_t, m_t, d_t | x_t)$. By definition of q^* and a^* , we have $\hat{C}(q_t^*, a_t^*, m_t^*, d_t | x_t) = C(q_t^*, a_t^*, m_t^*, d_t | x_t)$, and thus it is easy to verify that $q_t^\dagger = q_t^*$

and $a_t^\dagger = a_t^*$ achieve minimized long-term total cost by a similar construction and induction in the proof in Section A.1.1. \square

A.1.3 Proof of Proposition 2.2.3

The single period costs in $\tilde{C}(q_t, a_t, d_t|x_t)$ are monotonically non-decreasing in a_t and q_t for all t . The chance constraints are equivalent to $Pr[s_t + d_t > b_t^0 + b_t^1 + q_t] \leq \alpha_B$ and $Pr[s_t + d_t > r_t + a_t] \leq \alpha_R$. $Pr[s_t + d_t > b_t^0 + b_t^1 + q_t]$ and $Pr[s_t + d_t > r_t + a_t]$ are monotonically non-increasing in q_t and a_t , respectively. Thus, the solution of ABQ-c is to minimize a_t and q_t , subject to $a_t \geq 0$, $q_t \geq -(b_t^0 - m_t + b_t^1)$, and the chance constraints. Then a_t and q_t must satisfy

$$\begin{aligned} r_t + a_t &\geq F_d^{-1}(1 - \alpha_R) + s_t, \quad a_t \geq 0, \\ b_t^0 + b_t^1 + q_t &\geq F_d^{-1}(1 - \alpha_B) + s_t, \quad q_t \geq -(b_t^0 - m_t + b_t^1). \end{aligned}$$

Therefore, $a^*(x_t) = (s_t + F_d^{-1}(1 - \alpha_R) - r_t)^+$ and $q^*(x_t) = s_t + F_d^{-1}(1 - \alpha_B) - (b_t^0 + b_t^1)$. The optimality of $m^*(x_t)$ is similar with the argument on $m^p(x_t)$ in Section A.1.1. \square

A.1.4 Proof of Proposition 2.3.2

Let $u(x_1|q_1, a) = \mathbb{E}\{\tilde{C}(q_1, a(x_1), d_1|x_t) + \beta v_2(x_2|q_1, a(x_1))\}$. Suppose there exists another feasible policy (q'_1, a') such that $u(x_1|q'_1, a') < u_0(x_0|q^*, a^*)$. By Proposition 2.3.1 and the definition of q^* ,

$$u(x_1|q_1^*, a^*) > u(x_1|q'_1, a') \geq u(x_1|q'_1, a^*) \geq u(x_1|q_1^*, a^*),$$

which is a contradiction. \square

A.1.5 Proof of Proposition 2.3.3

- (i) Let $\mathbf{d} = \{d_1, \dots, d_T\}$ be a realization of demand over the planning horizon, and $\mathbf{d}_{(t)}$ be the first t elements of \mathbf{d} , i.e. $\mathbf{d}_{(t)} = \{d_1, \dots, d_t\}$.

Lemma A.1.2. *Pr* $[s_{t+1} > b_{t+1}^0 | x_1, \mathbf{d}_{(t)}, q_1]$ is non-increasing in q_1 .

Proof. (Lemma A.1.2) Given $x_1, \mathbf{d}_{(t)}$ and a^*, s_{t+1} and b_{t+1}^0 can be explicit as:

$$s_{t+1} = s_1 - m_1 + d_1 - \dots - m_t + d_t, b_{t+1}^0$$

$$= \begin{cases} q_1 + b_1^0 - m_1 + b_1^1 - \dots - m_t + b_1^t & t \leq T-1 \\ q_1 + b_1^0 - m_1 + b_1^1 - \dots - m_{T-1} + b_1^{T-1} \\ \quad -m_T + m_1 - \dots - m_t + m_{t-T+1} & t \geq T, \end{cases}$$

and hence the inequality $s_{t+1} > b_{t+1}^0$ can be rewritten as

$$s_1 + \sum_{i=1}^t d_i > \begin{cases} q_1 + \sum_{j=0}^t b_1^j & t \leq T-1 \\ q_1 + B_1 + \sum_{j=1}^{t-T+1} m_j & t \geq T. \end{cases} \quad (\text{A.3})$$

It requires to show that the right-hand side of (A.3) is non-decreasing in q_1 .

Case 1: $t \leq T-1$, $q_1 + \sum_{j=0}^t b_1^j$ is non-decreasing in q_1 ;

Case 2: $T \leq t \leq 2T-1$,

$$\sum_{j=1}^{t-T+1} m_j = \min \left\{ s_1 + \sum_{i=1}^{t-T+1} d_i, q_1 + \sum_{i=1}^{t-T+1} b_1^i, s_1 + \sum_{i=1}^{t-T} d_i + Z_R \right\}$$

is non-decreasing in q_1 ;

Case 3: $t \geq 2T$, we show $\sum_{j=1}^{t-T+1} m_j$ is increasing in q_1 by induction. For $t = 2T$,

$$\sum_{j=1}^{T+1} m_j = \min \left\{ s_1 + \sum_{i=1}^{t-T+1} d_i, q_1 + B_1 + m_1, s_1 + \sum_{i=0}^{t-T} d_i + Z_R \right\}$$

where m_1 is non-decreasing (shown in Case 2). Therefore, $\sum_{j=1}^{T+1} m_j$ is non-decreasing in q_1 . Suppose for a $t > 2T$, $\sum_{j=1}^{t'-T+1} m_j$ is non-decreasing in q_1 for all $t' \leq t$, remains to show $\sum_{j=1}^{t+1-T+1} m_j$ is non-decreasing in q_1 .

$$\sum_{j=1}^{t-T+2} m_j = \min \left\{ s_1 + \sum_{i=1}^{t-T+2} d_i, q_1 + B_1 + \sum_{i=1}^{t+2-2T} m_i, s_1 + \sum_{i=1}^{t-T+1} d_i + Z_R \right\}$$

where $\sum_{i=1}^{t+2-2T} m_i$ is non-decreasing in q_1 by hypothesis, thus $\sum_{j=1}^{t+1-T+1} m_j$ is non-decreasing in q_1 . \square

Therefore, $\Pr[s_{t+1} > b_{t+1}^0 | x_1, q_1] = \sum_{\mathbf{d}(t)} \Pr[\mathbf{d}(t)] \Pr[s_{t+1} > b_{t+1}^0 | x_1, \mathbf{d}(t), q_1]$. \square

(ii) By definition of \tilde{C} ,

$$\begin{aligned} \tilde{C}(q_1, a^*(x_1), d_1 | x_1) &= C_B q_1 + C_R a^*(x_1) + h_B(b_2^1 - s_2)^+ + h_R(r_2 - s_2)^+ \\ &= C_B q_1 + C_R(s_1 + Z_R - r_1) + h_B(b_2^0 - s_2)^+ + h_R(Z_R - d_1); \\ \tilde{C}(0, a^*(x_t), d_t | x_t) &= C_R a^*(x_t) + h_B(b_{t+1}^0 - s_{t+1})^+ + h_R(r_{t+1} - s_{t+1})^+ \\ &= C_R(s_t + Z_R - r_t) + h_B(b_{t+1}^0 - s_{t+1})^+ + h_R(Z_R - d_t), \end{aligned}$$

where $C_B q_1$ is trivially non-decreasing in q_1 , and $C_R(s_1 + Z_R - r_1)$, $h_R(Z_R - d_1)$, $C_R(s_t + Z_R - r_t)$ and $h_R(Z_R - d_t)$ are independent with q_1 . Thus it remains to show that $h_B(b_{t+1}^0 - s_{t+1})^+$ is non-decreasing in q_1 . With a similar derivation as in the proof of Lemma A.1.2,

$$b_{t+1}^0 - s_{t+1} = \begin{cases} q_1 + \sum_{j=1}^t b_1^j - (s_0 + \sum_{i=1}^t d_i) & t \leq T-1 \\ q_1 + B_1 + \sum_{j=1}^{t-T+1} m_j - (s_1 + \sum_{i=1}^t d_i) & t \geq T, \end{cases}$$

we can show $b_{t+1}^0 - s_{t+1}$ is non-decreasing in q_1 , which completes our proof. \square

A.1.6 Proof of Proposition 2.3.4

We first show that the lower bound $q_L(x_1, \alpha_B)$ ensures $Pr[s_{t+1} > b_{t+1}^0 | x_1, q_L(x_1, \alpha_B)] \leq \alpha_B$ for $t = 1, \dots, T$, by showing that for $k = 1, \dots, T$, $q_0^{(k)}$ is the smallest bioreactor adjustment such that $Pr[s_{k+1} > b_{k+1}^0 | x_1, q_1^{(k)}] \leq \alpha_B$ is ensured.

Case 1: For $k = 1, \dots, T - 1$, $s_{k+1} > b_{k+1}^0$ with a given x_1 implies $s_1 + \sum_{\tau=1}^k d_\tau > q_1 + \sum_{\tau=0}^k b_1^\tau$. To guarantee $Pr[s_{k+1} > b_{k+1}^0 | x_1, q_1^{(k)}] \leq \alpha_B$, we have $q_1 \geq s_1 + F_{(k)}^{-1}(1 - \alpha_B) - \sum_{\tau=1}^k b_1^\tau = q_1^{(k)}$;

Case 2: For $k = T$, $s_{T+1} > b_{T+1}^0$ with a given x_1 implies $s_1 + \sum_{\tau=1}^T d_\tau - m_1 > q_1 + \sum_{\tau=0}^{T-1} b_1^\tau$. To guarantee $Pr[s_{T+1} > b_{T+1}^0 | x_1, q_1^{(T)}] \leq \alpha_B$, we have $q_1 \geq s_1 + F_{(T)}^{-1}(1 - \alpha_B) - m_1 - B_1 = q_1^{(T)}$. To ensure feasibility of the first T bioreactor outage chance constraint, q_1 need to be no less than the maximum of $q_1^{(k)}(x_1, \alpha_B)$ for $k = 1, \dots, T$. Since $T \leq \mathcal{T}$, we claim $q_L(x_1, \alpha_B)$ is a lower bound on an optimal q_1 . \square

A.2 Chapter 3 Proofs

A.2.1 Proof of Lemma 3.2.1

Proof. Let $w_0 = 0$ and $w_{n+1} = \bar{H}w_n$ for all n , and $g_n(z_t, y_t | x_t) = \mathbb{E}\{G(z_t, y_t, d_t | x_t) + \beta w_n(x_{t+1})\}$. For $n = 0$, $g_0(z_t, y_t | x_t) = \mathbb{E}\{G(z_t, y_t, d_t | x_t)\}$ is convex in z_t and y_t . Suppose g_n is convex in z_t and y_t , then

$$w_{n+1}(x_t) = \begin{cases} g_n(z^*(x_t), r_t | x_t), & r_t \geq y^*(x_t) \\ g_n(z^*(x_t), y^*(x_t) | x_t), & y^*(x_t) - A_t \leq r_t \leq y^*(x_t) \\ g_n(z^*(x_t), r_t + A_t | x_t), & r_t \leq y^*(x_t) - A_t. \end{cases}$$

which is convex in r_t and irrelevant with \mathbf{b}_t . Since r_{t+1} can be expressed as a linear expression of y_t , w_{n+1} is convex in y_t . Therefore g_{n+1} is convex in z_t and y_t . Inductively, we show the convexity of g . \square

A.2.2 Proof of Proposition 3.2.5

Proof. It is sufficient to show that the objective functions of the two parts are not dependent with production quantity m . The reagent inventory level $y_{t+k} - s_{t+k} - d_{t+k}$ can be rewritten as

$$r_t + a_t + \cdots + a_{t+k} - s_t - d_t - \cdots - d_{t+k}$$

which is independent of m_t, \dots, m_{t+k} . A similar result can be shown for idle bioreactor inventory level. \square

A.2.3 Proof of Lemma 3.2.6

Proof. It is sufficient to show that $r_t \leq \bar{y}(s_t) \leq r_t + A_t$ implies $r_{t+1} \leq \bar{y}(s_{t+1}) \leq r_{t+1} + A_{t+1}$. Note $r_t \leq \bar{y}(s_t) \leq r_t + A_t$ is equivalent to

$$0 \leq F^{-1}(\bar{\rho}_R) + s_t - r_t \leq A_t$$

and $r_{t+1} = F^{-1}(\bar{\rho}_R) + s_t - m_t$. Therefore, substituting r_{t+1} , $r_{t+1} \leq \bar{y}(s_{t+1}) \leq r_{t+1} + A_{t+1}$ is equivalent to $0 \leq d_t \leq A_{t+1}$, which follows by the assumptions $Pr(d_t \leq A_t) = 1$ for all t and $\{d_t, t \geq 0\}$ is IID. \square

A.2.4 Proof of Proposition 3.2.9

Proof. (i) **Left \leq :** Define $g_n(z_t, y_t | x_t) = \mathbb{E}\{G(z_t, y_t, d_t | x_t) + \beta w_n(x_{t+1})\}$, and let $y_n^*(x_t)$ be a minimizer of g_n . For $n = 0$, $\bar{y}(s_t) = y_0^*(x_t)$. Suppose $\bar{y}(s_t) \leq y_n^*(x_t)$, it is sufficient to show $\bar{y}(s_t) \leq y_{n+1}^*(x_t)$. By definition of \bar{y}_t , $\left. \frac{\partial \mathbb{E}\{G(z_t, y_t, d_t | x_t)\}}{\partial y_t} \right|_{\bar{y}_t(s_t)} = 0$, and it is sufficient to show that $\left. \frac{\partial \mathbb{E}\{w_n(x_{t+1})\}}{\partial y_t} \right|_{\bar{y}_t(s_t)} \leq 0$. Let \bar{r}_{t+1} (\bar{s}_{t+1}) be the inventory level (specimen queue length) if \bar{y} is implemented at epoch t , then

$$\bar{r}_{t+1} = \bar{y}(s_t) - m_t \leq s_{t+1} + F^{-1}(\bar{\rho}_R) = \bar{y}(\bar{s}_{t+1}) \leq y_n^*(x_{t+1}),$$

hence $\left. \frac{\partial \mathbb{E}\{w_n(x_{t+1})\}}{\partial r_{t+1}} \right|_{\bar{r}_{t+1}} \leq 0$. Since \bar{r}_{t+1} is affine in \bar{y} , $\left. \frac{\partial \mathbb{E}\{w_n(x_{t+1})\}}{\partial y_t} \right|_{\bar{y}_t(s_t)} \leq 0$.

(ii) **Right \leq :** By contradiction, suppose $y^*(x_t) > \hat{y}(s_t)$ for any supplier disruption process, i.e. $\{A_\tau\}_{\tau \geq t+1}$. We can easily construct a counter example, e.g. a disruption process $A_{t+1} = \infty$ and $A_\tau = 0$ for $\tau > t + 1$. In this case, $y^*(x_t) = \bar{y}(s_t)$. Since $F_j(\bar{y}(s_t) - s_t) \leq \frac{p_R - (1-\beta)c_R}{p_R + h_R} \leq \frac{p_R}{p_R + h_R}$, then we have

$$(1 - \beta)c_R + \sum_{k=0}^{\infty} \beta^k [-p_R + (p_R + h_R)F_{(k+1)}(\bar{y}(s_t) - s_t)] \leq 0,$$

so that $\bar{y}(s_t) \leq \hat{y}(s_t)$. Thus we derive a contradiction,

$$\bar{y}(s_t) \leq \hat{y}(s_t) < y^*(x_t) = \bar{y}(s_t).$$

□

A.3 Chapter 4 Proofs

A.3.1 Proof of Theorem 4.3.1

Proof. The upper bound is obvious: since a Co-Net model is a relaxation of an Iso-Net model, the Iso-Net optimal cost is an upper bound of the Co-Net cost, i.e. $u^*(b_1^{1,0*}, \dots, b_1^{L,0*}) \leq u^\dagger(b_1^{1,0\dagger}, \dots, b_1^{L,0\dagger})$. We then focus on the lower bound.

Consider the following relaxation of a C-Net model:

(i) Let $K_S = K_R = K_B = 0$, so that

$$\tilde{C}(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) = \sum_l \tilde{C}^l(w_t^l, a_t^l, e_t^l, q_t^l, m_t^l, d_t^l | x_t^l),$$

where

$$\begin{aligned}\tilde{C}^l(w_t^l, a_t^l, e_t^l, q_t^l, m_t^l, d_t^l | x_t^l) &= c_R a_t^l + h_R (r_{t+1}^l - s_{t+1}^l)^+ + p_R (s_{t+1}^l - r_{t+1}^l)^+ \\ &\quad + h_B (b_{t+1}^{l,0} - s_{t+1}^l)^+ + p_B (s_{t+1}^l - b_{t+1}^{l,0})^+.\end{aligned}$$

Let $a_t = \sum_l a_t^l$, $s_t = \sum_l s_t^l$, $b_t^\tau = \sum_l b_t^{l,\tau}$ and $r_t = \sum_l r_t^l$, $m_t = \sum_l m_t^l$, $d_t = \sum_l d_t^l$

for all t and τ ,

$$\begin{aligned}
& \tilde{C}(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) \\
&= \sum_l [c_R a_t^l + h_R(r_{t+1}^l - s_{t+1}^l)^+ + p_R(s_{t+1}^l - r_{t+1}^l)^+ + h_B(b_{t+1}^{l,0} - s_{t+1}^l)^+ \\
&\quad + p_B(s_{t+1}^l - b_{t+1}^{l,0})^+] \\
&= \sum_l c_R a_t^l + \sum_l h_R(r_{t+1}^l - s_{t+1}^l)^+ + \sum_l (h_R + p_R)(s_{t+1}^l - r_{t+1}^l)^+ \\
&\quad + \sum_l h_B(b_{t+1}^{l,0} - s_{t+1}^l)^+ + \sum_l (h_B + p_B)(s_{t+1}^l - b_{t+1}^{l,0})^+ \\
&\geq \sum_l c_R a_t^l + \sum_l h_R(r_{t+1}^l - s_{t+1}^l)^+ + (h_R + p_R)(\sum_l (s_{t+1}^l - r_{t+1}^l)^+) \\
&\quad + \sum_l h_B(b_{t+1}^{l,0} - s_{t+1}^l)^+ + (h_B + p_B)(\sum_l (s_{t+1}^l - b_{t+1}^{l,0})^+) \\
&= c_R a_t + h_R(r_{t+1} - s_{t+1}) + (h_R + p_R)(s_{t+1} - r_{t+1})^+ + h_B(b_{t+1}^0 - s_{t+1}) \\
&\quad + (h_B + p_B)(s_{t+1} - b_{t+1}^0)^+ \\
&= c_R a_t + h_R(r_{t+1} - s_{t+1})^+ + p_R(s_{t+1} - r_{t+1})^+ + h_B(b_{t+1}^0 - s_{t+1})^+ \\
&\quad + p_B(s_{t+1} - b_{t+1}^0)^+ \\
&= c_R a_t + h_R(r_t - m_t + a_t + \sum e_t^l - s_t + m_t - d_t - \sum w_t^l)^+ \\
&\quad + p_R(s_t - m_t + d_t + \sum w_t^l - r_t + m_t - a_t - \sum e_t^l)^+ \\
&\quad + h_B(b_t^0 - m_t + b_t^1 + \sum q_t^l - s_t + m_t - d_t - \sum w_t^l)^+ \\
&\quad + p_B(s_t - m_t + d_t + \sum w_t^l - b_t^0 + m_t - b_t^1 - \sum q_t^l)^+ \\
&= c_R a_t + h_R(r_t + a_t - s_t - d_t)^+ + p_R(s_t + d_t - r_t - a_t)^+ \\
&\quad + h_B(b_t^0 + b_t^1 - s_t - d_t)^+ + p_B(s_t + d_t - b_t^0 - b_t^1)^+ = \ddot{C}(a_t, d_t | x_t)
\end{aligned}$$

(ii) Relax (4.1a) by $0 \leq \sum_{l \in [L]} a_t^l \leq \sum_{l \in [L]} A_t^l$

(iii) Relax (4.1b)-(4.1e)

(iv) Relax (4.1f) by $0 \leq \sum_{l \in [L]} m_t^l \leq \min\{s_t, b_t^0, r_t\}$.

Matching the term, we can verify that the relaxed model is essentially (L1-Apx-o) model, which will provide a lower bound expected total cost. \square

A.3.2 Proof of Theorem 4.3.2

Proof. $c_B \sum_{l \in [L]} b_1^{l,0}$ is linear in $b_1^{l,0}$'s, therefore, it remains to show that v_1^* is convex in $b_1^{l,0}$'s. Let

$$g_t(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t | \mathbf{x}_t) = \mathbb{E}[C(\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t, \mathbf{m}_t, \mathbf{d}_t | \mathbf{x}_t) + \beta v_{t+1}^*(\mathbf{x}_{t+1})],$$

if we can show that g_t is convex in $\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t$ and $b_t^{l,0}$'s, then we can show v_t^* is convex in $b_t^{l,0}$'s (so that v_1^* is convex in $b_1^{l,0}$'s) by the following lemma in [95].

Lemma A.3.1. *If $f(x, y) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ is convex in x and y , and $Y \in \mathbb{R}^m$ is a convex set, then $g(x) = \min_{y \in Y} f(x, y)$ is convex in x .*

Therefore, it remains to show that g_t is convex in $\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t$ and $b_t^{l,0}$'s. Prove by induction:

- $g_{\mathcal{T}}(\mathbf{w}_{\mathcal{T}}, \mathbf{a}_{\mathcal{T}}, \mathbf{e}_{\mathcal{T}}, \mathbf{q}_{\mathcal{T}}, \mathbf{m}_{\mathcal{T}} | \mathbf{x}_{\mathcal{T}}) = \mathbb{E}[C(\mathbf{w}_{\mathcal{T}}, \mathbf{a}_{\mathcal{T}}, \mathbf{e}_{\mathcal{T}}, \mathbf{q}_{\mathcal{T}}, \mathbf{m}_{\mathcal{T}}, \mathbf{d}_{\mathcal{T}} | \mathbf{x}_{\mathcal{T}})]$, which is obviously convex in $\mathbf{w}_{\mathcal{T}}, \mathbf{a}_{\mathcal{T}}, \mathbf{e}_{\mathcal{T}}, \mathbf{q}_{\mathcal{T}}, \mathbf{s}_{\mathcal{T}}, \mathbf{r}_{\mathcal{T}}$ and $b_{\mathcal{T}}^{l,0}$'s.
- Suppose g_{t+1} is convex in $\mathbf{w}_{t+1}, \mathbf{a}_{t+1}, \mathbf{e}_{t+1}, \mathbf{q}_{t+1}, \mathbf{s}_{t+1}, \mathbf{r}_{t+1}$ and $b_{t+1}^{l,0}$'s, then by Lemma A.3.1, v_{t+1}^* is convex in $\mathbf{s}_{t+1}, \mathbf{r}_{t+1}$ and $b_{t+1}^{l,0}$'s. We introduce another lemma from [95].

Lemma A.3.2. *If $f : \mathbb{R}^m$ is convex, then $g(x) = f(Ax + b)$ is convex in $x \in \mathbb{R}^n$ for any $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.*

Since $\mathbf{s}_{t+1}, \mathbf{r}_{t+1}$ and $b_{t+1}^{l,0}$'s are linear functions of $\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t$ and $b_t^{l,0}$'s, then by Lemma A.3.2, g_t is convex in $\mathbf{w}_t, \mathbf{a}_t, \mathbf{e}_t, \mathbf{q}_t$ and $b_t^{l,0}$'s. \square

APPENDIX B

COMPLIMENTARY ALGORITHMS

B.1 Chapter 2 Algorithm Complements

B.1.1 PST Feasibility Checking Algorithm

Algorithm 3: PST feasibility checking	
1	Initialization: set $N \leftarrow \inf\{N : \min\{N\alpha_B, N(1 - \alpha_B)\} \geq 5\}$, choose a confidence level θ , and calculate testing threshold $\Psi^{-1}(\theta)$;
2	Generate N \mathcal{T} -period simulation paths with bioreactor adjustment q_1 and base-stock reagent replenishment policy a^* . Collect incidents $\delta^{(i)}[s_{t+1} > b_{t+1}^0 x_1, q_1]$, for $i \in [N]$ and $\in [\mathcal{T}]$;
3	for $t \in [\mathcal{T}]$ do
4	Calculate test statistics $z_t = \frac{N\alpha_B - \sum_{i=1}^N \delta^{(i)}[s_{t+1} > b_{t+1}^0 x_0, q_0, a^*]}{N\sqrt{\alpha_B(1-\alpha_B)/n}}$;
5	if $z_t > \Psi^{-1}(\theta)$ then
6	Break for loop;
7	end
8	Set $k \leftarrow t$;
9	end
10	Output: 'Feasible' if $k = \mathcal{T}$, and 'Infeasible' if $k < \mathcal{T}$.

B.1.2 A Pracial Upper bound of q_1^*

Consider a finite horizon (ABQ-c) model:

$$\bar{v}_t(x_t) = \min_{q_t \geq -(b_t^0 - m_t + b_t^1), a_t \geq 0} \mathbb{E}\{\tilde{C}(q_t, a_t, d_t | x_t) + \beta \bar{v}_{t+1}(x_{t+1})\} \quad (\text{B.1})$$

$$s.t. \quad Pr[s_{t+1} > b_{t+1}^0 | x_t] \leq \alpha_B,$$

$$Pr[s_{t+1} > r_{t+1} | x_t] \leq \alpha_R.$$

and $\bar{v}_{\mathcal{T}+1}(x_{\mathcal{T}+1}) = 0$. For a given starting state x_1 and threshold α_B , for any demand realization scenario \mathbf{d} , the idle bioreactor adjustment and reagent replenishment policy yields a

dynamic total bioreactor quantity dynamics $B_t(x_1, \mathbf{d}_{(t)}, \alpha_B)$, where $\mathbf{d}_{(t)} = \{d_1, \dots, d_{t-1}\}$ is a truncated demand scenarios that is known at epoch t . Let $B(x_1, \mathbf{d}, \alpha_B) = \max_t B_t(x_1, \mathbf{d}_{(t)}, \alpha_B)$, which is a random variable. Let $F_{B(x_1, \alpha_B)}$ be the cdf of $B(x_1, \mathbf{d}, \alpha_B)$, define $B_U(x_1, \alpha_B) = F_{B(x_1, \alpha_B)}^{-1}(1 - \alpha_B)$. Then we obtain an practical upper bound $q_U(x_1, \alpha_B) := B_U(x_1, \alpha_B) - B_1$. Through our numerical tests on a large amount of input scenarios, $q_U(x_1, \alpha_B)$ is always a feasible bioreactor adjustment, while proofing this upper bound analytically could be one of our future research topics.

REFERENCES

- [1] A. Gunasekaran and E. W. Ngai, “Modeling and analysis of build-to-order supply chains,” *European Journal of Operational Research*, vol. 195, no. 2, pp. 319–334, 2009.
- [2] M. Holweg and F. K. Pil, “Successful build-to-order strategies start with the customer,” *MIT Sloan Management Review*, vol. 4, no. 1, pp. 74–83, 2001.
- [3] B. MacCarthy, P. G. Brabazon, and J. Bramham, “Fundamental modes of operation for mass customization,” *International Journal of Production Economics*, vol. 85, no. 3, pp. 289–304, 2003.
- [4] P. H. Zipkin, “Foundations of inventory management,” in. Boston: McGraw-Hill, 2000.
- [5] S. S. Malladi, A. L. Erera, and C. C. W. III, “A dynamic mobile production capacity and inventory control problem,” *IIE Transactions*, vol. 52, no. 8, pp. 926–943, 2020.
- [6] L. Faugere, B. Montreuil, and C. C. W. III, “Mobile access hub deployment for urban parcel logistics,” *Sustainability*, vol. 12, no. 17, p. 7213, 2020.
- [7] K. Wang, Y. Liu, J. Li, B. Wang, R. Bishop, C. White, A. Das, A. D. Levine, L. Ho, B. L. Levine, and A. D. Fesnak, “A multiscale simulation framework for the manufacturing facility and supply chain of autologous cell therapies,” *Cytotherapy*, vol. 21, no. 10, pp. 1081–1093, 2019.
- [8] A. Dolgui, D. Ivanov, and B. Sokolov, “Reconfigurable supply chain: The x-network,” *International Journal of Production Research*, vol. 58, no. 13, pp. 4138–4163, 2020. eprint: <https://doi.org/10.1080/00207543.2020.1774679>.
- [9] L. V. Snyder, Z. Atan, P. Peng, Y. Rong, A. J. Schmitt, and B. Sinoysal, “Or/rs models for supply chain disruptions: A review,” *IIE Transactions*, vol. 48, no. 2, pp. 89–109, 2015.
- [10] B. L. Levine, K. Wang, Y. Liu, J. Li, C. White, and B. Wang, “Modeling the effects of supply chain and operator disruptions on cell therapy manufacturing facility operations during the covid-19 pandemic,” *American Pharmaceutical Review*, 2020.
- [11] N. Delirezh, S. M. Moazzeni, F. Shokri, M. A. Shokrgozar, M. Atri, and P. Kokhaei, “Autologous dendritic cells loaded with apoptotic tumor cells induce t cell-mediated immune responses against breast cancer in vitro,” *Cellular Immunology*, vol. 257, pp. 23–31, 2009.

- [12] A. Gervais, F. Bouet-Toussaint, O. Toutirais, C. T. D. L. Pintiere, N. Genetet, and V. Catros-Quemener, “Ex vivo expansion of antitumor cytotoxic lymphocytes with tumor-associated antigen-loaded dendritic cells,” *Anticancer Research*, vol. 25, pp. 2177–2185, 2005.
- [13] R. L. Bartel, “Translational regenerative medicine,” in A. Atala and J. G. Allickson, Eds. Academic Press, 2015, ch. 8 - Stem Cells and Cell Therapy: Autologous Cell Manufacturing, pp. 107–112.
- [14] R. Harrison, E. Zylberberg, S. Ellison, and B. Levine, “Chimeric antigen receptor–t cell therapy manufacturing: Modelling the effect of offshore production on aggregate cost of goods,” *Cytotherapy*, vol. 21, Feb. 2019.
- [15] B. L. Levine, J. Miskin, K. Wonnacott, and C. Keir, “Global manufacturing of car t cell therapy,” *Molecular Therapy - Methods & Clinical Development*, vol. 4, pp. 92–101, 2017.
- [16] M. Papathanasiou, C. Stamatis, M. Lakelin, S. Farid, N. Titchener-Hooker, and N. Shah, “Autologous car t-cell therapies supply chain: Challenges and opportunities?” *Cancer Gene Therapy*, pp. 1–11, Jan. 2020.
- [17] G. D. Eppen, R. K. Martin, and L. Schrage, “Or practice—a scenario approach to capacity planning,” *Operations Research*, vol. 37, no. 4, pp. 517–527, 1989.
- [18] B. Fleischmann, S. Ferber, and P. Henrich, “Strategic planning of bmw’s global production network,” *Interfaces*, vol. 36, pp. 194–208, Jun. 2006.
- [19] R. Bihlmaier, A. Koberstein, and R. Obst, “Modeling and optimizing of strategic and tactical production planning in the automotive industry under uncertainty,” *OR Spectrum*, vol. 31, pp. 311–336, Apr. 2009.
- [20] G. Zhang, J. Shang, and W. Li, “Collaborative production planning of supply chain under price and demand uncertainty,” *European Journal of Operational Research*, vol. 215, no. 3, pp. 590–603, 2011.
- [21] H. Stephan, T. Gschwind, and S. Minner, “Manufacturing capacity planning and the value of multi-stage stochastic programming under markovian demand,” *Flexible Services and Manufacturing Journal*, vol. 22, pp. 143–162, Dec. 2011.
- [22] R. R. Inman and W. C. Jordan, “Integrated assembly line loading, design, and labor planning,” *Journal of Manufacturing Systems*, vol. 16, no. 5, pp. 315–322, 1997.
- [23] L. Benyoucef, V. Jain, N. Rezg, and X. Xie, “Production planning in an assembly system with social and industrial constraints: A dynamic programming approach,”

International Journal of Services Operations and Informatics, vol. 2, pp. 131–151, Jan. 2007.

- [24] T. Sillekens, A. Koberstein, and L. Suhl, “Aggregate production planning in the automotive industry with special consideration of workforce flexibility,” *International Journal of Production Research*, vol. 49, no. 17, pp. 5055–5078, 2011.
- [25] C. Lam, E. Meinert, A. Yang, and Z. Cui, “Tackling the unique challenges of capacity planning for autologous cell therapies,” *Cytotherapy*, vol. 22, no. 5, Supplement, S148–S149, 2020.
- [26] A. J. Askew, “Chance-constrained dynamic programming and the optimization of water resource systems,” *Water Resources Research*, vol. 10, pp. 1099–1106, 1974.
- [27] L. A. Rossman, “Reliability-constrained dynamic programming and randomized release rules in reservoir management,” *Water Resources Research*, vol. 13, no. 2, pp. 247–255, 1977.
- [28] E. Altman, *Constrained Markov Decision Processes*. Stochastic modeling. Chapman & Hall, 1999.
- [29] M. Ono, “Joint chance-constrained model predictive control with probabilistic resolvability,” in *2012 American Control Conference (ACC)*, 2012, pp. 435–441.
- [30] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram, “Chance-constrained dynamic programming with application to risk-aware robotic space exploration,” *Autonomous Robots*, vol. 39, Dec. 2015.
- [31] A. T. Schwarm and M. Nikolaou, “Chance-constrained model predictive control,” *AIChE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.
- [32] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. JOHN WILEY & SONS, INC, 1994.
- [33] C. C. White and H. A. Holloway, “Resolvability for imprecise multiattribute alternative selection,” *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 38, no. 1, pp. 162–169, 2008.
- [34] D. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1995.
- [35] SAS, *One sample test of proportions*, <https://sphweb.bumc.bu.edu/otlt/MPH-Modules/BS/SAS/SAS6-CategoricalData/SAS6-CategoricalData2.html>, 2016.

- [36] M. Schader and F. Schmid, “Two rules of thumb for the approximation of the binomial distribution by the normal distribution,” *The American Statistician*, vol. 43, no. 1, pp. 23–24, 1989.
- [37] L. M. Jarvis, “Hurricane maria’s lessons for the drug industry.,” *Chemical & Engineering News*, vol. 96, no. 37, 2017.
- [38] M. Wendelbo and C. C. Blackburn, “A saline shortage this flu season exposes a flaw in our medical supply chain,” *smithsonian.com*, 2018.
- [39] E. Palmer, “Lonza u.s. cell therapy plant slapped with fda warning letter,” *FiercePharma*, 2017.
- [40] V. Bachanova¹, M. R. Bishop, P. Dahi, B. Dholaria, S. A. Grupp, B. Hayes-Lattin, M. Janakiram, R. T. Maziarz, J. P. McGuirk, L. J. Nastoupil, O. O. Oluwole, M.-A. Perales, D. L. Porter, and P. A. Riedell, “Chimeric antigen receptor t cell therapy during the covid-19 pandemic,” *Biology of Blood and Marrow Transplantation*, no. 000, pp. 1–8, 2020.
- [41] J. Bell, “Covid-19 testing kits: What are they, who makes them, and why are there shortages?” *NS Medical Devices*, 2020.
- [42] S. Fitzpatrick, H. Przybyla, D. D. Luce, L. Strickler, and A. Kaplan, “Coronavirus testing must double or triple before u.s. can safely reopen,” *NBC News*, 2020.
- [43] M. Parlar, Y. Wang, and Y. Gerchak, “A periodic review inventory model with markovian supply availability,” *International Journal of Production Economics*, vol. 42, pp. 131–136, 1995.
- [44] S. Özekici and M. Parlar, “Inventory models with unreliable suppliers in a random environment,” *Annals of Operations Research*, vol. 91, pp. 123–136, 1999.
- [45] J. Song and P. H. Zipkin, “Inventory control with information about supply conditions,” *Management Science*, vol. 42, no. 10, pp. 1409–1419, Oct. 1996.
- [46] B. M. Lewis, A. L. Erera, M. A. Nowak, and C. C. White, “Managing inventory in global supply chains facing port-of-entry disruption risks,” *Transportation Science*, vol. 47, no. 2, pp. 162–180, 2013.
- [47] E. Li, S. H. Xu, and J. C. Hayya, “A periodic-review inventory system with supply interruptions,” *Probability in the Engineering and Informational Sciences*, vol. 18, no. 1, pp. 33–53, 2004.

- [48] A. Federgruen and P. Zipkin, “An inventory model with limited production capacity and uncertain demands i. the average-cost criterion,” *Mathematics of Operations Research*, vol. 11, no. 2, pp. 193–207, 1986.
- [49] ———, “An inventory model with limited production capacity and uncertain demands ii. the discounted-cost criterion,” *Mathematics of Operations Research*, vol. 11, no. 2, pp. 208–215, 1986.
- [50] C. Shaoxiang and M. Lambrecht, “X-y band and modified (s, s) policy,” *Operations Research*, vol. 44, no. 6, pp. 1013–1019, 1996.
- [51] C. Shaoxiang, “The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy,” *Operations Research*, vol. 52, no. 3, pp. 409–421, 2004.
- [52] G. Gallego and A. Scheller-Wolf, “Capacitated inventory problems with fixed order costs: Some optimal policy structure,” *European Journal of Operational Research*, vol. 126, no. 3, pp. 603–613, 2000.
- [53] F. W. Ciarallo, R. Akella, and T. E. Morton, “A periodic review, production planning model with uncertain capacity and uncertain demand—optimality of extended myopic policies,” *Management Science*, vol. 40, no. 3, pp. 320–332, 1994.
- [54] Y. Wang and Y. Gerchak, “Periodic review production models with variable capacity, random yield, and uncertain demand,” *Management Science*, vol. 42, no. 1, pp. 130–137, 1996.
- [55] C. A. Yano and H. L. Lee, “Lot sizing with random yields: A review,” *Operations Research*, vol. 43, no. 2, pp. 311–334, 1995.
- [56] M. Henig and Y. Gerchak, “The structure of periodic review policies in the presence of random yield,” *Operations Research*, vol. 38, no. 4, pp. 634–643, 1990.
- [57] H. L. Lee and C. A. Yano, “Production control in multistage systems with variable yield losses,” *Operations Research*, vol. 36, no. 2, pp. 269–278, 1988.
- [58] Q. Li, H. Xu, and S. Zheng, “Periodic-review inventory systems with random yield and demand: Bounds and heuristics,” *IIE Transactions*, vol. 40, no. 4, pp. 434–444, 2008.
- [59] W. T. Huh and M. Nagarajan, “Technical note—linear inflation rules for the random yield problem: Analysis and computations,” *Operations Research*, vol. 58, no. 1, pp. 244–251, 2010.

- [60] J. B. Mazzola, W. F. McCoy, and H. M. Wagner, “Algorithms and heuristics for variable-yield lot sizing,” *Naval Research Logistics (NRL)*, vol. 34, no. 1, pp. 67–86, 1987.
- [61] T. Cheong and C. White, “Inventory replenishment control under supply uncertainty,” *Annals of Operations Research*, vol. 208, pp. 581–592, Sep. 2011.
- [62] A. Y. Ng, D. Harada, and S. Russell, “Policy invariance under reward transformations: Theory and application to reward shaping,” in *In Proceedings of the Sixteenth International Conference on Machine Learning*, Morgan Kaufmann, 1999, pp. 278–287.
- [63] U. Karmarka, “The multiperiod multilocation inventory problem,” *Operations Research*, vol. 29, no. 2, pp. 215–228, 1981.
- [64] ———, “The multiperiod multilocation inventory problem: Bounds and approximations,” *Management Science*, vol. 33, no. 1, pp. 86–94, 1987.
- [65] N. Rudi, S. Kapur, and D. Pyke, “A two-location inventory model with transshipment and location decision making,” *Management Science*, vol. 47, no. 12, pp. 1668–1680, 2001.
- [66] Y. Herer, M. Tzur, and E. Yucesan, “Transshipments: An emerging inventory recourse to achieve supply chain leagility,” *International Journal of Production Economics*, vol. 80, no. 3, pp. 201–212, 2002.
- [67] ———, “The multilocation transshipment problem,” *IIE Transactions*, vol. 38, no. 3, pp. 201–212, 2006.
- [68] H. Zhao, J. K. Ryan, and V. Deshpande, “Optimal dynamic production and inventory transshipment policies for a two-location make-to-stock system,” *Operations Research*, vol. 56, no. 2, pp. 400–410, 2008.
- [69] R. Lien, S. Iravani, K. Smilowitz, and M. Tzur, “An efficient and robust design for transshipment networks,” *Production and Operations Management*, vol. 20, no. 5, pp. 699–713, 2011.
- [70] S. Axsater, J. Marklund, and E. Silver, “Heuristic methods for centralized control of one-warehouse, n-retailer inventory systems,” *Manufacturing & Service Operations Management*, vol. 4, no. 1, pp. 75–97, 2002.
- [71] K. Wee and M. Dada, “Optimal policies for transshipping inventory in a retail network,” *Management Science*, vol. 51, no. 10, pp. 1519–1533, 2005.

- [72] R. Halper and S. Raghavan, “The mobile facility routing problem,” *Transportation Science*, vol. 45, no. 3, pp. 413–434, 2011.
- [73] J. Qiu and T. Sharkey, “Integrated dynamic single-facility location and inventory planning problems,” *IIE Transactions*, vol. 45, no. 8, pp. 883–895, 2013.
- [74] G. Ghiani, F. Guerriero, and R. Musmanno, “The capacitated plant location problem with multiple facilities in the same site,” *Computers & Operations Research*, vol. 29, no. 13, pp. 1903–1912, 2002.
- [75] M. Melo, S. Nickel, and F. da Gama, “Dynamic multi-commodity capacitated facility location: A mathematical modeling framework for strategic supply chain planning,” *Computers & Operations Research*, vol. 33, no. 1, pp. 181–208, 2005.
- [76] S. Jena, J. Cordeau, and B. Gendron, “Dynamic facility location with generalized modular capacities,” *Transportation Science*, vol. 49, no. 3, pp. 489–499, 2015.
- [77] M. Dada, N. C. Petruzzi, and L. B. Schwarz, “A newsvendor’s procurement problem when suppliers are unreliable,” *Manufacturing & Service Operations Management*, vol. 9, no. 1, pp. 9–32, 2007.
- [78] Y. Wang, W. Gilland, and B. Tomlin, “Mitigating supply risk: Dual sourcing or process improvement?” *Manufacturing & Service Operations Management*, vol. 12, no. 3, pp. 489–510, 2010.
- [79] A. J. Schmitt, S. A. Sun, L. V. Snyder, and Z.-J. M. Shen, “Centralization versus decentralization: Risk pooling, risk diversification, and supply chain disruptions,” *Omega*, vol. 52, pp. 201–212, 2015.
- [80] H. Zhao, V. Deshpande, and J. K. Ryan, “Inventory sharing and rationing in decentralized dealer networks,” *Management Science*, vol. 51, no. 4, pp. 531–547, 2005.
- [81] G. Sobic, “Transshipment of inventories among retailers: Myopic vs. farsighted stability,” *Management Science*, vol. 52, no. 10, pp. 1493–1508, 2006.
- [82] X. Yan and H. Zhao, “Inventory sharing and coordination among n independent retailers,” *European Journal of Operational Research*, vol. 243, no. 2, pp. 576–587, 2015.
- [83] X. Hu, I. Duenyas, and R. Kapuscinski, “Optimal joint inventory and transshipment control under uncertain capacity,” *Operations Research*, vol. 56, no. 4, pp. 881–897, 2008.

- [84] U. Ozen, G. Susic, and M. Slikker, “A collaborative decentralized distribution system with demand forecast updates,” *European Journal of Operational Research*, vol. 216, no. 3, pp. 573–583, 2012.
- [85] S. Wright, “Coordinate descent algorithms,” *Mathematical Programming*, vol. 151, Feb. 2015.
- [86] D. Bertsekas, *Dynamic Programming and Optimal Control*. Jan. 2007, vol. 1.
- [87] D. Bertsekas and S. Shreve, “Stochastic optimal control: The discrete-time case,” in. 1978.
- [88] D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*. Jan. 1996.
- [89] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. A Bradford Book, 2018.
- [90] J. Benders, “Partitioning procedures for solving mixed-variables programming problems,” *Numerische Mathematik*, vol. 4, pp. 238–252, 1962.
- [91] D. Bertsimas and J. Tsitsiklis, *Introduction to Linear Optimization*. Athena Scientific, 1997.
- [92] L. Lasdon, *Optimization Theory For Large Systems*. Dover Publications, 2002.
- [93] J. Jones, M. Sapiano, A. Savinkina, K. Haass, M. Baker, R. Henry, J. Berger, and S. Basavaraju, “Slowing decline in blood collection and transfusion in the united states - 2017: Slowing decline in blood collection and use,” *Transfusion*, vol. 60 Suppl 2, Feb. 2020.
- [94] G. M. Association, “2008 joint industry unsaleables report: The real causes and actionable solutions,” 2008.
- [95] D. Bertsekas, A. Nedić, and A. Ozdaglar, *Convex Analysis and Optimization*. Athena Scientific, 2003.

VITA

Junxuan (Jun) Li is a senior data & applied scientist at Microsoft and a graduate researcher in Operations Research, with concentrations in sequential decision-making and dynamic control, at the H.Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology. Jun's research covers fundamental and application problems in dynamic programming and stochastic optimization, as well as their applications in supply chain, healthcare, production system, and e-commerce. Jun has contributed his expertise in dynamic programming, machine learning, statistics and simulation to a number of multi-disciplinary studies collaborating with not only academic research institutes, e.g. Georgia Tech CMaT Center and the University of Pennsylvania CVPF, but also companies, e.g. Microsoft, Roosterbio, BioLife Solutions, Dillard's, Ford Motor and etc.